Computer Architecture and Organization
Prof. Indranil Sengupta
Department of Computer Science and Engineering
Indian Institute of Technology, Kharagpur

Lecture - 43
Secondary Storage Devices

In this week we shall be starting our discussion on input output devices and the ways in which you can interface such devices to the computer system. Now when we talk about interfacing IO devices the most important kind of IO device interfacing that we encounter is that of the secondary memory, which traditionally has been the hard disk.

Well of course, very recently there is a new technology which has come up; the so called solid state disks that are based on something called flash memory technology. They are also very fast replacing the hard disk technology. The topic of our discussion in this lecture is secondary storage devices.

(Refer Slide Time: 01:15)

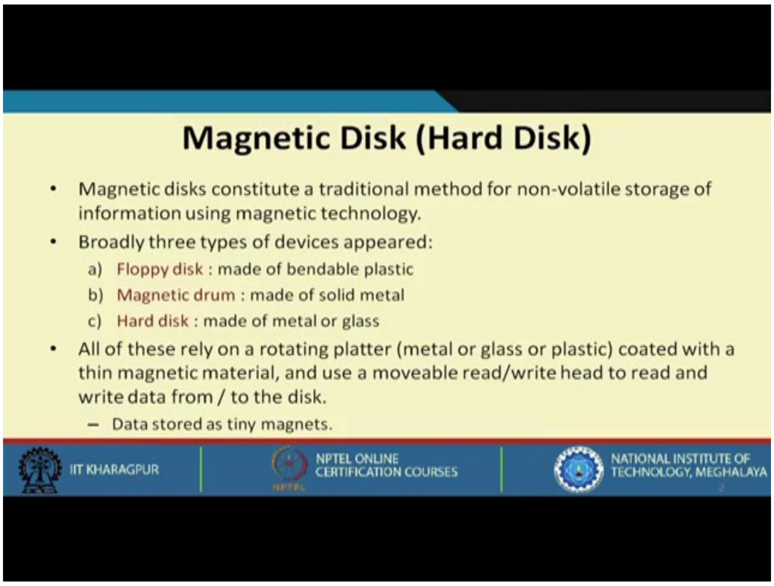

We start with our discussion on magnetic disks which has been around for several decades already. Now actually how does a magnetic disk work? Magnetic disks, as the name implies, rely on some kind of magnetic technology. There is some kind of ferromagnetic material coated on the surface of the disk, and with the help of some externally applied magnetic field we can change the orientation of tiny magnets on the surface.

Depending on the orientation we can say that we are storing either 0 or 1, this is the basic idea. Now because of the accuracy of positioning those electromagnets on the surface of the disk it is possible to store a very large number of zeroes and ones in the form of small magnets on the surface of the disk.

This is the traditional method that has been used since many decades. One characteristic of this kind of a storage media is that it is non volatile. Non volatile means whenever the power is withdrawn or the power is switched off, whatever you are storing will remain; that will not go off because it is based on magnetic technology, the magnets will remain even when we withdraw the power. So, this is non-volatile and it uses magnetic technology.

Now, over the years there are broadly three kinds of magnetic storage devices that have appeared. Magnetic drum was the first to appear, then floppy disks and hard disks. Magnetic drum traditionally was built out of solid metal on the top of which some ferromagnetic material was coated. Floppy disk was similar, but it was made of plastic, which was soft and you could bend the surface of the floppy disk, that is why the name was floppy.

In contrast the hard disk that you see today are really hard; they are made of metal or glass, you cannot bend the disks without breaking. Now all these different devices which I have talked about they are all relying on some kind of rotating platter. Either it is a wheel which is rotating or it is a drum which is rotating. So, the idea is that something is rotating and there is some kind of a sensor or an actuator we call it the read write head. The read write head is sitting in one place and the surface is rotating or moving under that read write head. So, you can read or write the bits on the surface.

The rotating platter as I had said are made of either metal or glass or plastic. They are coated with a very thin magnetic material and as I had said there is a read write head, for some of the devices the read write head is fixed like in magnetic drum, but for floppy and hard disks they can also move. And as I had said the data are stored as tiny magnets.

(Refer Slide Time: 05:06)

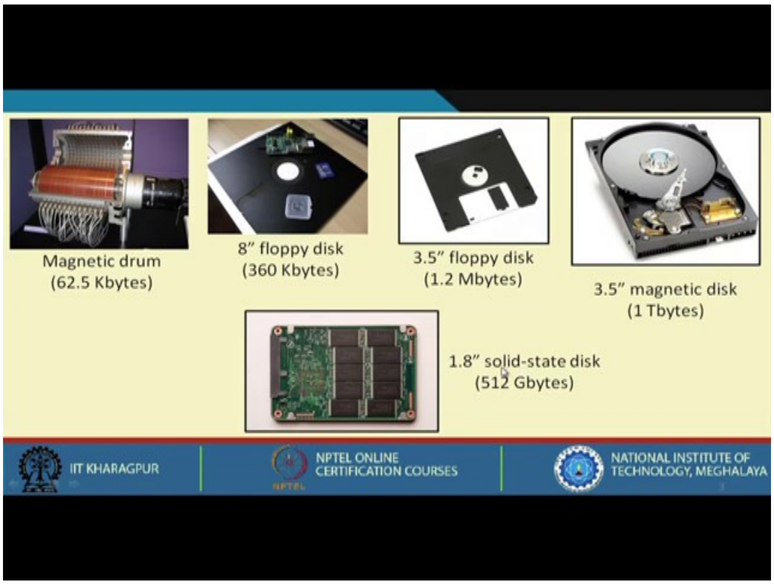

Here we have some pictures of the different kinds of devices which had appeared in the market over the years. You see this was the first kind of device to appear in the late 60s and early 70s; the magnetic drum. As you can see there was a cylindrical kind of a thing which was rotating and there were read write heads positioned along the axis. All these read write heads were able to read or write data bits on the surface, and in this kind of a device the total capacity was only about 62.5 kilobytes; it was pretty less in those days.

