Computer Components

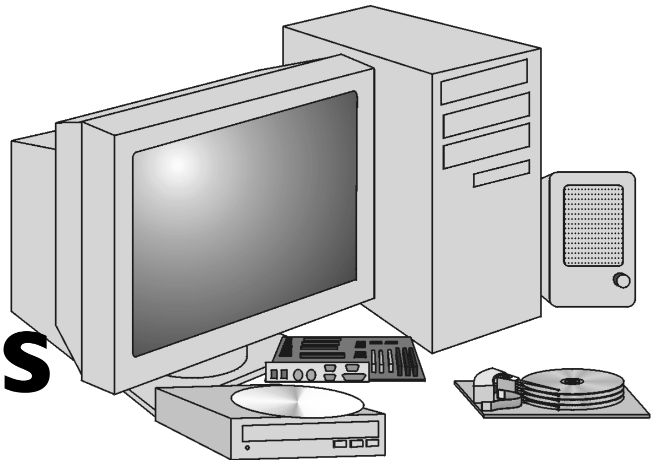

G. Snider

Input-Output Model

The heart of any computer system is the processor. The processor receives input from a variety of components and outputs it in many different forms.

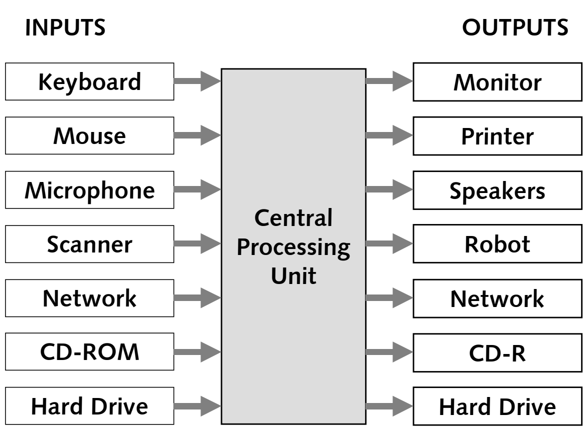

Electrical buses

Data is moved in and out of the computer on buses: groups of parallel conductors which can carry several bits of data at a time. The

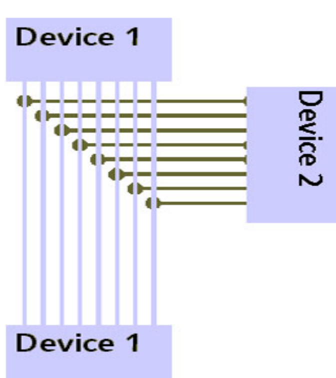

conductors in an electrical bus can be compared to the lanes of a superhighway. The more lanes of traffic, the faster data can be moved from one place to another.

The Processor

The processor is the part of the computer which interprets and carries out instructions from the computer programs. Modern desktop computers use microprocessors: complex integrated circuits containing millions of transistors and other electronic components. The CPU is mounted in a socket or slot on the mainboard (motherboard) so that it can be connected to other components.

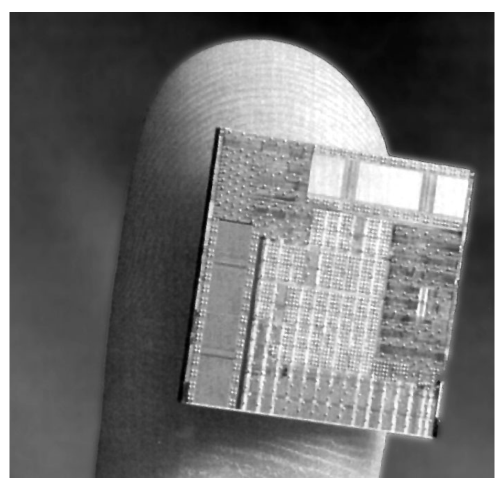

Microprocessor Specifications

- *Width of external data bus in bits*
- *Width of address bus in bits*
- *Width of internal data bus in bits*
- *Internal Cache Memory (Level 1 or L1)*
- *External Cache Memory (Level 2 or L2)*
- *Clock Speed in Megahertz (MHz)*

- *MIPS (Million Instructions per Second)*
- *Power Consumption (Watts)*

Width of Data bus: the earliest microprocessors could handle data only in bytes (8 bits). As the width of the data bus increases, the width of the data bus determines how much information can be moved in or out of the processor in one operation. It also determines the number and length of instructions which can be used.

Address bus: the width of the address bus determines how much memory can be addressed. The Intel 8086 had a 20 bit address bus. Since a 20 bit binary number can represent 2^{20} different numbers, the 8086 could address 1MB of memory. The 386 and 486 could address 4GB of memory using a 32-bit address bus, and the Pentium class processors have a 36-bit address bus capable of connecting 64GB of memory.

