Computer Architecture

Submitted by:

Awais Ghafar

ID: 15269

Assignment: 06

Submitted to:

M Amin

Semester : 04th

IQRA NATIONAL UNIVERSITY PESHAWAR

Q.1 Give detail answer to each of the following:

(a) What are the advantages of using a glass substrate for a magnetic disk?

The glass substrate has a number of benefits, including the following:

- Improvement in the uniformity of the magnetic film surface to increase disk

reliability

- A significant reduction in overall surface defects to help reduce read-write errors
- Ability to support lower fly heights (described subsequently)
- Better stiffness to reduce disk dynamics
- Greater ability to withstand shock and damage

(b) Define the terms track, cylinder, and sector.

Track - On a magnetic disk, data is organized on the platter in concentric sets of rings, called tracks. Or A **track** is that portion of a disk which passes under a single stationary head during a disk rotation, a ring 1 bit wide.

Cylinder - The set of all the tracks in the same relative position on the platter is referred to as a cylinder. Or A **cylinder** is comprised of the set of tracks described by all the heads (on separate platters) at a single seek position. Each cylinder is equidistant from the center of the disk.

Sector - Data are transferred to and from the disk in sectors. Or A track is divided into segments of **sectors**, which is the basic unit of storage.

(c) Define seek time, rotational delay, access time, and transfer time for magnetic disk.

Seek time - seek time Seek time is the time required to move the disk arm to the required track. It turns out that this is a difficult quantity to pin down. The seek time consists of two key components: the initial startup time, and the time taken to traverse the tracks that have to be crossed once the access arm is up to speed or

Time taken to position the head at the track.

Rotational delay - rotational delay Disks, other than floppy disks, rotate at speeds ranging from 3600 rpm (for handheld devices such as digital cameras) up to, as of this writing, 20,000 rpm; at this latter speed, there is one revolution per 3 ms. Thus, on the average, the rotational delay will be 1.5 ms

Or

Once the track is selected, the disk controller waits until the appropriate sector rotates to line up with the head. The time it takes for the beginning of the sector to reach the head is known as the rotational delay.

Access time - The sum of the seek time, if any, plus the rotational delay. The time it takes to get into position to read or write.

Transfer time - transfer time The transfer time to or from the disk depends on the rotation speed of the disk in the following fashion:

$T = b/ rN$ or

Time taken for data transfer. Once the head is in position, the read or write operation is performed as the sector moves under the head - data transfer portion of the operation.

(d) Briefly define the seven RAID levels.

RAID Level 0: RAID level 0 is not a true member of the RAID family because it does not include redundancy to improve performance. However, there are a few applications, such as some on supercomputers in which performance and capacity are primary concerns and low cost is more important than improved reliability. For RAID 0, the user and system data are distributed across all of the disks in the array. This has a notable advantage over the use of a single large disk: If twodifferent I/O requests are pending for two different blocks of data, then there is a good chance that the requested blocks are on different disks. Thus, the two requests can be issued in parallel, reducing the I/O queuing time.

RAID Level 1 RAID 1 differs from RAID levels 2 through 6 in the way in which redundancy is achieved. In RAID 1, redundancy is achieved by the simple expedient of duplicating all the data. in this case, each logical strip is mapped to two separate physical disks so that every disk in the array has a mirror disk that contains the same data. RAID 1 can also be implemented without data striping, though this is less common.

RAID Level 2 RAID levels 2 and 3 make use of a parallel access technique. In a parallel access array, all member disks participate in the execution of every I/O request. Typically, the spindles of the individual drives are synchronized so that each disk head is in the same position on each disk at any given time

RAID Level 3 RAID 3 is organized in a similar fashion to RAID 2. The difference is that RAID 3 requires only a single redundant disk, no matter how large the disk array. RAID 3 employs parallel access, with data distributed in small strips. Instead of an errorcorrecting code, a simple parity bit is computed for the set of individual bits in the same position on all of the data disks.

RAID Level 4 RAID levels 4 through 6 make use of an independent access technique. In an independent access array, each member disk operates independently, so that separate I/O requests can be satisfied in parallel. Because of this, independent access arrays are more suitable for applications that require high I/O request rates and are relatively less suited for applications that require high data transfer rates.