Later on in the 1980s floppy disks appeared. The first kind of floppy disks were of 8 inch size. The diameter was 8 inch; this kind of a disk was able to store 360 kilobytes of data. Subsequently we had a smaller version of the floppy disks 3.5 inch, which some of you may still be using today. Of course, these have become almost obsolete, their capacities are 1.2 megabytes, then came the hard disks.

Well here the picture that I am showing is the picture of a 3.5 inch magnetic disk, where the capacity can be as high as one terabytes or even more. But believe me the first hard disks which appeared in the market were pretty large in size. The diameter was more than a foot. So, it is more than 12 inches in diameter they were really big, but over the time it has become smaller; now we can have three and half inch hard disks.

Now, I am showing another kind of a device here. This is technically not a hard disk, but this is replacing hard disk, this is the so called solid state disk. Here there is something which is rotating, and there is a read write head which moves on the surface. So, there is some kind of mechanical movement here; because of mechanical movement the overall life of the device will be relatively less.

In contrast the solid state disk does not have any moving parts, they are based on integrated circuit technologies. They are chips, there are no mechanically moving parts and secondly, they are also smaller in size. This is an example which is an 1.8 inch disk, the capacity is 512 gigabytes.

(Refer Slide Time: 07:59)

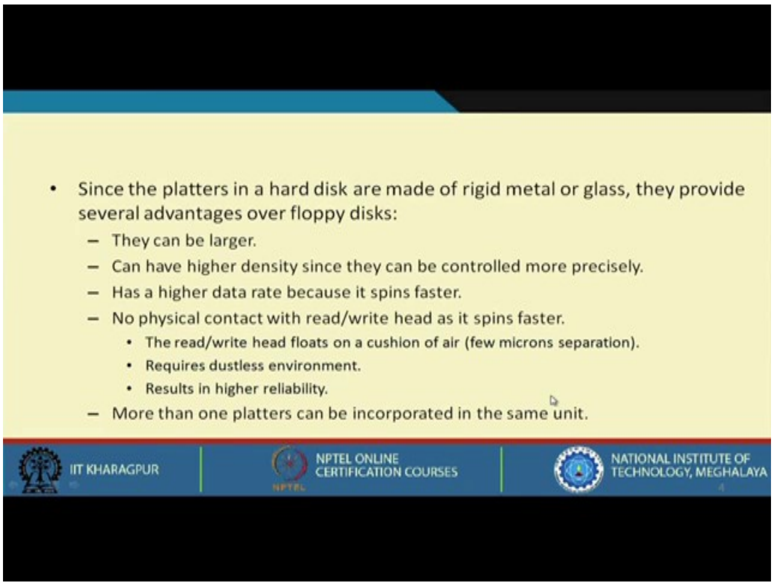

Now, let us come back to hard disk. In a hard disk as I had said the platters are made of metal or glass which are rigid. There are some advantages because in floppy disks the disk surface is soft it can bend. Because the hard disks are rigid they can be made larger because you can more accurately position the read write head on a solid surface. If it is soft it can bend, and a slight bend can lead to a misalignment of the read write head; you can read the wrong data by accident. Because of that hard disks can have higher density and they can be controlled more accurately or precisely.

And in a floppy disk the read write head was actually physically coming in contact with the disk and the disk was rotating. In contrast in a hard disk the disk is rotating at a much higher speed and the read write head is not touching the surface, there is a thin layer of air, the read write head floats on top of the surface. Because of that the wear and tear on the surface is also much less.

Because its spins faster the data rate can also be higher because you can read write at a higher rate, there is no physical contact as I had said. So, the read write head actually floats on a thin layer of air which is few microns thick. Because of this even if there is a single spec of dust on the surface of the disk the read write head will crash on the dust because the layer of air is even thinner than the spec of dust. So, the disk has to be put in a dust free enclosure.

This obviously results in higher reliability and for higher capacity you can have more than one disks in the same unit. There can be several disks, you can put them one on top of the other; several disk can rotate at the same time. They will all be having separate read write heads something like this.

(Refer Slide Time: 10:34)

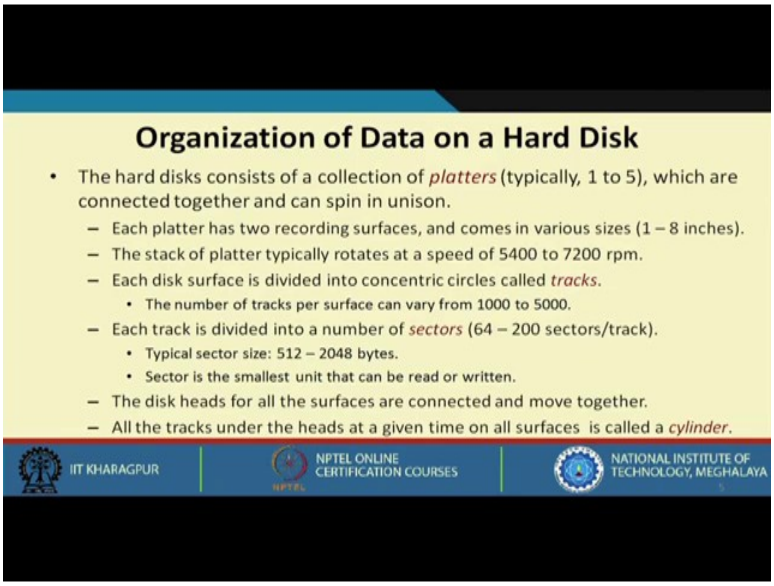

Let see how data are organized on a hard disk. As I had said each disk is called a platter, platter is a circular hard material which is magnetically coated typically on both the sides. Now a magnetic disk unit as I had said can contain more than one platter which rotates all simultaneously.

