UNIT 11: EXTERNAL MEMORY

UNIT STRUCTURE

11.1 Learning Objectives

11.2 Introduction

11.3 Magnetic Disk

11.4 Magnetic Read-Write Operation

11.5 Track

11.6 Sector

11.7 Inter Track and Inter Sector Gap

11.8 Cylinder

11.9 Fixed Head Disk

11.10 Seek Time

11.11 Access Time

11.12 Transfer Time

11.13 Rotational Delay

11.14 RAID

11.15 Optical Memory

11.16 Magnetic Tap

11.17 Let Us Sum Up

11.18 Further Reading

11.19 Answers to Check Your Progress

11.20 Model Questions

11.1 LEARNING OBJECTIVES

After going through this unit, you will be able to:

- learn about different external memory
- explain read-write operation
- describe different parts associated with external memory like track, sector, cylinder, etc.
- learn about seek time, access time and rotational delay
- learn about RAID, optical memory etc.

11.2 INTRODUCTION

In the previous unit we have learnt about the I/O processor and the interfacing of the I/O processors with the different devices. Also we have gone through the input and output devices.

In this unit, we will describe secondary storage devices, magnetic disk, optical memory, magnetic tape and various terms associated with these memories.

The speed and capacity of the memory impacts the performance of the computer. Magnetic disk is a random access medium. That is the data on the disk can be found almost immediately rather than sequentially going through the data. The data on the disk can be erased and reused virtually indefinitely. It is designed to reside in a protective case or cartridge to protect it from dust and other external interference.

The magnetic tape is another important storage device that is a sequential one. Also, the optical storage technology is important in all types of computer systems. The CD and the DVDs have also become increasingly important.

11.3 MAGNETIC DISK

Magnetic disks are the widely used as secondary storage device. They are mainly used for bulk storage. Unlike the magnetic tapes, the data on the magnetic disks can be read randomly. A magnetic disk is a thin, circular metal plate the surface is coated with a magnetic oxide layer usually on both the sides. These have a high storage capacity, reliable and have capability to access the stored data directly. The disks are mounted on a rotating shaft. Each magnetic disk has a head which consists of a magnetizing coil and a magnetic yoke. Information is stored on the disks which are represented as magnetic spots. Presence of a magnetic spot represents a 1 bit and the absence of a magnetic spot represents a 0 bit.

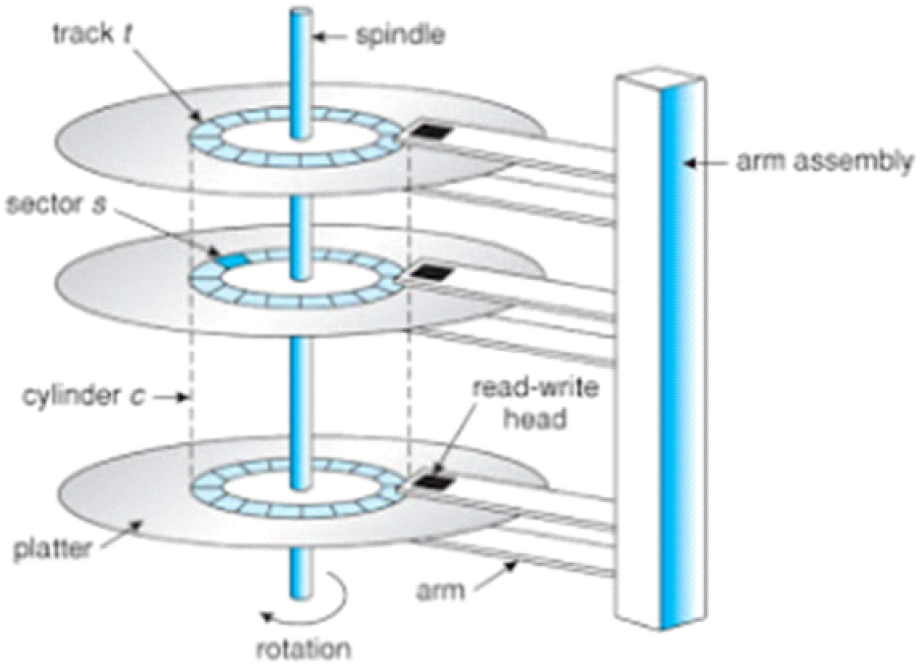


Figure 11.1

The individual disks are spaced on the shaft to allow room for a read/write mechanism to move between them.

Information is stored on the disk by magnetizing the surface in a particular direction as shown in the figure below.

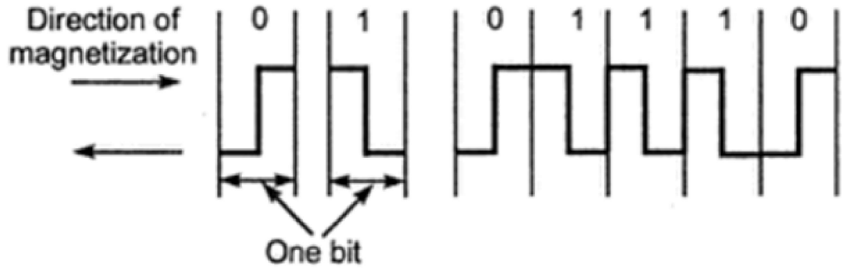


Figure 11.2: Bit Representation of Magnetic Disk

In most disk storage devices, several disks are assembled to form a disk pack as shown below:

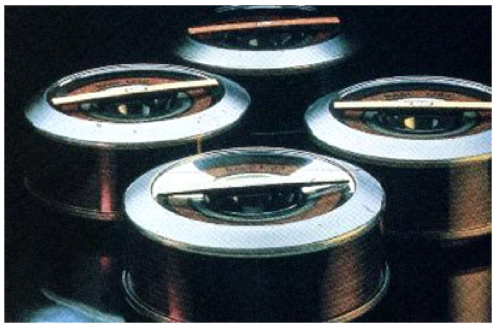


Figure 11.3

The top and bottom surfaces are not used for storing data because they are likely to become scratched or nicked. A disk pack may contain anywhere from five to one hundred disks.

The advantage of the magnetic disk over the magnetic tapes is that the access time is much faster.

11.4 MAGNETIC READ WRITE OPERATIONS

The data on a disk is read or written by the read/write heads located between the disks. Most disk units have one read/write head for each disk recording surface. All the heads are permanently connected to an access mechanism.

In some systems, there are two heads– a read head and a write head. During a Read/Write operation, the head is stationary and the platter rotates beneath it.

