

Know Your Personal Computer

2. The Personal Computer Hardware

S K Ghoshal



Siddhartha Kumar Ghoshal works with whatever goes on inside parallel computers. That includes hardware, system software, algorithms and applications. From his early childhood he has designed and built electronic gadgets. One of the most recent ones is a sixteen processor parallel computer with IBM PC motherboards.

This article surveys personal computing with the IBM personal computer as a platform to discuss details of hardware.

Introduction

The personal computer (PC) is the most widespread and popular computer because it runs all kinds of applications. In fact, if there is an application package that does not have a PC variant it is wise to not get familiar with it as its host may become obsolete and the package not survive. PCs will take a long time to get obsolete. They only get smaller, cheaper and more powerful as the years go by. Eventually, all the applications that people need will become available on PCs.

PCs can be connected to all types of peripherals and equipment. If any device does not interface with a personal computer, it may not be worth buying. In fact one should be sceptical of the device because a PC has an open architecture for which peripherals are easily developed. A device that does not plug into a PC does not live long. Such devices can be had from many vendors. They become cheaper and more powerful with time as their technology improves.

PCs come in all sizes, prices and capabilities. Some are large machines with a huge, fast, primary memory and a large backup disk which serves as secondary memory. Such PCs are capable of running an interactive operating system that hundreds of users can simultaneously use over a network. The tiny laptop PCs, barely larger than a pocket calculator in size, that run for hours on a rechargeable battery are also widely used. The base architecture and organization of system software remain the same across all of them and this is what

Address for correspondence

S K Ghoshal,
Supercomputer Education
and Research Centre,
Indian Institute of Science,
Bangalore 560 012, India.



A truly open system

All IBM PCs are compatible. Programs compiled in any one will run on any other. This is called *portability*. Almost any hardware component in the PC can be upgraded to a more powerful and modern variant and yet all the old programs run on the new system. This is called *scalability*. Old operating system, compilers and other system programs can be replaced with new ones without disturbing the way applications are run on the PC platform. This is called *interoperability*. These three properties make the IBM PC an open system. You can buy and use it in peace for years. You can upgrade it whenever you need it and can afford it. PCs do not get obsolete. Like a bicycle, all of it is never discarded. Components which are old are replaced with new and improved ones.

we will study in this article and the ones to follow.

Personal computers are a result of the tremendous advances made in very large scale integration (VLSI) design tools and implementation technology. This had two effects. First, powerful microprocessors were made in small areas of silicon. Second, system integration techniques allowed printed circuit boards to be mass-produced, populated with the semiconductor components and tested, resulting in highly reliable systems which could be sold at a low price. This trend continues and personal computers keep getting smaller, more powerful and more affordable. People who write system software have realized that programs should be written in such a way that they can be used, understood and run on many different machines. Thus system software and their applications are more standardized than they were a few years ago. Both these trends

PCs will take a long time to get obsolete. They only get smaller, cheaper and more powerful as the years go by.

Why IBM PC?

Personal computers are there because people want them (particularly the types that do not consume much electrical energy and don't need air-conditioning and other types of protected environments) at affordable prices and today's technology can mass-produce inexpensive but reasonably powerful ones. Most of them work reliably without much maintenance. The developments in personal computers were inevitable but the reason that the IBM PC is so popular is that its architecture is open and its standards are well-defined and published. There are therefore many vendors developing hardware for the PC and many billions of dollars worth of software runs on this platform.



If I did it again

Often pioneers assign bad names to what they invent and repent later in life. Ken Thompson was asked what his first step would be if he were to write Unix (Unix is a trademark of Unix Systems Laboratories) again. He replied: "I would call it create". He was referring to the *creat system call* of Unix. Unix is the most portable operating system today. That

System call is a request from a user's program to the operating system of a computer for a service such as creating a file, deleting a file etc.

is because Ken and others had designed the system calls and other internals of Unix so well that it has survived the test of time, has spread across virtually every computer system that exists today and has flourished among many different user communities. Unix also runs on IBM PCs. We will discuss Unix in subsequent articles in this series and also look at the system call in Unix that creates new files or rewrites old files. But we will have to call it *creat*.

People who write system software have realized that programs should be written in such a way that they can be used, understood and run on many different machines.

are good for computing in general and personal computing in particular.