So, there is a collection of platters, typically the number of platters will be 1 to 5 which are connected together, and as I had said they spin together. Each platter has two recording surfaces that are magnetically coated on both sides, and the platters can come in various sizes. There are many disks of about one inch diameter also available, but the larger ones are about 8 inches in diameter.

But the more commonly used ones today are typically three and half inch diameter, there are larger ones also. And these platters rotate at speed typically that range from 3600 revolutions per minute (RPM), to 5400 or 7200 RPM.

(Refer Slide Time: 12:17)

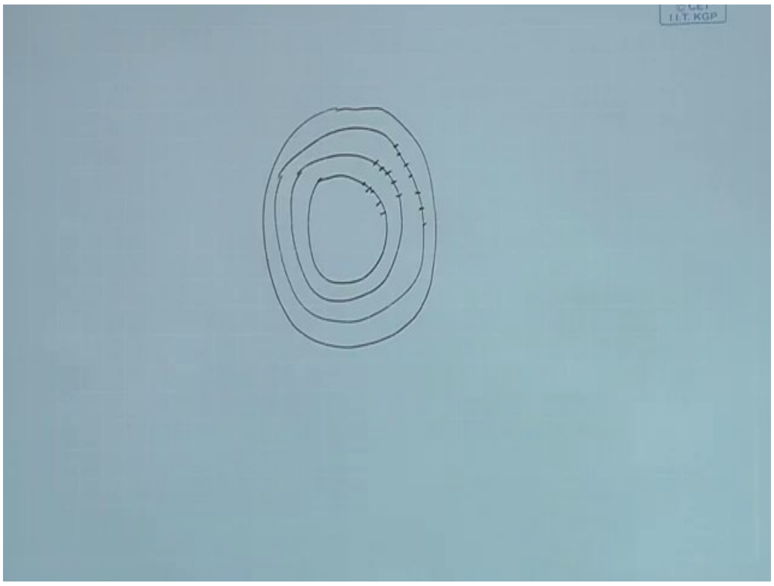

Each disks surface is divided into concentric circular regions called tracks. If this is the top view of a platter, then you can have the circular regions which are concentric circles. These are called tracks and we store the individual data bits along these tracks. This is unlike record player, for those of you have seen how a record player works. In a record player, you put the playback pin on one side and it is a helix kind of a thing, it slowly moves towards the centre; that is how a record player stores the data. But in a disk they are concentric circles called tracks, and typically the number of tracks can vary from 1000 to 5000.

And again each track is divided into a number of logical sectors. There can be 64 to 200 sectors per track, and each sector has size typically ranging from 512 to 2048 bytes. And with the disks the point to notice that sector is the smallest unit of data that you can read or write. This means you cannot read a single byte or you cannot write a single byte; you will have to read an entire sector or you will have to write an entire sector, this is something you have to remember.

Sector is the minimum unit of data that can be transferred to and from the disk, and the disk heads for all the surfaces as I had said are connected together and move. And if there are multiple platters all the tracks under the heads at a time is called a cylinder. So, there are several platters on each of the platter there is a track. Track number 1 on all the surfaces taken together form something called a cylinder. Similarly track 2 they form another cylinder, track 3 they form another cylinder, like that the same track on all the surfaces or the platters constitute a cylinder.

(Refer Slide Time: 15:01)

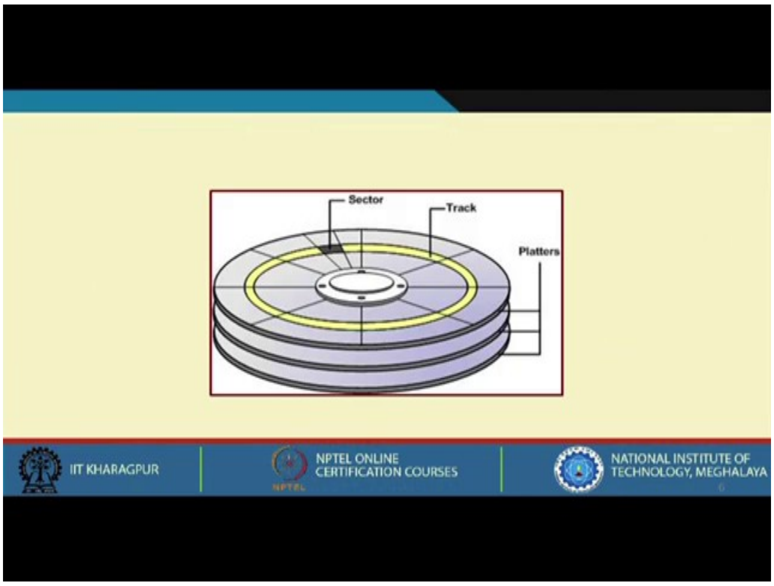

Pictorially you can see these are the platters. This yellow ring is the track, this is the concentric circle and the track is broken up into a number of small sectors, one of the sector is shown here. There are many sectors. And for individual disks or platters there are magnetic coatings on both sides. So, if there are three disks or platters, then there can be six recording surfaces.

(Refer Slide Time: 15:39)

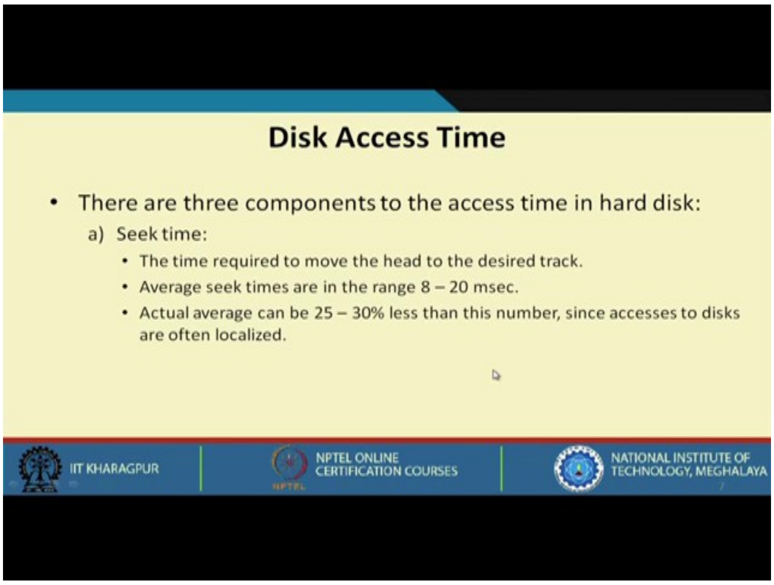

Now, when you talk about the access time of a disk there are three components to it. First is call the seek time. You see a disk contains a number of tracks, right now the read write head may be located somewhere. Suppose you want to read data from track number 10. The first thing is that you have to move the head to track number 10; that movement time is referred to as the seek time.