RAID Level 5 RAID 5 is organized in a similar fashion to RAID 4. The difference is that RAID 5 distributes the parity strips across all disks. A typical allocation is a round-robin scheme, as illustrated in Figure 6.6f. For an n-disk array, the parity strip is on a different disk for the first n stripes, and the pattern then repeats. The distribution of parity strips across all drives avoids the potential I/O bottle-neck found in RAID 4.

RAID Level 6 RAID 6 was introduced in a subsequent paper by the Berkeley researchers [KATZ89]. 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.

or

- RAID 0 - Non-redundant.
- RAID 1 - Mirrored, every disk has a mirror disk containing the same data.
- RAID 2 - Redundant via Hamming code; an error-correcting code is calculated across corresponding bits on each data disk, and the bits of the code are stored in the corresponding bit positions on multiple parity disks.
- RAID 3 - Bit-interleaved parity;
- RAID 4 - Block-interleaved parity;
- RAID 5 - Block-interleaved distributed parity;
- RAID 6 - Block-interleaved dual distributed parity;

(e) How is redundancy achieved in a RAID system?

RAID 1 differs from RAID levels 2 through 6 in the way in which redundancy is achieved. In these other RAID schemes, some form of parity calculation is used to introduce redundancy, whereas in RAID 1, redundancy is achieved by the simple expedient of duplicating all the data

(f) Discuss different optical Disk Products in detail.

CD

Compact Disk. A nonerasable disk that stores digitized audio information. The standard system uses 12-cm disks and can record more than 60 minutes of uninterrupted playing time.

CD-ROM

Compact Disk Read-Only Memory. A nonerasable disk used for storing computer data. The standard system uses 12-cm disks and can hold more than 650 Mbytes.

CD-R

CD Recordable. Similar to a CD-ROM. The user can write to the disk only once.

CD-RW

CD Rewritable. Similar to a CD-ROM. The user can erase and rewrite to the disk multiple times.

DVD

Digital Versatile Disk. A technology for producing digitized, compressed representation of video information, as well as large volumes of other digital data. Both 8 and 12 cm diameters are used, with a double-sided capacity of up to 17 Gbytes. The basic DVD is read-only (DVD-ROM).

DVD-R

DVD Recordable. Similar to a DVD-ROM. The user can write to the disk only once. Only one-sided disks can be used.

DVD-RW

DVD Rewritable. Similar to a DVD-ROM. The user can erase and rewrite to the disk multiple times. Only one-sided disks can be used.

Blu-ray DVD

High-definition video disk. Provides considerably greater data storage density than DVD, using a 405-nm (blue-violet) laser. A single layer on a single side can store 25 Gbytes.

(g) Discuss the CD read and write operation.

Read

Information is retrieved from a CD or CD-ROM by a low-powered laser housed in an optical-disk player, or drive unit. The laser shines through the clear polycarbonate while a motor spins the disk past it . The intensity of the reflected light of the laser changes as it encounters a pit. Specifically, if the laser beam falls on a pit, which has a somewhat rough surface, the light scatters and a low intensity is reflected back to the source. The areas between pits are called lands. A land is a smooth surface, which reflects back at higher intensity. The change between pits and lands is detected by a photo sensor and converted into a digital signal. The sensor tests the surface at regular intervals. The beginning or end of a pit represents a 1; when no change in elevation occurs between intervals, a 0 is recorded.

Write

Recall that on a magnetic disk, information is recorded in concentric tracks. With the simplest constant angular velocity (CAV) system, the number of bits per track is constant. An increase in density is achieved with multiple zoned recording, in which the surface is divided into a number of zones, with zones farther from the center containing more bits than zones closer to the center. Although this technique increases capacity, it is still not optimal.

(h) What differences between a CD and a DVD account for the larger capacity of the latter?

The DVD's greater capacity is due to three differences from CDs (Figure 6.11):

1. Bits are packed more closely on a DVD. The spacing between loops of a spiral on a CD is 1.6 *μm* and the minimum distance between pits along the spiral is *0.834 μm.* The DVD uses a laser with shorter wavelength and achieves a loop spacing of 0.74 *μm* and a minimum distance between pits of *0.4 μm.* The result of these two improvements is about a seven-fold increase in capacity, to about 4.7 GB.

