Construction and Operation of the Hard Disk

- Hard Disk Platters and Media

- Platter size, number of platters

- Platter Substrate Materials, Magnetic Media, Areal Density

Construction and Operation of the Hard Disk

- To many people, a hard disk is a "**black box**" of sorts--it is thought of as just a small device that "somehow" stores data. There is nothing wrong with this approach of course, as long as all you care about is that it stores data. If you use your hard disk as more than just a place to "keep stuff", then you want to know more about your hard disk. It is hard to really understand the factors that affect performance, reliability and interfacing without knowing how the drive works internally.

- **Fortunately, most hard disks are basically the same on the inside**.

- While the technology evolves, many of the basics are unchanged from the first PC hard disks in the early 1980s.

Construction and Operation of the Hard Disk

- **Photograph** of a **modern SCSI hard disk**, with major components annotated.The logic board is underneath the unit and not visible from this angle.*Original image* © *Western Digital Corporation*

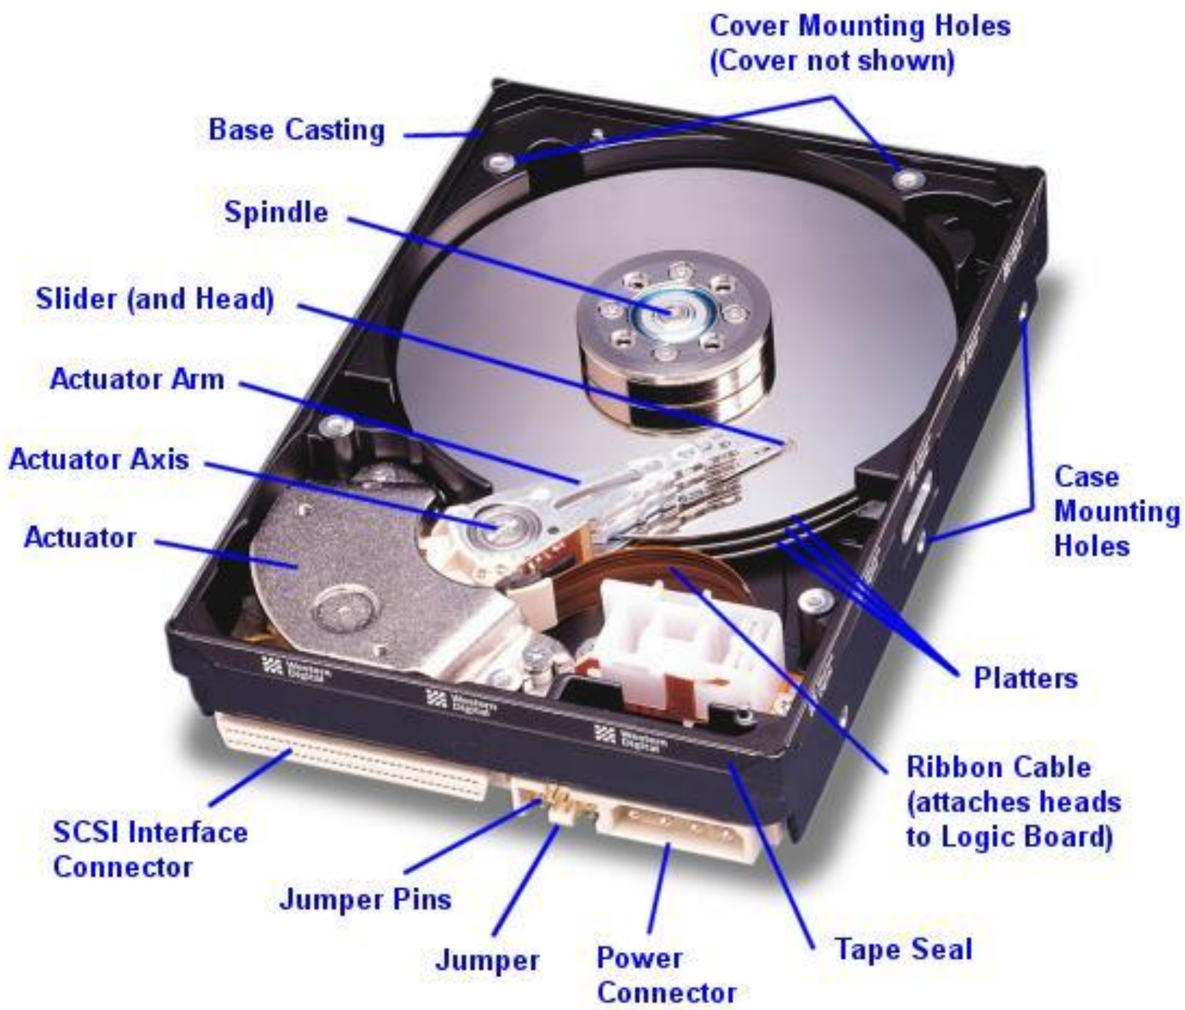

Hard Disk Operational Overview

- A hard disk uses round, flat disks called *platters*, coated on both sides with a special *media* material designed to store information in the form of magnetic patterns. The platters are mounted by cutting a hole in the center and stacking them onto a *spindle*. The platters rotate at high speed, driven by a special *spindle motor* connected to the spindle. Special electromagnetic read/write devices called *heads* are mounted onto *sliders* and used to either record information onto the disk or read information from it. The sliders are mounted onto *arms*, all of which are mechanically connected into a single assembly and positioned over the surface of the disk by a device called an *actuator*. A *logic board* controls the activity of the other components and communicates with the rest of the PC.

- Each *surface* of each platter on the disk can hold tens of billions of individual bits of data. These are organized into larger "chunks" for convenience, and to allow for easier and faster access to information. Each platter has two heads, one on the top of the platter and one on the bottom, so a hard disk with three platters (normally) has six surfaces and six total heads. Each platter has its information recorded in concentric circles called *tracks*. Each track is further broken down into smaller pieces called *sectors*, each of which holds 512 bytes of information.

- The **entire hard disk must be manufactured to** a **high degree of precision due to** the **extreme miniaturization of the components**, and the importance of the hard disk's role in the PC. The main part of the disk is isolated from outside air to ensure that no contaminants get onto the platters, which could cause damage to the read/write heads.

Hard Disk Operational Overview

- Exploded line drawing of a modern hard disk, showing the major components. Though the specifics vary greatly between different designs, the basic components you see above are typical of almost all PC hard disks.

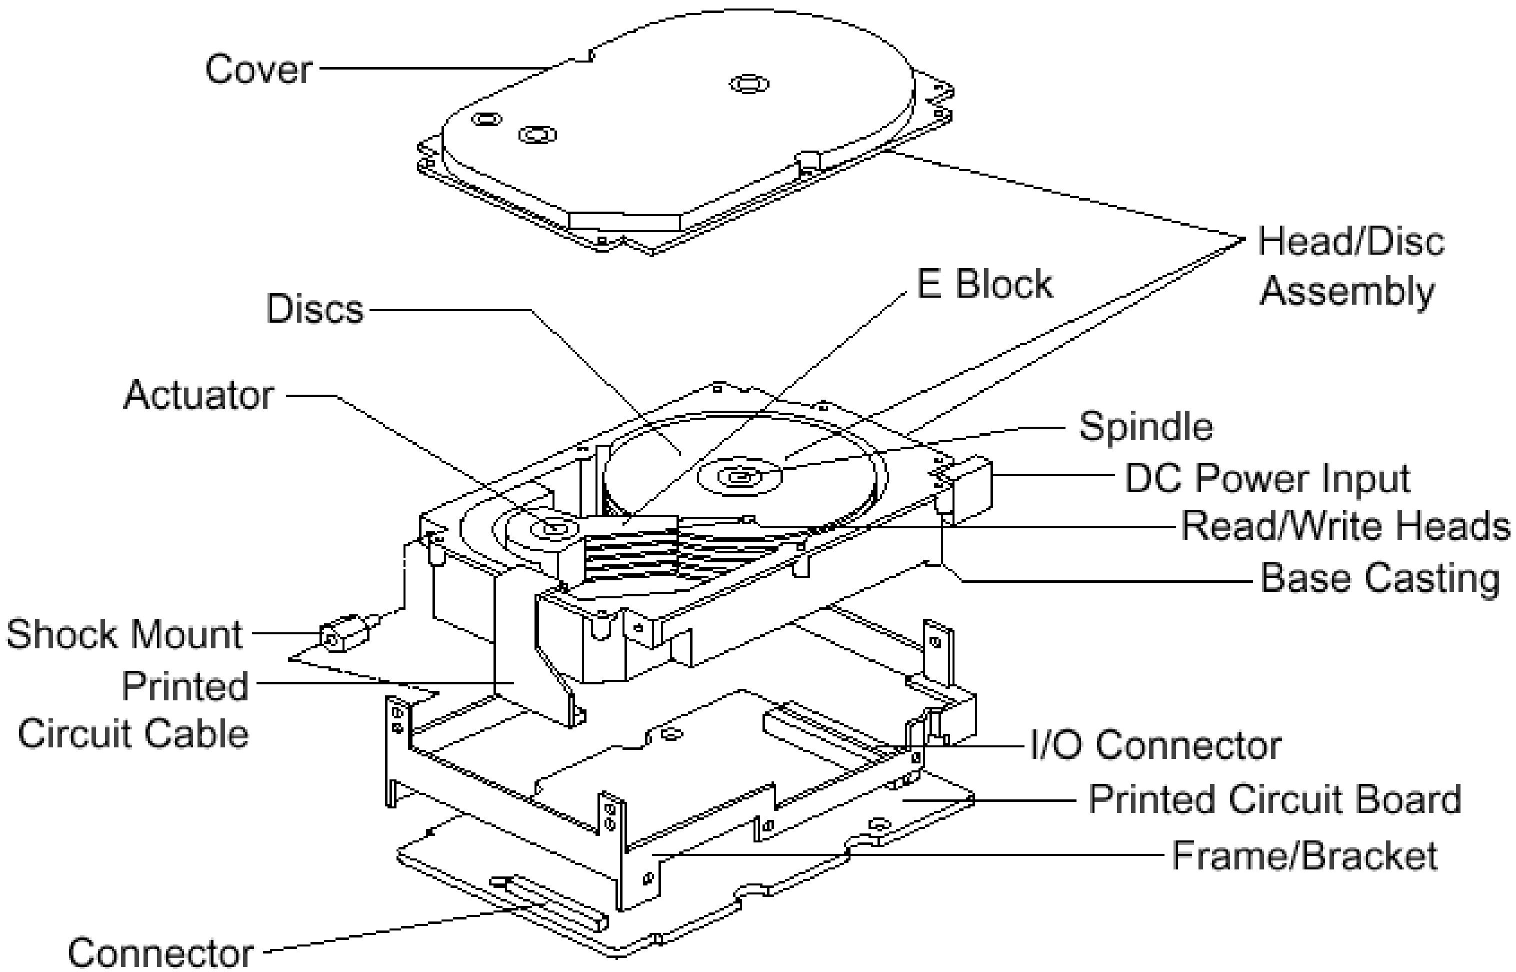

C

Hard Disk Operational Overview

- Here's an example case showing in brief what happens in the disk each time a piece of information needs to be read from it. This is a highly simplified example because it ignores factors such as disk caching, error correction, and many of the other special techniques that systems use today to increase performance and reliability. For example, sectors are not read individually on most PCs; they are grouped together into continuous chunks called clusters. A typical job, such as loading a file into a spreadsheet program, can involve thousands or even millions of individual disk accesses, and loading a 20 MB file 512 bytes at a time would be rather inefficient:

- The **first step** in accessing the disk is to figure out where on the **disk to look for the needed information**. Between them, the application, operating system, system BIOS and possibly any special driver software for the disk, do the job of determining what part of the disk to read.

- The location on the disk undergoes one or more translation steps until a final request can be made to the drive with an address expressed in terms of its geometry. The geometry of the drive is normally expressed in terms of the cylinder, head and sector that the system wants the drive to read. (A cylinder is equivalent to a track for addressing purposes). A request is sent to the drive over the disk drive interface giving it this address and asking for the sector to be read.

Hard Disk Operational Overview

- The hard disk's control program first checks to see if the information requested is already in the hard disk's own **internal buffer (or cache)**. It if is then the controller supplies the information immediately, without needing to look on the surface of the disk itself.
- In most cases the disk drive is already **spinning.** If it isn't (because power management has instructed the disk to "**spin down**" to save energy) then the drive's controller board will activate the spindle motor to "spin up" the drive to operating speed.
- The **controller board interprets** the address it received for the read, and performs any necessary additional translation steps that take into account the particular characteristics of the drive. The hard disk's logic program then looks at the final number of the cylinder requested. The cylinder number tells the disk which track to look at on the surface of the disk. The board instructs the actuator to move the read/write heads to the appropriate track.
- When the heads are in the correct position, the controller activates the head specified in the correct read location. The head begins reading the track looking for the sector that was asked for. It waits for the disk to rotate the correct sector number under itself, and then reads the contents of the sector.
- The controller board coordinates the flow of information from the hard disk into a temporary storage area (buffer). It then sends the information over the **hard disk interface**, usually to the system memory, satisfying the system's request for data.

Hard Disk Platters and Media

- **Every hard** disk contains one or more **flat disks** that are used to actually hold the data in the drive. These disks are called *platters* (sometimes also "disks" or "discs"). They are composed of two main substances: a ***substrate*** material that forms the bulk of the platter and gives it structure and rigidity, and a ***magnetic media coating*** which actually holds the magnetic impulses that represent the data.

- **Hard disks get their name from** the **rigidity** of the **platters used**, as compared to floppy disks and other media which use flexible "platters" (actually, they aren't usually even called platters when the material is flexible.)
- **platter** ['plaetÁ(r)] *n*
- 1. plitak tanjir (W also: tanjur); **on a silver* ~ na tanjiru
- 2. gramofonska ploča

- The **platters** are "where the action is"--this is where the data itself is recorded. For this **reason** the **quality** of the **platters and particularly**, their **media coating, is critical**.
- The surfaces of each platter are precision machined and treated to remove any imperfections, and the hard disk itself is assembled in a **clean room** to reduce the chances of any dirt or contamination getting onto the platters.

Platter Size

- The size of the platters in the hard disk
  - is the primary determinant of its overall physical dimensions
  - also generally called the drive's form factor
  - most drives are produced in one of the various standard hard disk form factors.
- Disks are sometimes referred to by a **size specification**;
  - for example,
  - someone will talk about having a "3.5-inch hard disk".
- When this terminology is used
  - it usually refers to the **disk's form factor**
  - and normally
  - the **form factor** is named based on the platter size.
- The platter size of the disk
  - is usually the same for all drives of a given form factor,
  - though not always,
  - especially with the newest drives, as we will see below.
- Every platter in any specific hard disk has the **same diameter**.

Platter Size

- The **first PCs** used hard disks that had a nominal size of 5.25".
- Today, by far the most common hard disk platter size in the PC world is 3.5".
- Actually, the platters of a 5.25" drive are 5.12" in diameter,
  - and those of a **3.5"** drive are **3.74"**
  - but habits are habits and the "approximate" names are what are commonly used.
- You will also notice that these numbers correspond to the common sizes for floppy disks because they were designed to be mounted into the same drive bays in the case.
- **Laptop** drives are usually smaller,
  - due to laptop manufacturers' never-ending quest for "lighter and smaller".
  - The platters on these drives are usually 2.5" in diameter or less;
  - **2.5"** is the standard form factor, but drives with 1.8"
  - and even 1.0" platters are
  - becoming more common in mobile equipment.

Platter Size

- A platter from a **5.25"** hard disk, with a platter from a **3.5"** hard disk placed on top of it for comparison. The quarter is included for scale (and strangely, fits right in the spindle hole for both platters... isn't that a strange coincidence?

Platter Size

- Traditionally, drives extend the platters to as much of the width of the physical drive package as possible, to maximize the amount of storage they can pack into the drive. However, as discussed in the section on hard disk historical trends, the trend overall is towards smaller platters. This might seem counter-intuitive; after all, larger platters mean there is more room to store data, so shouldn't it be more cost-effective for manufacturers to make platters as big as possible?

- There are several reasons **why platters are shrinking**, and they are primarily related to performance. The **areal density of disks is increasing so quickly** that the loss of capacity by going to smaller platters is viewed as not much of an issue--few people care when drives are doubling in size every year anyway!--while performance improvements continue to be at the top of nearly everyone's wish list.

- In fact, several hard disk manufacturers who were continuing to produce 5.25" drives for the "value segment" of the market as recently as 1999 have **now discontinued them**. (The very first hard disks were 24" in diameter, so you can see how far we have come in 40 or so years.)

Platter Size

- Here are the **main reasons** why companies are **going to smaller platters** even for desktop units:

- **Enhanced Rigidity**: The rigidity of a platter refers to how stiff it is. Stiff platters are more resistant to shock and vibration, and are better-suited for being mated with higher-speed spindles and other high-performance hardware. Reducing the hard disk platter's diameter by a factor of two approximately quadruples its rigidity.

- **Manufacturing Ease**: The flatness and uniformity of a platter is critical to its quality; an ideal platter is perfectly flat and consistent. Imperfect platters lead to low manufacturing yield and the potential for data loss due to the heads contacting uneven spots on the surface of a platter. Smaller platters are easier to make than larger ones.
- **Mass Reduction**: For performance reasons, hard disk spindles are increasing in speed. Smaller platters are easier to spin and require less-powerful motors. They are also faster to spin up to speed from a stopped position.

Platter Size

- **Power Conservation**: The amount of power used by PCs is becoming more and more of a concern, especially for portable computing but even on the desktop. Smaller drives generally use less power than larger ones.

- **Noise and Heat Reduction**: These benefits follow directly from the improvements enumerated above.

- **Improved Seek Performance**: Reducing the size of the platters reduces the distance that the head actuator must move the heads side-to-side to perform random seeks; this improves seek time and makes random reads and writes faster. Of course, this is done at the cost of capacity; you could theoretically achieve the same performance improvement on a larger disk by only filling the inner cylinders of each platter. In fact, some demanding customers used to partition hard disks and use only a small portion of the disk, for exactly this reason: so that seeks would be faster. Using a smaller platter size is more efficient, simpler and less wasteful than this sort of "hack".

Platter Size

- The trend towards smaller platter sizes in modern desktop and server drives began in earnest when:
  - some manufacturers "trimmed" the platters
  - in their **10,000 RPM** hard disk drives **from 3.74" to 3"**
  - (while keeping them as standard 3.5" form factor drives on the outside for compatibility.)
- Seagate's Cheetah X15 15,000 RPM drive goes even further,
  - **dropping the platter size down to 2.5"**
  - again trading performance for capacity
  - (it is "only" 18 GB, less than half the size of modern 3.5" platter-size drives.)
- This drive, despite having 2.5" platters:
  - still uses the common 3.5" form factor for external mounting
  - (to maintain compatibility with standard cases),
  - muddying the "size" waters to some extent
  - (it's a "3.5-inch drive" but it doesn't have 3.5" platters.)

Platter Size

- The smallest hard disk platter size
  - available on the market today
  - is a miniscule **1" in diameter!**
- IBM's amazing **Microdrive** has a single platter and
  - is designed to fit into
  - digital cameras
  - personal organizers
  - and other small equipment.
- The tiny size of the platters enables the Microdrive
  - to run off battery power,
  - spin down and back up again in less than a second,
  - and withstand shock that would destroy a normal hard disk.
- The downside? It's "only" 340 MB. :^)

Platter Size

- Internal view and dimensions of the amazing **IBM Microdrive.**

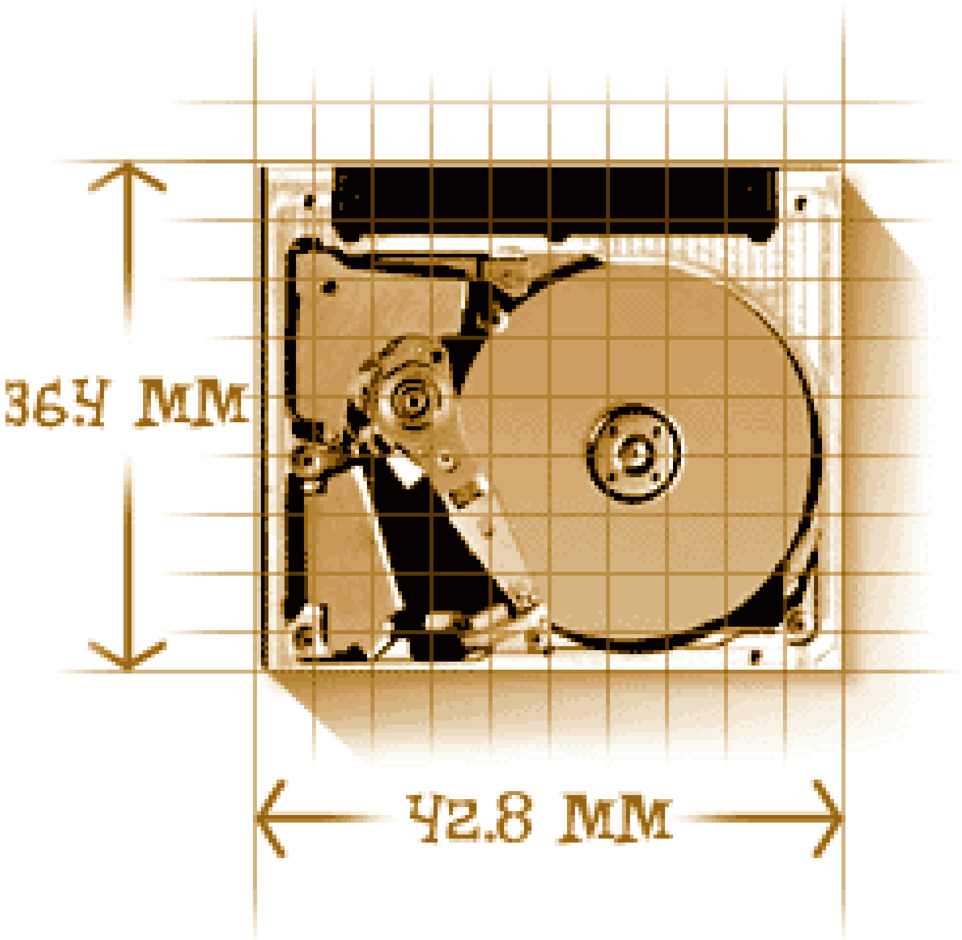

Platter Size

- Here's a **summary table showing** the most common platter sizes used in PCs, **in order of decreasing size** (which in most cases is also chronological order from their data of introduction, but not always) and also showing the most common **form factors** used by each technology:

Platter Diameter	Typical Form Factor	Application
5.12	5.25"	Oldest PCs, used in servers through the mid-1990s and some retail drives in the mid-to-late 1990s; now obsolete
3.74	3.5"	Standard platter size for the most common hard disk drives used in PCs
3.0	3.5"	High-end 10,000 RPM drives
2.5	2.5", 3.5"	Laptop drives (2.5" form factor); 15,000 RPM drives (3.5" form factor)
1.8	PC Card (PCMCIA)	PC Card (PCMCIA) drives for laptops
1.3	PC Card (PCMCIA)	Originally used on hand-held PCs (no longer made)
1.0	CompactFlash	Digital cameras, hand-held PCs and other consumer electronic devices

Number of Platters

- Hard disks **can have one platter**, or more, depending on the design.
- Standard consumer hard disks, the type probably in your PC right now, usually have between one and five platters in them.
- Some high-end drives--usually used in servers--have as many as a **dozen platters.**
- Some **very old** drives **had even more.**
- In every drive, all the platters are physically connected together on a **common central spindle**, to form a single assembly that spins as one unit, driven by the spindle motor.
- The platters are kept apart using spacer rings that fit over the spindle.
- The entire assembly is secured from the top using a cap or cover and several screws.