Seek time is the time required to move the read write head to the desired track, and this seek time on the average can range from 8 to 20 milliseconds. So, you see this times are of the order of milliseconds, but actually this number on the average can be 25 to 30 percent less, because for every stored data on the disks we normally do not have to move across many tracks, normally the related data are stored in nearby tracks. So, we have to move the read write head by smaller amounts that will be requiring less time. If you move the head across many tracks the time required will be more. So, this is seek time.

(Refer Slide Time: 17:08)

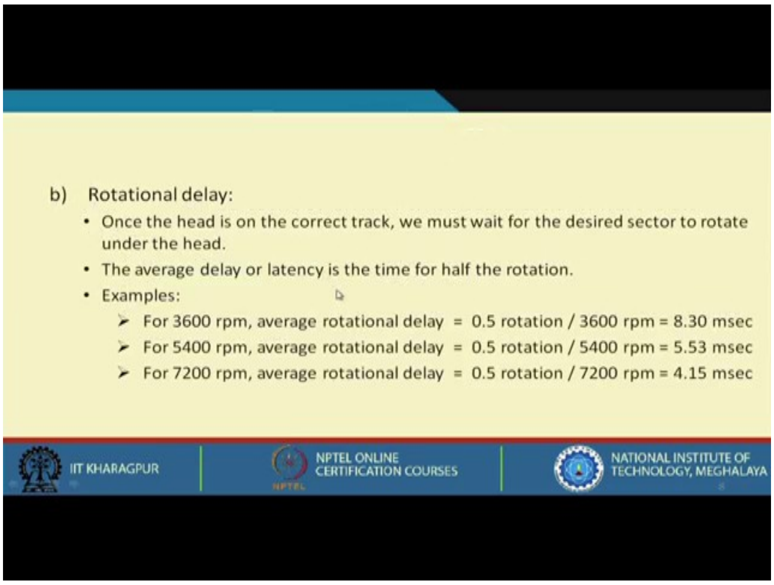

Next comes the rotational delay. Suppose we have reached the correct track. On the tracks suppose I want to read data from sector number 5. So, I will have to wait until sector number 5 rotates under the head. That is the rotational delay; I will have to wait for the disk to rotate, so that the data I want to read or write comes under the head. We must wait for the desired sector to rotate under the head.

On the average this delay will be half the rotational delay, because in the best case you may already be there, or in the worst case you may have to wait for an entire rotation. So, on the average it will be half. Some example value of rotational delay for 3600 rpm disk is shown. Half rotation per 3600 rpm, if you divide this and if you calculate you will get 8.3 milliseconds. For 5400 rpm the delay becomes 5.53 milliseconds, and for 7200 it becomes even less 4.15 milliseconds.

(Refer Slide Time: 18:55)

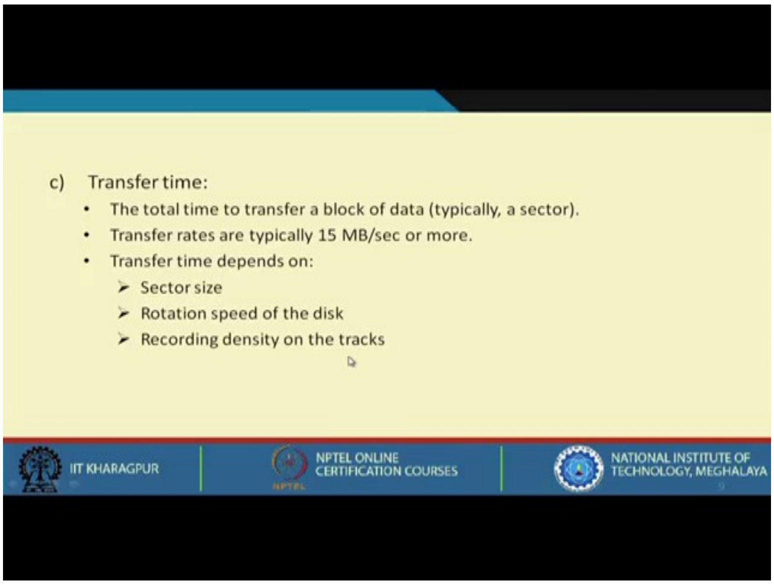

The data transfer rate in bytes per second or kilobytes per second whatever you say, is the total time required to transfer a block of data, say a sector or a track, in a disk. The typical transfer rates are typically of the order of 15 megabytes. The transfer time depends on several things, the sector size, rotation speed of the disk because faster the disk is rotating my speed will be higher, and also recording density on the tracks.

Are we recording the zero and ones very sparsely or very densely? If you are recording very densely then the reading and writing you can also do faster because the disk is rotating at the same speed. So, all these things will determine how fast we can read or write the data.