A read/write mechanism is used to achieve the access to read or write on the magnetic disk. When the device is used to read from a particular magnetic medium, it detects the individual magnetized areas and translates them into electrical pulses. In performing the write operations, the device magnetizes the appropriate areas on the magnetic medium thereby erasing the previously stored data, if any. Also, the electrical impulses from the

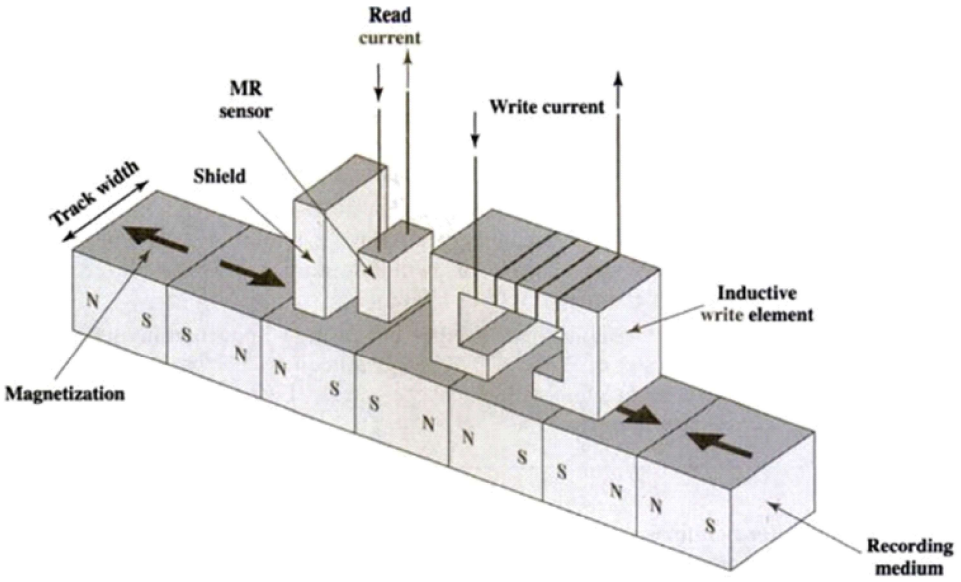


Figure 11.4

processor are converted into magnetic spots on the disks for writing to it. When electric pulses are passed through the write head, a magnetic field is produced. These magnetic patterns are recorded on the surface of the disk below creating different for the positive and the negative currents. If the direction of the current is reversed, the direction of the magnetization of the recording medium is also reversed.

The figure above shows the inductive Read/Write head. During a read operation, a magnetic field moving relative to the coil produces an electrical current in the coil. When the surface of the disk passes under the head, it produces the same current that was recorded.

The head consists of a partially shielded magnetoresistive (MR) sensor. By passing a current through the MR sensor, the resistive changes are detected as voltage signals.



Figure 11.5

11.5 TRACK

The data on the magnetic disk is organized in a concentric set of circles called track. There can be from a few hundred to a few thousand tracks on a single disk surface. Each of the tracks has the same width as the head and the adjacent tracks are separated by gaps. Data is stored serially on each track. The tracks are numbered from the outermost to the innermost circle starting with zero.

Each track on a disk can store the same amount of data even though the tracks get smaller toward the center of the disk. This results higher bit density in inner tracks than that of the outer tracks. Consider a disk pack with 4,000 usable tracks (20 surfaces x 200 tracks per surface) on which 7,294 characters can be stored on each track. The disk pack could conceivably store 29,176,000 characters of data (4,000 tracks x 7,294 characters per track).

11.6 SECTOR

Each track is divided into sectors. The sectors refers to the number of fixed size areas (imaginary pie slices) that can be accessed by one of the read/write heads, in one rotation of the disk, without the head having to change its position. The number of bytes stored in each sector is same. Thus, the size of the inner sectors decides the storage capacity for all other sectors on the disk. Sectors are typically 512 bytes in size, although there is an initiative to increase the sector size to 4096 bytes.

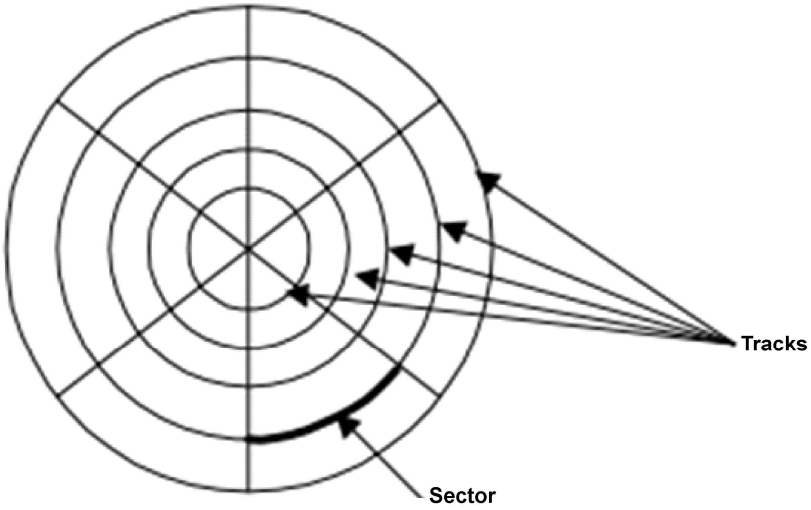


Figure 11.6

11.7 INTER TRACK AND INTER SECTOR GAP

The gap in between the tracks in a magnetic disk is called the Inter-track gap and the gap between the sectors in the disks is called inter-sector gap. The inter-track gaps help protects data integrity. To avoid, imposition

unreasonable precision requirements on the system, adjacent sectors are separated by inter-track (inter-sector) gaps.

Some means are needed to locate sector positions within a track. Clearly there must be some starting points on the track and a way of identifying the start and end of each sector.These requirements are handled by means of a control data recorded on the disk. Thus, the disk is formatted with some extra data used only by the disk drive and not accessible to the user.