However, while studying personal computers and the organization of their system and application software, it must be borne in mind that many developments that took place had a historic reason and many more were idiosyncratic in nature. The external circumstances and individual mind-sets that led to these changes no longer exist. So if one had to do it today all over again, the situation would be quite different.

How they missed the opportunity

It could have been the DEC PC. If they had only listened to engineer David Ahl at Digital Equipment Corporation (DEC). In 1974, he proposed that the company produce an inexpensive version of its PDP-8 minicomputer and sell it to computer enthusiasts for \$5000 a piece, at no loss for DEC (DEC and PDP are registered trademarks of the Digital Equipment Corporation). Top management shot down the idea. They thought that it would be foolish for an individual

to buy a computer. She would rather buy a dumb terminal that connected her to a powerful mainframe. That was the trend in computing those days. The VLSI revolution had just begun. Not many were aware of its potential. Few believed personal computers could be sold to the masses. Had Ahl's proposal been taken up, the PDP architecture would have played the same role in the history of personal computing as the 8088/80386/80486 series from Intel!



The First Personal Computer

About 20 years ago, only electronics hobbyists were seriously interested in personal computers and it existed only in the pages of magazines like *Popular Electronics* and in the laboratories of people who loved building them. Some companies sold kits (without power supply, keyboard and monitors) to hobbyists. Altair from MITS of Albuquerque, New Mexico, USA is the first personal computer which sold in volume. Introduced in 1975, it had the 8-bit Intel 8080 as its CPU and 256 bytes of memory. A young man wrote a BASIC language interpreter for this machine. His name: Bill Gates.

Wherever possible, I will try to explain why things were designed the way they were and how they changed as technology evolved. In the next section, we present an overview of the PC hardware.

PC Hardware

The most important component in the IBM PC is the *motherboard*. It is a rectangular printed circuit board (PCB) which has the 80386/87 (2 chips) or 80486 CPU, the primary memory, the controller chips that carry out operations such as direct memory access (DMA), interrupt handling and keeping track of the system time. As system integration technology improves, keeping pace with advances made in VLSI technology, all these controller chips get incorporated

Downsizing was hard to do those days

Single board microcomputers grew more powerful, sold well and finally edged out the mainframes. Mainframes did not take advantage of the VLSI revolution. They did not shrink their hardware in size and make their software more modular, tight and portable (this whole operation, done by computer system designers is called *downsizing* these days). All mainframe and supercomputer vendors failed in this aspect. IBM, Data General, Control Data, Texas Instruments and Cray Research attempted downsizing their popular models but without much success. They somehow did not understand the micros. Downsizing has been perfected by most organizations by now, as system software and applications must run on the PC platform, or else perish.

Personal Computing - the Apple way

One company that has refused to follow IBM and has created a niche for itself is Apple. Its founders, Steve Jobs and Steve Wozniak were pioneers in personal computing. They created personal computers before IBM got into the game. Apple uses the Motorola 68000/68020 series of microprocessors as their CPU. They are much easier to use than IBM PCs, thanks to their excellent icon-based graphical user interface. Apple did not make its architecture open. In India, they are popular only with desk top publishers.



What is DMA?

Copying data between memory devices (particularly between primary and secondary memory devices) is an operation very frequently used in a computer system. One can always use the CPU itself to copy data. But that would be wasting the hardware resources of the CPU and is also very slow. Thus every computer system has a number of channels bypassing the CPU, to directly access memory and copy data. This mechanism is called *direct memory access* or DMA. The channels are called DMA channels.

within complex VLSI chips called application specific integrated circuits (ASICs). A modern motherboard has only about 4-5 chips apart from the CPU. Therefore they consume less power and work more reliably than their ancestors. They even cost less since ASICs can be mass produced.

The electrical signals produced by the CPU are weak and they cannot traverse long lines without getting severely attenuated. So these signals are amplified (or *buffered*, in digital electronics parlance) using special chips called transceivers (which can buffer in both directions and are used for data lines) and address latches (which can retain an useful address output by the CPU and can thus be used to wait for slow memory to respond to a fast CPU) which are also present on the motherboard. In modern motherboards there are many sequential digital *finite state machines* (FSMs) which serve as an interface between CPU and memory and input/output (I/O) devices. They help the CPU work with a high throughput while working with slow devices. Just as a power transformer with multiple secondary windings steps down 220V and supplies 12V, 5V, 3V etc at different ampere ratings, the FSMs distribute the CPU's memory bus into a number of specialized buses with different speeds and widths. Thus maximum throughput is obtained from the system and a uniform view of the memory is projected to the CPU, even though it is physically heterogeneous with different types of devices with various bus speeds and widths. One of these buses which is relatively slow and only 8-bit wide, terminates on the motherboard itself and connects to an erasable programmable read-only memory (EPROM) device. This device has some programs that are run while initializing the motherboard and other programs which serve as a low-level operating system. These programs are collectively called *Basic Input Output System* (BIOS).