Number of Platters

- **Each platter** has **two surfaces** that are capable of holding data; each surface has a read/write head.
- Normally both surfaces of each platter are used, but that is not always the case.
- Some older drives that use dedicated servo positioning **reserve** one surface for holding **servo** information.
- Newer drives don't need to spend a surface on servo information, but sometimes leave a surface unused for marketing reasons--to create a drive of a particular capacity in a family of drives.
- With modern drives packing huge amounts of data on a single platter, using only one surface of a platter allows for increased "granularity".
- For example, IBM's Deskstar 40GV family sports an impressive 20 GB per platter data capacity. Since IBM wanted to make a 30 version of this drive, they used three surfaces (on two platters) for that drive.

Number of Platters

- Here's a good illustration of how **Western Digital** created five different capacities using three platters in their Caviar line of hard disk drives:

Model Number	Nominal Size (GB)	Data Sectors Per Drive	Platters	Surfaces
WD64AA	6.4	12,594,960	1	2
WD102AA	10.2	20,044,080	2	3
WD136AA	13.6	26,564,832	2	4
WD172AA	17.2	33,687,360	3	5
WD205AA	20.5	40,079,088	3	6

Number of Platters

- **Note:** In theory, using only one surface means manufacturing costs can be saved by making use of platters that have unacceptable defects on one surface, but I don't know if this optimizing is done in practice...

- From an **engineering standpoint** there are several factors that are related to the number of platters used in the disk.
- Drives with many platters are more difficult to engineer due to the increased mass of the spindle unit, the need to perfectly align all the drives, and the greater difficulty in keeping noise and vibration under control.
- **More platters also means more mass**, and therefore slower response to commands to start or stop the drive; this can be compensated for with a stronger spindle motor, but that leads to other tradeoffs. In fact, the trend recently has been towards drives with fewer head arms and platters, not more.
- **Areal density continues to increase**, allowing the creation of large drives without using a lot of platters. This enables manufacturers to reduce platter count to improve seek time without creating drives too small for the marketplace. See here for more on this trend.

Number of Platters

- This **Barracuda hard disk has 10 platters**

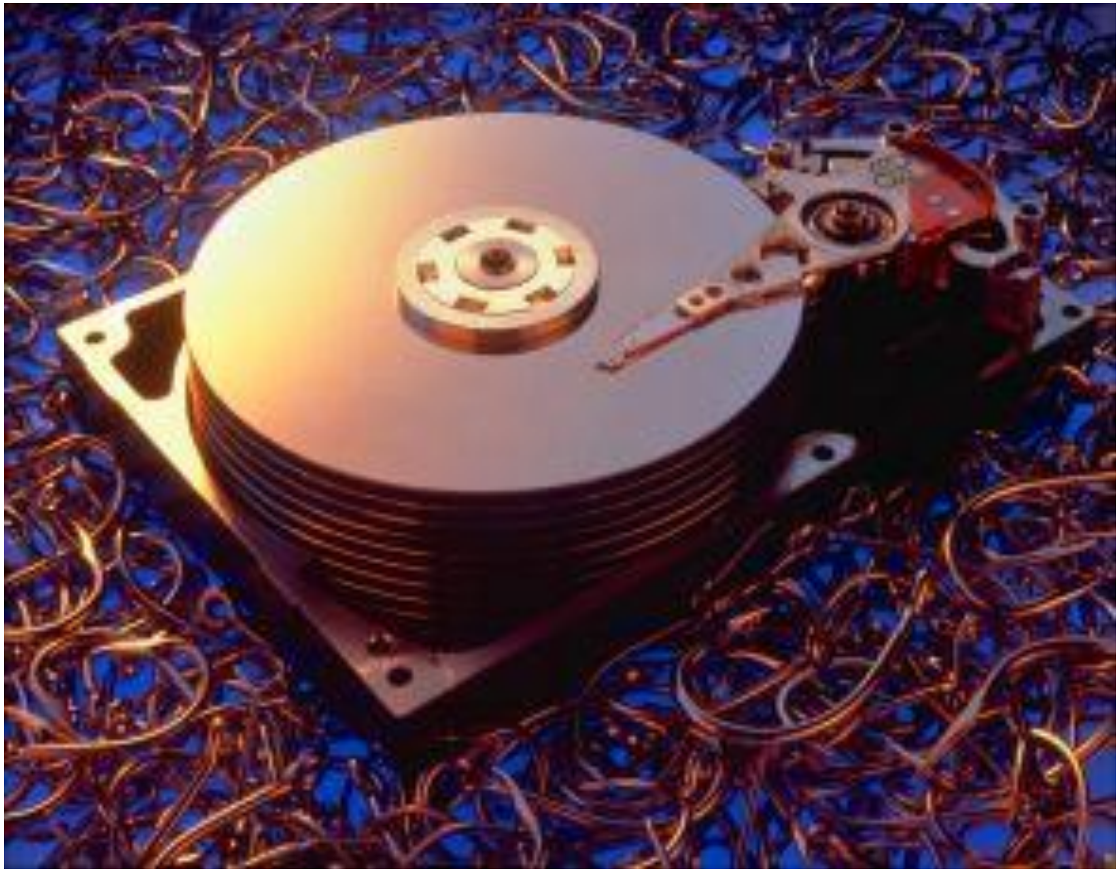

Number of Platters

- The form factor of the hard disk also has a great influence on the number of platters in a drive.
- Even if hard disk engineers wanted
  - to put lots of platters in a particular model,
  - the standard **PC "slimline"** hard disk form factor
  - is limited to **1 inch in height**,
  - which limits the number of platters that can be put in a single unit.
- Larger 1.6-inch "half height" drives
  - are often found in servers
  - and usually have many more platters than desktop PC drives.
- Of course, engineers are constantly **working to reduce**
  - the amount of clearance required **between platters,**
  - so they can increase the number of platters
  - in drives of a given height.

Platter Substrate Materials

- The magnetic patterns that comprise your data are recorded in a very thin media layer on the surfaces of the hard disk's platters; the bulk of the material of the platter is called the **substrate** and does nothing but support the media layer.
- To be suitable, a **substrate** material **must be**:
  - rigid
  - easy to work with, lightweight, stable, magnetically inert,
  - inexpensive and
  - readily available.

- The most commonly used material for making platters
- has traditionally been an **aluminum alloy**, (legura)
- which meets all of these criteria.

Platter Substrate Materials

- Due to the way the platters spin with the read/write heads floating just above them,
- the **platters must** be **extremely smooth (gladak)** and **flat**.
- With older, slower spindle drives and relatively high fly heights, the uniformity of the platter surface was less of an issue.
- Now, as technology advances,
  - the **gap** between the heads and the platter is **decreasing**,
  - and the **speed** that the **platters spin at is increasing**,
  - creating more demands on the platter material itself.
- Uneven platter surfaces on hard disks running at faster speeds with heads closer to the surface are more apt to lead to **head crashes**.
- For this reason many drive makers began several years ago
- to **look** at **alternatives to aluminum**, such as:
  - glass
  - glass composites
  - magnesium alloys

Platter Substrate Materials

- Hard disk platters are very smooth, right? Well, not to a scanning **electron microscope**!
  The image on the **left** is of the surface of an **aluminum alloy platter**; the one on the **right is** a **glass platter**. The images speak for themselves. The scale is in microns..

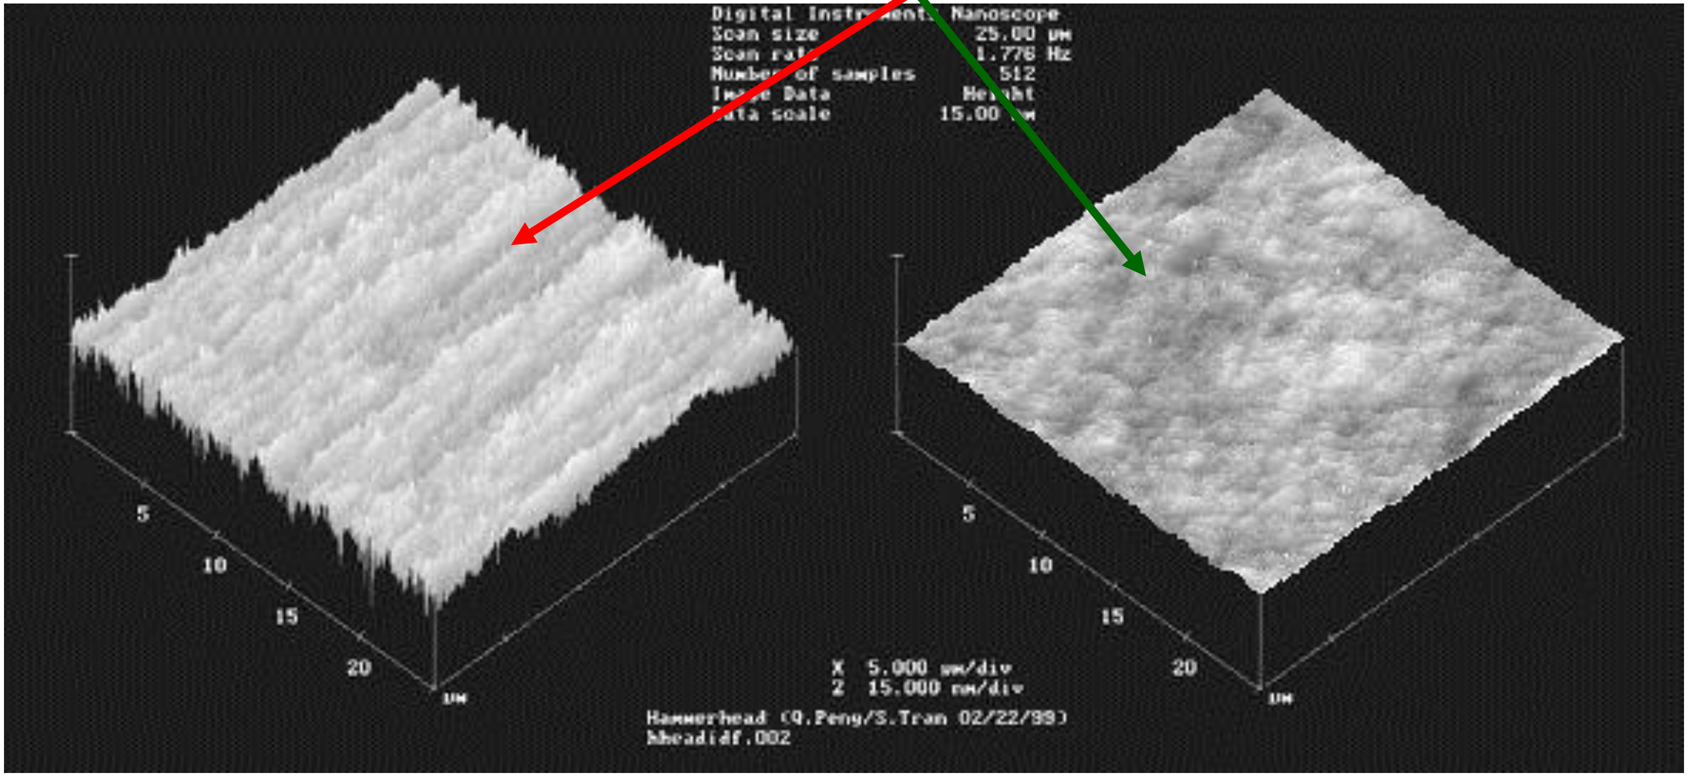

Platter Substrate Materials

- It now is looking increasingly likely that glass and composites made with **glass** will be the next standard for the platter substrate. IBM has been shipping drives with glass platters for several years and in 2000 is introducing them into the IDE/ATA consumer drive market.
- **Compared to aluminum platters, glass platters** have several **advantages**:
- **1. Better Quality**: The first and most important reason for going to glass is probably that glass platters can be made much smoother and flatter than aluminum, improving the reliability of the hard disk and making low flying heights and faster spindle speeds more feasible.
- **2. Improved Rigidity**: Another important consideration is that glass is more rigid than aluminum for the same weight of material. Improved rigidity, one of the reasons why platter sizes are also shrinking in size, is important for reducing noise and vibration with drives that spin at high speed.
- **3. Thinner Platters**: The enhanced rigidity of glass also allows platters to be made thinner than with aluminum, allowing more platters to be packed into the same drive dimensions. Thinner platters also weigh less, reducing spindle motor requirements and reducing start time when the drive is at rest.

Platter Substrate Materials

- **4. Thermal Stability**: When heated, glass expands much less than does aluminum. With some hard disk platters now containing 35,000 tracks per inch or more, even a small amount of expansion can causes these tracks to "move around". The drive's servo mechanism compensates for expansion and contraction, but it is still preferable to use materials that move as little as possible because this reduces the amount of adjusting the hard drive has to do, improving performance.

- One obvious **disadvantage** of glass compared to aluminum is **fragility**, particularly when made very thin. For this reason some companies are experimenting with **glass/ceramic composites**. One of these is a Dow Corning product called **MemCor**, which is a glass made with ceramic inserts to reduce the likelihood of cracking.
- Sometimes these composites are just called "glass", much the way **aluminum alloy platters**, which usually contain other metals, are just called "aluminum".

Magnetic Media – oxide media

- The substrate material of which the platters are made forms the base upon which the actual recording media is deposited. The **media layer** is a very thin coating of magnetic material which is where the actual data is stored; it is typically only a few millionths of an inch in thickness.
- iron(II) oxide, wüstite (FeO)
- iron(II,III) oxide, magnetite (Fe_3O_4)
- iron(III) oxide (Fe_2O_3)
- **Older hard disks used oxide media**. "Oxide" really means iron **oxide--rust**. Of course no high-tech company wants to say they use rust in their products, so they instead say something like **"high-performance oxide media layer**". :^) But in fact that's basically what oxide media is, particles of rust attached to the surface of the platter substrate using a binding agent. You can actually see this if you look at the surface of an older hard disk platter: it has the characteristic **light brown color**. This type of media is similar to what is used in audio cassette tape (which has a similar color.)
- **Oxide media is inexpensive to use**, but also has **several important shortcomings**.
- 1. The **first** is that it is a **soft material**, and **easily damaged from contact by a read/write head**.
- 2. The **second** is that it is only useful for **relatively low-density storage**. It worked fine for older hard disks with relatively low data density, but as manufacturers sought to pack more and more data into the same space, oxide was not up to the task: the **oxide particles became too large** for the small magnetic fields of newer designs.

Magnetic Media – thin film media

- Today's hard disks use **thin film media**.
- As the name suggests,
- thin film media consists of a very thin layer of magnetic material
- applied to the surface of the platters.
- (While oxide media certainly isn't thick by any reasonable use of the word, it was much thicker than this new media material; hence the name "thin film".)

- **Special manufacturing techniques** are employed to deposit the media material on the platters.

- 1. One method is **electroplating (oblagati)**, which deposits the material on the platters using a process similar to that used in electroplating jewelry.

- 2. Another is **sputtering (**puckati, puckettati), which uses a vapor-deposition process borrowed from the manufacture of semiconductors to deposit an extremely thin layer of magnetic material on the surface. Sputtered platters have the advantage of a more uniform and flat surface than plating. Due to the increased need for high quality on newer drives, sputtering is the primary method used on new disk drives, despite its higher cost.

Magnetic Media – thin film media

- **Compared to oxide media**:
- 1. thin film media is much **more uniform** and **smooth (gladak)**.

- 2. It also has greatly superior magnetic properties,
- allowing it to hold **much more data**
- **in** the **same amount of space**.

- 3. Finally, it's a much **harder** and **more durable material** than oxide, and therefore much less susceptible to damage.

oxide media v thin film media

- A thin film 5.25" platter (above) next to an oxide 5.25" platter (below). **Thin film platters are actually reflective**;

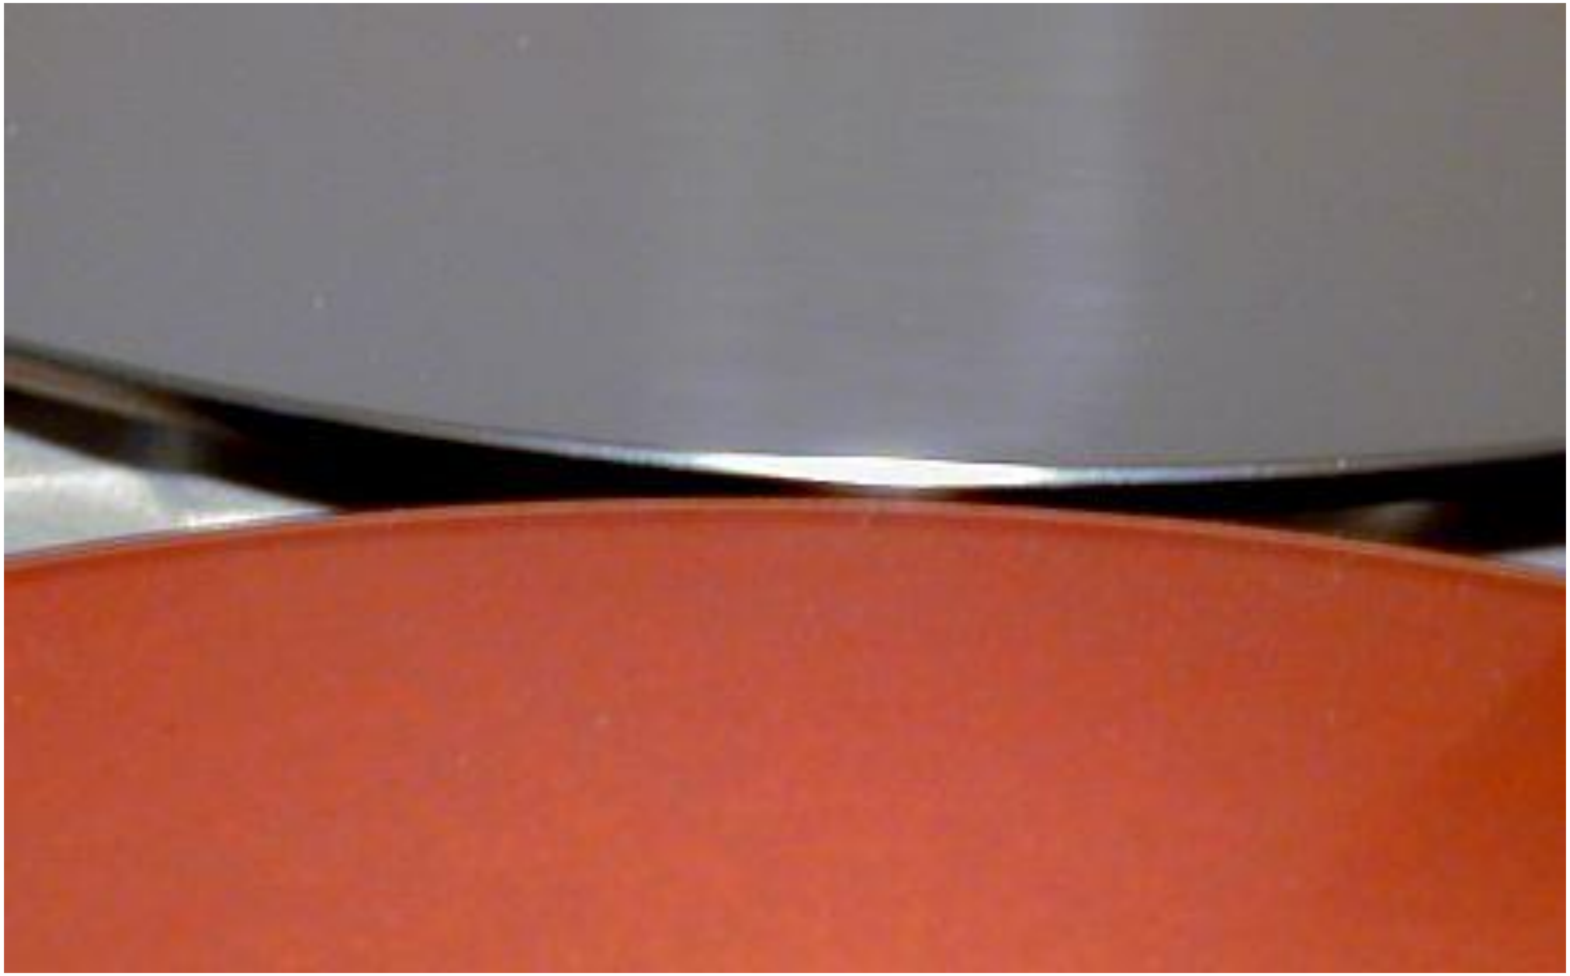

nanocrystal superlattice

- After applying the magnetic media, the surface of each platter is usually covered with a thin, protective, **layer made of carbon**. On top of this is added a **super-thin lubricating layer**. These material are used to protect the disk from damage caused by accidental contact from the heads or other foreign matter that might get into the drive.

- IBM's researchers are now working on a fascinating, experimental new substance that may replace thin film media in the years ahead.
- Rather than sputtering a metallic film onto the surface, a chemical solution containing **organic molecules** and **particles of iron** and **platinum is applied to the platters**.
- The solution is spread out and heated. When this is done, the iron and platinum particles arrange themselves naturally into a grid of crystals, with each crystal able to hold a magnetic charge. IBM is calling this structure a "**nanocrystal superlattice**". (super resetka)
- This technology has the potential to increase the areal density capability of the recording media of hard disks by as much as **10 or even 100 times**! Of course it is years away, and will need to be matched by advances in other areas of the hard disk (particularly read/write head capabilities) but it is still pretty amazing and shows that magnetic storage still has a long way to go before it runs out of room for improvement.

Tracks and Sectors

- Platters are organized into specific structures to enable the organized storage and retrieval of data. Each platter is broken into tracks--tens of thousands of them--which are tightly-packed concentric circles. These are similar in structure to the annual rings of a tree (but not similar to the grooves in a vinyl record album, which form a connected spiral and not concentric rings).

- A track holds too much information to be suitable as the smallest unit of storage on a disk, so each one is further broken down into sectors.
- A sector is normally the smallest individually-addressable unit of information stored on a hard disk, and normally holds 512 bytes of information.
- The **first PC hard disks** typically held **17 sectors per track**.
- Today's hard disks can **have thousands of sectors** in a single track, and make use of **zoned recording** to allow more sectors on the larger outer tracks of the disk.

