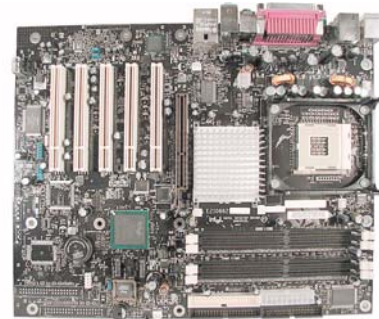# Chapter 2
Components of a Personal Computer

Every IT professional will be faced with preparing budgets, ordering hardware and software, and contracting for repair services. Thorough familiarity with hardware and software will be necessary to handle these tasks. The hardware portion will be covered in this chapter as well the Computer Organization and Architecture chapter.

The Motherboard

The Central Processing Unit (CPU) must be connected to memory and the input/output (I/O) modules. The most convenient way of doing this will be designing a printed circuit board having a slot for the CPU, memory and interface cards, and buses for the data, control and address signals. Such a circuit board is called a motherboard. We will be discussing more about these buses and interfaces in the Computer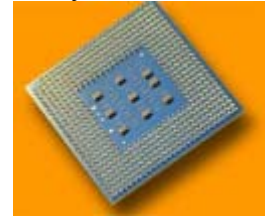 Organization and Architecture chapter. A motherboard is made to service a particular series of CPUs. Therefore, when selecting a motherboard we must assure that it will support the CPU in question. The buses are visible from the bottom side of the motherboard. These buses are specified electrically and mechanically by the number of pins on the connector, and by the voltage used and type of interface cards that plug into the motherboard.

On the motherboard in this figure only two types of cards can be plugged in, five PCI and one AGP. Motherboards today come with built-in PS/2 mouse, parallel, serial, USB and keyboard ports and floppy, hard drive (IDE) connectors. Just a few years ago a multi I/O card had to be installed to make some of these ports available. The 12V connector plugs into the motherboard that supplies power to the motherboard, CPU, interface cards and the ports. Additional power connectors are available on the power supply for the peripherals. The CPU socket is a surface mount, ZIF (zero insertion force) socket and accommodates the processors. Every motherboard has a chipset that coordinates the busses and the interfaces. Different options or features of the motherboard and the chipset can be changed or activated by either setting the DIP (dual inline package) switches or placing a jumper over certain pins, or by programming the flash EEPROM and setting the BIOS.

The motherboard is placed inside the case and screwed down, ensuring that no active components or wires of the motherboard are touching the case. Once the CPU, memory and the video card are installed the computer could be powered on to setup the features of the BIOS. After installing desired drives and peripherals the computer is ready to be used.

The Central Processing Unit (CPU)

There are many microprocessor chips on the market today manufactured by Intel, Motorola, IBM, Sun, and AMD to name a few. Intel sells the most microprocessors today. As of November of 2002 Intel introduced the Intel® Pentium® 4 Processor with Hyper-Threading (HT) Technology. HT Technology enables the processor to execute two threads of a program in parallel. Prior Pentium 4 chips did not have the Hyper-Threading. Intel Celeron processor is the lowest priced entry level processor they sell. Intel Itanium and Xeon processors are designed for the server market. AMD markets their Duron and Athlon processors to compete with Intel Celeron and Intel Pentium 4 chips. AMD's Opteron chip competes with Intel's Xeon and Itanium chips. Much of the computer use can be categorized either as CPU intensive or I/O intensive. If most of your computer time is CPU intensive then investing in the fastest chip available will be the way to go. In case of I/O intensive programs such as a database, other optimizations would be more worthwhile. There will be more architectural discussion about CPUs in later chapters. The heat produced by the CPU should be dissipated using a heat sink and a fan to avoid damage to the CPU and other components in the computer. Different types of heat sinks and fans are available to fit the various CPUs. Frequent computer freeze up is an indication of a failing CPU fan. It is a good idea to invest a few extra dollars to buy a good quality CPU fan to protect the most important part of your computer, the CPU.