The width of the internal data bus and the storage registers may differ from the external data bus. The 386SX processor, for example had the same internal 32-bit registers as other 386s, but only a 16-bit external bus. The Pentium class processors, have an external 64-bit data bus, but the internal registers are only 32-bit. For this reason the Pentium processors are referred to as 32-bit. Workstation processors like the SPARC and Alpha are 64-bit, as is the next generation of Intel chips, the Itanium.

Clock Speed: An oscillator mounted on the motherboard generates a series of electrical pulses which the computer uses to synchronize the operations of its many components. Each complete change in the signal, from positive to negative and back again is known as a cycle, and the number of cycles per second, or frequency, is measured in Hertz. 1000 HZ = 1KHz, 1000000 Hz = 1 MHz, 1000000000 Hz = 1 GHz. The speed of the processor is often a multiple of the external bus speed: for example a 500 MHz chip installed on a 100MHz mainboard will operate at 5x the bus speed.

MIPS: The clock speed does not relate directly to the speed at which the CPU processes instructions. Early microprocessors required as many as 10 clock cycles to complete a single instruction. Modern microprocessors with what is called 'superscalar" architecture have dual or multiple 'pipelines' so that more than one instruction can be executed at once. Therefore, a more accurate measure of processor speed is MIPS (Millions of Instructions per Second), although the number of actual instructions processed rarely reaches the theoretical maximum.

Power consumption is an often overlooked measure of microprocessor performance. Much of the power consumed is given off as heat, which must be dissipated, or it will cause malfunctions. Low power consumption is also a critical factor in extending the life of batteries in notebook computers. The formula for power consumption is Volts x Amps = Watts. However, a decrease in operating voltage also produces a drop in amperage (Ohm's Law). Older chips functioned at 5 volts, while Pentium class chips operate in the range of 2 volts. This results in a power saving of 84% without any other improvements in the circuitry.

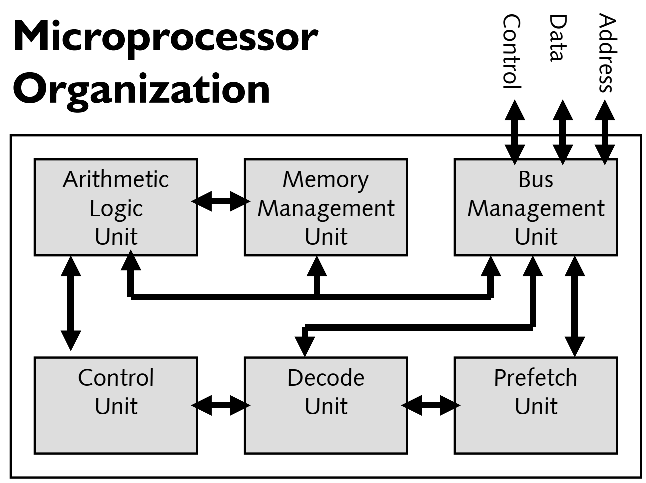

Arithmetic Logic Unit: executes all logic and arithmetic operations as specified by the instruction set

Control Unit: contains the microcode that tells the ALU how to function.

Decode Unit: translates instruction into control signals and microcode directions then queues them until requested by the Control Unit.

Prefetch Unit: queues instruction to assure that the microprocessor is in continuous operation

Memory Management Unit: converts internal logic addresses into external memory addresses.
Bus Management Unit: manages the flow of information into and out of the microprocessor with memory sources (main disk, CD-ROM, etc.) and peripherals (printer, display, etc.)