Tracks and Sectors

- A platter from a 5.25" hard disk, with **20 concentric tracks** drawn over the surface.
- This is far lower than the density of even the oldest hard disks; even if visible, the tracks on a modern hard disk would require high magnification to resolve.
- **Each track is divided into 16 imaginary sectors**.
- Older hard disks had the same number of sectors per track
- but new ones use zoned recording with a different number of sectors per track in different zones of tracks.

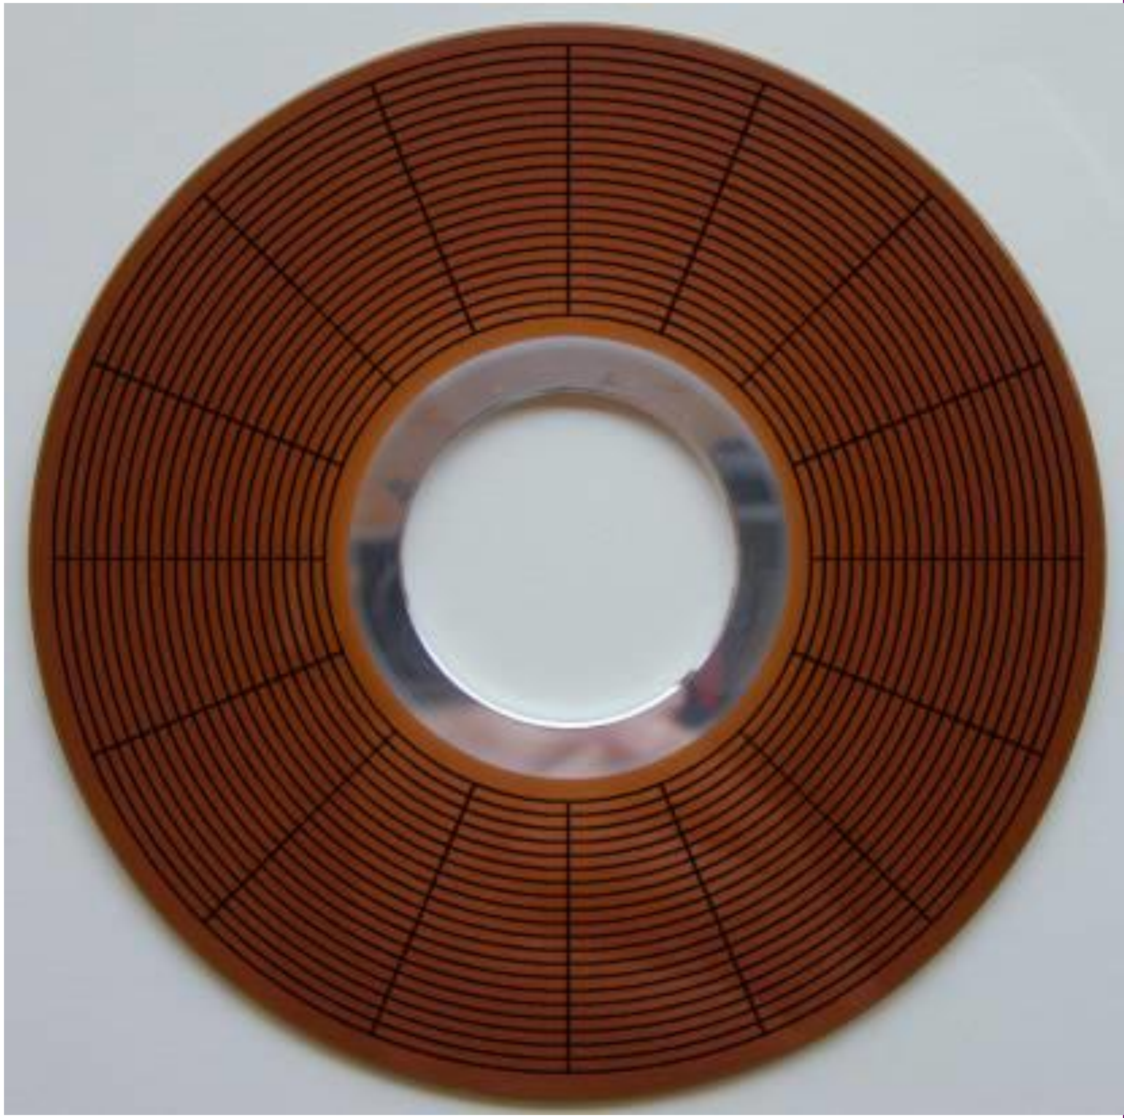

Areal Density 1 inch = 2.54 centimeters

- **1. Areal density,** also sometimes called **bit density**, refers to the amount of data that can be stored in a given amount of hard disk platter "real estate". Since disk platters surfaces are of course two-dimensional, **areal density is a measure of the number of bits that can be stored in a unit of area**. It is usually expressed in **bits per square inch (BPSI).**
- Being a two-dimensional measure, areal density is computed as the product of two other one-dimensional density measures:
- **2. Track Density:** This is a measure of how tightly the concentric tracks on the disk are packed: how many tracks can be placed down in inch of radius on the platters. For example, if we have a platter that is 3.74" in diameter, that's about 1.87 inches. Of course the inner portion of the platter is where the spindle is, and the very outside of the platter can't be used either. Let's say about 1.2 inches of length along the radius is usable for storage. If in that amount of space the hard disk has 22,000 tracks, the track density of the drive would be approximately 18,333 **tracks per inch** (**TPI**).
- **3. Linear or Recording Density:** This is a measure of **how tightly the bits are packed within** a **length of track**. If in a given inch of a track we can record 200,000 bits of information, then the linear density for that track is 200,000 **bits per inch per track (BPI).** Every track on the surface of a platter is a different length (because they are concentric circles), and not every track is written with the same density. Manufacturers usually quote the *maximum* linear density used on each drive.

Areal Density

- Taking the product of these two values yields the drive's areal density, measured in bits per square inch. If the maximum linear density of the drive above is 300,000 bits per inch of track, its maximum areal density would be 5,500,000,000 bits per square inch, or in more convenient notation, 5.5 Gbits/in2. The newest drives have areal densities exceeding **10 Gbits/in2**, and in the lab **IBM in 1999 reached 35.3 Gbits/in2-**-524,000 BPI linear density, and 67,300 TPI track density! In contrast, the **first PC hard disk had an areal density of about 0.004 Gbits/in2!**

- **Note:** Sometimes you will see areal density expressed in a different way: **gigabytes per platter (GB/platter).** This unit is often used when comparing drives, and is really a different way of saying the same thing--*as long as you are always clear about what the platter size is*. It's easier conceptually for many people to contrast two units by saying, for example: "Drive A has 10 GB/platter and drive B has 6 GB/platter, so A has higher density". Also, it's generally easier to compute when the true areal density numbers are not easy to find. As long as they both have the same platter size, the comparison is valid.
- Otherwise you are comparing apples and oranges.

Areal Density

- There are 2 ways to increase areal density:
  - **1. increase** the **linear density** by packing the bits on each track closer together so that each track holds more data;
  - or
  - **2. increase** the **track density** so that each platter holds more tracks.
- Typically new generation drives improve both measures.
- It's important to realize that increasing areal density leads to drives that are not just bigger, but also faster, all else being equal.
- The reason is that the **areal density** of the disk **impacts** both of the key hard disk performance factors: both **positioning speed** and **data transfer rate**.

Areal Density

- This illustration shows how areal density works.
- First, divide the circle in your mind into a left and right half, and then a top and bottom half.
- The **left half** shows **low track density** and the right half high track density.
- The **upper half shows low linear density**, and the bottom half high linear density.
- Combining them, the upper left quadrant has the lowest areal density; the upper right and lower left have greater density, and the bottom right quadrant of course has the highest density.

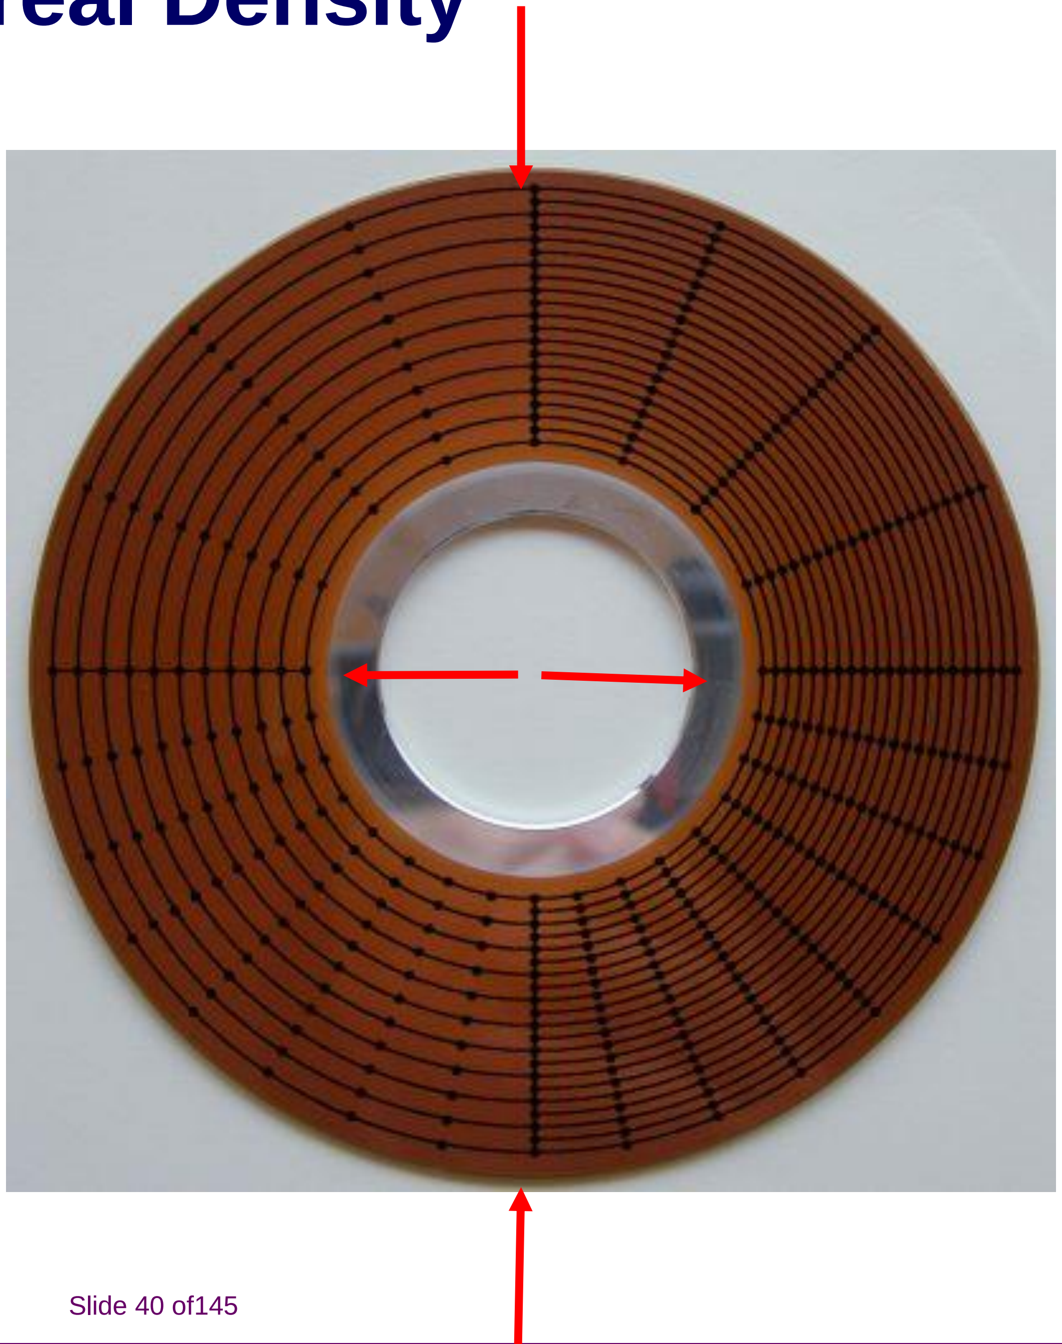

Areal Density

- **Increasing the areal density** of disks is a **difficult task** that requires many technological advances and changes to various components of the hard disk.

- As the data is packed closer and closer together, **problems result** with **interference between bits**. This is often dealt with by reducing the strength of the magnetic signals stored on the disk, but then this creates other problems such as ensuring that the signals are stable on the disk and that the **read/write heads are sensitive and close enough to the surface to pick them up**. In some cases the heads must be made to fly closer to the disk, which causes other engineering challenges, such as ensuring that the disks are flat enough to reduce the chance of a head crash.

- Changes to the media layer on the platters, actuators, control electronics and other components are made to continually improve areal density. This is especially true of the read/write heads. Every few years a read/write head technology breakthrough enables a significant jump in density, which is why hard disks have been doubling in size so frequently--in some cases it takes less than a year for the leading drives on the market to double in size.

Hard Disk Read/Write Heads

- The read/write heads of the hard disk are the interface between the magnetic physical media on which the data is stored and the electronic components that make up the rest of the hard disk (and the PC). The heads do the **work** of **converting bits** to **magnetic pulses** and **storing them** on the platters, and then reversing the process when the data needs to be read back.
- **Read/write heads are** an **extremely critical component** in determining the **overall performance of the hard disk**, since they play such an important role in the storage and retrieval of data.
- They are usually one of the **more expensive parts** of the hard disk, and to enable areal densities and disk spin speeds to increase, they have had to evolve from rather humble, clumsy beginnings to being extremely advanced and complicated technology.
- New head technologies are often the triggering point to increasing the speed and size of modern hard disks.

Hard Disk Read/Write Head Operation

- In many ways, the read/write heads
- are the most sophisticated part of the hard disk,
- which is itself a technological marvel.

- They don't get very much attention, perhaps in part because most people never see them.

- This section takes a look at how the heads work and discusses some of their key operating features and issues.

Function of the Read/Write Heads

- In concept, hard disk heads are relatively simple.
- They are **energy converters**:
  - they **transform**
  - **electrical signals to magnetic signals**,
  - and
  - **magnetic signals** back to **electrical** ones again.

- The heads on your VCR or home stereo tape deck perform a similar function, although using very different technology.

- The read/write heads are in essence **tiny electromagnets**
  - that perform this conversion
  - from electrical information to magnetic
  - and back again

- **Each bit of data** to be stored is recorded onto the hard disk
  - using a special encoding method
  - that translates zeros and ones into patterns of **magnetic flux reversals**.

Function of the Read/Write Heads

- **Older, conventional** (ferrite, metal-in-gap and thin film) hard disk heads work by making use of the **two main principles** of **electromagnetic force**.
- The **first** is that applying an electrical current through a coil produces a magnetic field; this is used when writing to the disk. The direction of the magnetic field produced depends on the direction that the current is flowing through the coil.
- The **second** is the opposite, that applying a magnetic field to a coil will cause an electrical current to flow; this is used when reading back the previously written information. (You can see a photograph showing this design on the page on ferrite heads.)
- Again here, the direction that the current flows depends on the direction of the magnetic field applied to the coil.
- **Newer (MR** and **GMR)** heads don't use the induced current in the coil to read back the information; they function instead by using the principle of **magnetoresistance**, where **certain materials change their resistance** when subjected to **different magnetic fields**.

Function of the Read/Write Heads

- The heads are usually called "read/write heads", and older ones did both writing and reading using the **same element**.

- Newer MR and GMR heads however, are in fact composites that include a **different element for writing and reading**.
- This design is **more complicated to manufacture**,
  - but is required
  - because the **magnetoresistance** effect
  - used in these heads **only functions** in the **read mode.**

- Having separate units for writing and reading
  - also allows each to be tuned to the particular function it does,
  - while a single head must be designed as a compromise between
  - fine-tuning for the write function or the read function.

- These dual heads are sometimes called "**merged heads**".

Function of the Read/Write Heads

- This graph shows how the bit size of hard disks is shrinking over time: dramatically.

- The width and length of each bit are shown for hard disks using varies areal densities.
- Current high-end hard disks have exceeded 10 Gbit/in2 in areal density, but it has been only a few years since 1 Gbit/in2 was state of the art.

- As the bit size drops and the bits are packed closer together, magnetic fields become weaker and more sensitive head electronics are required to properly detect and interpret the data signals.

Function of the Read/Write Heads

Crveni pravougaonik=1bit

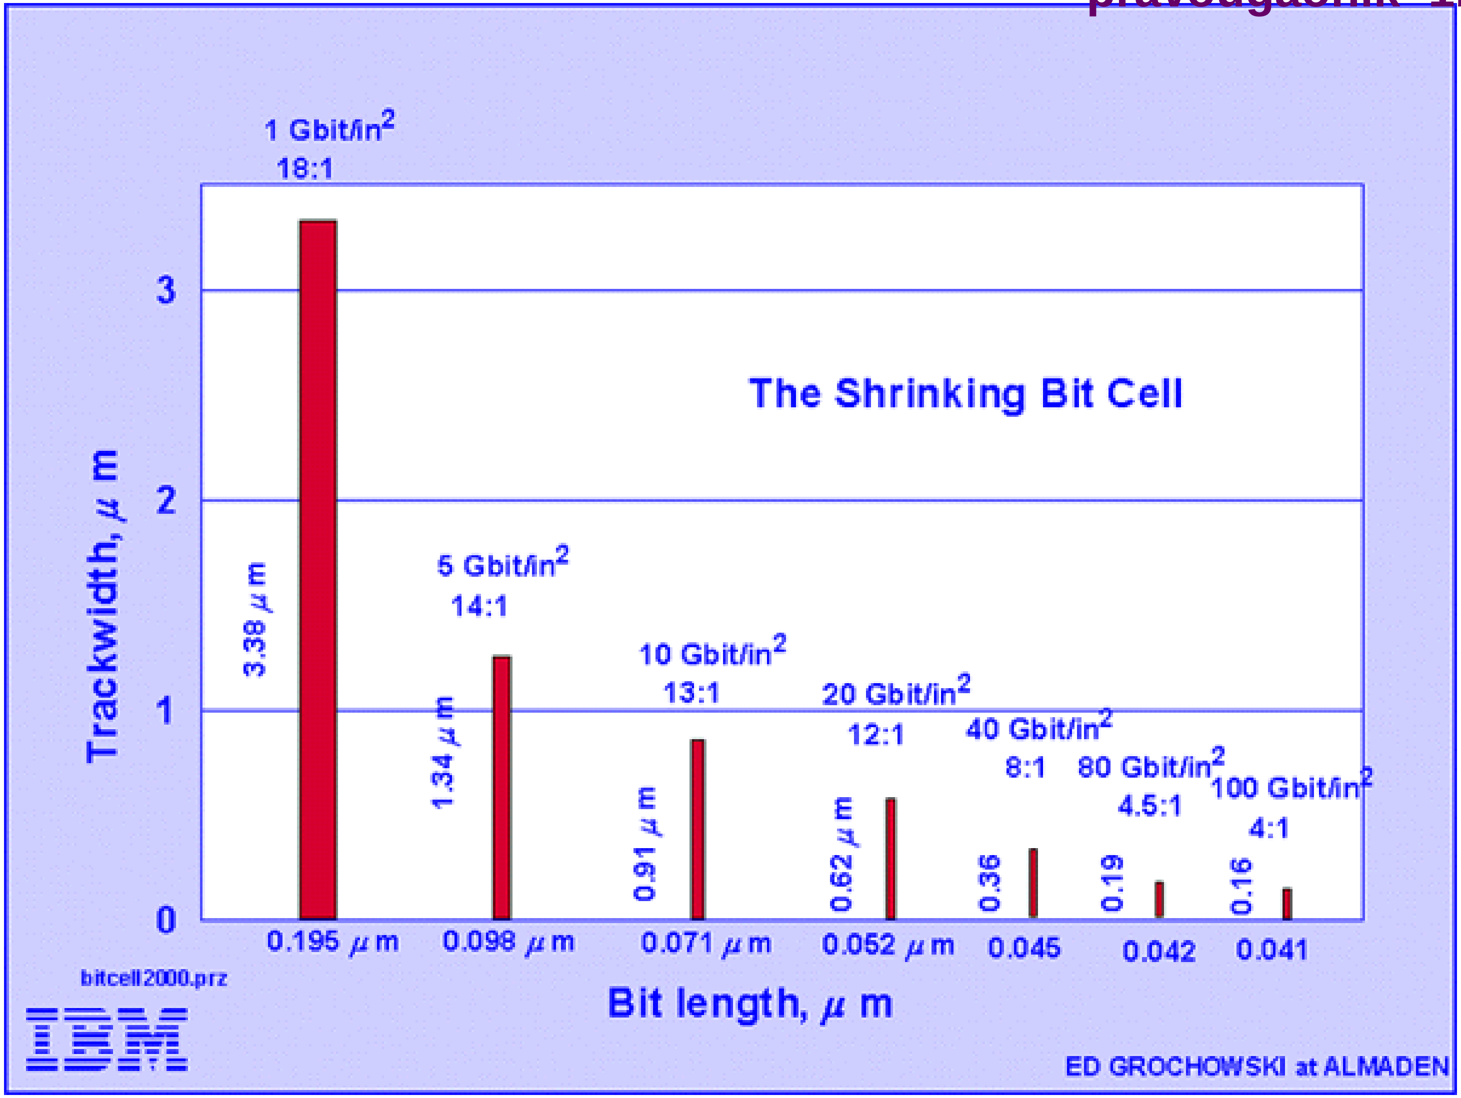

Function of the Read/Write Heads

- Because of the tight packing of data bits on the hard disk, it is important to make sure that the **magnetic fields don't interfere with one another**. To ensure that this does not happen, the stored fields are very small, and very weak. Increasing the density of the disk means that the fields must be made still weaker, which means the read/write heads must be faster and more sensitive so they can read the weaker fields and accurately figure out which bits are ones and which bits are zeroes.

- This is the reason why MR and GMR heads have taken over the market: **they are more sensitive and can be made very small** so as not read adjacent tracks on the disk. Special amplification circuits are used to convert the weak electrical pulses from the head into proper digital signals that represent the real data read from the hard disk.

- **Error detection and correction circuitry** must also be used to compensate for the increased likelihood of errors as the signals get weaker and weaker on the hard disk. In addition, some newer heads employ magnetic "shields" on either side of the read head to ensure that the head is not affected by any stray magnetic energy.

Number of Read/Write Heads

- Each hard **disk platter has two surfaces**, one on each side of the platter. More often than not, both surfaces of the platter are used for data on modern drives, but as described in the section discussing the number of platters in a drive, this is not always the case. There is normally one head for each surface used on the drive. In the case of a drive using dedicated servo, there will be one more head than there are data surfaces.

- Since most hard disks have one to four platters, most hard disks have between **two and eight heads**. **Some larger drives can have 20 heads or more.**

- Only one head can read from or write to the hard disk at a given time. Special circuitry is used to control which head is active at any given time.