The speed in each family of CPUs is given in terms of Megahertz or Gigahertz. An Intel Celeron 2.0 GHz CPU does not run at the same speed as a Pentium 4 2.0 GHz CPU. If it takes 0.01 seconds for a clock cycle, it is easier to say 100 hertz. Alternating current in the USA reverses polarity 60 times a second, or it takes 1/60 of a second for one cycle. 60 Hz, 1/60 of a second, or 60 cycles per second can be used as units of measurement, hertz being the standard unit of measurement of frequency. 1000 hertz is a kilohertz, 1000 kilohertz is a megahertz, and 1000 megahertz is a gigahertz. The occurrence of events in theCPU is determined by a clock that transmits regular sequences of alternating 1s and 0s of equal duration. A single transmission of the set (0,1) is called a clock cycle. The processor and the busses on the motherboard operate at different speeds. A clock multiplier is added to the system bus to obtain the CPU speed. For example, if the motherboard runs at 533 Megahertz and the CPU speed is 2 Gigahertz then a multiplier of 4 will be used. This needs to be set on the motherboard using jumpers provided; however, the newer CPUs and motherboards automatically configure the multiplier preventing "overclocking." Personal computers have come a long way from 4.77 kilohertz in 1984 to over 3 gigahertz in 2003.

A single operation may take multiple clock cycles. During each clock cycle a single event takes place such as fetch, decode, execute, memory access, and write back. Breaking down an operation to many discrete events like this provides for pipelining, which will be discussed in a later chapter.   The performance of a computer depends upon the clock speed of the CPU and buses, width of the buses, stages in pipelining, cache performance, addressable memory, miss rate, hit rate, instruction mix of the program, and many other such variables. Measuring and reporting performance are the focus of much

research. There are many benchmark suites such as the SPEC 2000 that can be used to compare performance of different computers.

Two laws come to mind while we are on the topic of CPUs: Moore's Law that predicts the means by which performance can be improved, and Amdahl's Law that calculates the improvement in speedup. Dr. Gordon E. Moore of Intel recalled "I first observed the *'doubling of transistor density on a manufactured die every year'* in 1965, just four years after the first planar integrated circuit was discovered. The press called this 'Moore's Law' and the name has stuck. To be honest, I did not expect this law to still be true some 30 years later, but I am now confident that it will be true for another 20 years. By the year 2012, Intel should have the ability to integrate 1 billion transistors onto a production die that will be operating at 10GHz. This could result in a performance of 100,000 MIPS, the same increase over the currently cutting edge Pentium® II processor as the Pentium II processor was to the 386! We see no fundamental barriers in our path to Micro 2012, and it's not until the year 2017 that we see the physical limitations of wafer fabrication technology being reached." Amdhal's Law states that the performance improvement to be gained from using some faster mode of execution is limited by the fraction of the time the faster mode can be used. For instance, if the floating point operation is enhanced such that a CPU can perform floating point operations twice as fast as the previous CPU, that does not mean the overall performance of the CPU is twice as fast. The performance improvement will depend upon the percentage of CPU time floating point operations are performed.

Example using Amdahl's Law: If a new generation of CPU can perform floating-point operations twice as fast as the previous one and 20% of all operations are floating point operations, what is the overall performance speedup?

Assume it takes 100 microseconds to complete a program using the old CPU.
To run the same program in the new improved CPU, it will take 80 microseconds for 80% of the program and the remaining 20% will run twice as fast, therefore will take 10 microseconds, a total of 90 microseconds.

