

Advanced Computer Architecture

Lecture No. 32

Reading Material

Vincent P. Heuring & Harry F. Jordan
Computer Systems Design and Architecture

Chapter 9
9.1

Summary

- Hard Disk
- Static and Dynamic Properties
- Examples
- Mechanical Delays and Flash Memory
- Semiconductor Memory vs. Hard Disk

Hard Disk

Peripheral devices connect the outside world with the central processing unit through the I/O modules. One important feature of these peripheral devices is the variable data rate. Peripheral devices are important because of the function they perform.

A hard disk is the most frequently used peripheral device. It consists of a set of platters. Each platter is divided into tracks. The track is subdivided into sectors. To identify each sector, we need to have an address. So, before the actual data, there is a header and this header consisting of few bytes like 10 bytes. Along with header there is a trailer. Every sector has three parts: a header, data section and a trailer.

Static Properties

The storage capacity can be determined from the number of platters and the number of tracks. In order to keep the density same for the entire surface, the trend is to use more number of sectors for outer tracks and lesser number of sectors for inner tracks.

Dynamic Properties

When it is required to read data from a particular location of the disk, the head moves towards the selected track and this process is called seek. The disk is constantly rotating at a fixed speed. After a short time, the selected sector moved under the head. This interval is called the rotational delay. On the average, the data may be available after half a revolution. Therefore, the rotational latency is half revolution.

The time required to seek a particular track is defined by the manufacturer. Maximum, minimum and average seek times are specified. Seek time depends upon the present position of the head and the position of the required sector. For the sake of calculations, we will use the average value of the seek time.

- **Transfer rate**

When a particular sector is found, the data is transferred to an I/O module. This would depend on the transfer rate. It would typically be between 30 and 60 Mbytes/sec defined by the manufacturer.

- **Overhead time**

Up till now, we have assumed that when a request is made by the CPU to read data, then hard disk is available. But this may not be the case. In such situation we have to face a queuing delay. There is also another important factor: the hard disk controller, which is the electronics present in the form of a printed circuit board on the hard disk. So the time taken by this controller is called over head time.

The following examples will clarify some of these concepts.

Example 1

Find the average rotational latency if the disk rotates at 20,000 rpm.

Solution

The average latency to the desired data is halfway round the disk so

$$\begin{aligned}\text{Average rotational latency} &= 0.5 / (20,000 / 60) \\ &= 1.5 \text{ms}\end{aligned}$$

Example 2

A magnetic disk has an average seek time of 5 ms. The transfer rate is 50 MB/sec. The disk rotates at 10,000 rpm and the controller overhead is 0.2 msec. Find the average time to read or write 1024 bytes.

Solution

$$\text{Average } T_{\text{seek}} = 5 \text{ms}$$

$$\text{Average } T_{\text{rot}} = 0.5 * 60 / 10,000 = 3 \text{ ms}$$

$$T_{\text{transfer}} = 1 \text{KB} / 50 \text{MB} = 0.02 \text{ms}$$

$$T_{\text{controller}} = 0.2 \text{ms}$$

$$\begin{aligned}\text{The total time taken} &= T_{\text{seek}} + T_{\text{rot}} + T_{\text{tsfr}} + T_{\text{ctr}} \\ &= 5 + 3 + 0.02 + 0.2 \\ &= 8.22 \text{ ms}\end{aligned}$$

Example 3

A hard disk with 5 platters has 1024 tracks per platter, 512 sectors per track and 512 bytes/sector. What is the total capacity of the disk?

Solution

$$512 \text{ bytes} \times 512$$

$$\text{sectors} = 0.2 \text{MB/track}$$

$$0.2 \text{MB} \times 1024 \text{ tracks} = 0.2 \text{GB/platter}$$

$$\text{Therefore the hard disk has the total capacity of } 5 \times 0.2 = 1 \text{GB}$$

Example 4

How many platters are required for a 40GB disk if there are 1024 bytes/sector, 2048 sectors per track and 4096 tracks per platter

Solution

The capacity of one platter

$$= 1024 \times 2048 \times 4096$$

$$= 8\text{GB}$$

For a 40GB hard disk, we need $40/8$

$$= 5 \text{ such platters.}$$

Example 5

Consider a hard disk that rotates at 3000 rpm. The seek time to move the head between adjacent tracks is 1 ms. There are 64 sectors per track stored in linear order.