- Warning: Most IDE/ATA hard disks come with "setup parameters" intended for use when configuring the disk in the BIOS setup program. Don't be fooled by these numbers, which sometimes bear confusing names like "physical geometry" even though they are not actually the physical geometry at all. For today's drives these numbers have nothing to do with what is inside the hard disk itself. Most new IDE/ATA disks these days are set up as having 16 heads, even though all have far fewer, and new drives over 8.4 GB are always specified as having 16 heads per the ATA-4 specification. BIOS translation issues can make this even more confusing. See here for more details.

Floating Height / Flying Height / Head Gap

- One distinguishing characteristic of hard disk technology
  - that makes it different from how floppy disks,
  - VCRs and tape decks work,
  - is that the **read/write heads do not make contact with the media**.
- The reason for this is that due to
  - the high speed that the hard disk spins,
  - and the need for the heads to frequently scan
  - from side to side to different tracks,
  - allowing the heads to contact the disk
  - would result in unacceptable wear to both
  - the delicate heads and the media.
- In fact the earliest hard disks
  - did have their heads in contact with the media,
  - and this design was changed due to the wear that contact caused.

Floating Height / Flying Height / Head Gap

- Modern drive heads float over the surface of the disk and
  - do all of their work without ever physically touching
  - the platters they are magnetizing.
- The **amount of space between** the **heads** and the platters is called
  - the floating height or
  - flying height.
- It is also sometimes called the head gap, and some hard disk manufacturers refer to the heads as riding on an "air bearing".
- The read/write head assemblies are spring-loaded—
  - using the spring steel of the head arms—
  - which causes the sliders to press against the platters
  - when the disk is stationary.
- (This is done to ensure that the heads **don't drift away from the platters**; maintaining an exact floating height is essential for correct operation.) When the disk spins up to operating speed, the high speed causes air to flow under the sliders and lift them off the surface of the disk--the same principle of lift that operates on **aircraft wings** and **enables them to fly**.

Floating Height / Flying Height / Head Gap

- A pair of mated head sliders with their platter removed.
- You can see that the tension of the head arms has caused them to press against each other.

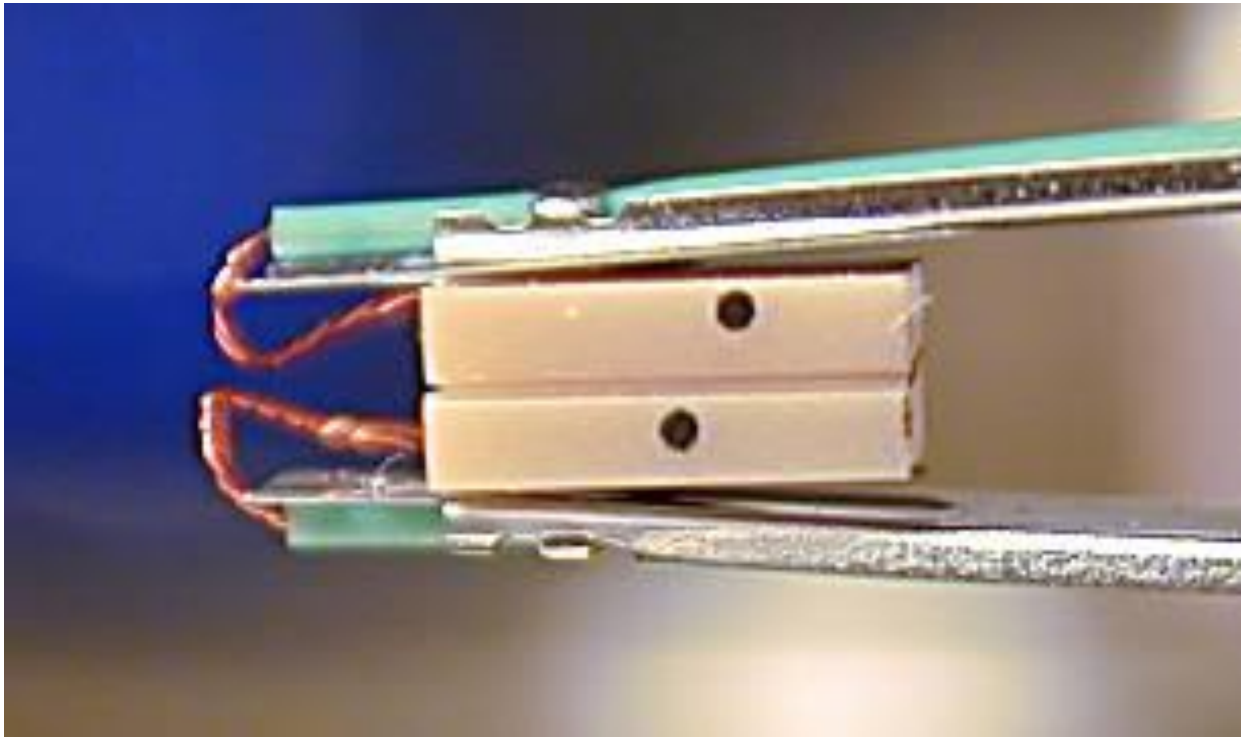

Floating Height / Flying Height / Head Gap

- Due to the very small distance from the heads to the platters--normally measured in millionths of an inch--the hard disk is assembled in a *clean room* containing air specially filtered to remove all but the tiniest particles.

- Air however is *required* for the heads to function. Whenever someone suggests that the inside of a hard disk is maintained under a vacuum--and it always happens--just ask them how exactly the heads can float on the surface of the disk if there is **no air**. :^) You will also hear people say that the drive's interior is "sealed" (including, I must admit, myself at one point).
- This is also generally untrue: while the disk's internal environment is separate from the outside air to keep it clean, air exchange is permitted between the outside and inside of the drive to allow the drive to adjust to changes in air pressure. A special **"breather" filter** is installed to prevent foreign matter from contaminating the drive.

- **Note:** If a drive is used at too high an altitude, the air will become too thin to support the heads at their proper operating height and **failure will result**; special industrial drives that truly *are* sealed from the outside are made for these special applications.

Floating Height / Flying Height / Head Gap

- The **distance from** the **platters** to the heads is a specific design parameter that is tightly controlled by the engineers that create the drive.
- By adjusting the **strength** of the springs to match the other drive parameters (such as the speed the disks are spinning and the size and shape of the heads) the float height can be precisely maintained.
- If the **height is too great**, the heads **can't properly read and write the platter**. If it is **too small**, there is increased chance of a **head crash** (ouch.)
- As mentioned in the section on operation, increasing areal density means that weaker magnetic fields must be used in storing data on the disks. When this is done the heads must be allowed to ride closer and closer to the platter surface to pick up the weaker signals, which requires other quality improvements to the drive to make sure that there is no chance of a head crash (ouch. :^))

- Tip: Some modern drives **include sensors that monitor** the flying height of the heads and signal a warning if the parameter falls out of the acceptable range.

Floating Height / Flying Height / Head Gap

- It's actually quite amazing how close to the surface of the disks the **heads fly without touching**.
- To put it into perspective, a modern hard disk has a floating height of an amazing **0.5 microinches**.
- A human hair has a thickness of over 2,000 microinches! You can see why keeping dirt out of the hard disk is so important!
- In fact, the floating height of a hard disk is smaller than the circuit size of a microprocessor. What's even more amazing is how much abuse these hard disks can take when they are placed in laptop PCs, for example, given these facts, and how many people take this technology for granted every day...

dust

- This illustration gives you some idea of just how small the flying height of a modern hard disk is (and today's hard disks have flying heights significantly lower than 3-7 millionths of an inch!

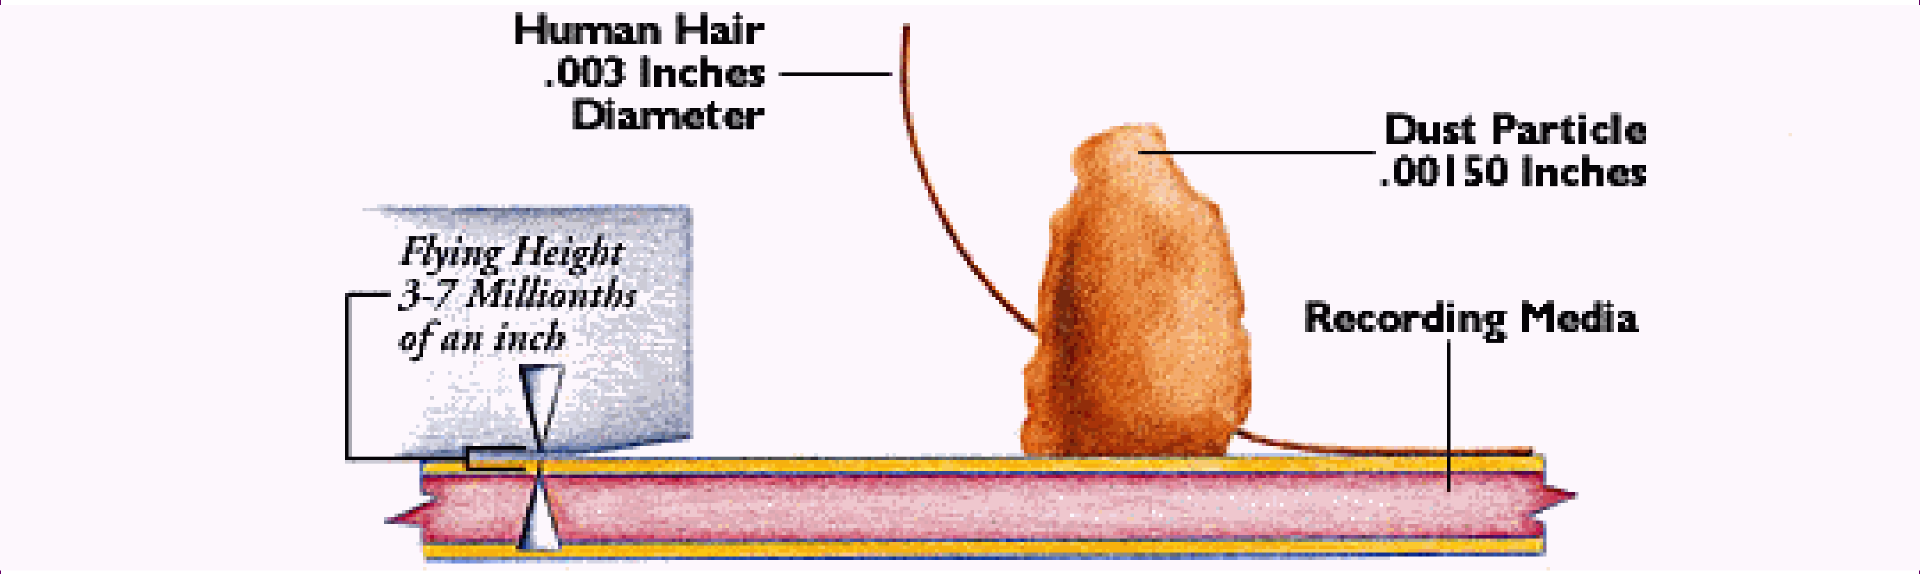

Floating Height / Flying Height / Head Gap

- When the areal density of a drive is increased to improve capacity and performance, the **magnetic fields are made smaller and weaker**.
- **To compensate**, either the heads **must be made more sensitive**, or the floating height must be decreased.
- Each time the floating height is decreased, the mechanical aspects of the disk must be adjusted to make sure that the platters are flatter, the alignment of the platter assembly and the read/write heads is perfect, and there is no dust or dirt on the surface of the platters.
- **Vibration** and shock also become more of a concern, and must be compensated for. This is one reason why manufacturers are turning to smaller platters, as well as the use of glass platter substrates.
- **Newer heads** such as **GMR** are preferred because they allow a higher flying height than older, less sensitive heads, all else being equal.

Floating Height / Flying Height / Head Gap

- As the flying height of drives continues to decrease,
  - hard disk engineers are recognizing
  - that we may soon reach the point where
  - it cannot be made any smaller
  - without touching the surfaces of the platters.
- There is actually talk about the possibility of going back to the concept of **contact disks**, where the head gap is intentionally made zero. This would allow even smaller magnetic fields than is possible in today's drives. Of course, this brings us full circle to the first hard disk experiments in the 1950s!
- The difference of course is almost 50 years of advances in technology. For example, thin film media is much tougher than the oxide media used on contact disks half a century ago, and lubricating agents are much more advanced as well.
- Even so, it will probably be several years before we know if this technology will be feasible from both an engineering and manufacturing standpoint.

Head Crashes

- Since the read/write heads of a hard disk
  - are floating on a microscopic layer of air
  - above the disk platters themselves,
  - it is possible that the heads can make contact with the media on the hard disk under certain circumstances.
- **Normally**, the heads only contact the surface when the drive is either **starting up** or **stopping**. Considering that a modern hard disk is turning over 100 times a second, this is not a good thing. :^)
- **If** the **heads contact** the surface of the disk while it is at **operational speed**, the result can be:
  - loss of data
  - damage to the heads
  - damage to the surface of the disk
  - or all three.
- This is usually called a **head crash**, two of the most frightening words to any computer user. :^) The most common causes of head crashes are **contamination** getting stuck in the thin gap between the head and the disk, and **shock applied** to the hard disk while it is in operation.

Head Crashes

- Despite the lower floating height of **modern hard disks**, they are in many ways less susceptible to head crashes than older devices.
- The reason is:
  - the superior design of hard disk enclosures
  - to eliminate contamination,
  - more rigid internal structures and
  - special mounting techniques designed to eliminate vibration and shock.
- The platters themselves usually have a **protective layer** on their surface that can tolerate a certain amount of abuse before it becomes a problem.
- Taking precautions to avoid head crashes
  - especially not abusing the drive physically
  - is obviously still common sense
- Be especially careful with portable computers; I try to never move the unit while the hard disk is active.

Read/Write Head Technologies

- There are several different technology families that have been employed to make hard disk read/write heads.
- Usually,
  - to enable hard disk speed and capacity to progress
  - to higher levels,
  - adjustments must be made to the way the critical read/write head operation works.
- In some cases
  - this amounts to minor tweaking of existing technologies,
  - but major leaps forward usually require
  - a breakthrough of some sort,
  - once the existing technologies have been pushed to their limits.

Read/Write Head Technologies

- Summary chart showing the basic design characteristics of most of the read/write head designs used in PC hard disks.

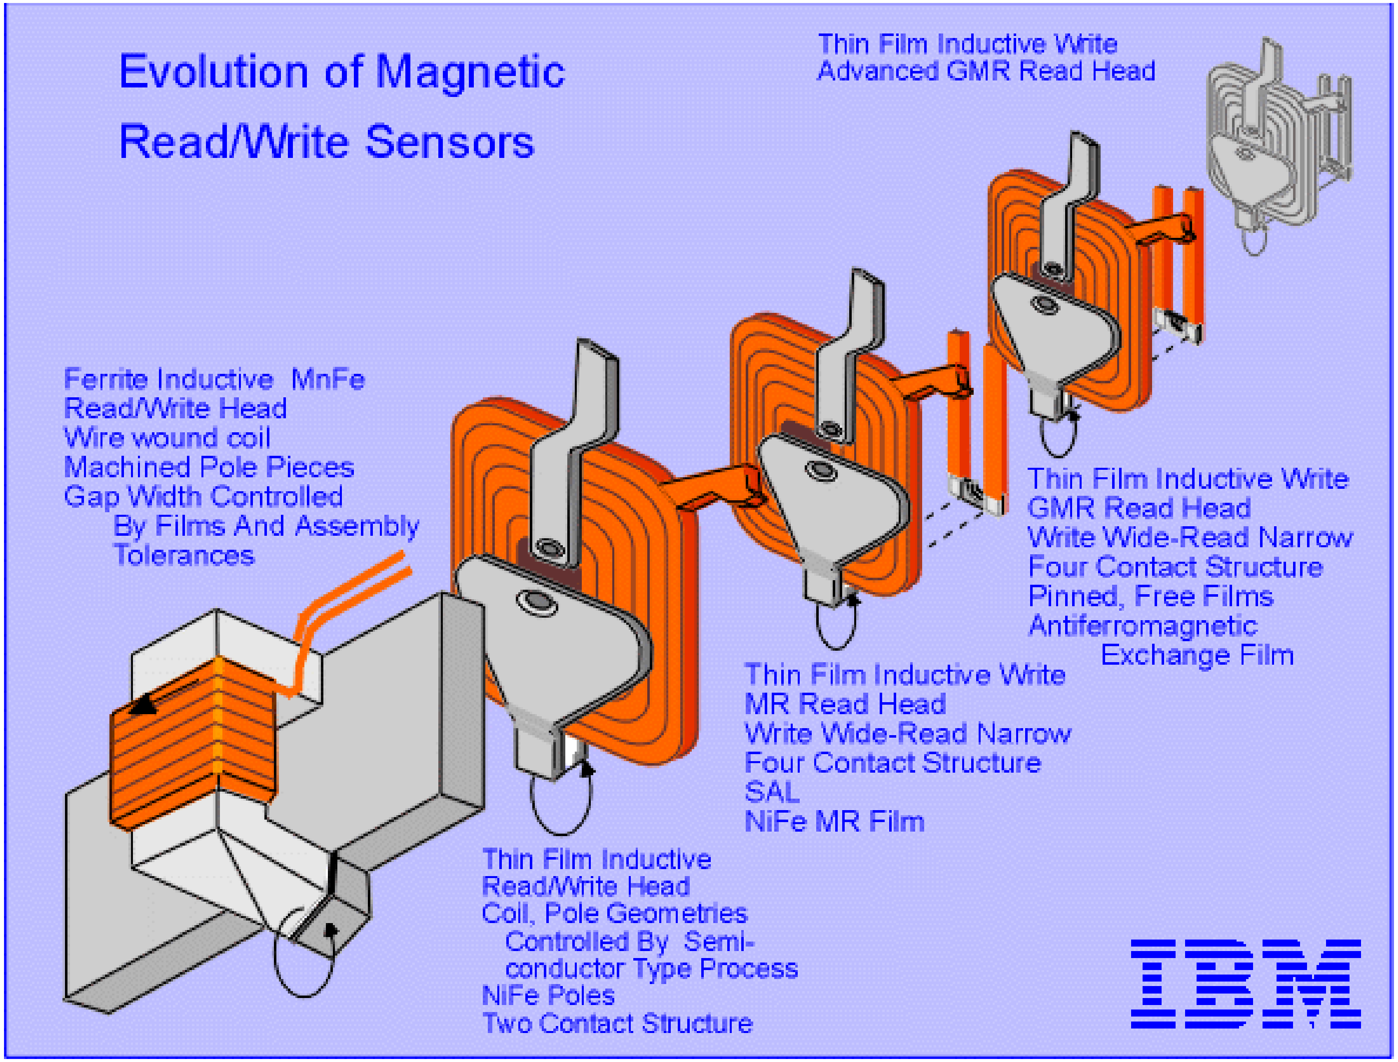

Ferrite Heads

- The oldest head design is also the simplest conceptually.
- A **ferrite head** is a
  - **U-shaped iron core**
  - wrapped with electrical windings
  - to create the read/write head
  - almost a classical electromagnet, but very small.
  - (The name "ferrite" comes from the iron of the core.)
  - The result of this design is much like a child's U-shaped magnet,
  - with each end representing one of the poles, north and south.
- **When writing**, the current in the coil creates a polarized magnetic field in the gap between the poles of the core, which magnetizes the surface of the platter where the head is located. When the direction of the current is reversed, the opposite polarity magnetic field is created.
- **For reading**, the process is reversed: the head is passed over the magnetic fields and a current of one direction or another is **induced** in the windings, depending on the polarity of the magnetic field.

Ferrite Heads

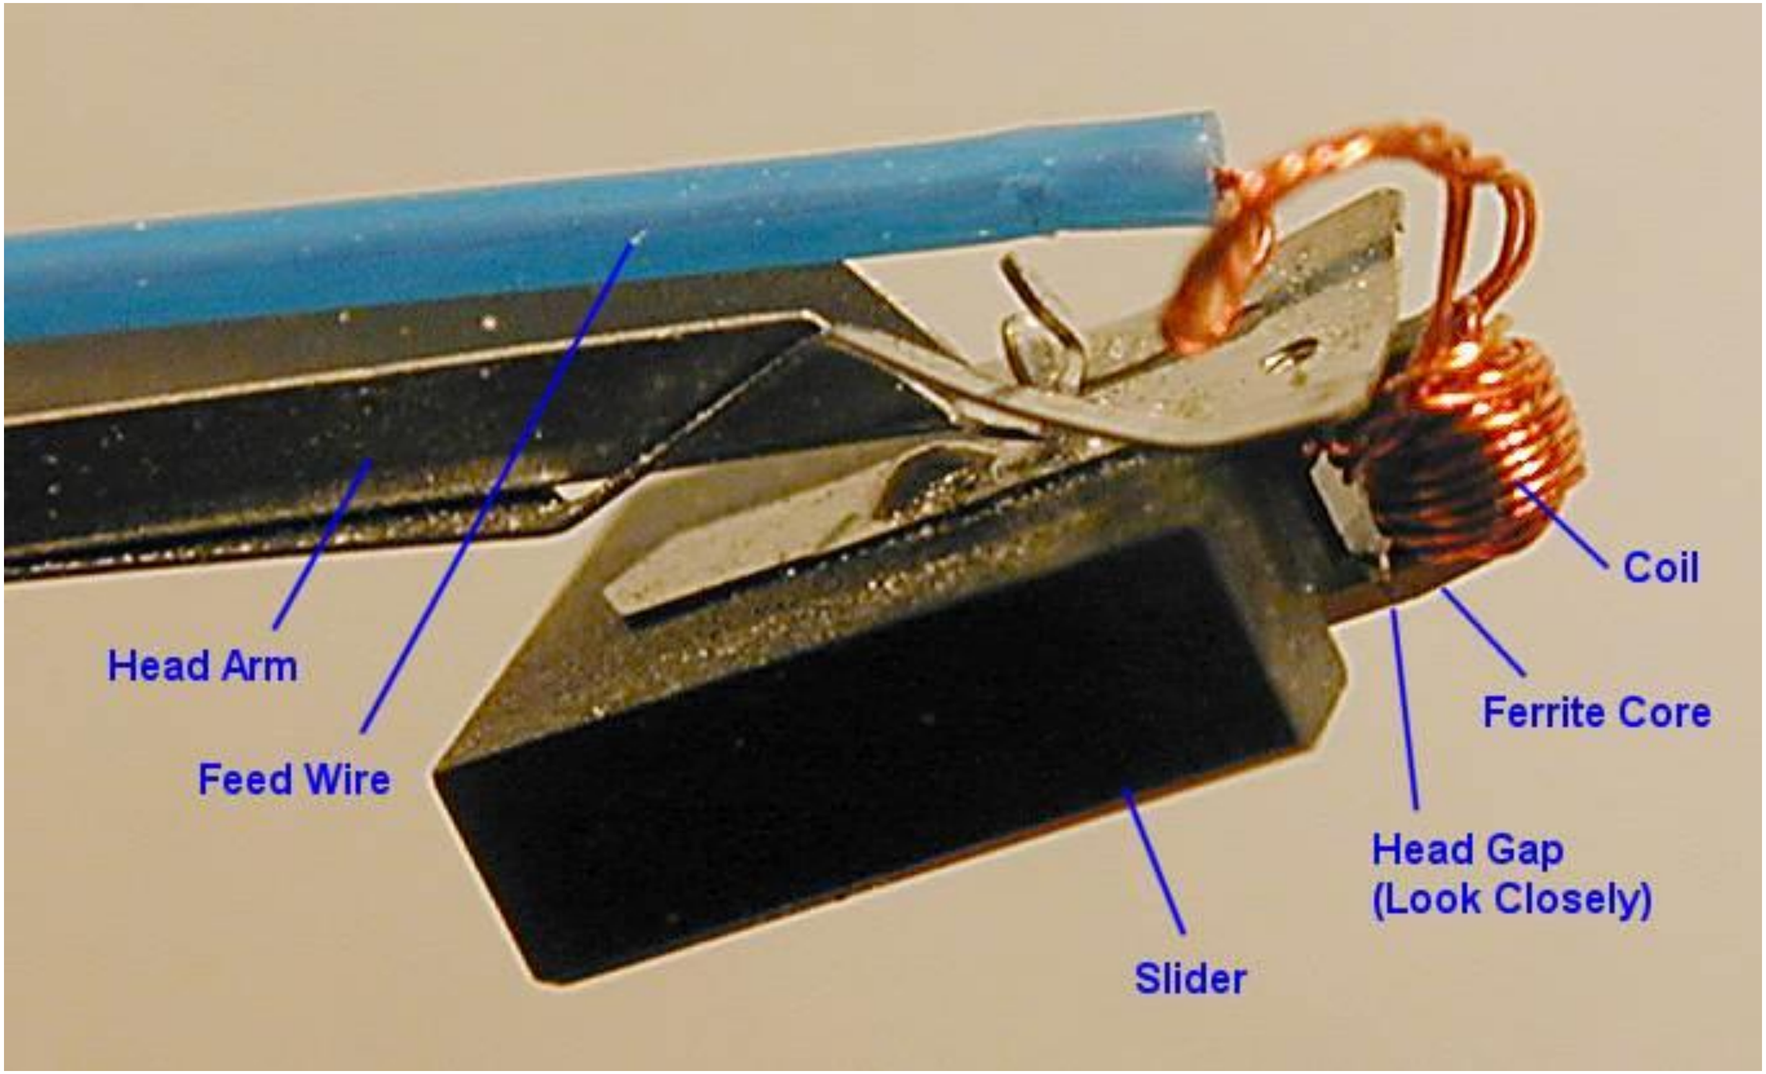

Ferrite Heads

- Extreme closeup view of a ferrite read/write head from a mid-1980s Seagate ST-251, one of the most popular drives of its era.
- The **big black** object is not actually the head, but the slider.
- The head is at the end of the slider, wrapped with the coil that magnetizes it for writing, or is magnetized during a read.
- If you look closely you can actually see the gap in the core, though it is very small.
- The blue feed wire runs back to the circuits that control the head.