Overall Speedup = Execution time of the unimproved CPU/Execution time of the improved CPU = 100 microseconds/90 microseconds = 1.1

Cache Memory

All modern CPUs contain varying amounts of Level 1 cache memory and all motherboards contain varying amounts of level 2 cache memory. The latest Intel Pentium 4 CPUs come with 8K of Level 1 cache and 512 K of Level 2 cache in the chip itself and the cache on the motherboard will now be referred to L3 cache. The latest PowerMac (Apple) has three levels of cache inside the CPU. Each level of memory is larger than the previous one and retrieves data from the immediate lower level. The cache memory works on the principles of spatial locality and temporal locality. Spatial locality states that data that are physically close together are accessed close together, and temporal locality states that recently accessed data will be accessed again in the near

future. Based on these principles, when one datum is requested by the CPU, reading data all around it and keeping it closer to the CPU for fast access makes sense. Also, once a datum is used by the CPU, keep it closer to the CPU for further use. The principle of locality is also true for program code. A program spends 90% of the execution time in only 10% of the code such as in case of loops and function calls. After the CPU registers that take one-half clock cycle to read or write, the internal cache is the fastest memory available and transfers data into and out of it with each clock cycle, same speed as a CPU cycle.

In order to examine how cache memory works, let us take a computer configuration with 256KB cache and 256MB of RAM, a 1 to 1024 ratio. The RAM is divided into blocks of certain bytes each (say 64 bytes) and cache will be divided into lines of same number of bytes (64 bytes) as in each block. In this scenario, there will be 4096 lines of cache and 4,194,304 blocks of RAM. Each line of cache should hold data from 1024 blocks of memory, a physical impossibility. However, it is possible to hold one out of the 1024 blocks of data in a single cache line at any given time. There are three different ways to map the blocks of memory to lines of cache: direct, fully associative, and set-associative mapping.

Direct associative mapping is the simplest and the most inexpensive to implement of the three. Continuing with the example of 256KB cache, 256MB RAM and 64 Byte block/line, the direct associative maps every 1024^{th} block into a particular line. This can be easily determined by modulo arithmetic. If a certain block of memory needs to be accessed, the predetermined line can be accessed to verify if that line contains the required block. To make this verifying process easy and fast, the memory address issued by the CPU is divided into three fields, tag, index, and offset. The tag and the index fields together make up the block address and the offset indicates the byte number in the block. The tag of every cache line contains information to check if that line contains the block address requested by the CPU. The memory address issued by the CPU (to address 256 Meg (268,435,455K) of RAM, bin 1111,1111,1111,1111,1111,1111,1111) requires 28 bits to represent it if the architecture is byte addressable. In order to indicate which of the 64 (0 to 63) bytes of a block is desired we need 6 bits ($2^6 = 63$), and the offset field requires 6 bits. The remaining 22 bits are required for the block address ($2^{22} = 4,194,304$). Since we have 4096 lines of cache (0-4095 or 1111,1111,1111), we would need 12 bits to indicate which line will hold the data. The final 10 bits are used for the tag field, ($2^{10} = 1024$), which is checked to see if the appropriate block is found in that line.

Example of Direct Mapped Cache 256 Meg RAM, 256 K Cache, 64 byte block	
Number of bits needed to Address 256M RAM	28 $2^{28} = 268,435,456$
Number of Blocks in 256M RAM (64 bytes = 2^6 per block)	$2^{28}/2^6 = 2^{22} = 4,194,304$
Number of bits required to address 4,194,304 memory blocks	22
Number of bits required to address each byte in a block	6
Number of lines of cache available for 256K, 64 bytes/line	$2^{18}/2^6 = 2^{12} = 4,096$
Number of bits required to address the 4096 cache lines	12

Number of bits for tag	$28 - (6+12) = 10$
Number of blocks of memory represented by each line of cache	$2^{10} = 1024$

In the case of fully associative cache, any block from the memory may reside in any line in the cache, and all tags searched in parallel to find the desired block address. The set associative cache provides for a number of lines (2, 4 or 8) to be in a set, and a block can only be placed within a specified set. The lines in that set must be searched to see if desired data is found there. Intel Pentium 4 implements a 8-way associative L2 cache.