(Refer Slide Time: 20:00)

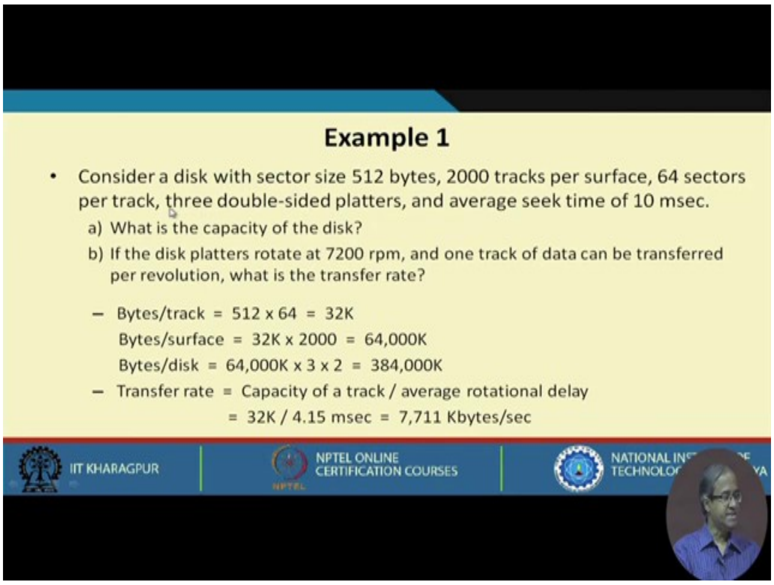

Let us work out a simple example. Let us consider a disk, where the sector size is 512 bytes, there are 2000 tracks per surface 64 sectors per track, and there are 3 platters double-sided, so there are 6 surfaces, and average seek time is 10 millisecond. The first question is; what is the capacity of the disk and the second question is if the disk rotates at 7200 rpm and we assume that one track of data can be transferred per revolution what is the transfer rate?

In the first question, capacity of the disk calculation is shown.

The transfer rate calculation is also shown.

(Refer Slide Time: 22:26)

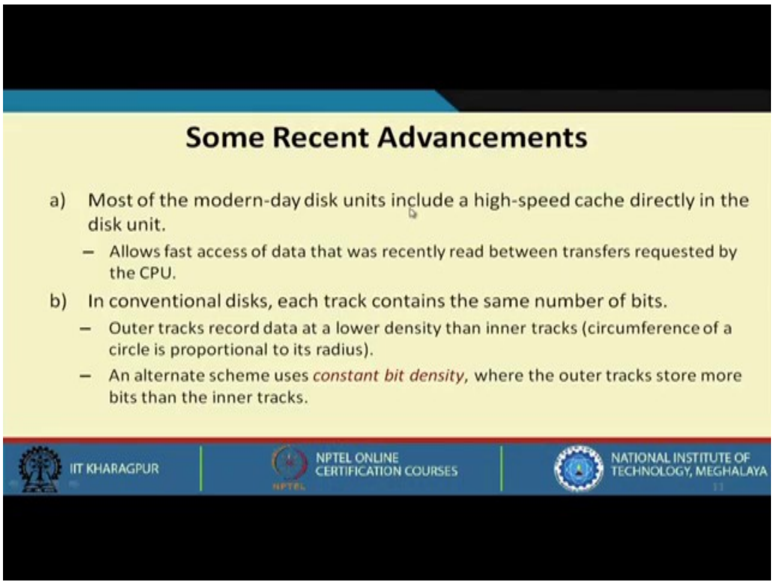

Some recent advancements in hard disk technology they allow two things. Firstly, if you break open a disk sometime you will see that other than the disk surface there are lot of circuits inside. Among other things they include a high speed memory that is called as disk cache directly inside the disk unit. Whenever data are transferred from the disk, they are first stored in the cache, and whenever cpu requests they can be directly accessed from the disk cache. So, it will much faster that way. This use of the disk cache speeds up the data transfer of the average.

And the second advancement is that in the conventional earlier disks. So, each of the tracks has the same capacity means if one track content 32 kilobytes of data, the innermost track which is much smaller will also contain 32 kilobytes of data. So, irrespective of the circumference of the tracks the capacity was the same. But, imagine the tracks which are outside the circumference is much larger. So, potentially they can store more data.

In the earlier disks it was not done that way. Now you can have a scheme called constant bit density scheme where the outer tracks can store more bits than the inner tracks, this is in contrast with the traditional or conventional way where the outer and inner tracks all stored the same amount of data, which means the outer tracks recorded data at a lower density.

(Refer Slide Time: 24:34)

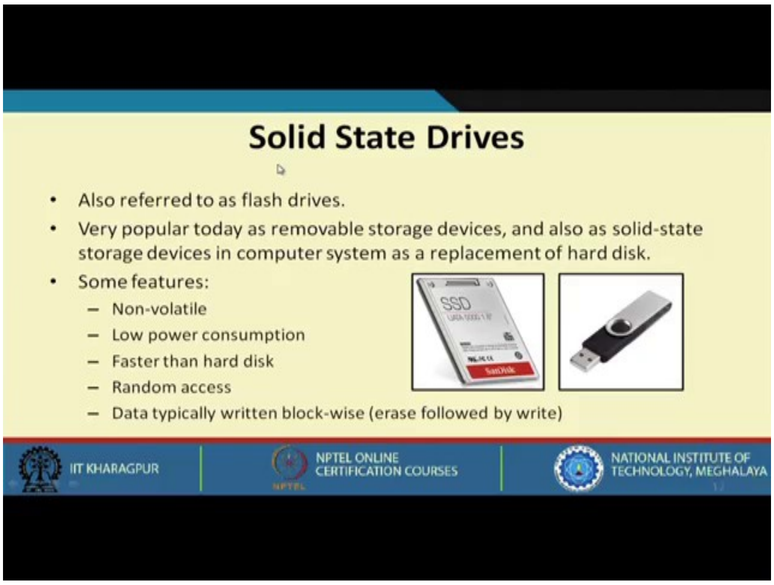

Now, let us come to the other competing technology which is now threatening to replace the hard disks, the so called solid state drives. We are already using solid state drives in our daily life in the form of the pen drives, pen drives are everywhere. Pen drive is one form of solid state drive that is storage where there is no moving part, it is electronic, it is non-volatile, and it is very convenient. Also it is faster than hard disk this is another very important characteristics.