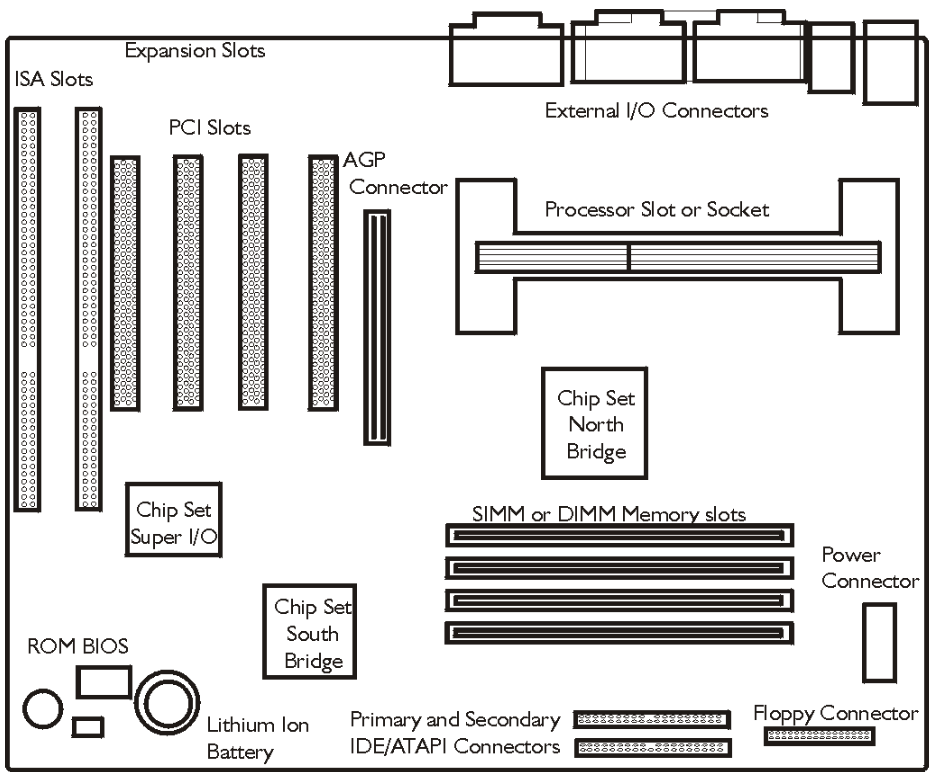

Components found on motherboard:
Memory modules
System clock
Expansion buses
Real-time clock
Disk Interfaces
Chip Set
BIOS
Input/Output Ports

Mainboard

The mainboard or motherboard connects the microprocessor to all the other components of the computer.

System Clock

•All operations in the computer must be precisely synchronized.
•The system clock generates a series of pulses known as a square wave.
•The frequency of the pulses is controlled by the oscillations of a tiny sliver of quartz.
•Quartz crystals are used because they vibrate at a constant speed, which depends on the size and shape of the crystal.

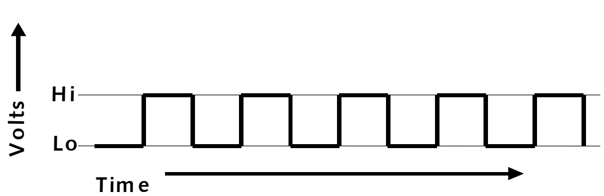

Frequency Units

The unit of Frequency is the Hertz, named after a 19th century German physicist.

1 cycle per second = 1 Hertz
1 000 Hz = 1 kilohertz (kHz)
1 000 000 Hz = 1 Megahertz (MHz)
1000 000 000 Hz = 1 Gigahertz (GHz)

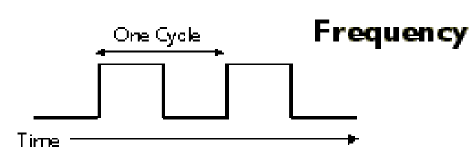

Timing

It is essential that operations within the computer be synchronized. Some events may be triggered by the leading edge of the clock pulse, others are triggered by the trailing edge. Clock

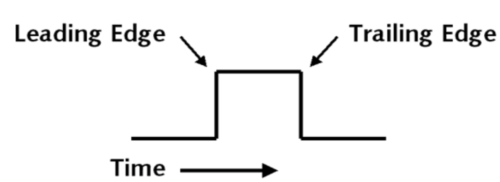

doubling occurs when events are triggered twice in one clock cycle.

BIOS (Basic Input/Output System)

The BIOS originated when Gary Kildall sought a method of adapting CP/M to run on systems made by different manufacturers. The BIOS is a set of device drivers, some stored in ROM on the motherboard or expansion card, others loaded by the operating system. The BIOS is customized for each machine so that the operating system can complete the same tasks on different hardware.

Hardware/Software Interface

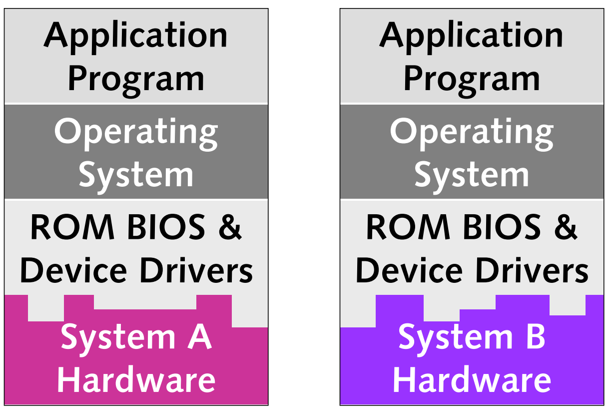

Types of ROM

•ROM - Read-Only Memory
•PROM - Programmable Read-Only Memory
•EPROM - Erasable PROM
•EEPROM - Electrically Erasable PROM (often called Flash ROM)
•Motherboards with Flash BIOS can be upgraded

RAM Memory

RAM (Random Access Memory) is used as the temporary memory space, or workspace of the computer. It is where the operating system, application programs and documents currently in use are stored. RAM is semiconductor memory, and as such, it needs continuous power: it is erased when the power is shut off. (Computers use hard disks, floppy disks, CD-ROMS and other media for permanent storage.) RAM cards or modules are installed in slots on the mainboard.