Ferrite Heads

- Ferrite heads suffer from **being large** and **cumbersome (glomazan)**,
- which means
- they must ride at a relatively great distance from the platter
- and must have reasonably large and strong magnetic fields.

- Their design prevents their use with modern,
- very-high-density hard disk media,
- and they are now obsolete and no longer used.

- They are typically encountered in PC hard disks under 50 MB in size.

Metal-In-Gap (MIG) Heads

- An evolutionary improvement to the standard ferrite head design was the invention of **Metal-In-Gap** heads.
- These heads are essentially of the same design as ferrite core heads, but **add** a **special metallic alloy in** the **head**.
- This change greatly increases its magnetization capabilities, allowing MIG heads to be used **with higher density media**, increasing capacity.
- While an improvement over ferrite designs, MIG heads themselves have been supplanted by thin film heads and magnetoresistive technologies.
- They are usually found in PC hard disks of about 50 MB to 100 MB.
- **Note:** The word "gap" in the name of this technology refers to the gap between the poles of the magnet used in the core of the read/write head, not the gap between the head and the platter.

Thin Film (TF) Heads

- Thin Film (TF) heads--also called thin film inductive (TFI)—
- are a totally different design from ferrite or MIG heads.

- They are so named because of how they are manufactured.

- TF heads are made using a photolithographic process
- similar to how processors are made.

- This is the same technique
  - used to make modern thin film platter media,
  - which bears the same name;
  - see here for more details on this technology.

Thin Film (TF) Heads

- In this design,
  - developed during the 1960s but not deployed until 1979
  - the iron core of earlier heads, large, bulky and imprecise
  - is done away entirely.
- A substrate wafer is coated with a **very thin layer** of **alloy material** in specific patterns.
- This produces a **very small, precise head** whose characteristics can be carefully controlled, and allows the bulky ferrite head design to be completely eliminated.
- Thin film heads are capable of being used on much higher-density drives and with much smaller floating heights than the older technologies.
- They were used in many PC hard disk drives in the late 1980s to mid 1990s, usually in the 100 to 1000 MB capacity range.
- As hard disk areal densities increased, however, thin film heads soon reached their design limits. They were eventually replaced by magnetoresistive (MR) heads.

TF head

- A pair of mated thin film head assemblies, greatly magnified. The heads are gray slivers with coils wrapped around them, embedded at the end of each slider (large beige objects). One feed line (with green insulation) is visible.

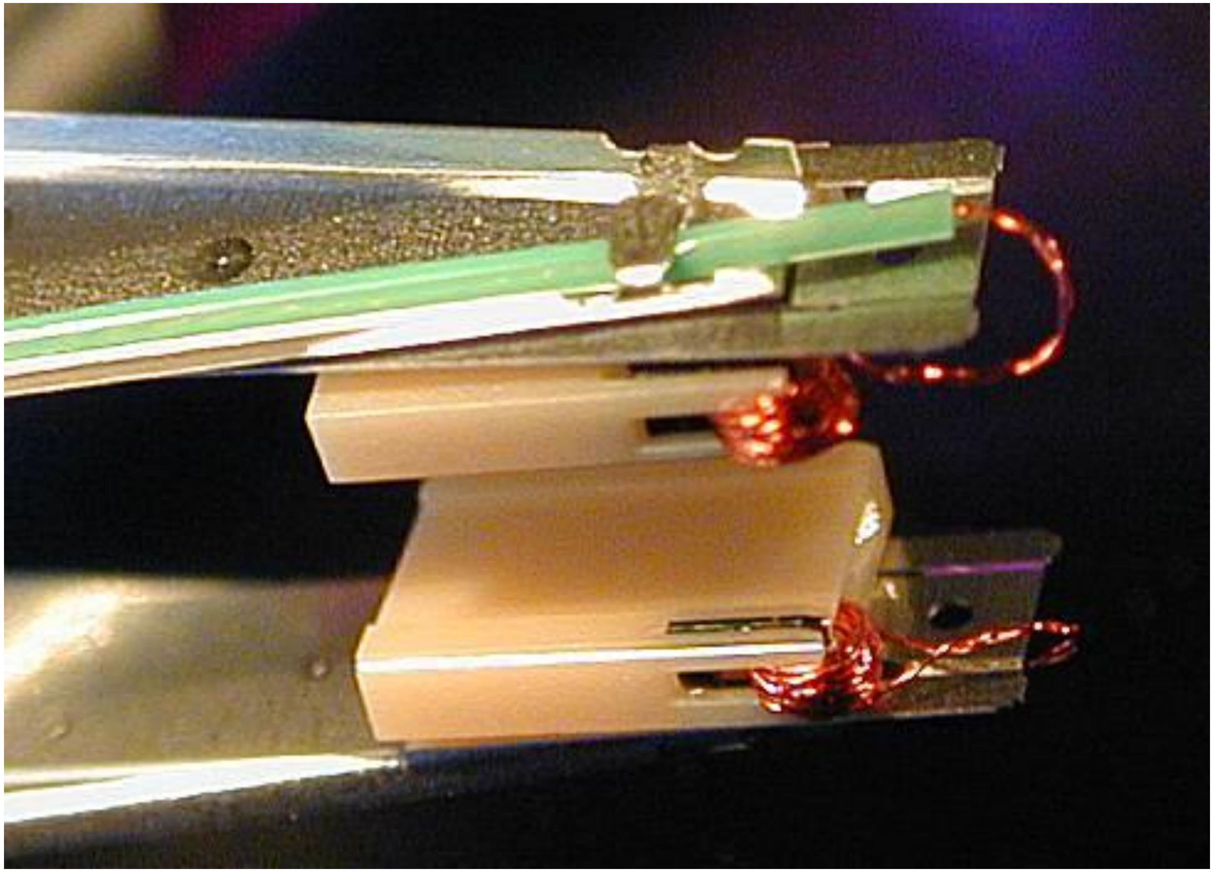

(Anisotropic) Magnetoresistive (MR/AMR) Heads

- The **newest type** of t**echnology** commonly
  - used in read/write heads
  - is much more of a radical change to the way the read/write head works internally than the earlier advances,
  - which were much more evolutionary and more related to how the head was made than how it worked.
  - While conventional ferrite or thin film heads work
  - on the basis of inducing a current in the wire of the read head
  - in the presence of a magnetic field,
  - ***magnetoresistive* (*MR*)** heads use a **different principle entirely to read the disk**.
- **Note:** The correct technical name for first-generation MR heads
  - is ***anisotropic magnetoresistive* (AMR),**
  - but traditionally they have just been called "magnetoresistive" (MR).
  - With GMR heads now on the market, there is the potential for confusion between the terms "magnetoresistive" and "giant magnetoresistive".
- Therefore, some companies have now gone back to calling the older MR heads "AMR" heads to distinguish them from GMR ones. Normally though, if you are told a drive has "MR heads", this means the older technology described here

(Anisotropic) Magnetoresistive (MR/AMR) Heads

- An **MR head employs**
  - a **special conductive material**
  - that changes its resistance
  - in the presence of a **magnetic field**.

- As the head passes over the surface of the disk
  - this material changes resistance
  - as the magnetic fields change
  - corresponding to the stored patterns on the disk.
  - A sensor is used to detect these changes in resistance,
  - which allows the bits on the platter to be read.

Magnetoresistive (MR/AMR) Heads

- The use of MR heads
  - allows **much higher areal densities** to be used on the platters
  - than is possible with older designs,
  - greatly increasing the storage capacity
  - and (to a lesser extent) the speed of the drive.
- Because the MR head
  - **is not generating a current**
  - directly the way standard heads do,
  - it is several times **more sensitive** to magnetic flux changes in the media.
- This allows the use of weaker written signals
  - which lets the bits be spaced closer together
  - without interfering with each other
  - improving capacity by a large amount.

Magnetoresistive (MR/AMR) Heads

- **MR technology is used for reading** the **disk only**.
- **For writing**, a **separate standard thin-film head is used**.
- This splitting of chores into one head for reading and another for writing has **additional advantages**.

- Traditional heads that do both reading and writing are an exercise in tradeoffs, because many of the improvements that would make the head read more efficiently would make it write less efficiently, and vice-versa.

- For example, if you increase the **number of windings of wire** around the core of a standard read/write head, you increase the sensitivity of the head when reading, but you make it much more difficult to **write at high speed**.

- Also, for best results we want to write a wider data track (to ensure the media is properly magnetized) but read a narrower one (to make sure we don't accidentally pick up signals from adjacent bits). In an MR design the MR head does the reading, so the thin film write head can be optimized solely for writing without worrying about these sorts of compromises.

Magnetoresistive (MR/AMR) Heads

- Closeup view of an MR head assembly. Note that the separate copper lead wire of older head designs is gone, replaced by thin circuit-board-like traces. The slider is smaller and has a distinctive shape.
  The actual head is too small to be seen without a microscope.

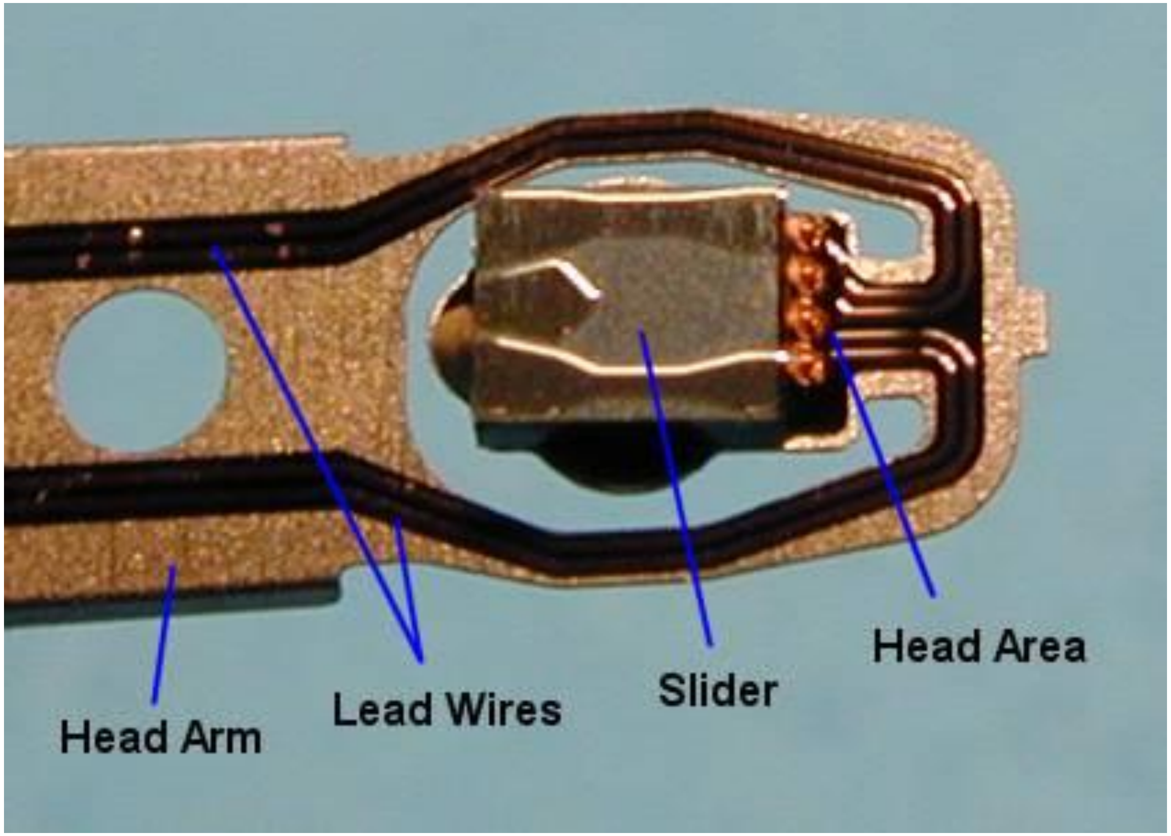

Magnetoresistive (MR/AMR) Heads

- **First introduced** in 1991 by IBM--who else—
  - but not used widely until several years later
  - MR heads were one of the key inventions
  - that led to the creation of hard disks over 1 GB in size
  - and the subsequent explosive growth in size since then.

- Despite the increased cost of MR heads
  - they have now totally replaced thin film heads
  - which just are not up to the challenge of hard disks
  - in the tens of gigabytes
  - MR heads are commonly found in hard disks
  - from about **1 GB to about 30 GB** in size.

Magnetoresistive (MR/AMR) Heads

- Even MR heads however have a limit in terms of how much areal density they can handle.
- Successive generations of MR heads were reduced in size to allow still greater areal density.
- Sometimes these more advanced designs were dubbed *MRX* for *Extended Magnetoresistive* heads.
- The successor to MR now appears to be GMR heads, named for the *giant magnetoresistive effect*.
- They are similar in basic concept to MR heads but are more advanced.

Giant Magnetoresistive (GMR) Heads

- From the beginning in the 1950s and its first hard disk drive,
- IBM has been at the forefront of hard disk technological advances.

- IBM's groundbreaking research and development efforts continue, nearly 50 years later, with the introduction of giant magnetoresistive (GMR) hard disk read heads.

- These heads work on the same general principles
  - as the original (anisotropic) magnetoresistive heads,
  - but use a **somewhat different design**
  - that makes them superior in several ways.

Giant Magnetoresistive (GMR) Heads

- GMR heads are not named "**giant**" because of **their size**;
- they are actually smaller than the regular (A)MR heads
- developed by IBM many years earlier.

- Rather, they are named after the **giant magnetoresistive effect**,
  - first discovered in the late 1980s by two European researchers,
  - Peter Gruenberg and Albert Fert, who were working independently.

- Working with **large magnetic fields**
  - and thin layers of **various magnetic materials**,
  - they noticed **very large resistance changes**
  - when these materials were subjected to magnetic fields.
  - These early experiments used techniques and materials
  - that were not suitable to manufacturing,
  - but they formed the basis for the technology.

Giant Magnetoresistive (GMR) Heads

- From there, engineers and scientists at IBM's Almaden Research Center took over.
- IBM developed GMR into a commercial product
- by experimenting with thousands of different materials and methods.

- A key advance was the discovery that the GMR effect would work
  - on multilayers of materials deposited by sputtering
  - the same technique used to make thin film media
  - and thin film read/write heads.

- By December 1997, IBM had introduced its first hard disk product using GMR heads, and areal density records have been looking over their shoulders nervously ever since.

Giant Magnetoresistive (GMR) Heads

- Detailed structure diagram of a GMR head assembly. Notice how small these devices are: the scale is magnified twice to get to the detail level of the GMR sensor itself, and the arm/slider/head structure at top is itself only about a quarter of an inch long (see the illustration on the page covering MR heads to see what one of those assemblies looks like in real life.)

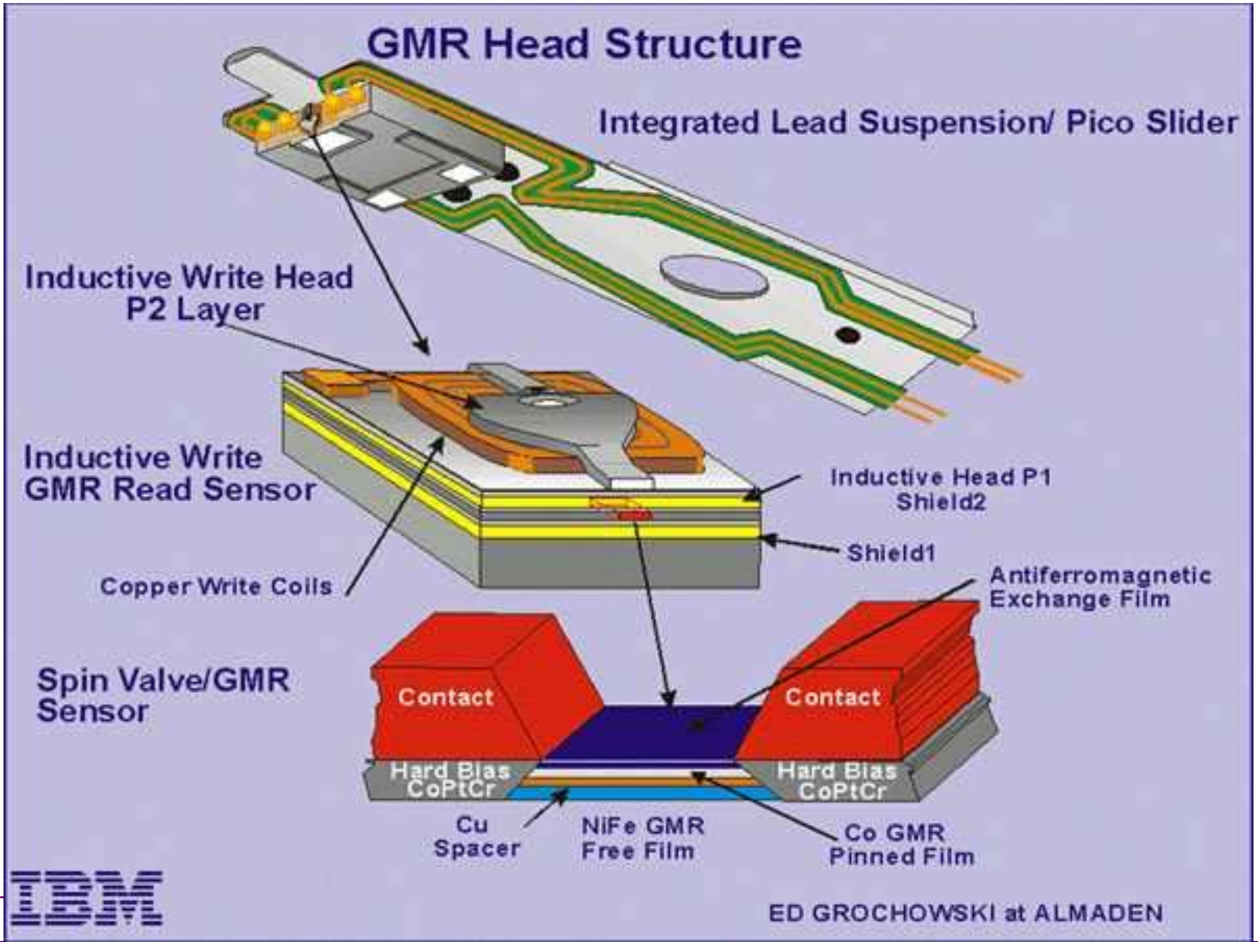

Giant Magnetoresistive (GMR) Heads

- GMR heads are comprised of **four layers** of thin material sandwiched together into a single structure:
- **Free Layer:** This is the **sensing layer**, made of a **nickel-iron alloy**, and is passed over the surface of the data bits to be read. As its name implies, it is free to rotate in response to the magnetic patterns on the disk.
- **Spacer:** This layer is nonmagnetic, typically made from copper, and is placed between the free and pinned layers to separate them magnetically.
- **Pinned Layer:** This layer of cobalt material is held in a fixed magnetic orientation by virtue of its adjacency to the exchange layer.
- **Exchange Layer:** This layer is made of an "antiferromagnetic" material, typically constructed from iron and manganese, and fixes the pinned layer's magnetic orientation.

Giant Magnetoresistive (GMR) Heads

- Here's how the GMR head works in a nutshell (without getting into quantum physics that would make your brain melt--and mine. :^))
- When the head passes over a magnetic field of one polarity
  - (say, a "0" on the disk),
  - the free layer has its electrons turn to be aligned
  - with those of the pinned layer;
  - **this creates a lower resistance** in the **entire head structure.**
- When the head passes over a magnetic field of the **opposite polarity** ("1"),
  - the electrons in the free layer rotate
  - so that they are not aligned with those of the pinned layer.
- This causes an **increase** in the **resistance of the overall** structure.
  - The resistance changes are caused by changes to the spin characteristics of electrons in the free layer,
  - and for this reason, IBM has named these structures *spin valves*.
  - If you imagine a plumbing pipe with a rotatable shut-off valve, that's the general concept behind the name.

Giant Magnetoresistive (GMR) Heads

- Conceptual operation of a GMR head, showing the four layers.
- As the head moves across a bit, electrons in the free layer rotate, increasing the resistance of the overall structure.