Solid state drives are very popular today as removable storage devices. You can see a picture of a traditional solid state hard disk, and a pen drive which you all know. These devices are already replacing hard disks in many computer systems like laptops. Some of the features are unlike hard disks, they are non volatile, they have extremely low power consumption, they are faster than hard disk; at least 10 times faster, random access and data are written blockwise. One thing I did not mention; this device does not have any moving parts like a hard disk, that is another advantage. So, the life is also higher.

(Refer Slide Time: 26:08)

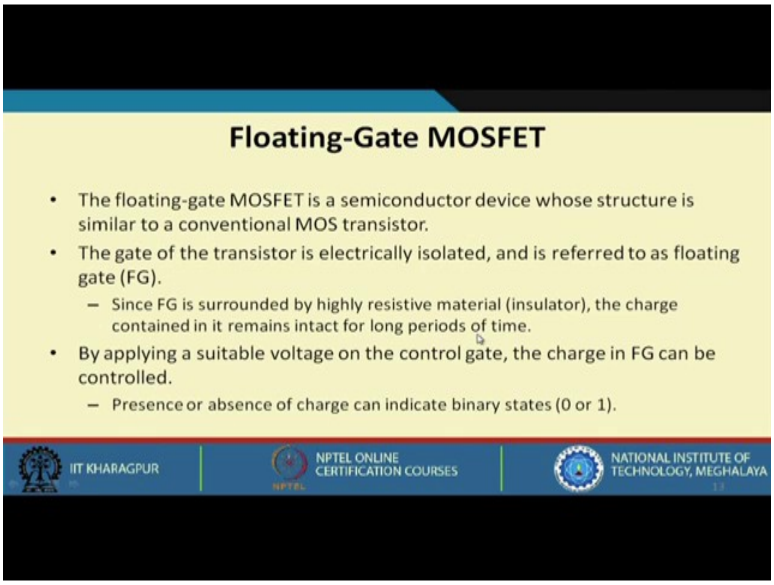

The basic technology that is used to built this kind of flash or solid state memory devices is called a floating gate MOS transistor. The floating gate MOSFET is a device whose structure is somewhat similar to a conventional MOS transistor, which is used in CMOS technology for building chip today.

The difference is that the gate of the transistor is electrically isolated, and this is referred to as floating gate. Electrically isolated means there is an insulating material and inside that insulating material the gate is residing or is hidden inside. If you can somehow put some charge on the gate there is no path for the charge to leak out, the charge will remain. This is the basic idea how it can memorize some information; you put or force some charge in the gate, and because of the insulation the charge cannot leak out and it can be retained for quite long amount of time.

(Refer Slide Time: 27:45)

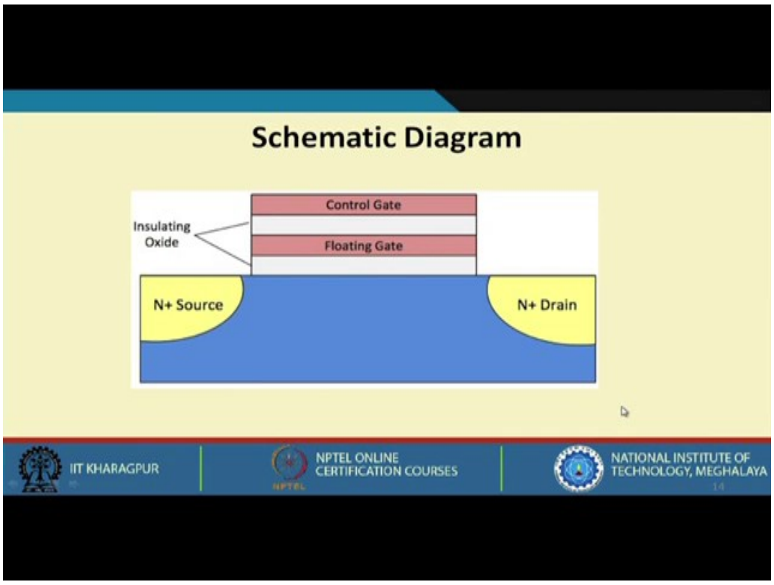

By applying a suitable voltage on a control gate you can control the charge and the presence or absence of the charge can be representative of the binary state 0 or 1. This is how you can store. Pictorially the floating gate transistor looks like this. Just like a normal MOS transistor we have the n-type regions. One is the source other is drain; there is a gate, but in between there is another floating gate which is covered by insulating oxide.

(Refer Slide Time: 28:18)

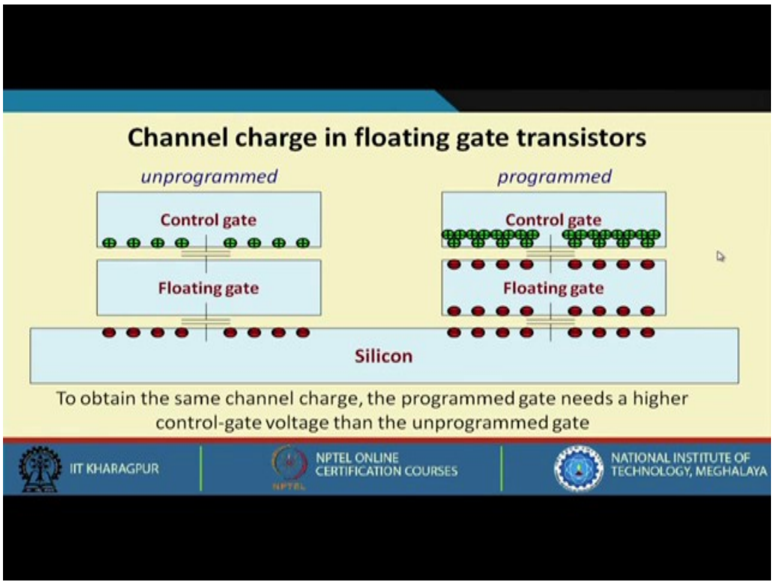

By applying some voltage on the control gate you can inject some electrons or some charges within the floating gate like this. Suppose when there is no charge on the floating gate, when you apply some normal voltage on the control gate some electrons will be attracted on the channel just like a normal transistor and between the source and drain a current will be flowing.