2. The DVD employs a second layer of pits and lands on top of the first layer. A dual-layer DVD has a semi-reflective layer on top of the reflective layer, and by adjusting focus, the lasers in DVD drives can read each layer separately. This technique almost doubles the capacity of the disk, to about 8.5 GB. The lower reflectivity of the second layer limits its storage capacity so that a full doubling is not achieved.

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.

Q.2 Write note on each of the following:

(a) Physical Characteristics of magnetic disk system

Fixed-head disk

- One read-write head per track
- Heads are mounted on a fixed ridged arm that extends across all tracks

Movable-head disk

- One read-write head
- Head is mounted on an arm
- The arm can be extended or retracted

Non-removable disk

- Permanently mounted in the disk drive
- The hard disk in a personal computer is a non-removable disk

Removable disk

- Can be removed and replaced with another disk
- Advantages:
    - Unlimited amounts of data are available with a limited number of disk systems
    - A disk may be moved from one computer system to another
- Floppy disks and ZIP cartridge disks are examples of removable disks

Double sided disk

- Magnetizable coating is applied to both sides of the platter

OR

Fixed-head disk

In a fixed-head disk, there is one read-write head per track. All of the heads are mounted on a rigid arm that extends across all tracks; such systems are rare today.

Movable-head disk

In a movable-head disk, there is only one read-write head. Again, the head is mounted on an arm. Because the head must be able to be positioned above any track, the arm can be extended or retracted for this purpose

Non removable

A non removable disk is permanently mounted in the disk drive; the hard disk in a personal computer is a non removable disk.

Removable disk

A removable disk can be removed and replaced with another disk. The advantage of the latter type is that unlimited amounts of data are available with a limited number of disk systems. Furthermore, such a disk may be moved from one computer system to another.

Floppy disks and ZIP cartridge disks are examples of removable disks.

Double sided disk

the magnetizable coating is applied to both sides of the platter, which is then referred to as double sided. Some less expensive disk systems use single-sided disks

(b) Solid state drives (SSD)

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

(c) Parallel access and independent access RAID schemes

Parallel access - All member disks participate in the execution of every I/O request. Typically, the spindles of the individual drives are synchronized so that each disk head is in the same position on each disk at any given time.

Independent access - Each member disk operates independently, so that separate I/O requests can be satisfied in parallel.

(d) Reasons for DVD's greater capacity over CD

The DVD's greater capacity is due to three differences from CDs :

1. Bits are packed more closely on a DVD. The spacing between loops of a spiral on a

CD is 1.6 μm and the minimum distance between pits along the spiral is $0.834\ \mu m$.

The DVD uses a laser with shorter wavelength and achieves a loop spacing of

0.74 *µm* and a minimum distance between pits of *0.4 µm.* The result of these two improvements is about a seven-fold increase in capacity, to about 4.7 GB.

2. The DVD employs a second layer of pits and lands on top of the first layer. A dual-layer DVD has a semi-reflective layer on top of the reflective layer, and by adjusting focus, the lasers in DVD drives can read each layer separately. This technique almost doubles the capacity of the disk, to about 8.5 GB. The lower reflectivity of the second layer limits its storage capacity so that a full doubling is not achieved.

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.

(e) Solid State Drive Architecture

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 6.8 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.

In addition to the interface to the host system, the SSD contains the following components:

■ Controller: Provides SSD device level interfacing and firmware execution.

■ Addressing: Logic that performs the selection function across the flash memory components.

■ Data buffer/cache: High speed RAM memory components used for speed matching and to increased data throughput.

■ Error correction: Logic for error detection and correction.

■ Flash memory components: Individual NAND flash chips.

(f) Practical Issues peculiar to SSDs

There are two practical issues peculiar to SSDs that are not faced by HDDs:

- **SDD performance has a tendency to slow down as the device is used**
  - The entire block must be read from the flash memory and placed in a RAM buffer
  - Before the block can be written back to flash memory, the entire block of flash memory must be erased
  - The entire block from the buffer is now written back to the flash memory
- **Flash memory becomes unusable after a certain number of writes**
  - Techniques for prolonging life:
    - Front-ending the flash with a cache to delay and group write operations
    - Using wear-leveling algorithms that evenly distribute writes across block of cells
    - Bad-block management techniques
  - Most flash devices estimate their own remaining lifetimes so systems can anticipate failure and take preemptive action

Q.3 Differentiate each of the following:

(a) Magnetic disk read and write Mechanisms

READ:

The traditional read mechanism exploits the fact that a magnetic field moving relative to a coil produces an electrical current in the coil. When the surface of the disk passes under the head, it generates a current of the same polarity as the one already recorded. The structure of the head for reading is in this case essentially the same as for writing and therefore the same head can be used for both. Such single heads are used in floppy disk systems and in older rigid disk systems. Contemporary rigid disk systems use a different read mechanism, requiring a separate read head, positioned for convenience close to the write head. The read head consists of a partially shielded **magnetoresistive (MR) sensor. The MR material** has an electrical resistance that depends on the direction of the magnetization of the medium moving under it. By passing a current through the MR sensor, resistance changes are detected as voltage signals. The MR design allows higher-frequency operation, which equates to greater storage densities and operating speeds.

WRITE:

The write mechanism exploits the fact that electricity flowing through a coil produces a magnetic field. Electric pulses are sent to the write head, and the resulting magnetic patterns are recorded on the surface below, with different patterns for positive and negative currents. The write head itself is made of easily magnetizable material and is in the shape of a rectangular doughnut with a gap along one side and a few turns of conducting wire along the opposite side . An electric current in the wire induces a magnetic field across the gap, which in turn magnetizes a small area of the recording medium. Reversing the direction of the current reverses the direction of the magnetization on the recording medium.

(b) CAV system and multiple zoned recording system

CAV:

A bit near the center of a rotating disk travels past a fixed point (such as a read–

write head) 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. 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, known as the constant angular velocity (CAV)

Multiple zone recording,:

It is a technique in which the surface is divided into a number of concentric zones (16 is typical). Within a zone, the number of bits per track is constant. Zones farther from the center contain more bits (more sectors) than zones closer to the center. This allows for greater overall storage capacity at the expense of somewhat more complex circuitry. As the disk head moves from one zone to another, the length (along the track) of individual bits changes, causing a change in the timing for reads and writes.

(c) Solid-state drives and hard-disk drives

SSD:

1. They have a copy/write speed of 200–550 Mbps.
2. They darw less power averages about 2-3 watts resulting in an increase of 30+ minutes of battery life.
3. Typically they are not larger than 512 GB for notebook size drives and are maximum of 1 TB for desktops.
4. They cost Approx. $0.50 per GB for a 1-TB drive.

HHD:

1. They have a copy/write speed of 50-120 Mbps.
2. They draw more power averages about 6-7 watts and therefore uses more battery

3. Typically they are around 500 GB and 2 TB for notebook size drives and are maximum of 4 TB for desktops.
4. They cost Approx. $0.15 per GB for a 4-TB drive.

(d) CD and DVD

CD

Compact Disk. A nonerasable disk that stores digitized audio information. The standard system uses 12-cm disks and can record more than 60 minutes of uninterrupted playing time.they have max size of 680 mb

DVD

Digital Versatile Disk. A technology for producing digitized, compressed representation of video information, as well as large volumes of other digital data. Both 8 and 12 cm diameters are used, with a double-sided capacity of up to 17 Gbytes. The basic DVD is read-only (DVD-ROM).

OR

Definition of CD: CD is also known as Compact Disc; it is small, round & portable in size. It is made up of molded polymer & is used to store data.

Definition of DVD: DVD is known as Digital Video Disc or Digital Versatile Disc. It stores information but uses a bit different technique then CD.

Now let us discuss few of the differences between CD & DVDs.

Difference between CD/DVD:

- **Storage:**

A CD can store up to 700 Mb of data whereas; a DVD can store 17 GB of data in it.

- **Popular:**

A DVD is more popular in comparison to CD as it can store more data in it.

- **Metal Layers:**

Metal layer is the centre layer in a DVD whereas; in a CD it is beneath the labeling layer.

- **Pit & Land Layer:**

A CD can either have pit or land layer; whereas a DVD has both.

(e) HD DVD and Blu-ray DVD

HD DVD players have been much cheaper than Blu-ray machines,

but Blu-ray discs have more storage space and more advanced protections against piracy.

Both versions deliver sharp resolution.

Blu-Ray has 25 GB capacity (50 GB for dual-layer) and is more expensive.

HD-DVD has 15 GB (30GB for dual layer) and is cheaper than Blu-Ray.