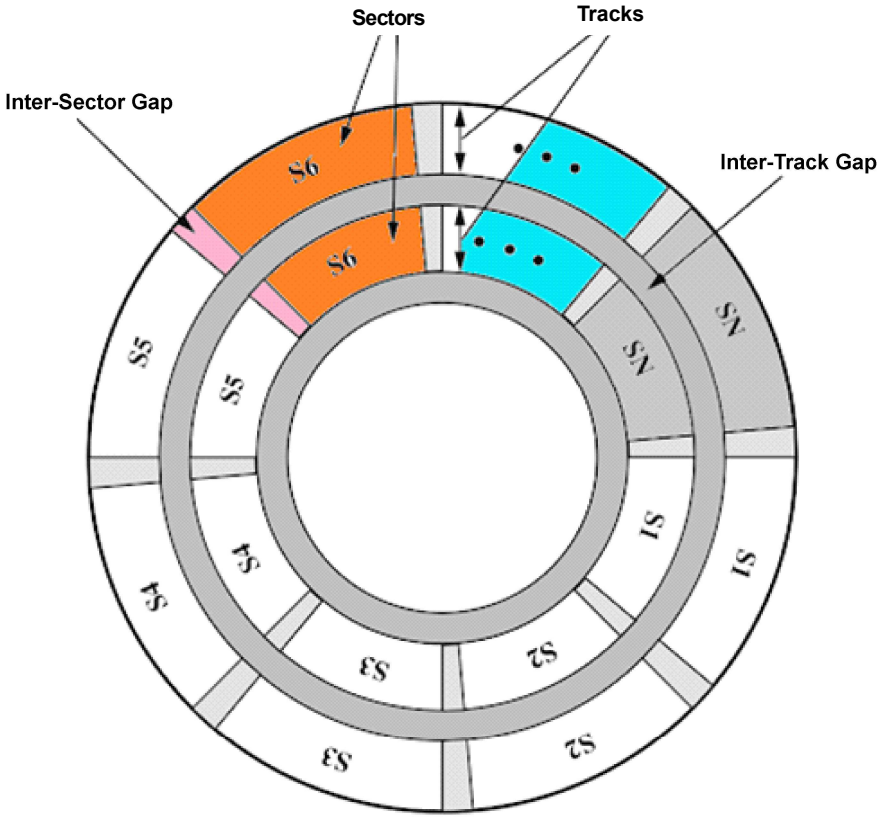


Figure 11.7

11.8 CYLINDER

To have a large storage system in the computer, multiple magnetic disks are maintained and used together. The disks or the platters are stacked on top of each other with a common spindle. The spindle turns or rotates the disks of the pack.

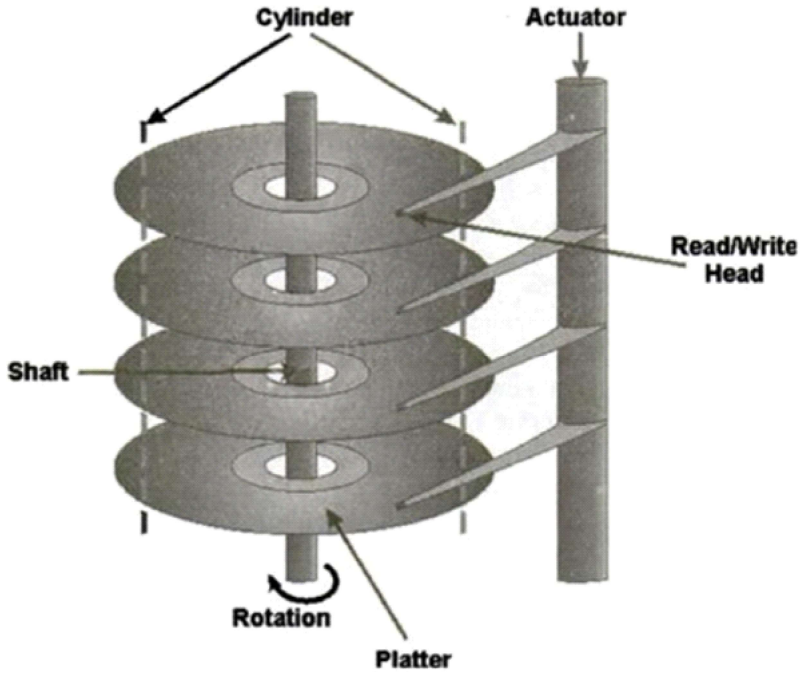


Figure 11.8

There is a gap between the platters so that there is space for the magnetic read/write head. For each side of each platter, there is a magnetic read/write head. The head is mounted on arms which can move the heads towards the central spindle or towards the edge. This concept of stacking the disks is called cylinder. A cylinder is made up of all the tracks of the same number from all the metal disks. Thus cylinder is an imaginary view of all the matching tracks on top of each other as shown in the figure above. Thus the storage capacity of a magnetic disk can be defined as the following:

Storage capacity of a magnetic disk = Number of recording surfaces
x Number of tracks per surface
x Number of sectors per track
x Number of bytes per sector

11.9 FIXED HEAD DISK

There are two types of disks depending upon the two methods of accessing data. One is moving head and the other is fixed head.

A fixed head disk has many read/write heads on each disk surface as the number of tracks which are mounted on the access arm as shown in the figure below.

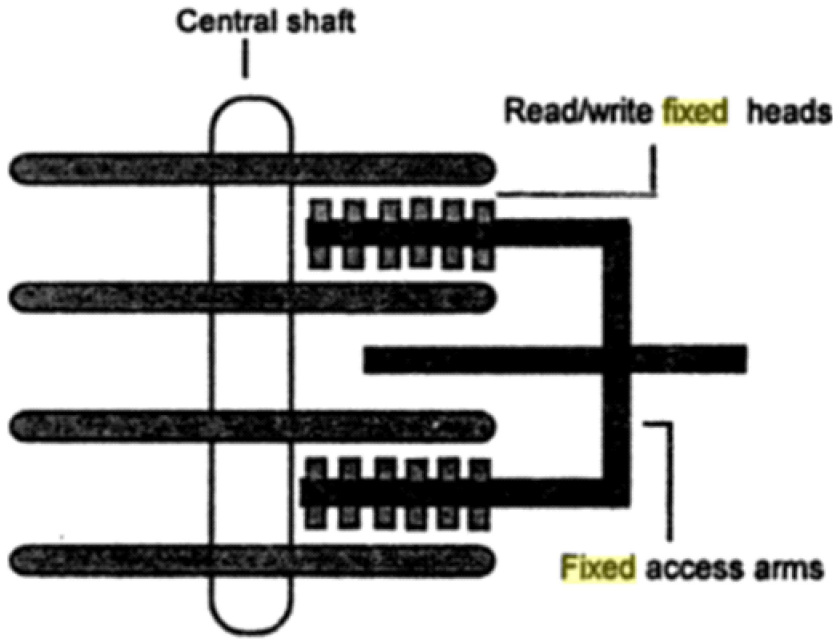


Figure 11.9

The head does not move as each track has a separate read/write head. The moving head disk systems on the other hand have one read/ write head for each disk surface which is mounted on the access arm. The head in such systems moves horizontally on the disk surface to read or write data.