RAM Modules

The earliest microcomputers had RAM mounted directly on the mainboard or on expansion cards. To make RAM easier to install or replace, it is now mounted on small printed circuit boards and installed in dedicated slots on the motherboard.

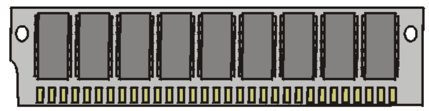

30 pin Single In-Line Memory Module (SIMM)

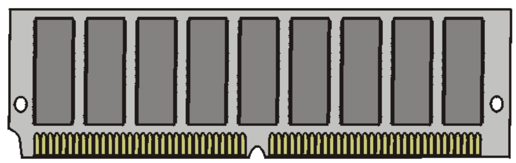

72 pin Single In-Line Memory Module (SIMM)

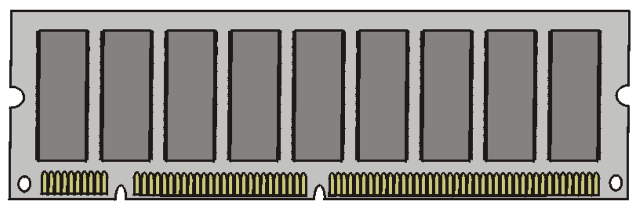

168 pin Dual (double-sided) In-Line Memory Module (DIMM)

Bits, Bytes and Megabytes

RAM memory is measured in bytes.
1 byte = 8 bits
1 kilobyte = 1024 bytes (210)
1Megabyte = 1 048 576 bytes (220)
1Gigabyte = 1 073 741 824 bytes (230)
1Terabyte = 1 099 511 627 776 bytes (240)

Types of RAM

DRAM - Dynamic RAM uses one transistor and one tiny capacitor to store each bit as a voltage. The capacitor will store the voltage only for a few microseconds unless it is recharged regularly. SRAM - Static RAM uses a circuit of six transistors to store each bit of data. It is much faster, because it does not have to be refreshed, but also more expensive. It is used for cache memory operating at the same speed as the processor.

Transfer Rates

The data transfer rate is a measure of the data transferred in one second. The transfer rate is the product of frequency and bus width.
Example:
100 MHz x 32 bits = 3200 bits/second = 400 bytes/second

Access Time

Access time is the length of time that it takes for the system to find and retrieve data in RAM or on disk. RAM access times typically range from 60 nanoseconds (DRAM) to 10 nanoseconds (SDRAM).Hard disk access times are measured in milliseconds. Access time is inversely related to the clock speed at which RAM can operate.
Example:
1 second/60ns = 16.67 MHz

Types of DRAM

Several improvements have been made to increase the speed of DRAM.
•Fast Page Mode RAM organizes data into consecutive groups or pages to shorten access time.
•EDO (Extended Data Out) RAM speeds access times by reading the next address while the previous data is being transferred.
•Synchronous DRAM operates at the same speed as the processor bus. It is rated by the speed it can maintain, for example PC/100 SDRAM can be used with 100 MHz bus.

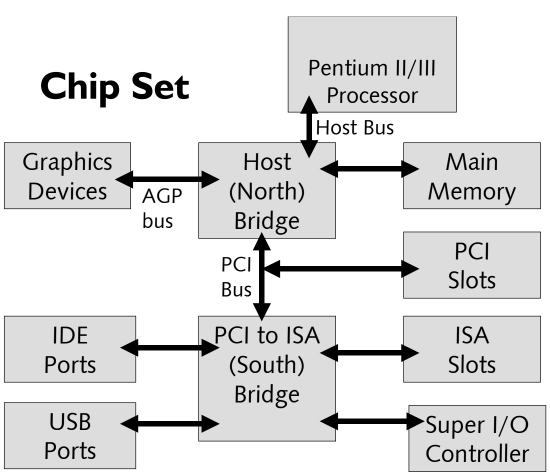

RDRAM

Rambus DRAM uses a combination of techniques to triple the effective transfer speed. SDRAM transfers and stores 64 bits of data at one time. Rambus DRAM has a 16-bit external interface, but it can operate at 800 MHz.
Data is transferred on both the leading edge and trailing edge of the clock pulse.

Chip Set

The chip set is a group of chips which act as the interface between the microprocessor and the rest of the devices connected to the computer. The devices found in these chips include the clock generator, system timer, interrupt handler, DMA controller, CMOS RAM, Real-time clock, and bus controller.
In today's computers, the functions are contained in two chips known as the North Bridge and the South Bridge. The North Bridge connects the processor to main memory, the video display card, and the PCI bus. The second chip provides a bridge from the 33MHz PCI bus to the slower devices such as ISA, IDE and USB. The South Bridge also connects to the Super Input Output Controller chip, which interfaces with the basic input and output ports, for attaching external devices, including the floppy drive, keyboard, mouse, modem, and printer.