A high-speed 32-bit wide bus connects the CPU, through the cache controller FSM, to up to 512 Kilobytes of SRAM cache memory that can be present on the motherboard. Yet another medium-speed 32-bit bus connects the CPU, through DRAM controllers, to banks of



Is BIOS Hardware or Software?

The original IBM PC BIOS was a trademark of IBM. IBM also retains copyright of that BIOS. Now BIOS is written and sold by many vendors. When you buy a motherboard, you get BIOS along with it. BIOS is closely associated with the hardware and needs to be changed or augmented if the hardware is changed in a major way. It remains physically with the hardware in that it is kept in EPROMs that are plugged into the motherboard. It does not have to be loaded from an external device every time the computer is turned on. Some people regard BIOS as part of the hardware. We however, take a middle path and denote BIOS by a term "*firmware*", to designate that it is in between hardware and software.

dynamic RAMs integrated as single in-line memory modules (SIMMs) which constitute the primary memory of the computing system. SIMMs are removable from the motherboard. Many other buses connect the CPU to different peripheral support and other controller chips that are present on the motherboard.

A number of buses are brought out of the motherboard as edge-connectors. They are used to plug in add-on cards and extend the architecture. IBM PC supports the world's largest variety of add-on cards available from a wide range of vendors at affordable prices. The old and slow *Industry Standard Architecture* (ISA) bus which had two varieties 8-bit and 16-bit wide respectively, the 32-bit *Extended Industry Standard Architecture* (EISA) bus and the 32-bit high-speed *Peripheral Component Interconnect* (PCI) bus are some popular standard buses. A couple of these edge connectors are taken up by the graphics adapter and the hard disk controller. The rest are left free for the user to plug in adapters of his choice and requirements. The motherboard has a DIN socket into which a 5-core cable from the keyboard plugs in. It also has a polarized power supply connector (so that one cannot connect it in reverse and blow up all the hardware in the motherboard) to feed it power from a *switched mode power supply* (SMPS). Once the keyboard, power supply, disk adaptor

Personal computers are a result of the tremendous advances made in very large scale integration (VLSI) design tools and implementation technology. This had two effects. First, powerful microprocessors were made in small areas of silicon. Second, system integration techniques allowed printed circuit boards to be mass-produced, populated with the semiconductor components and successfully tested.




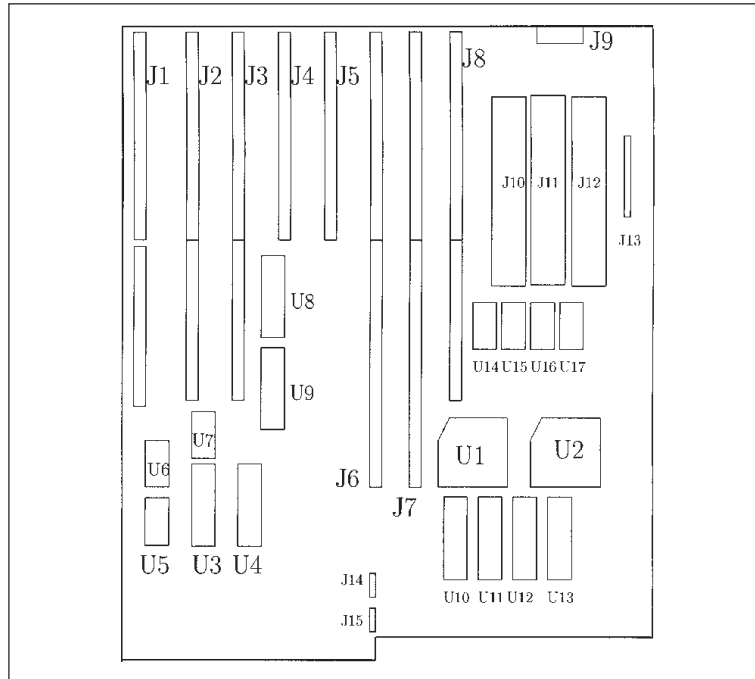


Figure 1 Photograph of an Intel iSBC386ATZ motherboard..