The access speed of the fixed head disks are more as compared to the moving head disks. In a fixed head disk, the track or the cylinder is selected by electronically switching to the appropriate read/write head rather than by mechanically moving it. This makes it faster.

However, as the cost of the additional read/write heads is quite high, so the fixed head disks are not commonly used.



CHECK YOUR PROGRESS

Q.1: Fill In the blanks:

i) Magnetic disk has a head which consists of a ..... and a .....

ii) The data on a disk is read or written by the ..... heads located between the disks.

iii) During a read operation, a magnetic field moving relative to the coil produces an ..... in the coil.

iv) The head consists of a partially shielded ..... sensor.

v) The data on the magnetic disk is organized in a concentric set of circles called .....

vi) Each track on a disk can store the same amount of data even though the tracks get ..... toward the center of the disk.

vii) Tracks are divided into .....

viii) The gap in between the tracks in a magnetic disk is called the ..... and the gap between the sectors in the disks is called .....

11.10 SEEK TIME

The time required to position the read/write head from the current track or cylinder to the desired track or cylinder is called seek time. The seek time varies depending upon the position of the head when a data request is made. It is zero when the head is already positioned on the specified cylinder.

The seek time is maximum when the read/write has to move from the outermost track to the innermost track or vice-versa.

11.11 ACCESS TIME

Access time is the time taken to read a record from the magnetic disk. In other words, it is the time elapsed between the initiation of a particular request for data and the receipt of the data.

The sum of the seek time and the rotational latency is the access time. The average disk access time is equal to seek time + average rotational delay+ transfer time + controller overhead.

11.12 TRANSFER TIME

Once the read/write head is in position, the read or write operation can be performed as the sector moves under the head. This is the data

transfer operation of the disk. The transfer time is the time taken to transfer a block of bits. It is a function of the sector size, rotation speed, and the recording density of a track.

The transfer time to or from the disk depends on the rotation speed of the disk. The transfer time, T can be calculated as

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

Where, T = transfer time

b = number of bytes to be transferred

N = number of bytes on a track

r = rotation speed

Thus, the Total average access time can be expressed as

$$T_a = T_s + \frac{1}{2r} + \frac{b}{rN}$$

here, T_S is the average seek time.

11.11 ROTATIONAL DELAY

Seek time as we have seen is the time taken to position the head at the track. After the track is selected the disk controller waits until the appropriate sector rotates to line up with the head. This time which is taken for the beginning of the sector to reach the head is called rotational delay or rotational latency.

11.12 RAID

RAID is ***Redundant Array of Independent Disks***. Originally it was known as Redundant Array of Inexpensive Disks. It is a way of storing data in different places on multiple disks to protect the data in case of a drive failure. It 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. The data recording is distributed on the different drives which appear as a single drive to the application program.

The access speed of the magnetic disks is considerably slow as compared to the speed of the processor and the main memory. By improving

the disk performance the overall performance of the computer can be increased. This led to the development of arrays of disks that operate independently and in parallel. The same data is kept in multiple disks. With the multiple disks, separate I/O requests can be handled in parallel provided the data that is required is on the separate disks. Also, a single I/O request can be handled in parallel as long as the data to be accessed is in the multiple disks.

In a RAID, a single large file is stored in several different disks by breaking up the file into a number of smaller pieces and storing the pieces on the separate disks. This is called ***data stripping.*** Thus the file modules can be accessed simultaneously. This increases the access speed and the data transfer speed. Thus the transfer time would be the time that would be required in a single disk divided by the number of disks used in the array. A data stripe can be any size between 512 bytes up to several megabytes.

There are a number of ways in which the data can be organized with the use of multiple disks. These are called as the RAID levels. These levels do not imply a hierarchical relationship. They are different design architectures. The RAID levels use these stripes and the multiple disks in different ways.

RAID 0: The RAID level 0 uses the data stripping as shown in the figure above. A number of disks are connected together to form a single disk.

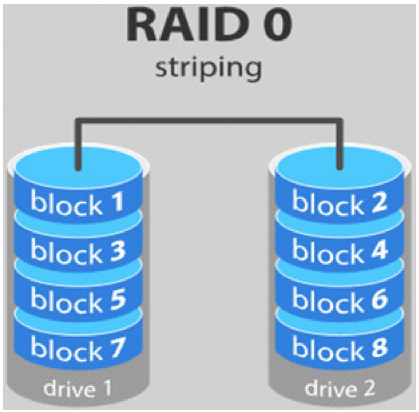


Figure 11.10

In this level, the data are split up into blocks that are written across all the drives in the array. The multiple disks here offer better I/O performance and also a disk system with higher capacity is realized.

RAID 0 does not provide fault handling capacity because if one disk fails, the entire disk system fails.

RAID 1: RAID 1 uses mirroring. It uses pairs of stripped disks and the data is written to both the disk in the pair at once.

If one disk controller is used for both the disk in each pair, the process is called ***mirroring***, and if a separate disk controller is used for each disk in the pair it is called ***duplexing***.

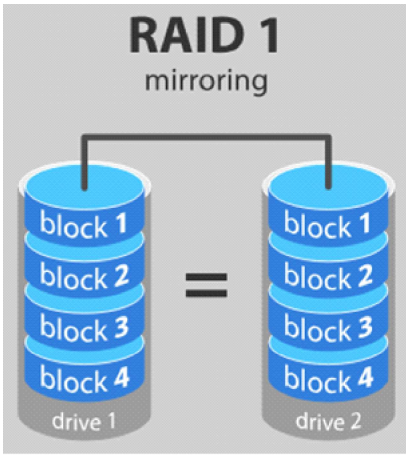


Figure 11.11

RAID 1 is a highly reliable system. If one of the drive fails, then the controller uses either the data drive or the mirror drive for data recovery.

RAID 2: In RAID 2, the data striping is used as in RAID 0 and RAID 1. The data is stripped across two or more disks and one or more extra disk are used to hold error correcting code for the stored data. Thus, if one of the disk crashes, the data of that disk can be regenerated using the appropriate error correcting code held on another disk.

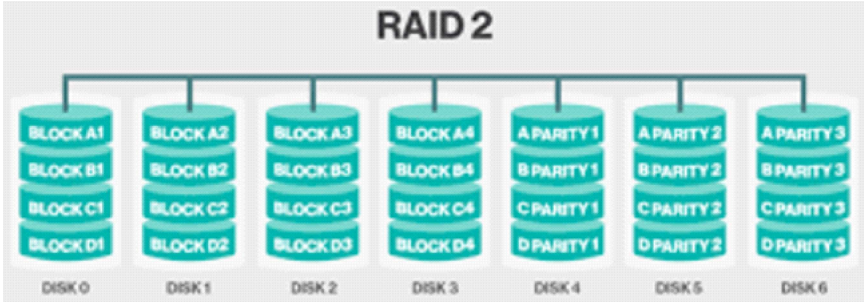


Figure 11.12

These systems are expensive mainly because of the error correcting code that has to be stored. They are rarely used.