Super Input/Output Controller

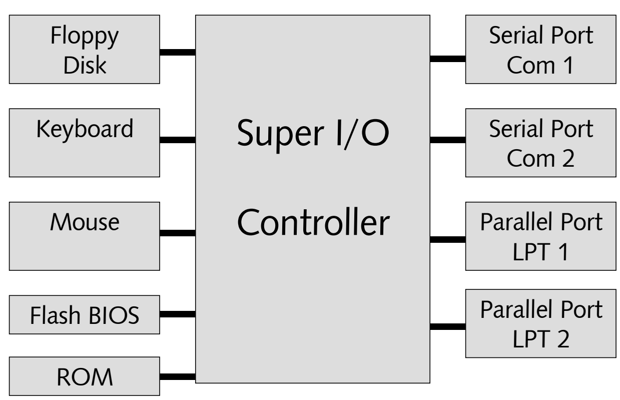

Common Expansion Buses

The constantly increasing power and speed of the microprocessor has created a need for improved buses to connect external components. Some, such as IBM's MCA and the VESA local bus, proved to be dead ends. The most common buses still in use are:
•ISA (Industry Standard Architecture)
16-bit, 8MHz (soon to be history)
•PCI (Peripheral Component Interconnect)
32-bit, 33MHz
•AGP (Accelerated Graphics Port)
32-bit, 66/133/266 MHz
•Processor Bus
64-bit, 66/100/133 MHz

IDE Interface

•Stands for Integrated Drive Electronics - meaning the hard-disk controller is built in to the drive.
•Official ANSI standard is called ATA, for AT Attachment.
•Enhancements include:
•ATA-2 or EIDE •ATA-3
•ATA-4 (UDMA 33) •ATA-5 (UDMA 66)

IDE Enhancements

•faster transfer using Direct Memory Access and Ultra DMA
•33MB/second transfer rate with UDMA 33, 66MB/second with UDMA 66 (requires special 80-conductor cable)
•BIOS support for drives over 528MB and 8.4GB limits.
•PC Card support
•ATAPI - ATA Packet Interface for CD-ROM and tape drives.

IDE Connections

•Hard-disks and CD-ROM drives are connected using a 40-conductor ribbon cable which plugs into 40-pin connectors
•Most main boards have a primary and secondary IDE connector.
•A total of 4 drives may be connected, two on each of two cables.

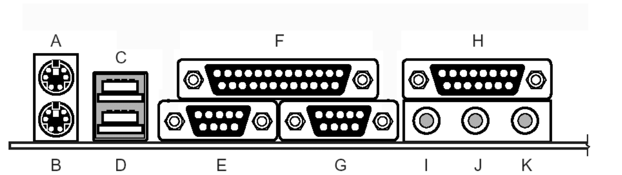

Standard arrangement of input/output ports at the rear of Celeron/Pentium III/Pentium 4

•Each drive must be designated a master or slave using jumpers on the drive case.

SCSI

The Small Computer System Interface (Scuzzy) was originally adopted by Apple for the Macintosh. It can support up to 15 devices on a single bus. Each device must have a binary ID, set by jumpers or software, so that it can be recognized by the controller. SCSI provides data transfer rates up to 80MB (Ultra2/Wide SCSI) All SCSI cables require a terminator, which may be a separate device or built into the last device in the chain.

SCSI Standards

	MHz	Bits	MB/Sec	Devices	Pins
Standard	5	8	5	7	50
Wide	5	16	10	15	68
Fast	10	8	10	7	50
Fast/Wide	10	16	20	15	68
Ultra	20	8	20	7	50
Ultra/Wide	20	16	40	15	68
Ultra2	40	8	40	7	50
Ultra2/Wide	40	16	80	15	68
Ultra3	40x2	8	120	7	50
Ultra3/Wide	40x2	16	120	15	68

System Resources

Hardware Interrupt Channels (IRQs)
are used by most devices to signal the CPU when they have data to transmit. Only 15 are available, of which several are reserved for system timer, keyboard controller, floppy disk

controller, real-time clock, co-processor, and IDE controller. Interrupt conflicts occur when two devices attempt to use the same IRQ at the same time. The PCI bus reduces interrupt conflicts by using IRQ steering to share interrupts. DMA (Direct Memory Access) channels are used by high-speed devices such as disk drives and sound cards. I/O Port Addresses are the addresses assigned to the devices connected to the bus; 64 000 addresses are available.