Figure 2 Floor-plan of the Intel iSBC386ATZ motherboard.



Compatibility

The first PC from IBM, released in 1981, had an Intel 8088 CPU working at 4.77 MHz clock speed, in a motherboard which did not have any FSM or cache memory and only had 8-bit ISA slots. Any program that ran on that platform, can run on any modern motherboard. It uses all the resources present both on the motherboard and across PCI or EISA buses, to deliver, on an average, more than 20 times the performance the original one did.

with drives, graphics adaptor and monitors are connected, the PC is a complete computer ready to switch on and be used. *Figure 1* shows an old-generation 80386-based motherboard from Intel. *Figure 2* depicts its floor-plan. U1 is the 80386 CPU. U2 is the 80387 numeric co-processor. U3 and U4 are the DMA controllers. U5 and U6 are the interrupt controllers. U7 is the programmable interval timer chip that helps keep track of the system time and generates many other kinds of timing interrupts that are used by the system software. U8 and U9 are EPROMs that contain BIOS. U10, U11, U12 and U13 are *programmable array logic* (PAL) chips that implement the different FSMs. U14, U15, U16 and U17 are transceivers. J1, J2 and J3 are 16-bit wide ISA slots. J4 and J5 are 8-bit wide ISA slots. J6 and J7 are 32-bit wide Intel’s own slots. (which never became very popular). J8 is a 16-bit wide ISA slot. J9 is the keyboard connector. J10, J11 and J12 are the SIMM module sockets. J13 is the power supply connector. J14 is the speaker connector. J15 is the key-lock connector.

In subsequent articles we will learn more about the hardware and software of the IBM PC. Write to me if you want more information or have any questions or comments about this article.



What is a Bus?

The CPU accesses memory and exchanges data with input/output devices using a set of wires called a *bus*. Each wire conveys a bit. If it is at +5V potential with respect to the ground line, then it represents a 1. If it is less than 0.6V potential with respect to the ground, it conveys a 0. The number of lines in a bus is called its width (see *Figure 3*). A number of ground lines are also run along with the signal lines in order to nullify any effect of difference of ground voltages at the two ends of a bus. The ground lines are not counted in the width of the bus. If a bus conveys the address sent out by a CPU, it is called an *address bus*. It is unidirectional. Its direction of signal flow is always from the CPU to the memory, and never the other way. If a bus carries data that is exchanged between the CPU and the memory, then it is called the *data bus*. Data buses

are bi-directional. A collection of data and address buses is simply called a *bus*. *Figure 4* shows a CPU accessing a memory with 2^{32} bytes using an address bus that is 32 bits wide and a data bus which is 8 bits (i.e. one byte) wide. The **READ**, **WRITE** and **WAIT** lines are together called the *control signals*. Often power supply lines carrying many amperes of +5V, important control signals and other signal lines like clock and interrupt lines (lines by which devices can send signals to the CPU) are included in a bus. Buses follow a specification that lays down the geometrical arrangement of the wires, the maximum permissible length, the maximum number of devices that can be connected to it and other details. By common usage of the term, a *N*-bit wide bus means a complete bus, whose data-bus width is *N*-bits. Thus, *Figure 3* shows a 4-bit bus.

Figure 3 A unidirectional bus of width 4. The four wires and the ground wire as shown in (a) form a bus; it is normally represented as in (b). The digit 4 above the / symbol indicates that the bus width is 4 and the arrow indicates that the bus carries signals from left to right..

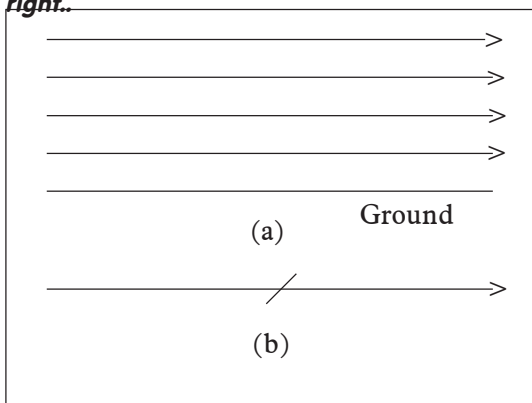


Figure 4 CPU accessing memory using a bus.