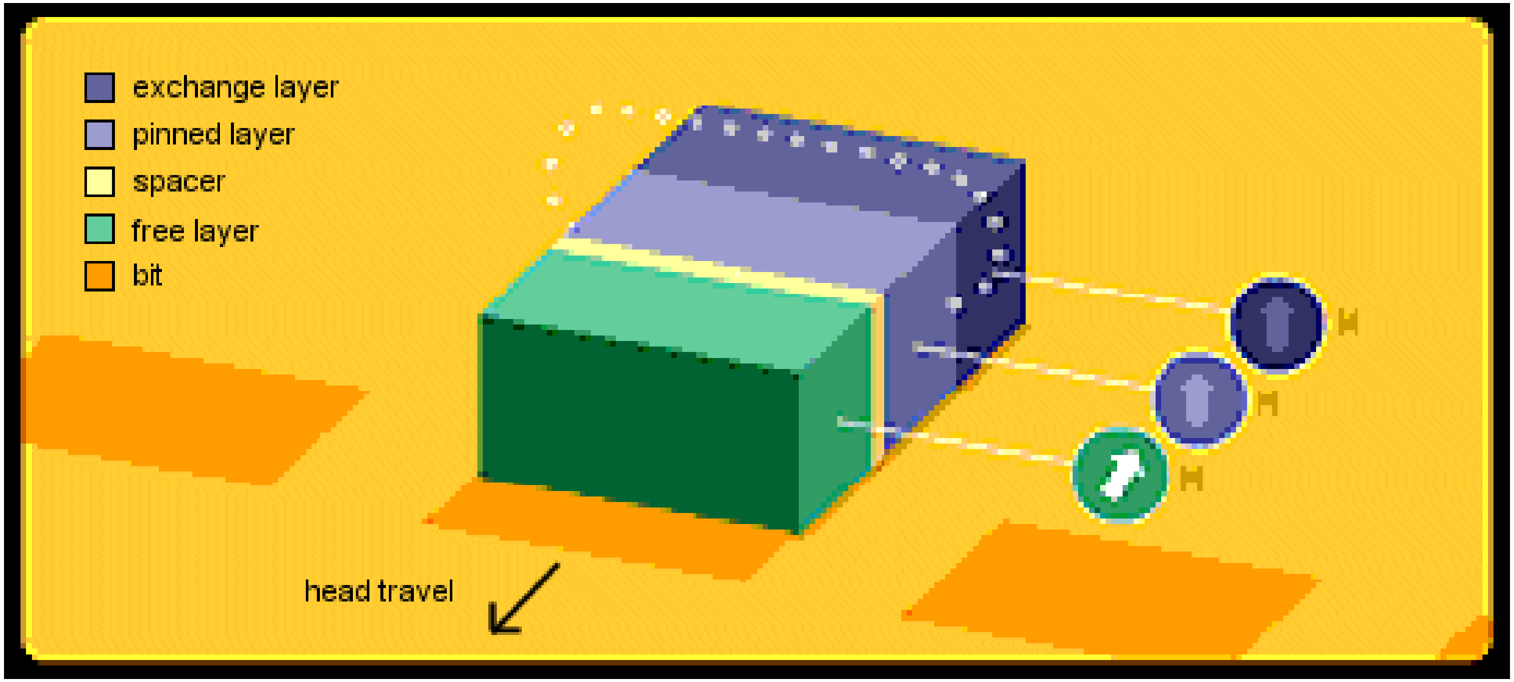

Giant Magnetoresistive (GMR) Heads

- GMR heads are superior to conventional MR heads because they are more sensitive. **While older MR heads** typically exhibit a resistance change when passing from one magnetic polarity to another of about 2%, for GMR heads this is anywhere from 5% to 8%.

- This means GMR heads can detect much weaker and smaller signals, which is the key to increasing areal density, and thus capacity and performance.

- They are also much less subject to noise and interference because of their increased sensitivity, and they can be made smaller and lighter than MR heads.

- GMR heads are also typically fitted with "shields" that prevent them from being affected by stray magnetic fields, i.e., anything but the bit directly below the head.

Giant Magnetoresistive (GMR) Heads

- GMR heads are used in the latest technology drives,
  - which currently have capacities of up to 75 GB
  - and areal densities of approximately 10 to 15 Gbits/in2.

- As of early 2000, IBM has already produced GMR heads
  - in the lab capable of 35 Gbits/in2,
  - suggesting that standard form factor hard disks of 200 GB
  - and up are just around the corner.

Colossal Magnetoresistive (CMR) Heads

- Now that new giant magnetoresistive (GMR) heads
  - are taking over the market,
  - engineers have turned their attentions to the next breakthrough
  - that will take hard disk read heads to the next level of performance.
- One of the designs being researched
  - is an evolutionary advance based on GMR
  - that has been dubbed by some
  - as *colossal magnetoresistive* or *CMR*.
- At this stage this technology is in the early stages of investigation, and I have not been able to find much information on it. It would be logical to assume that it involves using slightly different materials than GMR in an effort to further increase sensitivity.

Hard Disk Head Sliders, Arms and Actuator

- The hard disk platters are accessed for read and write operations using the read/write heads mounted on the top and bottom surfaces of each platter.
- Obviously, the read/write heads don't just float in space;
  - they must be held in an exact position relative to the surfaces
  - they are reading, and furthermore,
  - they must be moved from track to track
  - to allow access to the entire surface of the disk.
- The heads are mounted onto a structure that facilitates this process.
- Often called the
- **head assembly**
- or
- **actuator assembly**
- (or even the **head-actuator assembly**),
- it is comprised of several different parts.

Hard Disk Head Sliders, Arms and Actuator

- The heads themselves are mounted on **head sliders**.

- The sliders are suspended over the surface of the disk
- at the ends of the **head arms**.

- The head arms are all mechanically fused into a single structure that is moved around the surface of the disk by the **actuator.**

Actuator assembly (montaža)

- Annotated illustration of a typical PC actuator assembly, showing the major components. The platters have been removed from the drive to provide a better view of the actuator arms and heads. There are four sliders but only one of each pair is visible. The spindle motor is visible at the top right.

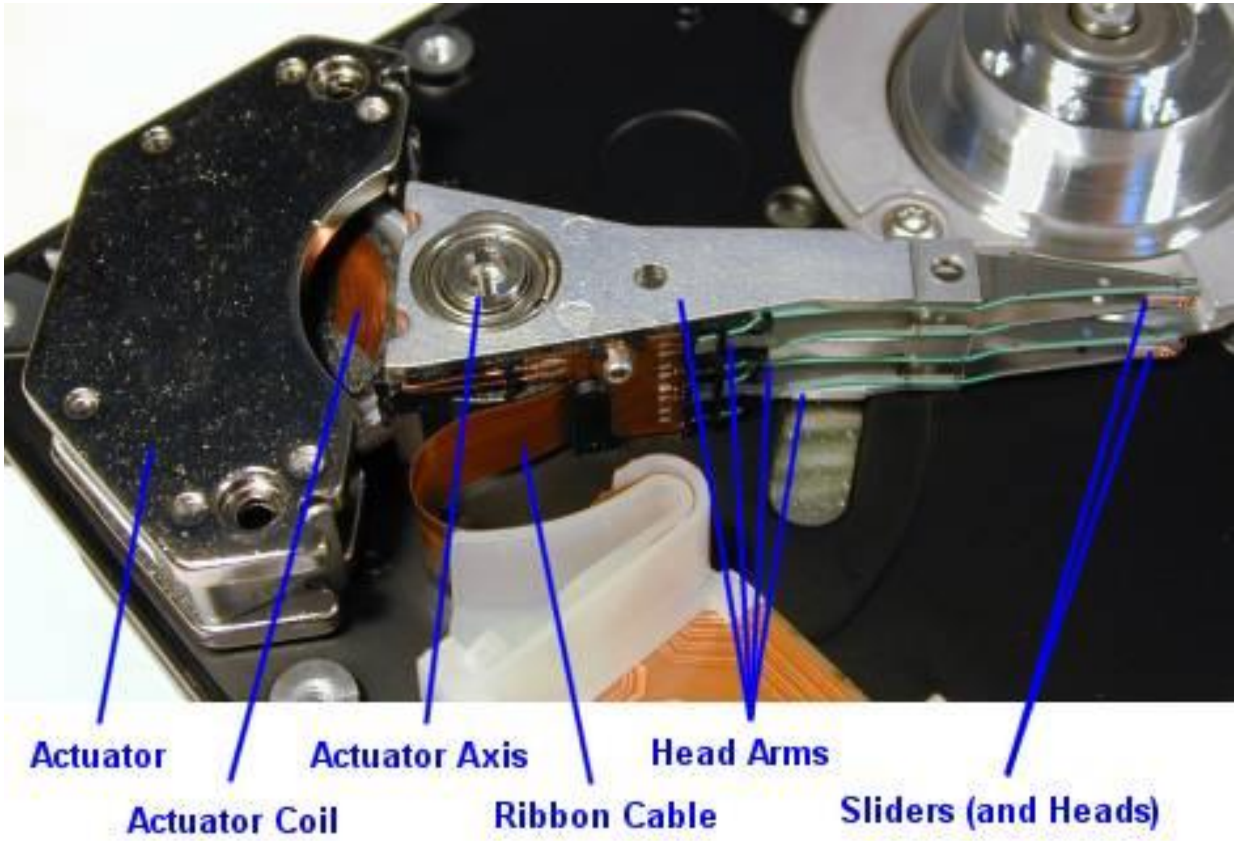

Head Sliders

- Hard disk read/write **heads are too small** to be used without attaching them **to a larger unit**. This is especially true of modern hard disk heads. Each hard disk head is therefore mounted to a special device called a **head slider** or just *slider* for short. The function of the slider is to physically support the head and hold it in the correct position relative to the platter as the head floats over its surface.

- **Sliders are given** a **special shape** to allow them to ride precisely **over the platter**. Usually they are shaped somewhat like a sled (sanke); there are two rails or runners on the outside that support the slider at the correct flying height over the surface of the disk, and in the middle the read/write head itself is mounted, possibly on another rail.

- **As hard disk read/write heads have been shrinking in size**, so have the sliders that **carry them**. The main advantage of using small sliders is that it reduces the weight that must be yanked around the surface of the platters, improving both positioning speed and accuracy. Smaller sliders also have less surface area to potentially contact the surface of the disk.

Head Sliders

- A graphic illustration of what approximately 15 years' worth of technological evolution has done to hard disk head sliders.
  At **left**, a **slider from a 40 MB 5.25" ferrite-head drive**;
  at **right**, the **slider from a 3.2 GB, 3.5" MR-head drive**.

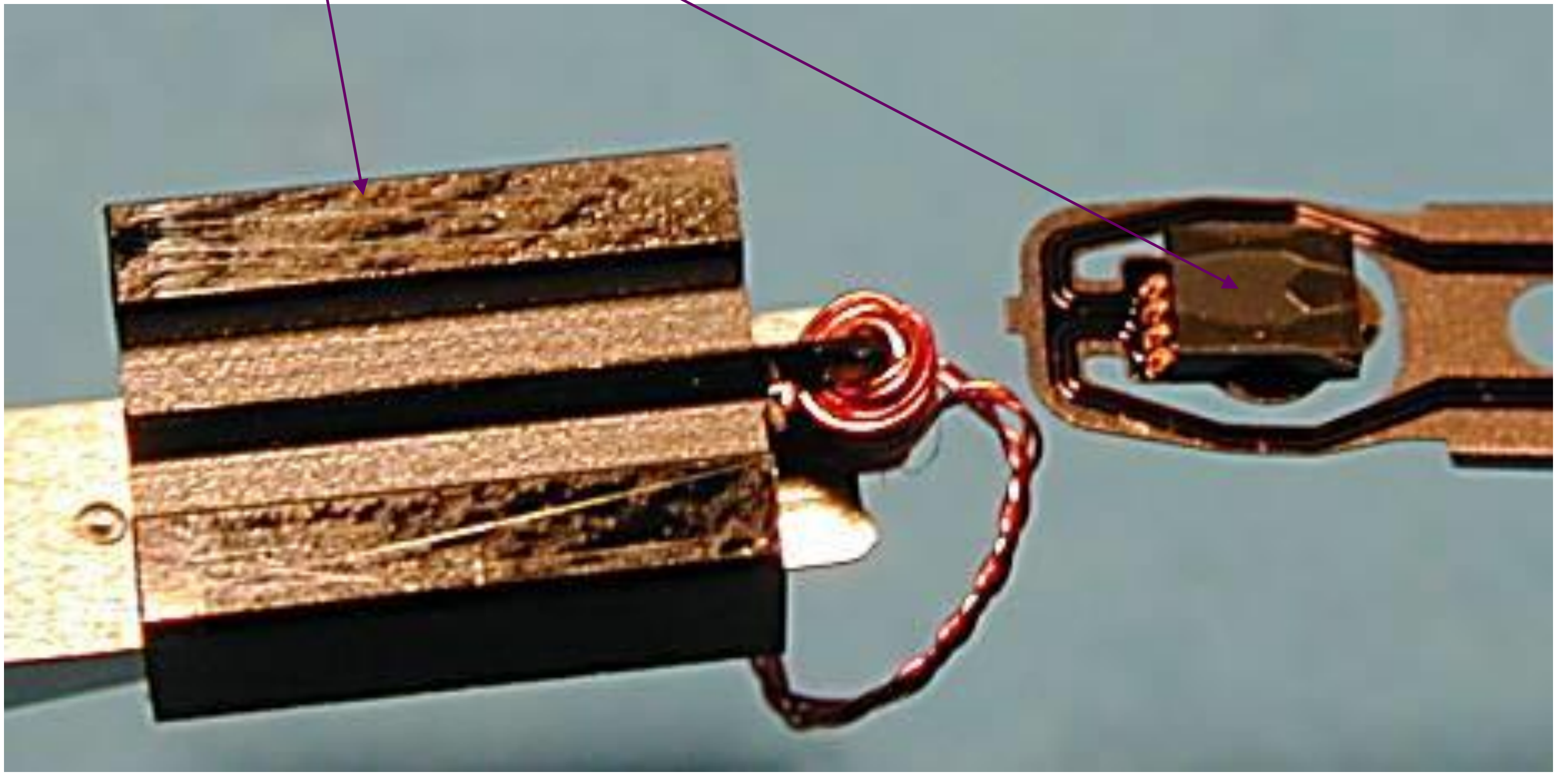

Head Arms

- The top head arm of a typical recent-design hard disk drive. Note that the arm is not solid, but rather has a structural triangular shape. This is done to reduce weight while maintaining rigidity.

Head Actuator

- The **actuator is the device**
  - **used to position** the **head arms**
  - to different tracks on the surface of the platter
  - (actually, to different cylinders, since all head arms are moved as a synchronous unit, so each arm moves to the same track number of its respective surface).
- The actuator is a very important part of the hard disk, because changing from track to track is the only operation on the hard disk that requires active movement: changing heads is an electronic function, and changing sectors involves waiting for the right sector number to spin around and come under the head (passive movement).
- Changing tracks means the heads must be shifted, and so making sure this movement can be done quickly and accurately is of paramount importance. This is especially so because physical motion is so *slow* compared to anything electronic--typically a factor of 1,000 times slower or more.

Head Actuator

- Head actuators come in two general varieties:
- **1. Stepper Motors**: Originally, hard disk drives used a stepper motor to control the movement of the heads over the surface of the platters.
- A regular motor turns in a rotary fashion continuously; it can stop at any point in its rotation as it spins around, kind of like the second hand on a wind-up wristwatch.
- A stepper motor can only stop at predefined "steps" as it turns around, much the way the second hand turns on an electronic, quartz wristwatch.
- A hard drive using a stepper motor for an actuator attaches the arms to the motor, and each time the motor steps one position clockwise or counterclockwise, the arms move in or out one position.
- **Each position defines** a **track on** the **surface of the disk**.
- Stepper motors are also commonly used for both turning the spindle and positioning the head on floppy disk drives. If you have a floppy drive, find one of its motors and turn it slowly with your hand; you will feel the discrete step-wise nature of its motion.

Head Actuator

- **2. Voice Coils**:
- The actuator in a modern hard disk
  - uses a device called a voice coil
  - to move the head arms in and out over the surface of the platters,
  - and a closed-loop feedback system called a servo system to dynamically position the heads directly over the data tracks.
- The voice coil works using electromagnetic attraction and repulsion.
  - A coil is wrapped around a metal protrusion
  - on the end of the set of head arms.
  - This is mounted within an assembly containing a strong permanent magnet
  - When current is fed to the coil, an electromagnetic field is generated that causes the heads to move in one direction or the other based on attraction or repulsion relative to the permanent magnet.
- By controlling the current, the heads can be told to move in or out much more precisely than using a stepper motor.
  - The name "voice coil" comes from the resemblance of this technology
  - to that used to drive audio speakers,
  - which are also basically electromagnets.
  - All PC hard disk voice coil actuators are rotary, meaning that the actuator changes position by rotating on an axis.

Head Actuator

- stepper motor actuator. The motor moves in steps, which you can feel if you move the motor shaft by hand. The shaft has two thin strips of metal wrapped around it, which are connected to a pivot that is rigidly attached to the actuator arms. As the motor shaft turns, one half of this "split band“ coils onto the shaft and the other half uncoils. When the motor turns in the opposite direction the process reverses. As this occurs the pivot moves and in doing so, moves the actuator arms and the hard disk heads.

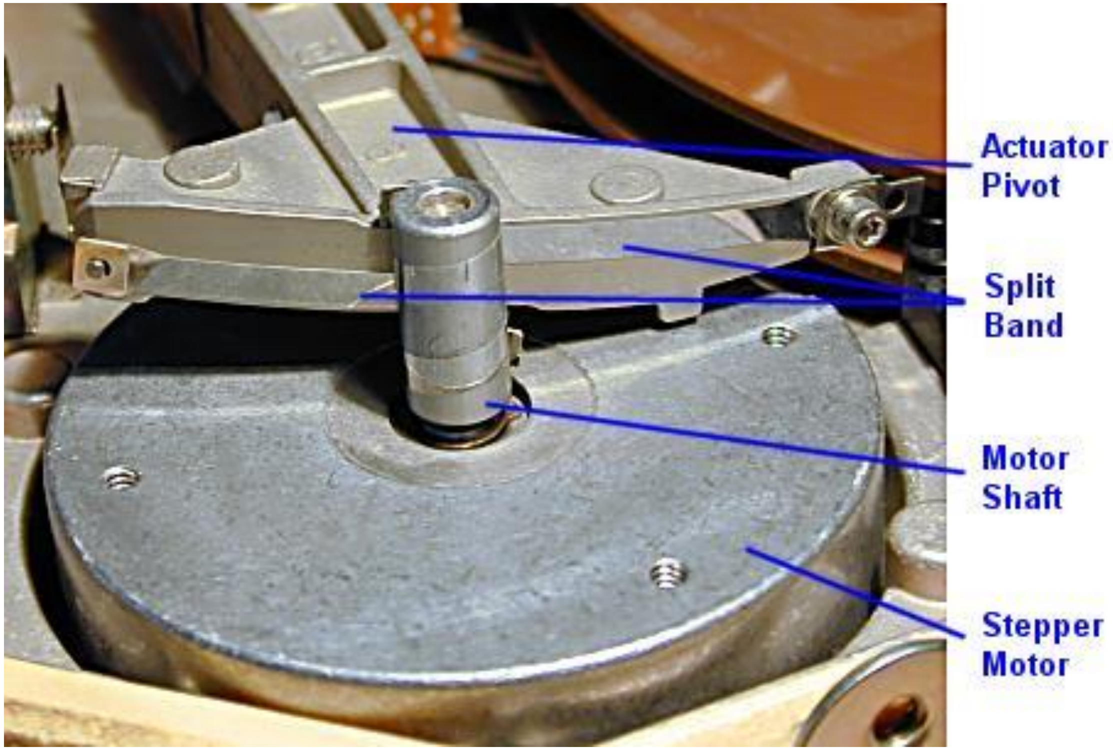

Voice coil actuator

- A partially-disassembled voice coil actuator.
- The magnet assembly has been unscrewed from its mounting and pulled to the left to expose the coil.
- The magnet assembly consists of two metal plates (top one easily visible above, and part of the bottom one visible.)
- The magnet itself is mounted on the underside of the top plate, and spacers used between the plates to create the gap for the coil assembly.
- Being non-ferrous the coil moves freely between the plates, rotating the actuator on its axis as its magnetic polarity is changed. (Incidentally, the magnet is strong enough that after removing the spacers between the plates, the bottom plate got "stuck" on the magnet and required considerable effort to remove!)