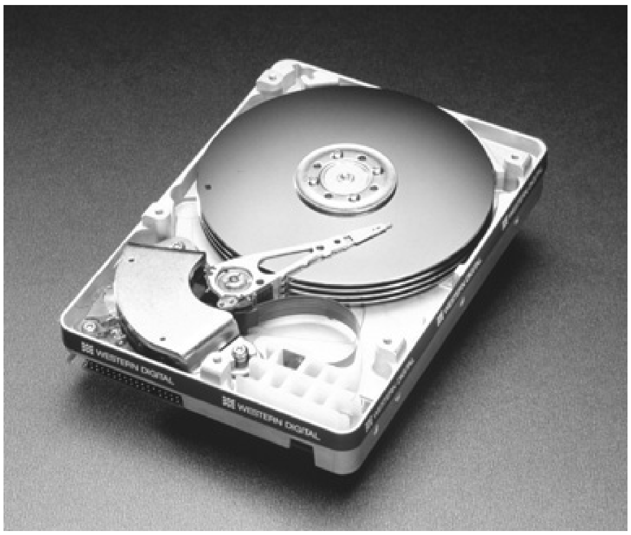

Magnetic Storage

Secondary or mass storage of programs and data requires systems which can retain information when the power is off. The most common methods of storing data, since the earliest days of computing, have involved electromagnetism.

Electromagnetism

Magnetic data storage relies on the relationship between electricity and magnetism. When an electrical current flows through a wire, it produces a magnetic field around it. If the wire is wound into a coil, each turn of wire adds to the strength of the magnetic field. This principle is used to create electromagnets and electric motors.

This relationship works in reverse as well. When a wire—or coil of wire— is moved through a magnetic field an electrical current is created in the wire. An electrical current is also generated when the magnetic field surrounding a conductor changes in strength or polarity (direction).

Floppy disks and hard drives are similar in the way in which data is stored. The disks are made from a substrate (underlying layer) of a nonmagnetic material—Mylar plastic for floppy disks, aluminum or glass for hard disks—coated with a layer of iron oxide. The magnetic particles in the coating media are initially arranged in a random fashion; think of playing a

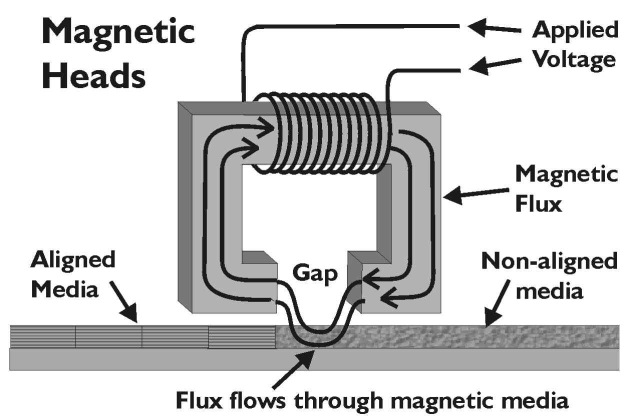

blank videotape on which nothing has yet been recorded. The result is only snow on the screen. As the recording head passes over the disk, a tiny electromagnet creates a magnetic field which causes the magnetic particles in the media to align themselves, an remain in that pattern after the head has passed.

Write Current

The direction of the current through the coil determines the direction of the magnetic flux. Reversing the polarity of the electrical current reverses the direction of the magnetic field on the disk surface.

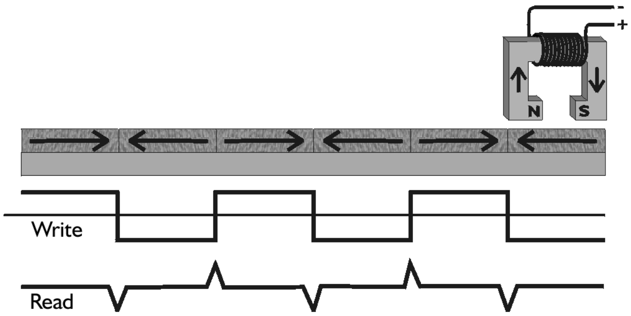

Read Current

Current is generated in the coil only when the magnetic field changes. A current spike is produced each time the head passes over a transition from one magnetic polarity to another.

Data Synchronization

Synchronization between the rotating disk and the heads is critical to reading and writing data. Imagine asking a band to hold a long sustained note and then hit the next note at exactly the same time, without a conductor to keep time.

Data Encoding

To ensure synchronization, a clock signal must be encoded with the data. Encoding is based on flux transitions rather than on the polarity of the magnetic field. Run-Length Limited (RLL) encoding is the most common for modern hard disks.