But when you program this gate like you inject some charge in the floating gate; that means, already some electrons are injected then to have this channel you have to put in more voltage on the control gate to create the channel. So, just like checking this that how much voltage is required for a current to flow you can check whether this gate is charged or not charged, whether it is 1 or 0.

(Refer Slide Time: 29:12)

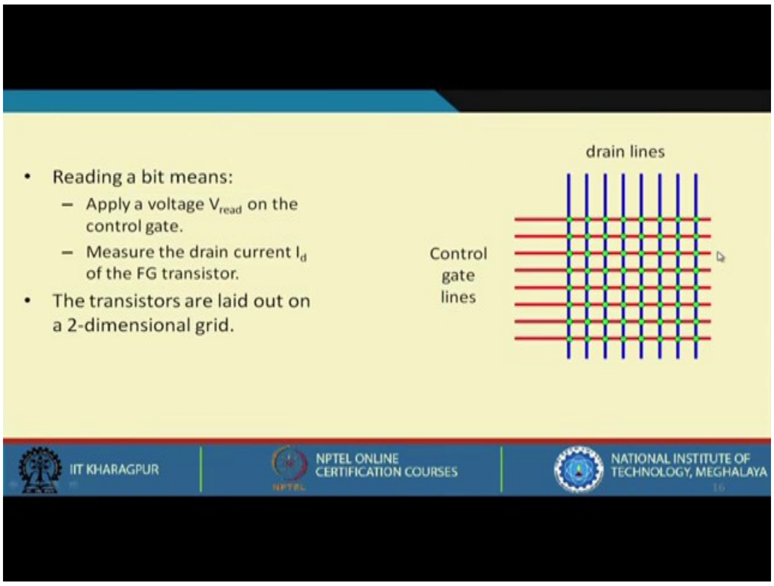

Now, when you talk about reading a bit it means you apply a voltage on the control gate and measure the drain current whether the current is flowing or not; and typically this transistors are laid out on a two dimensional grid. These red and blue lines are on two different levels they are not touching, and the transistors are fabricated at the junctions; one is the drain other is the source, this is the gate and transistors connected like this.

(Refer Slide Time: 29:45)

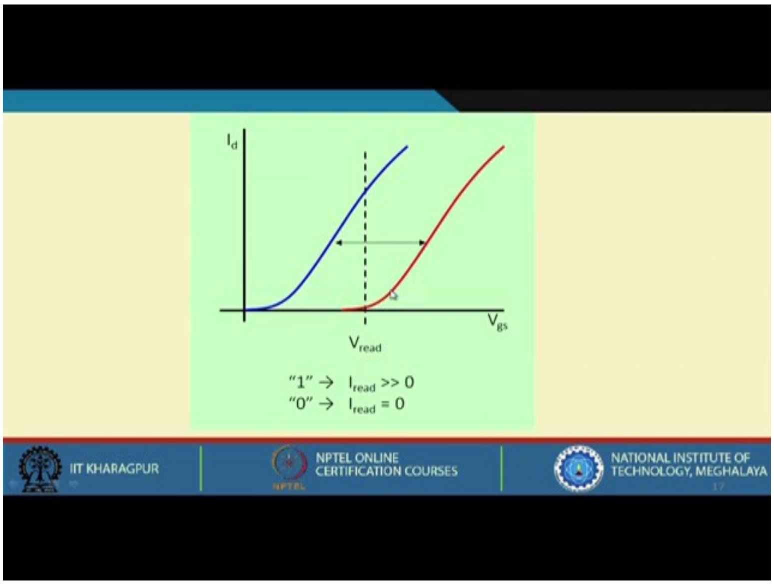

These are the characteristic curves. As you can see for one state the current will be very less, and for the other state the current will be quite high. So, just by measuring the current you can sense that in which state your device is in, 1 or 0.

(Refer Slide Time: 30:10)

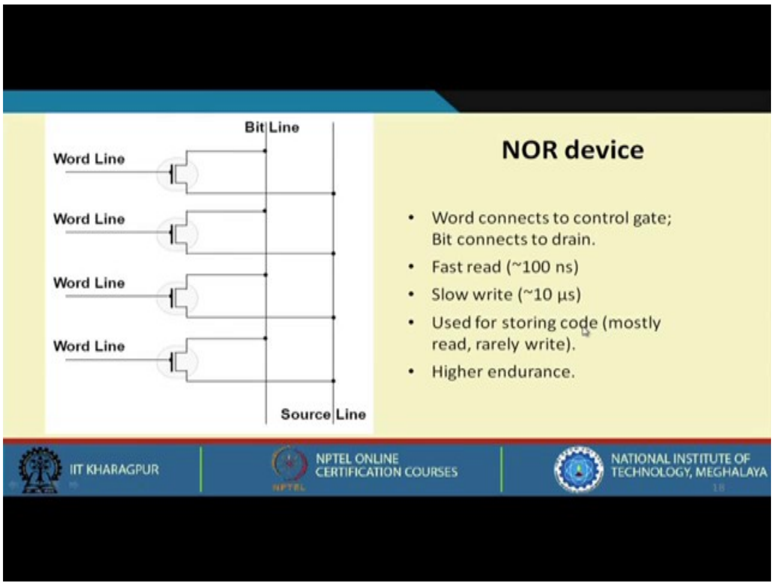

Now, the solid state device is can be fabricated in two ways, one is called NOR flash or a NOR device, other is called NAND device. A NOR device looks like this. Here the floating gate transistors are fabricated like this, so all the sources are made common and the drains are connected to the bit line and the word lines are connected to the gate. So, when you enable a transistor, whatever you apply you can measure the current between source and bit line, and you can check whether you are storing 1 or 0. By applying a high voltage you can inject some charge on the floating gate; this is how NOR devices are laid out.