RAID 3: This is similar to RAID 2. The data is stripped across a number of disks and an extra disk is used to hold the parity information. Instead of an error correcting code, a simple parity bit is computed for the set of individual bits in the same position. In the event of a disk failure, the data on the failed disk can be revived from the surviving data and parity information.

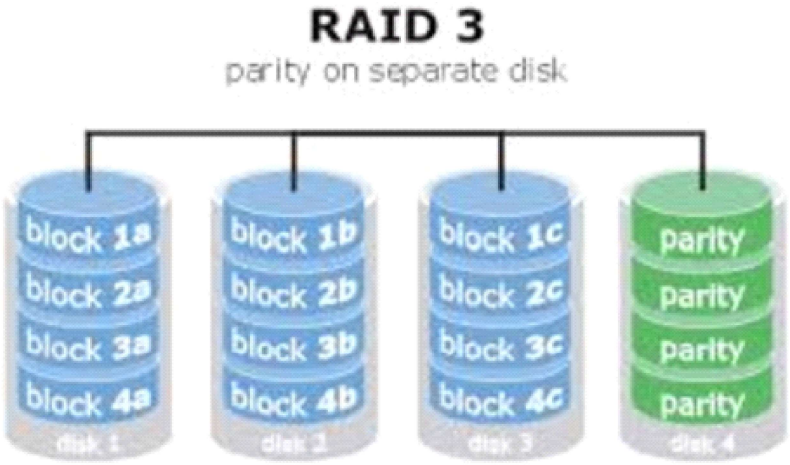


Figure 11.13

RAID 4: RAID 4 consists of a data disks and a parity disk. This level is similar to RAID level 3 because similar stripping method is used. But it employs stripped data in much larger blocks or segments.

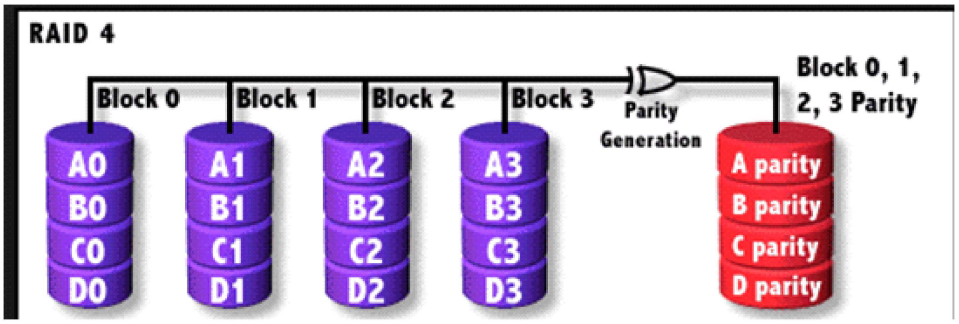


Figure 11.14

As all the parity data is written in to one single disk, so RAID level 4 suffers from a write bottleneck and is not eneally used.

RAID 5: RAID 5 also uses disk stripping with parity. In this configuration, the data and the parity are stripped evenly across all of the disks. So no single disk is a bottleneck. Striping also allows users to reconstruct data in case of a disk failure.

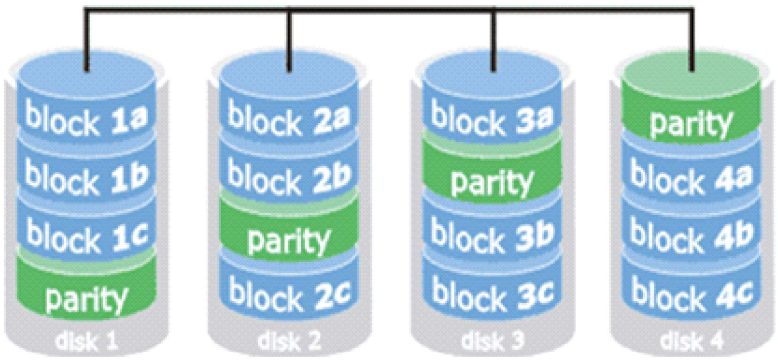


Figure 11.15

RAID 5 used parity instead of mirroring for data redundancy. When data is written to a RAID 5 drive, the system calculates parity and writes that parity into the drive. While mirroring maintains multiple copies of data in each volume to use in case of failure, RAID 5 can rebuild a failed drive using the parity data, which is not kept on a fixed single drive.

RAID 6: As like RAID 5, the RAID level 6 does block level striping. However, it uses dual parity. It creates two parity blocks for each data block. With RAID 6, the RAID array will even survive that second failure.

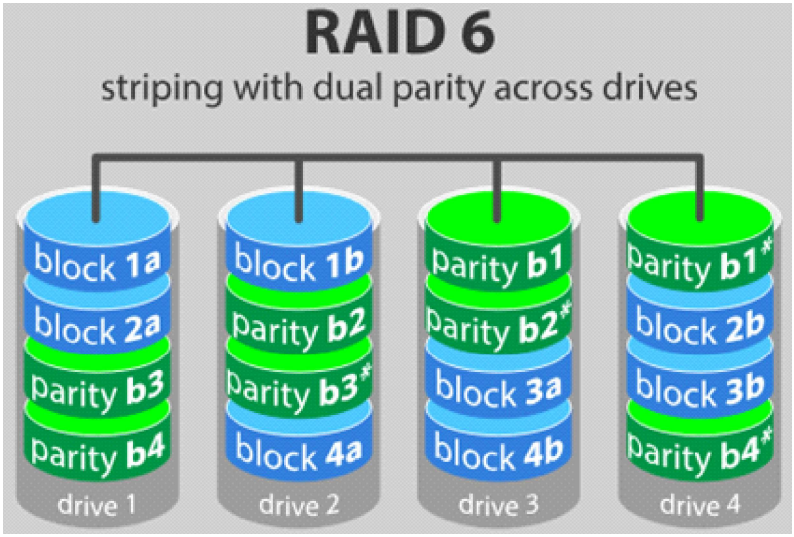


Figure 11.16

If two drives fail, the data can still be accessed, even while the failed drives are being replaced. So RAID 6 is more secure than RAID 5. RAID 6 is a good all-round system that combines efficient storage with excellent security and decent performance. It is preferable over RAID 5 in file and application servers that use many large drives for data storage.