RLL 2,7 Encoding

Data Bits	Flux Encoding
10	NTNN
11	TNNN
000	NNNTNN
010	TNNTNN
011	NNTNNN
0010	NNTNNTNN
0011	NNNNTNNN
T = Transition N = No Transition	

T = Transition N = No Transition

The code is designed so that there is a minimum length and maximum limit to the number of cells between transitions. For example, RLL 2,7 encoding has a minimum of two spaces between transitions, and a maximum of 7 spaces between transitions. The string 000000000 would be encoded NNNTNNNNNTNNNNNTNN, ensuring that the there are no long gaps between transitions.

Hard Disks

Hard disk technology was pioneered by IBM researchers, beginning with the RAMAC disk drive in the fifties, and including the Winchester drive in the sixties, and magneto-resistive heads in the nineties. Floppy disk drives were a later derivative of hard disk technology. A typical hard disk consists of two to eight rigid disks with data stored on both sides of each disk. Each side has its own read-write head mounted on an armature so that it can be positioned anywhere over the surface of the disk. The heads rotate around a central spindle at speeds of 3 600 to 12 000 RPM.

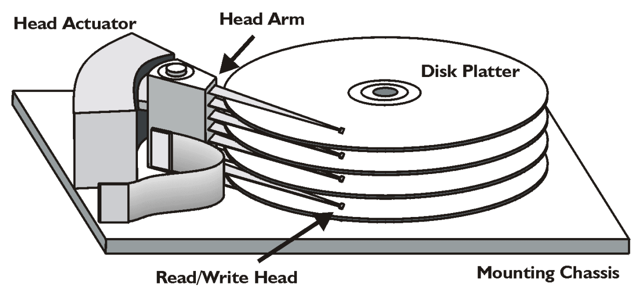

The position of the head arm in hard drives is controlled by a device known as a *voice coil actuator*. The actuator is an electromagnetic coil which moves the head arm magnetically. The armature controller receives feedback from positional data written on the disk when it is manufactured. In some systems an entire side of one platter is used for positional information, but in most disks, the positional code is written in wedges between the data sectors.

The positional information is encoded using a variation of binary known as the Gray code which is frequently used in applications such as robotics and automation, where constant feedback on the position of parts of the machine is necessary.

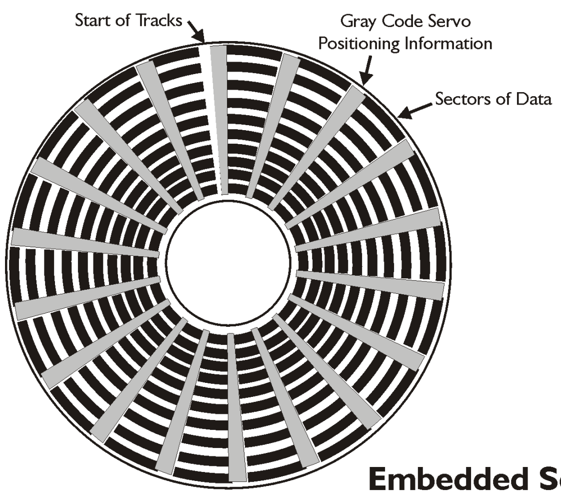

Embedded Servo

In binary code, the number following 0111 is 1000: every bit must change, although the difference is only 1. The numbers in Gray code are chosen so that only one bit changes between each adjacent number in a sequence. That reduces the likelihood of errors in the transmission of information. The four-bit Gray code is shown on the table below.

Decimal	Gray code	Decimal	Gray code
0	0000	8	1100
1	0001	9	1101
2	0011	10	1111
3	0010	11	1110
4	0110	12	1010
5	0111	13	1011
6	0101	14	1001
7	0100	15	1000

Head Technology

The heads are aerodynamically designed so that a cushion of air separates the heads from the surface of the disk at all times.

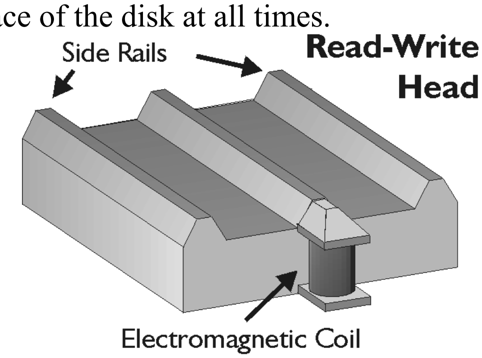

Dimensions: 2mm x 1.5mm x 0.5mm

In order to reduce the size of the heads and increase the density of data on the disk, the original ferrite core/wire coil design has been replaced. Currently, hard disks have separate read and write heads on each head arm. Thin Film heads are manufactured in a manner similar to ICs. The core is a nickel-iron alloy and the gap is sealed with aluminum. They are now used primarily for writing data. Magneto-resistive heads give exceptional read sensitivity. A small sense current flows constantly through the head. The conductivity of the head changes in response to changes in the magnetic field, in turn altering the sense current.