Now, the characteristics of NOR devices is that reading of data is very fast, of the order 100 nanosecond. So, you can compare the time with that of hard disk where the seek time itself takes around 10 milliseconds plus rotational delay another 10 milliseconds. But write is relatively slow or because in write you have to apply a higher voltage and have to wait for a longer time for the charge to leak inside the floating gate; that takes about 10 microsecond.

So, this is a memory which can be read fast, but writing is slower. So, you can use it in applications where you are storing some program code which you are mostly reading, but rarely writing. This kind of NOR devices have higher endurance, but the capacity is smaller because layout wise it occupies more area.

(Refer Slide Time: 32:17)

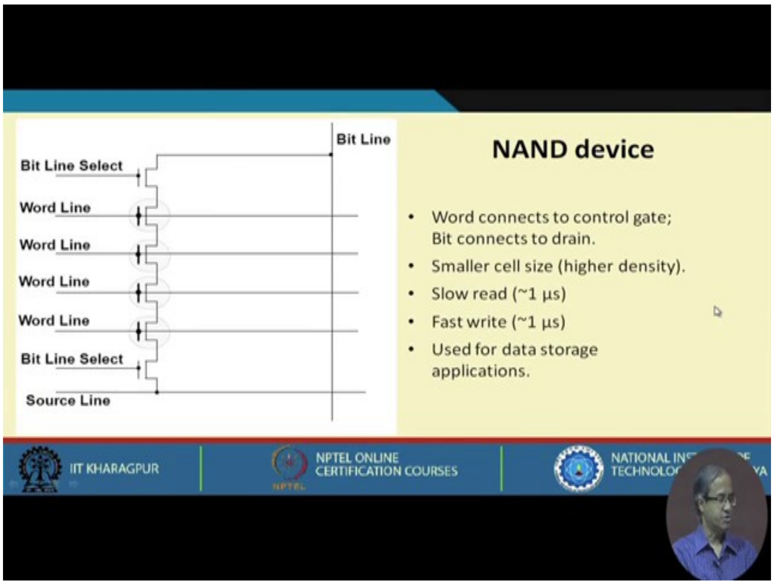

The other device is NAND device. You can see here there are two wires coming out of each transistor that is why it takes more area, but for a NOR device the transistors are connected in series like this, there is a bit select line on two side I am not going to detail. The source is here and bit line if they are activated, when you are selecting this line there will be a current flowing you can sense it.

In case of NAND device the cell size is smaller resulting in higher density. All these solid state drives that you see today are built using this kind of NAND technology, because it has higher density. Now here reading and writing times are almost similar, slower read as compared to NOR, but as compared to NOR writing is faster. So, read and write are almost the same time and this is used for data storage applications like pen drives, solid state drives etc.

(Refer Slide Time: 33:24)

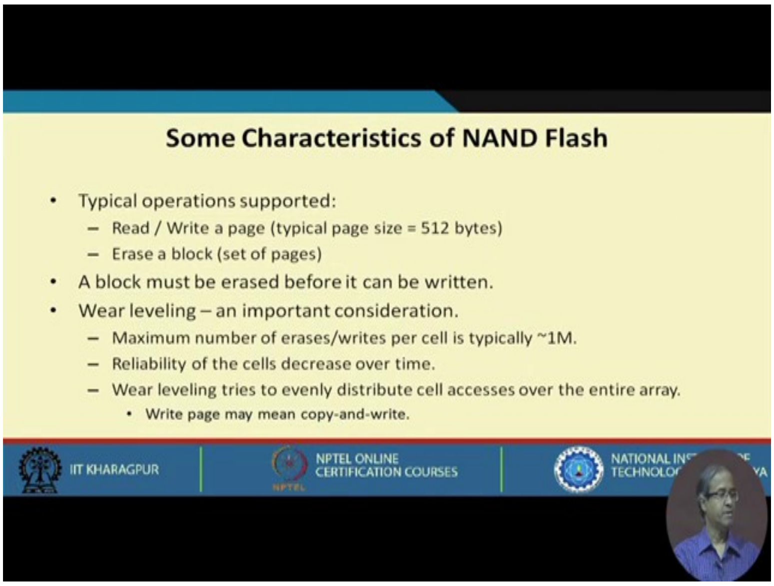

Some characteristics of NAND flash are as follows. Some typical operations that are supported are read or write a page. Again you cannot read or write a byte or word. You can write an entire page which is of the order 512 bytes, because you see in a NAND so many things are connected in series. So, when you do an operation you have to operate on all of them, and by applying a voltage you can erase all the blocks in a page; means all the cells in a block. A block is defined as a set of pages and another thing is that when you write into such a device first you have to erase the block only then you can write.

Another important consideration in this kind of flash is something called wear levelling. You see the life of the devices specified by the manufacture as 1 million writes, that the manufacture tells you, but if by accident you are writing most frequently in a few of the cells, those few cells will go bad very quickly because they will reach 1 million, but the others may be still fresh.

Wear levelling means you consciously try to distribute the writes across all the cells. If you see that you are writing into one cell many times you will explicitly make a copy of that cell to some other cell and write it there. So, lot of copies and writes are required due to wear levelling considerations. Wear levelling tries to evenly distribute cell accesses, write page may mean copy to some other position and then write.

With this we come to the end of this lecture where we tried to give you a brief idea about the technologies that are used to built a hard disk or a solid state disk. Due to their characteristics and differences whenever you are designing or building a memory hierarchy or a computer system you will be able to take a more informed decision based.

Thank you.