Voice coil actuator

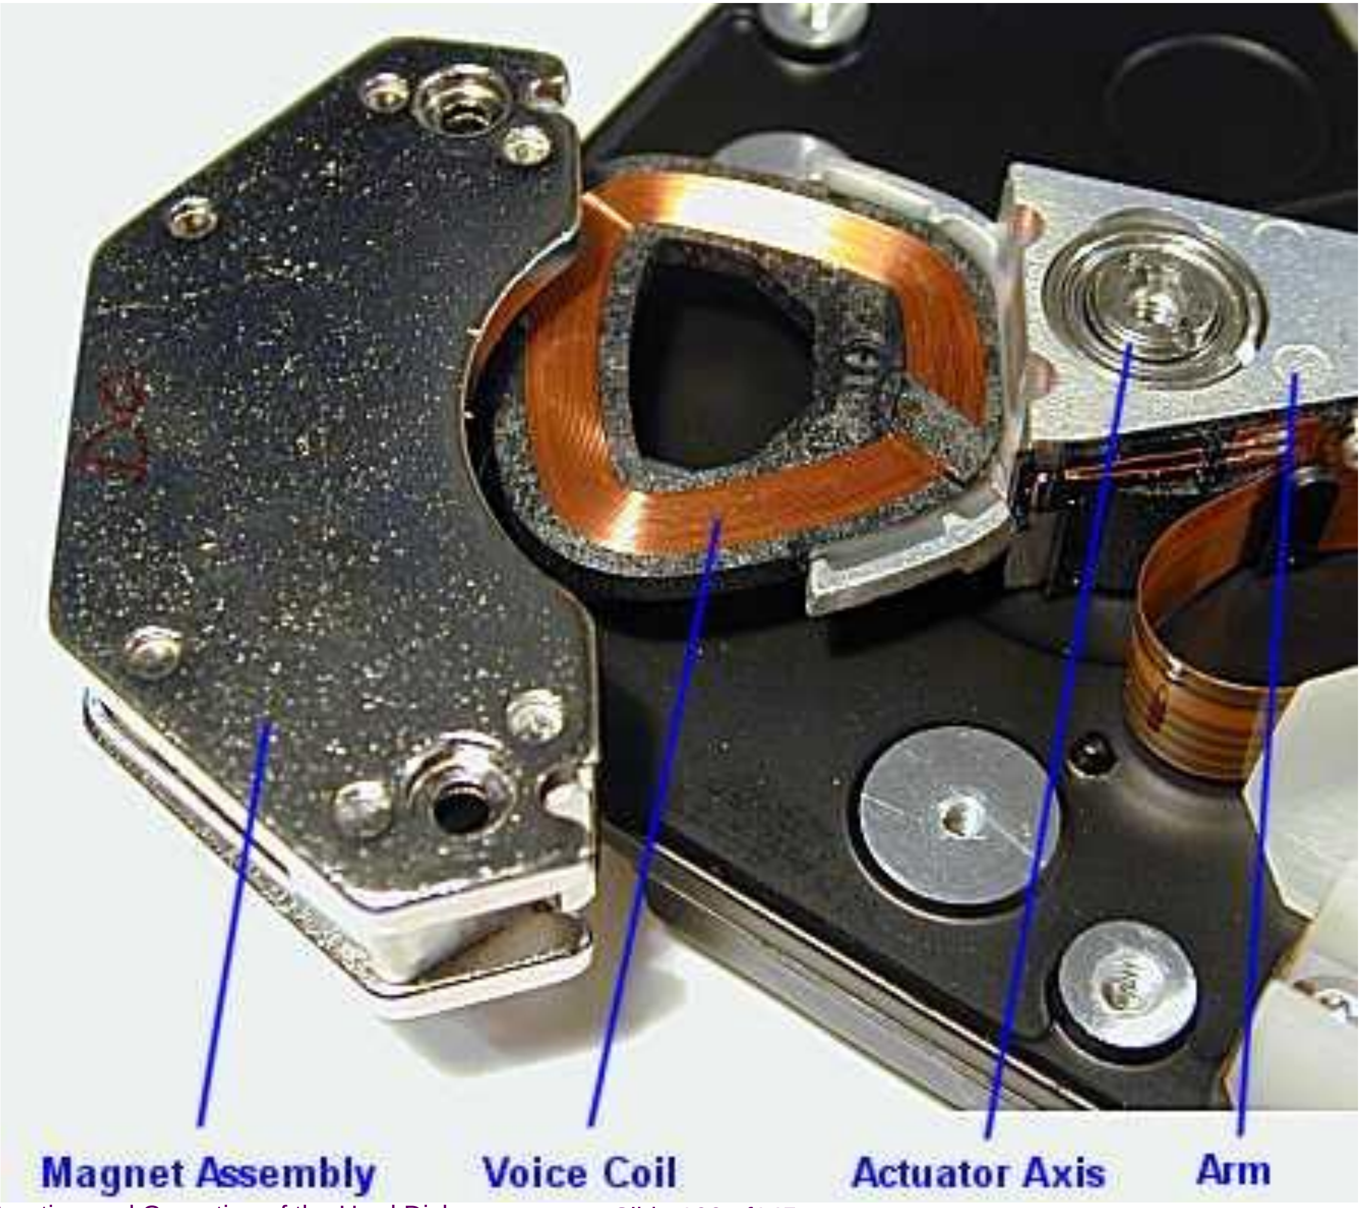

Head Actuator

- The primary distinction between the two designs is that
  - the **stepper motor is an *absolute* positioning system**,
  - while the voice coil is a *relative* positioning system.
- Commands given to a stepper motor actuator are generally of the form "Go in this direction to position A, where you'll find item B".
- Commands to a voice coil actuator are of the form "Go in this direction until you find item B".

- Consider this analogy. In your backyard you have buried a "secret treasure" and want to tell a friend where to find it. When you buried it, you walked down a path 139 paces to the fourth oak tree, and buried it at the edge of the path. The stepper motor analog would be to tell your friend to walk 139 paces down the path, and start digging.
- The voice coil analog would be to tell him to look for the fourth oak tree and dig there. Obviously, using the "139 paces" method, your friend has a problem: his paces aren't likely to be the same length as yours. In fact, even if you yourself walked out 139 paces twice, you'd probably end up in very different spots, since a "pace" isn't an accurate or repeatable measure. On the other hand, the fourth oak tree will always be the fourth oak tree (barring disastrous chain-saw activity :^)).

Head Actuator

- Now hard disks of course don't have to use inaccurate measures like "paces", and it's always the same stepper motor accessing the disk, not a "friend", so why is saying "track #139" a big problem?

- For starters, motors change their characteristics over time,
- and **after a year or two position** #139 might not be where it was when the drive was first formatted.

- However, they have an even **more serious problem:**
  - ☞ disk components (the platters and the head arms themselves especially) **expand** and **contract** with **heat**.

- Even if a stepper motor was perfect, it could not properly account for the fact that the disks are changing in size, and therefore, the tracks are literally moving around.

- If you consider our backyard analogy and think about what it would be like if the oak tree moved a few feet closer to or further from the house based on the day's temperature, you start to realize how inadequate absolute positioning of this form can be.

Head Actuator

- A stepper motor has **no way**
  - to compensate for expansion or contraction of the disk:
  - all it can do is go to where "track #139" is supposed to be,
  - and hope it finds it there!
- If it doesn't find it because
  - the motor and the disk have become out of sync,
  - errors and data loss result.
- This is why older disks were so sensitive to temperature,
  - and normally had to be low-level formatted periodically
  - to make sure the tracks lined up with the heads properly.
- This is also why many drives would fail
  - when first powered up after a weekend,
  - but would work properly
  - after the drive had been allowed to warm up.

Head Actuator

- The shortcomings of stepper motors were unfortunate
  - but acceptable with old hard disks,
  - because of their relatively **low track density**.
- To compensate, tracks could be written fairly wide
  - so that the head would find them
  - even if it was a bit misaligned.
- The first PC hard disks in 1982 had a track density of only two or three hundred tracks per inch 200-300 (TPI).
- Even In 1986, the year Conner Peripherals introduced
  - the **first voice coil PC** hard disk,
  - density had increased to only about 1,000 TPI.
- Stepper motors **are still used to drive floppy disks**, for example, because the accuracy demands for floppies are much lower: a 1.44 MB floppy disk has a track density of 135 tracks per inch.
- In contrast, today's hard disks have densities as high as *30,000* tracks per inch. With data packed this densely, tracks are extremely thin, and a stepper motor lacks the accuracy and stability necessary for proper operation.

Head Actuator

- All modern hard disks use voice coil actuators.
- The voice coil actuator is not only
  - far **more adaptable** and insensitive to **thermal issues**
  - it is much faster and more reliable than a stepper motor.
- The actuator's positioning is dynamic
  - and is **based on feedback**
  - from **examining** the **actual position** of the tracks.
- This closed-loop feedback system is also sometimes called
  - a servo motor
  - or servo positioning system
  - and is commonly used in thousands of different applications
  - where precise positioning is important.
- There are several different ways that the servo positioning system is implemented in PCs; the servo's operation is discussed in its own section.

Servo Techniques and Operation

- Modern hard disks use voice coil actuators
  - to position the heads
  - on the surface of the hard disk's platters.
- This actuator is one instance of what
  - is commonly called a servo system,
  - which is a type of closed-loop feedback system.
- In this sort of positioning system,
  - a device is controlled by doing something,
  - measuring the result
  - seeing how far off the device is from its target
  - making an adjustment
  - and repeating
- This enables the device
  - to reach its target intelligently
  - instead of just taking a guess and hoping it is correct.

Servo Techniques and Operation

- One example of a **closed-loop feedback system** is
  - a modern heating system
  - that uses a **thermostat:**
  - when the temperature gets too low relative to the "target",
  - the heat turns on
  - and stays on until the target temperature is reached.
- Another example is a driver steering through a curve in the road:
  - when the car starts to veer off course,
  - the driver turns the wheel
  - and looks to see if the turn was enough
  - to keep the card on the road.
  - If not, he or she turns the wheel more.
  - When the curve straightens out, the steering wheel is returned to the normal position.
- The *feedback* is what makes the control system "closed-loop".
- In contrast a "send it and hope it finds what is supposed to be there" system such as a stepper motor actuator is called an *open-loop system*.

Servo Techniques and Operation

- A **key element** of any closed-loop feedback system
  - is a measuring device,
  - to provide the feedback.
  - In the case of the thermostat it is a thermometer,
  - and for the driver, it is his or her eyes viewing the road.
- For the hard disk, the feedback device is
  - the **read/write head itself**, and **special codes** written on the disk
  - that let the hard disk know
  - where the heads are when the actuator moves.
- These codes are, unsurprisingly, typically called ***servo codes*.**
- They are read by the heads and
  - fed back to the actuator control logic (at very high speed of course)
  - to guide the actuator to the correct track.
- By putting **different codes** on each track of the disk,
  - the actuator can always figure out
  - which track it is looking at.

Servo Techniques and Operation

- There are **three different ways** that the hard disk servo mechanism has been implemented.
- Each **uses a different way** of
  - recording and
  - reading the servo information from the disk:
- **1. Wedge Servo:** klin
  - In this implementation used in older drives
  - the servo information is recorded in a "wedge" of each platter
  - sort of like a "slice" out of a pie
  - The remainder of the "pie" contains data
  - This design has an important flaw:
  - the servo information is **only in one location** on the hard disk
  - which means that to position the heads a lot of waiting
  - must be done for the servo wedge to rotate around to where the heads are.
  - All this waiting makes the positioning performance of drives that use this method **painfully slow.**
- Obsolete, this technique is no longer used.

Servo Techniques and Operation

- **2. Dedicated Servo:** In this technique:
  - an entire surface of one disk platter is "dedicated"
  - *just* for servo information,
  - and no servo information is recorded on the other surfaces.
- **One head** is constantly reading servo information,
  - allowing very fast servo feedback,
  - and eliminating the delays associated with wedge servo designs.
- Unfortunately, an entire surface of the disk is "wasted"
  - because it can contain no data.
- Also, there is **another problem:**
  - the **heads where data** is recorded may not always line up exactly
  - with the head that is reading the servo information,
  - so adjustments must be made to compensate,
  - and since the servo platter may be warmer or cooler than the data platters,
  - these drives are notorious for needing frequent thermal recalibration.
- Because one platter surface is used for servo information and not data,
  - dedicated servo drives usually have an **odd number of heads**
  - (though there are also marketing reasons why this can happen.)
  - They were found in many drives through the mid-1990s.

Servo Techniques and Operation

- **Embedded Servo:**
- The newest servo technique intersperses
  - servo information with data
  - across the entire surface of all of the hard disk platter surfaces.
  - The servo information and data are read by the same heads,
  - and the heads never have to wait for the disk to rotate the servo information into place as with wedge servo.
- This method doesn't provide
  - the constant access to positioning information
  - that is available with dedicated servo,
  - but it also doesn't require an entire surface
  - to be expended on overhead.
- Also, the need for constant thermal recalibration
  - is greatly reduced
  - since the the **servo information and data**
  - are the same distance from the center of the disk
  - and will expand or contract together.
- All modern hard disks use embedded servo.

Servo Techniques and Operation

- A simple illustration of the difference between dedicated servo and embedded servo.
- On the left, dedicated servo: one platter surface contains nothing but servo information, and the others nothing but data.
- On the right, embedded servo, with data and servo information together.
- (Note that for clarity **only one track on each platter** (one cylinder) is **shown in this illustration**; in fact every track of the servo surface has servo information in the dedicated servo design, and every track of *every* surface has interspersed servo information in the embedded design.

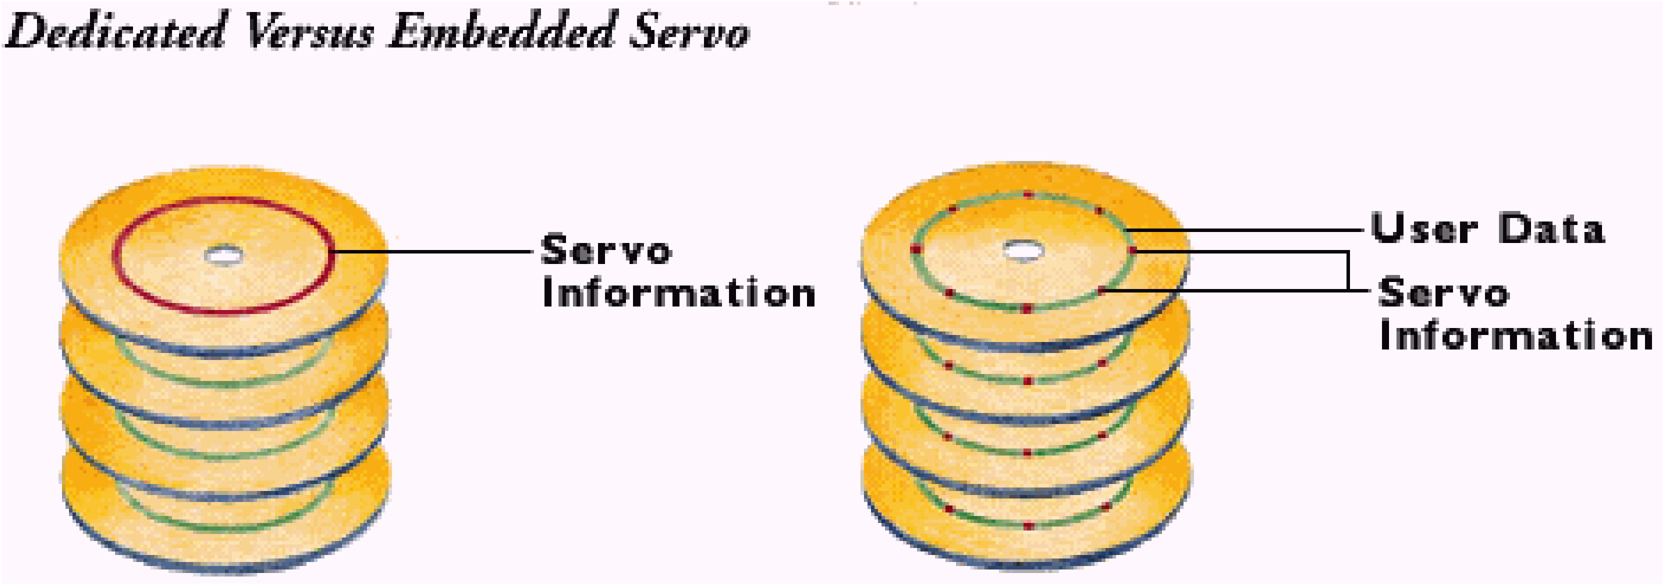

Servo Techniques and Operation

- The **servo codes** are written to the disk surfaces
- at the time the hard disk is manufactured.
- Special, complex and expensive equipment is employed
  - to record this information,
  - which as you can imagine must be placed very precisely
  - on each surface.

- The machines that do this are called ... wait for it... ***servowriters***. :^)

- The servo codes are put in place for the life of the drive
  - and cannot be rewritten without returning the drive to the factory
  - (which never happens because it would be *way* too expensive).
- The hard disk heads themselves
  - are locked out at the hardware level
  - by the drive's controller from writing to the areas
  - where servo information is written.

Servo Techniques and Operation

- The creation of this precise pre-written information
  - ☞ is part of the low-level formatting of a modern drive
  - ☞ and the need for the fancy machine is one reason
  - ☞ why modern disks **cannot be low-level-formatted outside the factory**

- There is nothing a user can do with a drive
- to touch the servo information
- (well, short of using a screwdriver, which is not recommended... :^))

Thermal Recalibration

- All modern drives use **voice coil actuators** instead of the older stepper motors, which makes them far less sensitive to thermal effects than older hard disks were.
- The built in servo mechanism automatically adjusts for shifts in the position of the media due to temperature variations. However, even the newest drives have issues related to **thermal stability**.
- Since engineers continue to decrease the width of each track of data, and increase the number of tracks per inch on the surface of the disk, even the highest-quality electronic servo motors can have, well, "issues" when the various pieces of metal in the hard disk start **expanding and contracting at different rates**.
- This is especially the case with drives using **dedicated servo**, because the servo information and data are on different physical surfaces.
- Since the surfaces can easily be at different temperatures and can therefore expand or contract at different rates, there is the potential that the servo data and user data **might become misaligned**.

Thermal Recalibration

- To combat this problem,
  - most drives manufactured in the mid-1990s
  - include a feature
  - called **thermal recalibration**.

- Every few minutes,
  - the heads are moved
  - and the distance between tracks measured.
  - This information is recorded in the servo
  - used to aid in positioning the heads
  - when reading or writing needs to be done.

- When the recalibration occurs
  - you can hear the disk operate
  - as if you were reading or writing to it,
  - even if you are not.

Thermal Recalibration

- Thermal recalibration produces one **unfortunate side-effect**:
  - if you attempt to access the disk
  - while a recalibration cycle is taking effect,
  - there is a slight pause until it completes.
  - This does not cause any read or write requests to be lost,
  - but it can cause a problem if you are performing a read or write task that is operating in "real time" and might be sensitive to the delay.
- Common examples include real-time video playback,
  - audio file editing, or burning a recordable CD-ROM disk.
  - For users working with these applications,
  - thermal recalibration represents an **unacceptable problem,**
  - so the hard disk manufacturers created special drives that work around the recalibration "feature".
  - Typically by using special buffering techniques or by intelligently "scheduling" recalibration activity to avoid platters being accessed by the user, these drives essentially "hide" recalibration from the PC operator.
- They were frequently marketed as being "**audio/visual" or "A/V" drives during the 1990s.**

Thermal Recalibration

- Today, recalibration has become largely a **moot point.**
- The need for constant recalibration was greatly diminished with the creation of **embedded servo technology, which is now** the **standard for hard** disk drives.
- Recalibration had become a real "black eye" of sorts to hard drive manufacturers over the years,
- so they were happy to announce to the world that they had "eliminated" thermal recalibration in their newest units.

- This is sort of true and not true; it would be more accurate to say that it has **been greatly reduced**. **Some degree of recalibration is still required even with the newest drives**, but it does not cause the potential impact on performance that users of the old style drives had to deal with.

- **Tip:** You may still find companies trying to sell "**special A/V drives**" at higher prices than regular drives. These days, this is more likely to be an attempt to get you to part with a greater percentage of your wallet than anything else. It's not necessary to look for such designations on modern drives.

Head Parking and the Landing Zone

- When the platters **are not spinning**,
  - the **heads rest on the surface of the disk**.
  - When the platters spin up, the heads rub along the surface of the platters
  - until sufficient speed is gained for them
  - to "lift off" and float on their cushion of air.
- When the drive is spun down, the process is repeated in reverse. In each case, for a period of time the heads make contact with the surface of the disk--while in motion, in fact.
- While the platters and heads are designed with the knowledge in mind
  - that this contact will occur,
  - it still makes sense to avoid having this happen
  - over an area of disk where there is data!
- For this reason, most disks **set aside** a **special track** that is designated to be where the heads will be placed for **takeoffs** and **landings**.
- Appropriately, this area is called the **landing zone**, and no data is placed there. The process of moving the heads to this designated area is called **head parking**.

Head Parking and the Landing Zone

- Most early hard drives that **used stepper motors** did not **automatically park the heads of the drive**.

- As a safety precaution, small utilities were written that the user would run before shutting down the PC.

- The utility would instruct the disk to move the heads to the landing zone, and then the PC could be shut off safely.

- A parameter in the BIOS setup for the hard disk told the system which track was the landing zone for the particular model of hard disk.

- Usually,
- it was the next consecutive-numbered track
- above the largest-numbered one actually used for data.

Head Parking and the Landing Zone

- Modern **voice-coil actuated** hard disk drives are **all auto-parking**.
- On some disks,
  - a weak spring (opruga) is attached to the head assembly
  - that tries to pull the heads to the landing zone.
  - When power is applied the actuator is able to overpower the spring and position the heads normally.
  - When the power is shut off,
  - the electromagnetic force from the voice coil abates,
  - and the spring yanks the heads to the landing zone
  - before the platters can spin down;
  - this can sometimes be heard on older drives as an audible clunk
  - when you turn the power off.
- Other disks use a different mechanical or electronic scheme to achieve the same goal.
- Some even make use of the **rotational energy remaining** in the spindle motor to move the heads off the **data surface when the power is cut off**!
- This means that modern hard disks will automatically park their heads--even in the event of a power failure--and no utilities are required. The BIOS landing zone parameter for modern drives is ignored.

Head Parking and the Landing Zone

- Some people still think that it is necessary to **manually park** the heads of modern hard disks, **but this is not true.**

- I sometimes think of head parking utilities as the disk drive's equivalent of a screen saver. In both cases, the software was invented as a preventative measure, and one that made sense for use with the technology that prevailed at the time it was thought up.

- And in both cases, the technology has evolved to the point where utility is no longer necessary, yet many people still think it is. :^)

Head Parking and the Landing Zone

- IBM has developed an alternative to conventional head parking that I think is really a great idea.
- Instead of letting the heads **fall down** to the **surface of** the **disk**
  - when the **disk's motor is stopped**,
  - the **heads are lifted completely** off the surface of the disk
  - while the drive is still spinning,
  - using a **special ramp**.
- Only then are the disks allowed to spin down.
- When the power is reapplied to the spindle motor, the process is reversed: the disks spin up, and once they are going fast enough to let the heads fly **without contacting the disk surface**, the heads are moved off the **"ramp"** and **back onto the surface of the platters**.
- IBM calls this **load/unload technology**.
- In theory it should improve the reliability of the h eliminating most contact between the heads an unaware of any other drive manufacturers usin more about it here.
- Another feature related to reducing damage to t wear from the heads is wear leveling, which mo surface of the drive to avoid "wearing out" one s discussed in this quality and reliability section

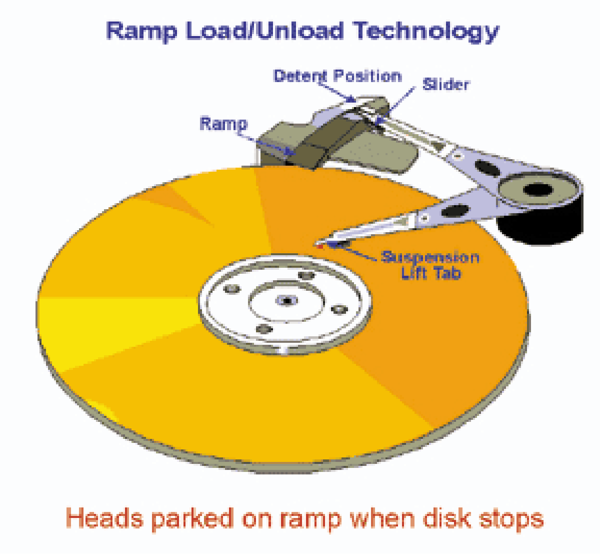

Single vs. Multiple Actuators

- One common question that folks studying hard drives often have is something like the following:
- "If the mechanical motion of the heads across the surface of the hard disk platters is so slow--relative to other components in the hard drive and the rest of the computer--why don't the hard disk manufacturers **just make hard disks with more than one actuator**?"
- It's a very good question.
- Putting a **second set of heads** in the drive
  - would allow **two areas** of the disk to be accessed simultaneously,
  - greatly improving performance,
  - particularly on random seeks.
- **Sounds great**... in fact, why not put four actuators and sets of heads on the drive, while we're at it? :^)

Single vs. Multiple Actuators

- In fact, such hard disks *have* been built.
- Conner Peripherals, which was an innovator in the hard disk field
  - in the late 1980s and early 1990s
  - (they later went bankrupt and their product line and technology were purchased by Seagate)
  - had a drive model called the **Chinook**
  - that had two complete head-actuator assemblies:
  - two sets of heads, sliders and arms and two actuators.
- They also duplicated the control circuitry to allow them to run independently. For its time, this drive was a great performer.
- But the drive **never gained wide acceptance**, and the design was dropped.
- Nobody to my knowledge has tried to repeat the experiment in the last several years.