11.11 OPTICAL MEMORY

In 1983, the first optical memory device CD (Compact Disc) was successfully launched. The huge success of CD encouraged the development of optical storage technology.

CD (Compact Disk) and DVD (Digital Versatile Disk) are common examples of optical memory.

Compact Disk: Information is digitally recorded on the compact disk. The recording is done in the form of microscopic dots on the surface of the polycarbonate. It is done by focusing a laser beam on the surface of the disk. The pitted surface is then coated with a highly reflexive surface usually aluminum or gold. The shiny surface protects against dust and scratches by a top coat of clear acrylic.

To record data in a compact disk, the high-intensity laser beam is turned On and Off according to a digital signal. This causes tiny holes or pits are burnt into the disk. This is shown in the figure below:

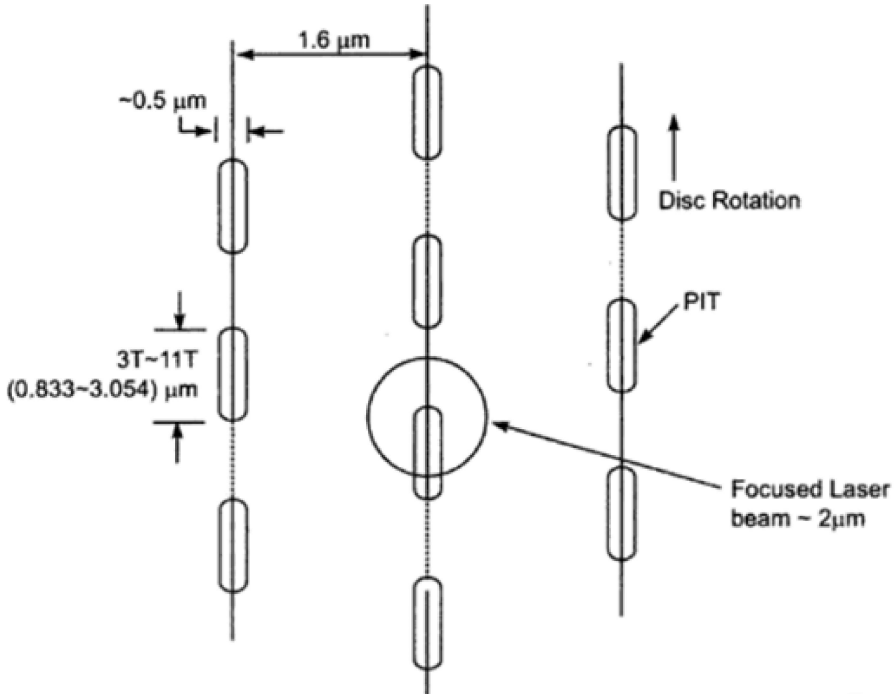


Figure 11.17: Recorded Portion of CD Showing Pits along its Tracks

The recorded data is recovered from the optical disk with an optical pickup, which moves across the surface of the rotating disk. The pickup system is shown in the figure below:

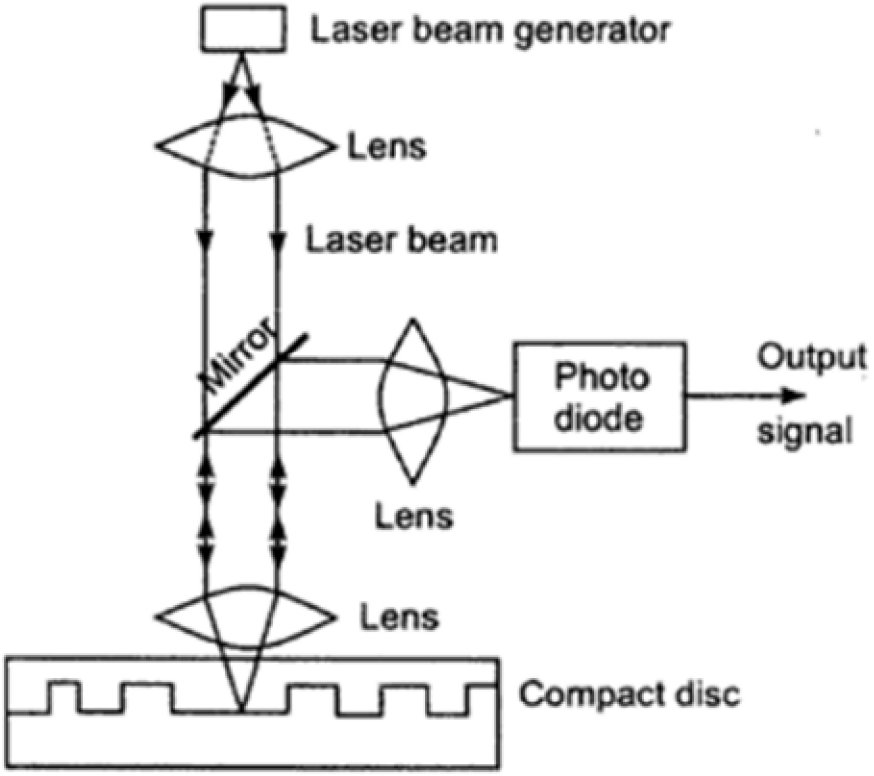


Figure 11.18: Optical Pickup System of a Compact Disk

The area in between the pits is called ***lands***. A land is a smooth surface that reflects back at a higher intensity. A photosensor detects the change between the change between the pits and the lands and converts them to a digital signal.

The information on the compact disks is not stored in tracks. Rather, the disk contains a single spiral track, beginning near the center and spiraling out to the outer edge of the disk. The sectors inside and the outside of the spirals are of the same length. So the information is evenly spread across the disk. The magnetic spots or pits are read by the laser at a *Constant Linear Velocity*. The disk rotates more slowly to access the data near the outer edge.

When data is read from a CD, a laser is directed to the surface of the disk. The laser follows the stream of pits and lands from the inside of the disk to the outer in a spiral direction. As the laser shines on the data track, it reflects one pattern of light off of a pit and a different pattern off of a land area. The resulting reflections equate to a series of ones and zeros by a photo sensor. It is shown in the figure below.