Other important cache related information to consider are hit rate, miss rate, replacement policy in case of miss, and write policy. When a requested data is found in the cache it is called a hit, otherwise it is called a cache miss. A cache miss causes the CPU to stall; the number of stalled cycles (miss penalty) depends upon memory speed, bus speed and bus width. When a miss occurs the new requested data must be brought into the cache, but to which line? In direct mapped the answer is straightforward since the new block can only be placed in one pre-ordained line. In case of the fully associative and set associative, a decision has to be made as to which line of data to replace. Some algorithms that could be used are first-in, first-out and least-recently used.

In case of a write to the memory, three options exist: only update the RAM, only update the cache, and update both memory and cache. The first two options make the other unusable, writing only to the RAM makes the corresponding cache line unusable and writing only to the cache makes the corresponding RAM block unusable. Since reads are more important for performance of the CPU, updating RAM without updating cache is simply not done. Cache can be updated without updating the RAM as long as when that cache line is replaced the contents are written to the corresponding block. But how can it be determined if a line has been changed? A bit can be kept to indicate if the line is changed (the memory is dirty) or not; this bit is called a dirty bit. A line with a dirty bit should be written back to the corresponding block and this is called *write back*. When both the cache and the memory are updated simultaneously it is referred to as *write through*.

Primary Storage (Main Memory)

There is a performance gap between the CPU and RAM technologies. Memory speed has only increased by a factor of 2 over past 15 years where as the CPU speed has increased almost 1,000 times. CPU speed does not amount to much unless it can access instructions and data to manipulate. It takes more than 60 clock cycles for the CPU to access the data from RAM. This is why cache plays such an important role in performance. With increased cache availability, what is needed is high bandwidth for the RAM.

In order to install 64KB of memory, the 8088 PCs required 8 Dual Inline Pin Package (DIPP) chips and one extra chip for parity checking, all installed in a bank, making sure none of pins bent. Plug in 36 chips and you have 256 KB

memory. The speed of these chips was 150 nanoseconds (ns). In order to make installation of these chips easier, Single Inline Pin Package (SIPP) memory modules were introduced in 1982. Each SIPP had 9 DIPP chips including parity and had 30 pins. Perhaps it was easier to install one SIPP than 9 individual chips.   Removing and reinserting them were very difficult as the pins were easily breakable. The next improvement in memory design was the SIMM (Single Inline Memory Module). SIMM was very similar to SIPP except for the edge connector and only allowed installation in proper way through the aid of wholes and notches. The 30-pin SIMMs came in 64K, 256K and 1 Meg capacities with an average speed of 80 nanoseconds. As the data width of the CPU and motherboards increased to 32 bits, 30 pins were not enough and the new generation of 72 pin SIMMs were introduced. The speed of memory on the SIMM only improved moderately, 70 nanoseconds on the average. As the technology improved, parity chips are no longer necessary. With introduction 64-bit Intel Pentium CPUs came the need for wider memory chips and the 168-pin dual inline memory module (DIMM) was introduced.

As previously mentioned, the parity chip was used to detect errors. The parity could be set as either even or odd, and in case of an even parity setup, all sending data bits with the value of 1 were counted. If the count yielded an odd number the parity bit would be set to 1 to make the total to be an even number. The receiver also would count number of 1s received, and if an odd number was received an error in transmission will be assumed. In 1994 SIMMs incorporated EDO (extended data out) in which read or write cycles were batched in bursts of four and incorporated fast page mode. EDO RAM had its own error correction techniques and kept track of most recent read/write location for rapid repeated access. The DIMM uses Synchronous Dynamic RAM (SDRAM) and allows for error detection and correction using the ECC (error correcting code). Examples of error correcting codes are the Reed-Solomon algorithm and the Hamming Code. SDRAM synchronizes with the system clock and read time is between 8 to10 ns. Another innovation is the DDR-DRAM that increases bandwidth by transferring data on both the rising edge and falling edge of the clock signal thereby yielding double data rate (DDR).