Hard Disk Controller

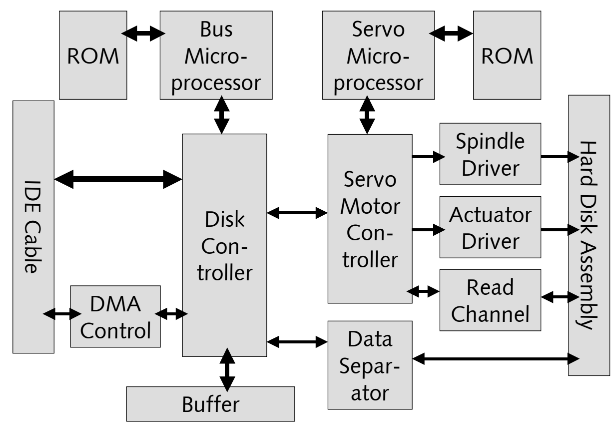

Hard Disk Controller

Each hard disk has a printed circuit board attached which interfaces between the disk and the remainder of the system. The tasks which must be performed by the controller are so complex that it usually has two dedicated microprocessors on the board. The Servo controller must process information from the FAT and the positional data from the embedded servos and send signals to the head actuator to position the heads correctly on the disk. During a write operation, the second microprocessor receives binary data through the IDE cable, encodes it using RLL or a similar method, and then sends it to be recorded on a blank portion of the disk.

During the read cycle, the data must be changed back to binary format before it is transmitted over the IDE cable to the motherboard. Hard disk controllers now also include substantial RAM memory on the board to serve as a buffer, or temporary storage area.

Installation

When installing an IDE hard drive, it must be connected to the mainboard through a data cable. The first hard disk should be connected to the primary IDE connector. A power cable must be connected which provides 5-volt power to the controller and 12-volt power to the spindle motor. There are jumpers on the rear of the drive next to the power and data connectors, which must be correctly set. Two drives can be connected to each IDE connector on the mainboard using a double-headed cable. One drive must be configured using the jumpers to be the "master", while the second is the "slave". As a rule, CD-ROM drives should be connected to the secondary IDE connector, since they use a different protocol (ATAPI) for data transfer.

Formatting

In order to store the data on the disk in such a way that it can be retrieved, the disk must be divided into sections. The path that the heads follow as the disk rotates are known as tracks. The tracks are divided into sectors during a process known as low-level formatting.
The operating system performs a high-level format of the disk. In this process, the disk is first scanned for bad sectors, which are marked so that they will not be used for storing data. The first sector on the first disk side is reserved for the boot sector, which is where the information regarding the size and partitioning of the disk is stored. If this data is damaged, the disk is rendered unusable. An index of the file locations on the disk—known as the *File Allocation Table*, or FAT—is created and stored in the second sector of the disk. The FAT enables the operating system to create files and directories on the disk. If the disk is to be used as the primary (C:\) drive, then the first files created on the disk must be the operating system files. For DOS, there are three essential files: IO.SYS, MSDOS.SYS, and COMMAND.COM. Since the file system is created by the operating system, it is unreadable by incompatible operating systems. UNIX, LINUX, MAC/OS, each use a unique file system. FAT32, which makes more efficient use of disk space, and can be used with larger hard drives, was introduced as an add-on to Windows 95, and is used in Windows 98 and SE. Windows NT, Windows 2000 and Windows XP use NTFS.

Partitioning

A hard drive may be subdivided into smaller sections known as *partitions*. Usually this is done using FDISK before the disk is formatted, because FDISK erases the FAT, effectively erasing any data stored on the disk. Programs such as Partition Magic, however, can change the number and size of partitions on a disk while preserving the data. The subdivisions on the physical disk, which appear to the operating system as separate disks, are known as logical drives. Generally, only one drive—the primary drive—can be bootable. However, boot manager utilities can circumvent this limitation and create two or more bootable drives by hiding the other drive[s] from the one which is in use.

Drive Specifications

The most common specification for hard disks is the *storage capacity* in Megabytes or Gigabytes. (Some manufacturers, when reporting the size of the drive, use 1GB = 1 000 000 000 bytes. When it is formatted by the operating system, the size is reported using binary, where 1GB = 2^{30} = 1 073 741 824. As a result the size reported by Windows will be less than the advertised capacity.) The time it takes the processor to find data on the disk—the *seek time*—is usually in the range of 8ms, far longer than it is for RAM. Once the data is located, however, it can be moved to RAM or to another device at up to 30MB/sec or more—this is called the *data transfer rate*. The factor which has been creating larger and larger drives in recent years is *areal density*: the amount of data in a given area. Hard disk density is now greater than 10 Gigabits per square inch.