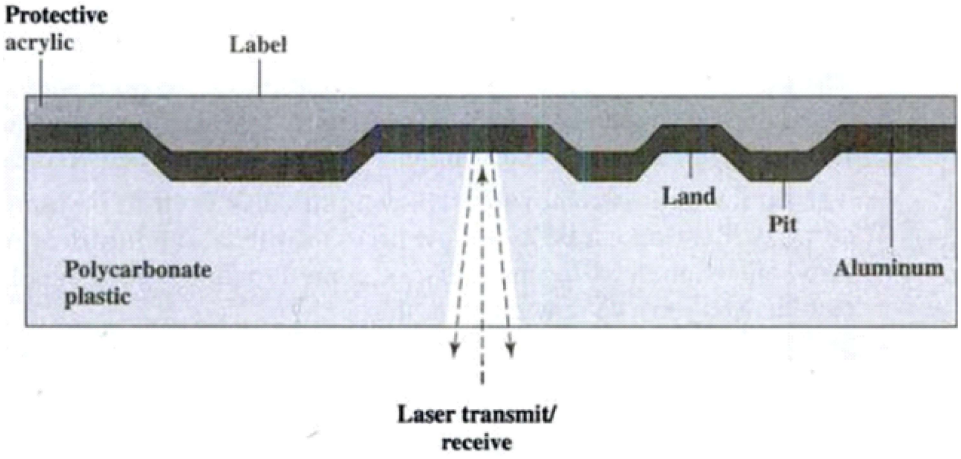


Figure 11.19: CD Operations

Data in CD ROM is stored as a sequence of blocks. It consists of the following fields.

Sync– it identifies the beginning of a block. It has a byte of all 0s, 10 bytes of all 1s and then a byte of all 0s.

Header– it contains the block address and mode byte.

Mode 0 = Blank data field

Mode 1 = 2048 data bytes + error correcting code

Mode 2 = 2336 data bytes

The block format is shown in the figure below.

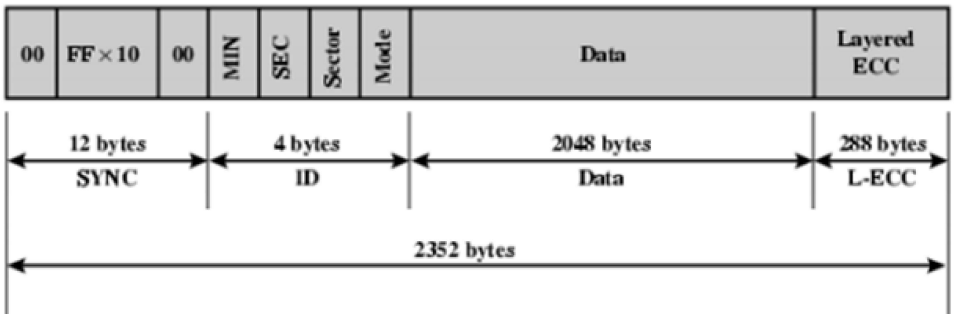


Figure 11.20

The advantages of a compact disk are mainly two. These are

1) The disk together with its information can be replicated inexpensively.
2) It is removable. Thus, the disk can be used for archival storage.

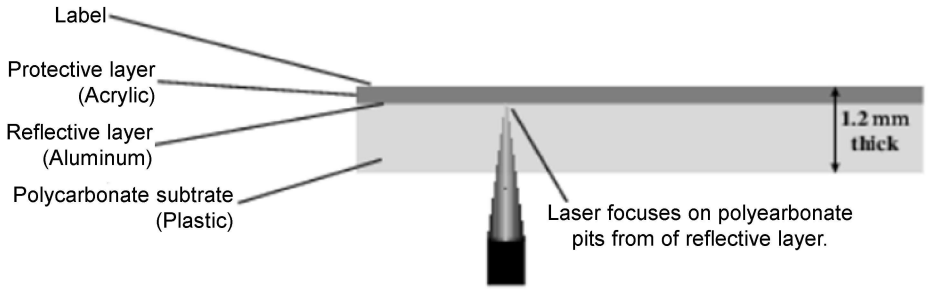
The disadvantages are:

1) The CDs are Read only and cannot be updated.
2) The access time is more than that of a magnetic disk.

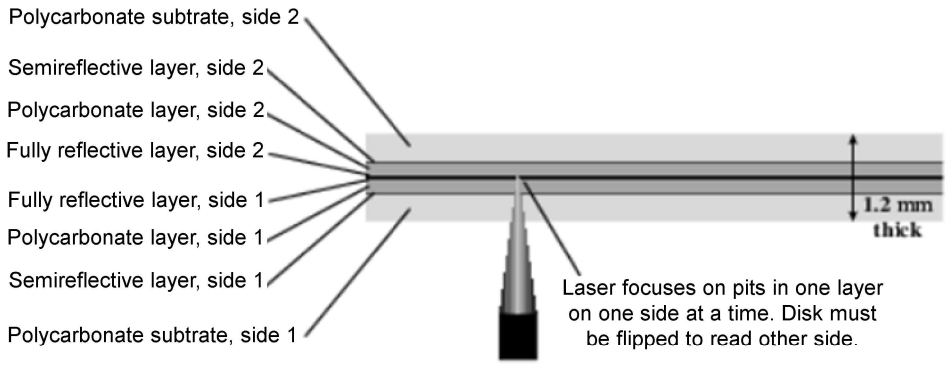
DVD (Digital Versatile Disk): The DVDs took the video to the digital age. It is based on the CD technology with increased storage. It has much larger storage capacity. They can be single or double sided. A two layered DVD can hold up to 17 GB of data, video or audio or other information. The DVD has a multiple layers of plastic pieces with a total thickness of 1.2 millimeters. Polycarbonate plastic is basically used for making these layers.

The greater capacity of the DVDs than that of the CDs are due to three reasons as given in the figure below.

1) The bits in a DVD are more closely packed as that of a CD.
2) The DVD has a second layer of pits and lands on top of the first layer. A dual layer DVD has a semi-reflexive layer on top of the reflexive layer. Each layer of the DVD can be read by the lasers.
3) DVDs can be two-sided, whereas the data can be recorded on only one side of a CD. This increases the capacity of the DVD.



(a) CD-ROM – Capacity 682 MB



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

Figure 11.21

11.16 MAGNETIC TAP

Magnetic tape is one of the most popular storage medium for large data that are sequentially accessed and processed. It is a flexible polyester

tape coated with a magnetizable material. Data in the magnetic tape is stored in the form of magnetized and de-magnetized portion on the surface of the tape. The magnetized portion on the surface refers to '1' bit value and the demagnetized portion refers to '0' bit value. The amount of data that can be stored on a linear inch of tape is known as the tape's recording density.