The Dynamic RAM (DDRAM) is cheaper than the static RAM (SRAM) and therefore most widely used, but, it is slower than the SRAM. The less expensive construction of DRAM looses its content due to capacitor leakage and upon a read; therefore, the data DRAM must be refreshed periodically. Each bit of DRAM has one capacitor and one transistor. The capacitor holds the value of that bit, while the transistor reads it, amplifies it and transfers the data to the data bus. The SRAM uses flip-flops (an electronic circuit that can assume either of the two states, on or off) and does not need to be refreshed. Each flip-flop is made of 6 to18 transistors and stores one bit. The flip-flop retains its status until it is changed or the power is turned off. Since it is faster than DRAM, it is used for cache memory.

No amount of memory in a multitasking computer is enough; each process must have its own address space. In order to accommodate multiple processes, particularly with limited memory available, only a portion of the needed memory is allocated to each process. Only code and data that are being used are brought into the RAM, and rest remains in the hard drive. The primary storage and secondary storage is seen as contiguous memory. This technique is referred to as virtual memory. In the days of 8088 PCs, it was the responsibility of the programmer to keep the programs small or segment them by overlay techniques and swap in and out as needed. Now, the operating system takes care of virtual memory and uses paging techniques. When the CPU requests data that is not available in the cache and not available in the memory, then a page has to be swapped out. This causes the CPU to swap that process out for another and waits for an interrupt once the data is available. We will be discussing this further in the chapter dealing with operating systems.

Secondary Storage (Auxiliary Memory)

A computer that cannot store data permanently is just a toy. When Radio Shack first introduced their TRS-80 microcomputer in 1977, the only secondary storage available was a very unreliable cassette recorder. Most programs written at that time had to typed in over and over again. The TRS-80 series of computers became very popular once Radio Shack introduced a 5.25" floppy disk drive with 90KB capacity. Floppy disks have come a long way since then, a 3.5" high-density, double sided disk of today has a capacity of 1.44 MB. The 5.25" high density double-sided floppy disk had a capacity of 1.2 MB, and drives using them are not installed in computers made later than 1995. These diskettes can be formatted with or without system files on them. Those with system files can be used to boot a computer. It is a common mistake to turn on a computer with a non-system floppy disk inserted, and if the computer is configured to boot first from the floppy disk drive, the system will hang until the problem is rectified.

In selecting a disk drive, it is important to compare them based on some criteria. As a read/write request is issued, the arm with the read/write head moves to position above the desired track; the time to do this is called the seek time (example: 10 milliseconds). The next step is to wait for the rotation of the disk to bring the sector under the head; this wait is called the rotational delay time and is inversely proportional to the RPM of the drive (example: 5 milliseconds). The sum of these two measurements is called the access time (example: 15 milliseconds). Now, the read/write operation can begin at the data transfer rate. Reliability of a drive is measured by the mean time to failure (MTF) given in hours.

Hard disk drives are the most popular form of secondary storage. Early IBM PCs did not come standard with a hard drive. When hard drives became a standard item, the computer was called the IBM-XT. IBM used the Seagate controller card and ST-506 MFM (Modified Frequency Modulation) hard drive. MFM is the encoding method used for recording bits onto the media of the drive. The Run Length Limited (RLL) encoding

was used for a brief period of time. This gave fifty percent more data capacity on the same size drive over MFM encoding, but required a better quality recording surface. The RLL encoding gave too many errors when disks intended for MFM were used. Another Seagate encoding method and interface, The Enhanced System Device Interface (ESDI), did not remain on the market very long. All these drives required low level formatting as well as high level formatting. The controller card's ROM chip contained the low-level format routine and this routine was accessed using the DOS debug utility. During the low level format the interleave factor (which sectors to read during a rotation) and bad sectors (printed on the drive) were manually entered, and additional bad sectors were added during the surface scan. Today's hard drives come already low-level formatted.