Single vs. Multiple Actuators

- There are several reasons why it is not practical to make a drive with more than one actuator.
- Some are technical; for starters, **it is very difficult to engineer**.
- Having multiple arms moving around on a platter makes the design complex, especially in small form factors.
- There are more issues related to thermal expansion and contraction.
- The heat generated inside the hard drive is increased.
- The logic required to coordinate and optimize the seeks going on with the two sets of heads requires a great deal of work.
- And with hard disk designs and materials changing so quickly, this work would have to be re-done fairly often.

Single vs. Multiple Actuators

- However, the biggest reasons why multiple actuators designs aren't practical are related to marketing.
- The added expense in writing specialty electronics and duplicating most of the internal control components in the drive would make it **very expensive**, and most people just don't care enough about performance to pay the difference.
- Hard disks are complex technology that can only be manufactured economically if they are mass-produced, and the market for those who would appreciate the extra actuators isn't large enough to amortize the development costs inherent in these fancy designs.
- It makes more sense instead to standardize on mass-produced drives with a single actuator stack, and build RAID arrays from these for those who need the added performance. Compare a single 36 GB drive to an array of four 9 GB drives: in effect, the array is a 36 GB drive with four sets of everything. It would in most cases yield performance and reliability superior to a single 36 GB drive with four actuators, and can be made from standard components without special engineering.

Hard Disk Spindle Motor

- The spindle motor, also sometimes called the spindle shaft,
- is responsible for turning the hard disk platters,
- allowing the hard drive to operate.

- The spindle motor is sort of a "work horse" of the hard disk.

- It's not flashy, but it must provide stable, reliable and consistent turning power for thousands of hours of often continuous use, to allow the hard disk to function properly.

- In fact, many drive failures are actually failures with the spindle motor, not the data storage systems.

Hard Disk Spindle Motor

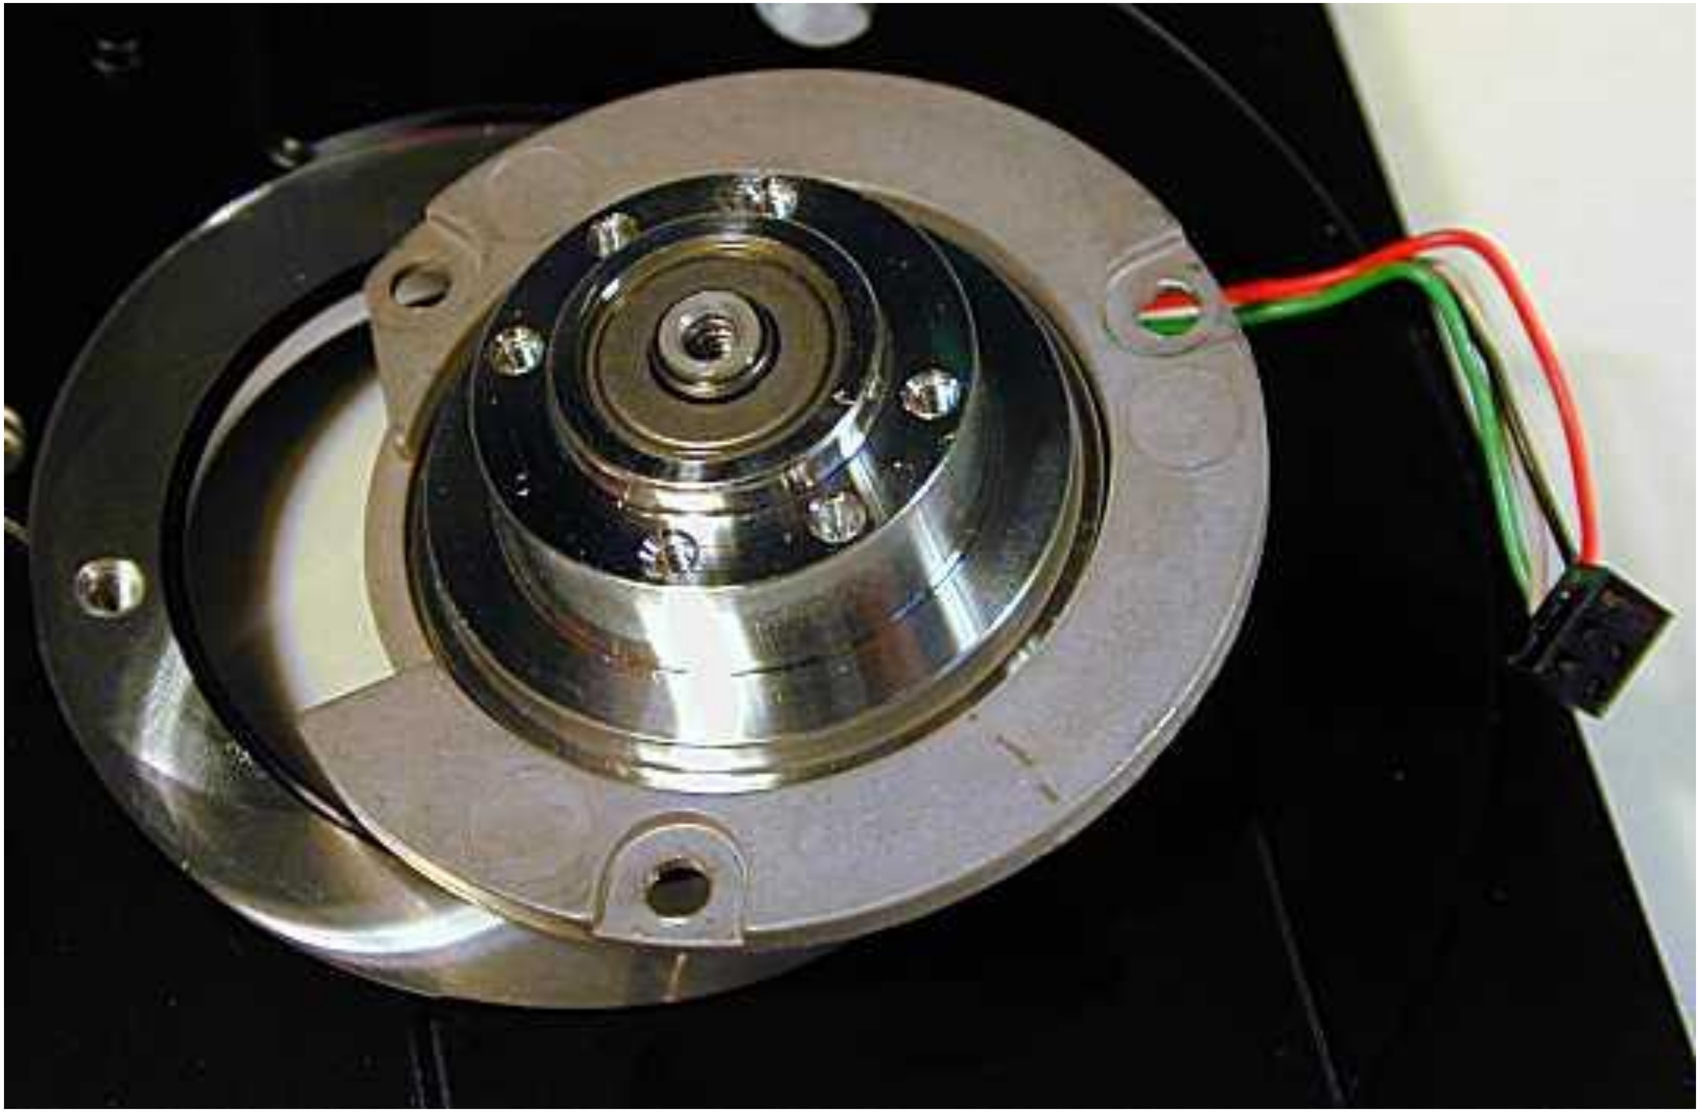

Hard Disk Spindle Motor

- A hard disk spindle motor, stripped of its platters and other components, and detached from the drive's base casting.
- You can see that it attaches with three screws around its perimeter.
- The shiny metal is the shaft, which rotates; the dull metal is the base of the motor.
- The six small screw holes on the top of the shaft are for securing the platters.
- You can also see a large screw hole in the center top of the shaft, which is used to attach the top cover to the spindle shaft for added stability. the four wire connector attaches to the hard disk logic board.

Hard Disk Spindle Motor

- For many years hard disks all spun at the **same speed.**
- In the interests of performance, manufacturers have been steadily ratcheting up their products' spin speeds over the last few years.
- These higher-speed spindles often have issues related to the amount of **heat** and **vibration they generate**.
- The **increased performance** and also the new potential issues related to the spindle motor have given it renewed attention in the last few years.

Spindle Motor Operation

- It will not surprise you, given the precision involved in every facet of the construction of the hard disk drive, that the spindle motor has several important demands placed upon it.
- **First**, the motor must be of high quality, so it can run for thousands of hours, and tolerate thousands of start and stop cycles, without failing.
- **Second**, it must be run smoothly and with a minimum of vibration, due to the tight tolerances of the platters and heads inside the drive.
- **Third,** it must not generate excessive amounts of heat or noise.
- **Fourth,** it should not draw too much power. And finally, it must have its speed managed so that it turns at the proper speed.

- To meet these demands, all PC hard disks use **servo-controlled DC spindle motors**. A **servo system** is a **closed-loop feedback system**; this is the exact same technology as is used in modern voice coil actuators, and I discuss how servo systems work in detail in that section. In the case of the spindle motor, the feedback for the closed-loop system comes in the form of a **speed sensor**. This provides the feedback information to the motor that allows it to **spin at exactly** the **right speed**.

Spindle Motor Operation

- All hard disk spindle motors are configured for **direct connection**; there are no belts or gears that are used to connect them to the hard disk platter spindle.
- The spindle onto which the platters are mounted is attached directly to the shaft of the motor.
- The platters are machined with a hole the exact size of the spindle, and are placed onto the spindle with separator rings (spacers) between them to maintain the correct distance and provide room for the head arms.
- The entire assembly is secured with a head cap and usually, lots of small Torx screws.

Spindle Motor Operation

- Components of the spindle motor assembly. Platters and spacer rings have the same inside diameter and alternated over the spindle motor axis to build the platter stack. The top cap goes, well, on top, and is secured using those annoying teeny weeny Torx screws. :^) Note that this particular drive does not have a screw hole in the center of the spindle motor shaft to secure it to the drive cover.

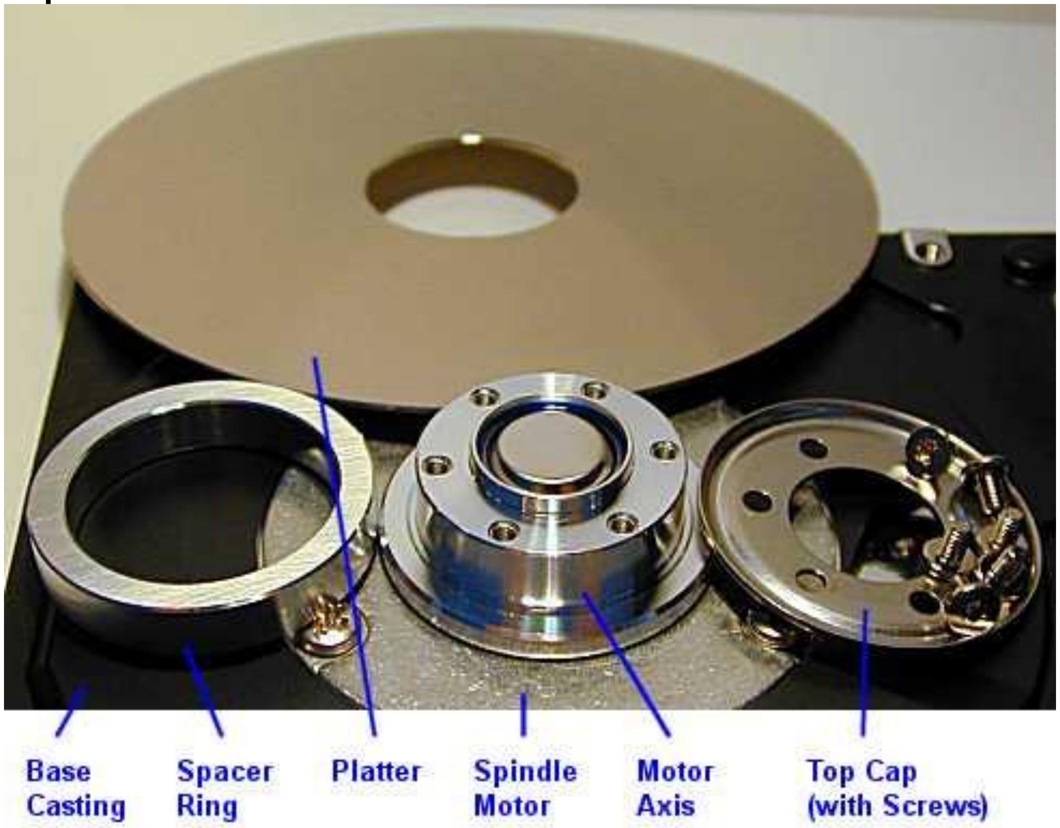

Spindle Motor Operation

- The amount of work that the spindle motor has to do is dependent on a **number** of **factors**.
- The first is the **size** and **number of platters** that it must turn.

- **Larger platters** and **more platters** in a drive mean
  - **more mass** for the motor to turn,
  - **so more powerful motors** are required.

- The same is true of higher-speed drives.

- Finally:
  - with **power management** becoming more of a concern today
  - users increasingly want hard disks that will spin up
  - from a stopped position to operating speed quickly
  - which also requires faster or more powerful motors

Spindle Motor Operation

- One important quality issue that has become a **focus of attention with newer** hard disks is the:
  - amount of noise
  - heat
  - vibration they generate.
- The reason for this becoming more of an issue is the increase in spindle speed in most drives.
- On older hard disks that typically spun at 3600 RPM, this was much less of a problem. Some newer drives, especially 7200 and 10,000 RPM models, can make a lot of noise when they are running.
- If possible, it's a good idea to check out a hard disk in operation before you buy it, to assess its noise level and see if it bothers you; this varies greatly from individual to individual. The noise produced also varies to some extent depending on the individual drive even in the same family. Heat created by the spindle motor can eventually cause damage to the hard disk, which is why newer drives need more attention paid to their cooling.
- **Tip:** Newer high-speed drives almost always run cooler and quieter than the first generation of drives at any new spindle speed. It can be painful to be a pioneer. :^)

Spindle Motor Operation

- A critical component of the hard disk's spindle motor
- that has received much attention recently due to concerns
- over noise, vibration and reliability
- is the set of **spindle motor bearings** *(leziste)*.

- Bearing are precision components that are placed around the shaft of the motor to support them and ensure that the spindle turns smoothly with no wobbling or vibration.
- **As hard disk speeds increase**, the **demands placed on** the **bearings increase dramatically**.
- Many of the noise and heat issues
  - created by fast motors
  - are related to the bearings
  - so engineers are constantly trying to improve them.

Spindle Motor Operation

- Most hard disk motors use **ball bearings**.
  - These are **small metal balls**
  - that are placed in a ring around the spindle motor shaft;
  - you have no doubt seen them used in many different applications outside the PC.
- They are also used elsewhere inside the PC,
  - such as higher-quality power supply fans.
- Some hard disks use special **fluid-dynamic bearings**
  - instead of ball bearings.
  - Here, the metal balls are replaced with a thick oil,
  - which reduces noise significantly
  - because the metal-to-metal of ball bearings is removed.
  - It also theoretically greatly increases bearing life,
  - though ball bearings should have a life exceeding the hard disk's normal service life anyway.

Continuous Power vs. Peak Power at Spin-Up

- Most of the power used by a modern hard disk drive
- is taken up by the spindle motor.

- Due to their smaller,
  - more efficient designs,
  - newer drives require relatively little power
  - to keep their platters spinning continuously.

Continuous Power vs. Peak Power at Spin-Up

- Even drives with faster spindle speeds take less power than the large, inefficient motors of the drives of a decade ago.
- However, when the hard disk is first started up,
  - the motor can draw a peak level of power
  - **that is more than two times**
  - what it takes to keep the disks spinning.
  - (Most things require more energy to start them in motion than to keep them in motion, but in the case of hard disk motors this is especially so because of their electrical characteristics.)
- While in most cases even the peak start-up power usage isn't all **that much**, there can be an issue when you are **using multiple hard disks** that all try to spin up at once when you turn on the machine.
- The ratings of your power supply, particularly for the +12 V level, must be sufficient to take this initial demand into account; these matters are discussed in this section of the power supply reference

Continuous Power vs. Peak Power at Spin-Up

- 2V power profile (current vs. time) of an IDE/ATA hard disk **at startup**. You can see that the peak power draw is over quadruple the steady-state operating requirement. The graph appears "noisy" due to frequent oscillations in current requirements

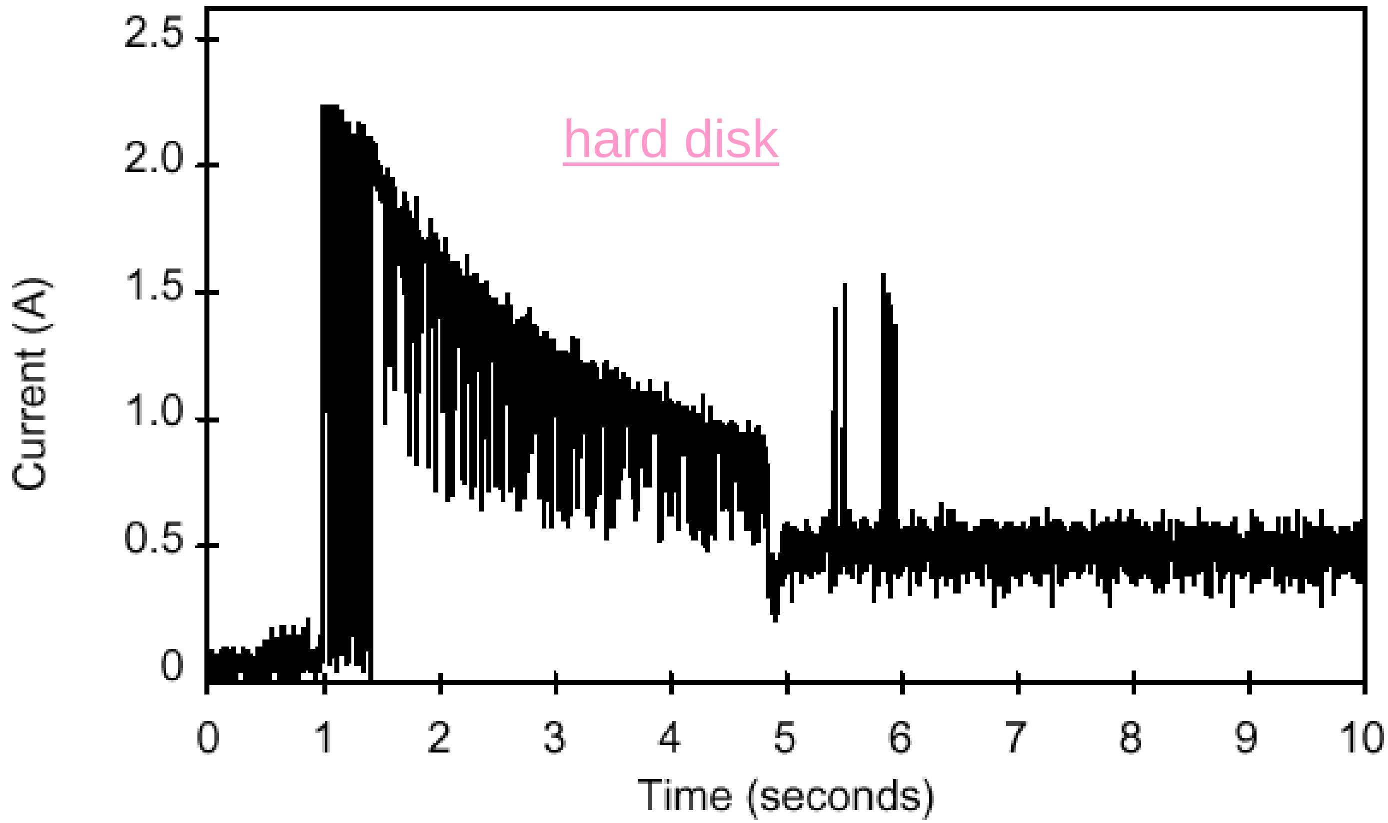

Continuous Power vs. Peak Power at Spin-Up

- Hard drive manufacturers generally program their drives
  - so that when two are used on the same IDE channel
  - as master and slave,
  - the slave drive delays its spin up by several seconds
  - to offset its start from that of the master drive
  - and reduce the peak demand on the power supply.

- Similarly, many SCSI drives
  - can be programmed to delay their startups,
  - using the "Remote Start" command or a special jumper.
  - These features offset the time
  - that the drives spin up to reduce the peak demand on the power supply.

Continuous Power vs. Peak Power at Spin-Up

- Another issue with spin-up relates to boot speed.
- In order for the system to boot up properly when it is turned on,
  - the **platters** must be brought from a **standstill to operating speed**
  - in a relatively short period of time.

- Modern PCs often have "expedited" BIOSes that go from power on to attempting to boot the hard disk in only a **few seconds.**

- If the disk isn't ready, often the system will hang, even though it will boot if you press the reset button and start the boot process over again.

- Some system BIOSes include a delay feature that can be used to slow down the boot process to avoid this problem.

- Otherwise, slowing down the boot process by turning on the memory test feature, for example, can sometimes cure this problem.

Power Management

- Power management is a set of protocols
- used to allow PCs to reduce the amount of power
- that they consume, especially when they lie idle.

- Hard disk drives are always spinning during operation,
- and the spindle motor is what is using the bulk of the power consumed by the drive as a whole.

- Therefore, many power management schemes
  - include a mechanism
  - by which the hard disk's spindle motor can be "spun down"
  - after a certain amount of inactivity
  - and then can be "spun up" (reactivated)
  - when the system needs them again

Power Management

- There is some controversy surrounding this feature.
- First, there is some debate
  - as to whether or not dramatically increasing
  - the number of times the drive spins up and down
  - (which power management does) is detrimental to the drive
  - or could cause it to fail prematurely.
- Second, there is some skepticism about
  - how much energy is really saved,
  - given that the hard disk doesn't use that much power
  - under normal operation to begin with.
- See here for a discussion of pros and cons of power management.
- Certainly for laptop users power management is an important feature; all laptop hard disks support power management via BIOS settings and/or operating system controls.