The data in the magnetic tapes are stored in the tracks that run parallel lengthwise. The recording of the data in the parallel tracks is called parallel recording. If the data is laid out as a sequence of bits along each track, the recording is called serial recording. Data are written in blocks which are called the physical records on the tape. The blocks are separated by gaps called the inter-record gaps.

The magnetic tape drive is an example of a sequential storage device. The tape must be rewind or fast forwarded to get the correct block under the read/write head. The magnetic disks, on the other hand, are direct access devices.

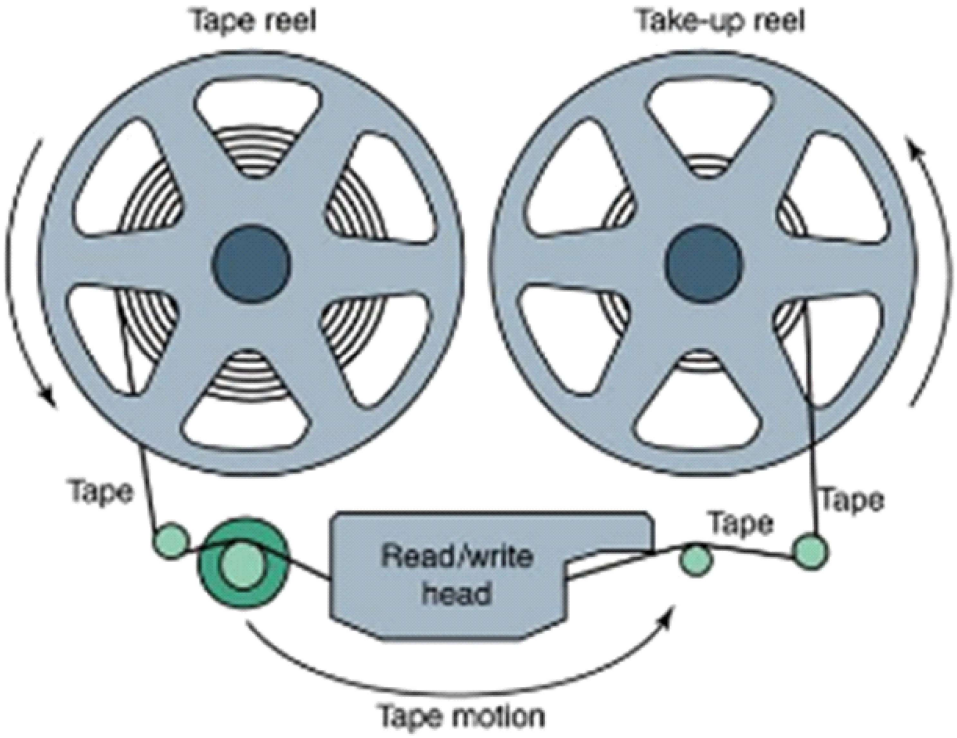


Figure 11.22: A Magnetic Tape Storage Medium



CHECK YOUR PROGRESS

Q.2: Fill In the blanks:

i) The time required to position the read/write head from the current track or cylinder to the desired track or cylinder is called .....

ii) ..... is the time elapsed between the initiation of a particular request for data and the receipt of the data.

iii) The transfer time to or from the disk depends on the ..... of the disk.

iv) The time which is taken for the beginning of the sector to reach the head is called .....

v) RAID is .....

vi) The access speed of the magnetic disks is considerably ..... as compared to the speed of the processor and the main memory.

vii) ..... 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.

viii) Examples of optical memory are ..... and .....

ix) A ..... detects the change between the change between the pits and the lands and converts them to a digital signal.

x) The magnetic spots or pits are read by the laser at a .....



11.11 LET US SUM UP

- A magnetic disk is a thin, circular metal plate the surface is coated with a magnetic oxide layer usually on both the sides.

- Information is stored on the disks which are represented as magnetic spots. Presence of a magnetic spot represents a 1 bit and the absence of a magnetic spot represents a 0 bit.
- A disk pack may contain anywhere from five to one hundred disks.
- A read/write mechanism is used to achieve the access to read or write on the magnetic disk.
- The data on the magnetic disk is organized in a concentric set of circles called track
- Each of the tracks has the same width as the head and the adjacent tracks are separated by gaps.
- The tracks are numbered from the outermost to the innermost circle starting with zero.
- Each track is divided into sectors.
- The gap in between the tracks in a magnetic disk is called the Inter-track gap and the gap between the sectors in the disks is called inter-sector gap.
- A cylinder is made up of all the tracks of the same number from all the metal disks. Thus cylinder is an imaginary view of all the matching tracks on top of each other.
- There are two types of disks depending upon the two methods of accessing data. One is moving head and the other is fixed head.
- A fixed head disk has many read/write heads on each disk surface as the number of tracks which are mounted on the access arm
- The moving head disk systems have one read/write head for each disk surface which is mounted on the access arm.
- The time required to position the read/write head from the current track or cylinder to the desired track or cylinder is called seek time.
- Access time is the time taken to read a record from the magnetic disk.
- The transfer time is the time taken to transfer a block of bits.
- After the track is selected the disk controller waits until the appropriate sector rotates to line up with the head. This time which is taken for the beginning of the sector to reach the head is called rotational delay or rotational latency.

- RAID is a way of storing data in different places on multiple disks to protect the data in case of a drive failure
- In a RAID, a single large file is stored in several different disks by breaking up the file into a number of smaller pieces and storing the pieces on the separate disks. This is called ***data stripping.***
- CD (Compact Disk) and DVD (Digital Versatile Disk) are common examples of optical memory.



11.18 FURTHER READING

1) Chaudhuri, P. Pal. *Computer Organization and Design*. PHI.
2) Godse, A. P., Godse, D. A. *Computer Architecture*. Technical Publications Pune.
3) Stallings, William. *Computer Organization and Architecture: Designing for Performance*. PHI.



11.19 ANSWERS TO CHECK YOUR PROGRESS

Ans. to Q. No. 1: i) Magnetizing coil, magnetic yoke; ii) read/write; iii) electrical current; iv) magnetoresistive; v) tracks; vi) smaller; vii) sectors; viii) Inter-track gap, inter-sector gap

Ans. to Q. No. 2: i) seek time; ii) access time; iii) rotation speed; iv) rotational delay; v) Redundant Array of Independent Disks; vi) Slow; vii) RAID; viii) CD, DVD; ix) Photosensor; x) Constant Linear Velocity



11.1 MODEL QUESTIONS

Q.1: Explain the organization of the magnetic disk.

Q.2: Compare magnetic disk with that of magnetic tape

Q.3: What is RAID technology? Explain the different levels of RAID.

Q.4: What is optical memory? Explain with examples.

*** ***** ***