Nearly all hard drives in PCs today are either Integrated Drive Electronics (IDE) or Small Computer System Interface (SCSI) drives. The IDE (also known as Advanced Technology Attachment, ATA) drives are found in low-end systems, while the SCSI drives are found in high end PCs, workstations, and servers. The IDE (integrated drive electronics) drive was originally developed in 1981 for Compaq Corporation by Western Digital, a drive manufacturer.   The device controller is built directly on the drive itself, giving flexibility to manufacturers to make the drives larger and faster using proprietary methods while abiding within the ANSI standards known as the ATA standards. Over the years many improvements have been incorporated into the ATA drives such as Direct Memory Access (DMA), synchronous mode, and self-monitoring, and analysis and reporting technology. And the drive capacity have increased from 60 megabytes to over 100 gigabytes. For ATA drives larger than 528 megabytes the LBA (Logical Block Addressing) should be enabled. The drive is connected to the system board through a host bus adapter. Most motherboards today have two host adapters for IDE devices built in, each capable of handling two drives. If two drives are connected to one host bus adapter, they should be configured as master and slave using the jumpers provided on the drives.

Every effort is being made to use IDE (ATA) drives for servers and other high-end computers as evidenced by the development of IDE RAID (Redundant Array of Inexpensive Disks) technology which previously was available only for more expensive SCSI drives. IDE RAID controller cards support up to 30 drives and provide for 64 to 256 MB of unbuffered PC133 ECC SDRAM memory onboard.   When configuring RAID, we must choose a RAID level. There are several RAID levels available today and the list is expanding. Multiple disks can appear as one interleaved large disk if we select RAID level 0. If maximum reliability and recoverability are our concern then we must select RAID level 1 that mirrors all data on duplicate drives. Other RAID series give reliability through error detection and error correction techniques. RAID level 10 is similar to RAID 1 (mirroring) and achieves high I/O rates by stripping the RAID 1 segments.

Another improvement in IDE technology is its the serial version. The days of "parallel is faster than serial" is gone. Newer serial technologies such as the Firewire and the serial ATA are examples. The serial ATA yields double the bandwidth of the parallel ATA transferring about 150MB of

data per second and expected to increase to 600MB soon. It uses thinner, more flexible cables and lower pin counts allowing computer makers to design systems with cables that are simpler to route and install. The data transfer in serial ATA is accomplished by packing bits into packets and transmitting these packets.

Even though similar technologies have been around since 1960, the SCSI (Small Computer Systems Interface) technology was formally adopted by ANSI (American National Standards Institute) in 1986.   The SCSI, pronounced, "Skuzzy", is provides for a standardized way of connecting peripherals to a computer. Three ANSI standards exist for SCSI and they are SCSI-1, SCSI-2, and SCSI-3. The SCSI-1 used an 8-bit wide bus, could connect seven devices, and ran at the rate of 5 MHz. Each device must be assigned an address value from 0-7 (7 is usually assigned to the SCSI card itself). The clock rate was increased to 10 MHz on the SCSI-2 and it was called the Fast SCSI. Then the data width was increased to 16-bit for the SCSI-2, called the Wide SCSI, and was able to connect 15 devices. The clock rate was further improved on SCSI-3 to 20, 30 and 40 MHz are called Ultra 1, 2 and 3 respectively.   The cable connectors used for SCSI are very different and could lead to confusion. All types of adapters are available to convert from one connector to another. Since SCSI cables use bus topology it should be terminated at the last device. Most SCSI devices have a terminator that can be enabled, otherwise an external active terminator should be plugged in. SCSI controllers come with a utility program that can be accessed at boot up by pressing a key combination. Adaptec Corporation is the leading manufacturer of SCSI controllers and most of their new controllers allow for RAID. The SCSI technology is used in most fileservers, workstations and other high-end systems.
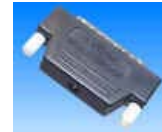
Here I need to talk about CDs, CDRWs, DVDs and tape drives.