Assume that the read/write head is initially at the start of sector 1 on track 7.

- How long will it take to transfer sector 1 on track 7 to sector 1 on track 9?
- How long will it take to transfer all the sectors on track 12 to corresponding sectors on track 13?

Solution

Time for one revolution $= 60/3000 = 20\text{ms}$

- Total transfer time = sector read time + head movement time + rotational delay + sector write time

$$\text{Time to read or write on sector} = 20/64 = 0.31\text{ms/sector}$$

$$\text{Head movement time from track 7 to track 9} = 1\text{ms} \times 2 = 2\text{ms}$$

After reading sector 1 on track 7, which takes 0.31ms, an additional 19.7 ms of rotational delay is needed for the head to line up with sector 1 again.

The head movement time of 2 ms gets included in the 19.7 ms. Total transfer time $= 0.31\text{ms} + 19.7\text{ms} + 0.31\text{ms} = 20.3\text{ms}$

- The time to transfer all the sectors of track 12 to track 13 can be computed in the similar way. Assume that the memory buffer can hold an entire track. So the time to read or write an entire track is simply the rotational delay for a track, which is 20 ms. The head movement time is 1ms, which is also the time for $1/0.3 = 3.3 \approx 4$ sectors to pass under the head. Thus after reading a track and repositioning the head, it is now on track 13, at four sectors past the initial sector that was read on track 12. (Assuming track 13 is written starting at sector 5) therefore total transfer time $= 20 + 1 + 20 = 41\text{ms}$. If writing of track 13 start at the first sector, an additional 19 ms should be added, giving a total transfer time $= 60\text{ms}$

Example 6

Calculate time to read 64 KB (128 sectors) for the following disk parameters.

- 180 GB, 3.5 inch disk
- 12 platters, 24 surfaces
- 7,200 RPM; (4 ms avg. latency)
- 6 ms avg. seek (r/w)
- 64 to 35 MB/s (internal)
- 0.1 ms controller time

Solution

Disk latency = average seek time + average rotational delay + transfer time + controller overhead
 $= 6 \text{ ms} + 0.5 \times 1/(7200 \text{ RPM}) / (60000 \text{ ms/M}) + 64 \text{ KB} / (64 \text{ MB/s}) + 0.1 \text{ ms}$
 $= 6 + 4.2 + 1.0 + 0.1 \text{ ms} = 11.3 \text{ ms}$

Mechanical Delay and Flash Memory

Mechanical movement is involved in data transfer and causes mechanical delays which are not desirable in embedded systems. To overcome this problem in embedded systems, flash memory is used. Flash memory can be thought of a type of electrically erasable PROM. Each cell consists of two MOSFET and in between these two transistors, we have a control gate and the presence/absence of charge tells us that it is a zero or one in that location of memory.

The basic idea is to reduce the control overheads, and for a FLASH chip, this control overhead is low. Furthermore flash memory has low power dissipation. For embedded devices, flash is a better choice as compared to hard disk. Another important feature is that read time is small for flash. However the write time may be significant. The reason is that we first have to erase the memory and then write it. However in embedded system, number of write operations is less so flash is still a good choice.

Example 7

Calculate the time to read 64 KB for the previous disk, this time using 1/3 of quoted seek time, 3/4 of internal outer track bandwidth

Solution

Disk latency = average seek time + average rotational delay + transfer time + controller overhead
 $= (0.33 \times 6 \text{ ms}) + 0.5 \times 1/(7200 \text{ RPM})$
 $+ 64 \text{ KB} / (0.75 \times 64 \text{ MB/s}) + 0.1 \text{ ms}$
 $= 2 \text{ ms} + 0.5 / (7200 \text{ RPM} / (60000 \text{ ms/M}))$
 $+ 64 \text{ KB} / (48 \text{ KB/ms}) + 0.1 \text{ ms}$
 $= 2 + 4.2 + 1.3 + 0.1 \text{ ms} = 7.6 \text{ ms}$

Semiconductor Memory vs. Hard Disk

At one time developers thought that development of semiconductor memory would completely wipe out the hard disk. There are two important features that need to be kept in mind in this regard:

1. Cost

It is low for hard disk as compared to semi-conductor memory.

2. Latency

Typically latency of a hard disk is in milliseconds. For SRAM, it is 10^5 times lower as compared to hard disk.