Input/Output Interfaces

While a computer does not require any input or output devices for its operation, these are needed to communicate with the users. The keyboard and monitor are the most common input and output devices to which users are accustomed. If the I/O devices are connected directly to the system bus the CPU will have to generate signals to control them that will slow down the CPU tremendously because the I/O devices are very slow compared to the speed of the processor. Other reasons why the CPU should not control the peripherals are that there a wide variety of devices with differing methods of operation, and the data formats and block transfer modes are different. Therefore, for the input/output devices to be connected to the buses on the motherboard, we need to install the appropriate interface cards (I/O modules). The I/O module has a processor on it that handles all the operations of the attached devices, satisfies requests of the CPU, and may have local memory for buffering the I/O stream. An example of such an I/O module is the video card.

The 8088 PCs came with a monochrome text-only video adapter that displayed 25 lines of 80 characters each. Even though each character was formed with dots, these dots were not individually accessible, and only characters and certain blocks could be displayed. When Lotus 123 (a spreadsheet program) implemented program capability to produce graphs, dot addressable monochrome video cards became a necessity. In the mid 1980s the color graphics card (CGA) was introduced that supported 16 colors at the resolution of 640 x 200. The resolution was so poor that dots were highly visible in characters. The Enhanced Graphics Adapter (EGA) introduced next, and it had a resolution of 640 x 350. When IBM introduced their PS/2 a new video adapter card, the Video Graphics Array (VGA) with resolution of 640 X 480 also was introduced that became the basis of all video cards in use today. Various improvements to VGA are available in the market such as SVGA (Super VGA), XGA (Extended Graphics Array), and others; they give 1024 x 768 resolution or more. Further improvements have to do with how they communicate with the CPU through different ways of connecting to the data bus on the motherboard. The different types of buses will be discussed in a later chapter dealing with computer architecture. Here are some calculations demonstrating the difference in throughput using different types of interfaces.

Bus transfer rate or Throughput in MBps = Speed of the bus in MHz * (Width of data transferred in bytes/ Cycles to transfer these bytes)

First let us calculate the throughput of 8088 machines with 8.33 MHz ISA bus speed. These machines were 16 bit machines, but the bus width was only 8 bits, therefore it took 2 cycles to transfer the 16 bits.

ISA Throughput = 8.33 * (2/2) = 8.33 MB/sec

Here is a chart showing some calculations.

Bus Type	Bus Speed in MHz	Data Width in bytes	Cycles to transfer	Calculation	Throughput in MB/sec
ISA	8.33	2	2	8.33*(2/2)	8.33
EISA	8.33	4	1	8.33*(4/2)	33.32
MCA	10.0	4	1	10.0*(4/1)	40
VESA	33	4	1	33*(4/1)	132
SUN	50	4	1	50*(4/1)	200
PCI	33	8	1	33*(8/1)	264
AGP	66	8	1	66*(8/1)	528

The Case and the Power Supply

Every computer manufacturer will market different models of a desktop computer. The models differ based on the case, power supply and motherboard used. The design of the case will determine the number of bays for removable and non-removable drives, type of motherboard, the type and wattage of the power supply, and number of fans used for air circulation. We could avoid cases altogether for each computer and go with a rack mount system.

The power supply converts the AC to DC and steps down the voltage from 120 or 220 volts to 12, 5 and 3.5 volts. The power supply converts AC to DC using full wave rectifiers and filters. Rectifiers converts AC to DC by using power diodes or by controlling the firing angles of thyristors. At the component level a rectifier uses diodes and capacitors. Today's power supplies first convert the 60 Hz AC to a much higher frequency in order to make the stepping down much easier and more accurate. Such a power supply is called a switching power supply and provides a more uniform voltage to the computer. Voltage is stepped up or down by varying number of secondary windings in the transformer. In a transformer there are primary and secondary windings. If there are half as many windings in the secondary as the primary windings the induced voltage will be approximately half on the secondary windings. The wattage rating of the power supply should be sufficient to handle all components plugged into it. For example, drives generally consume 5 to 15 watts, motherboard 30 watts, CPU 50 watts and an AGP video card 30 watts. These specifications should be obtained from the manufacturer of the components.