Computer Systems

Authorship: Francesc Aulí, Joaquin Garcia-Alfaro, Oriol Pujol

Computer Foundations

[. . . . . .]

Introduction

The history of modern computers has less than a century of existence. From the very beginning, the race to get more and more powerful designs has followed a frenetic pace where, in a matter of months, top equipment was rendered obsolete by new, ever smaller, and faster products. Moore's law predicts that the power and storage capacity of computers doubles every 18 months, resulting in exponential growth. In the first section of this chapter, we briefly report the dazzling history of computer discoveries and their foundational nature. Then, we move to some more recent topics with respect to the hardware and software building blocks underlying current designs.

Moore's law

Named after Gordon Moore, co-founder and chairman emeritus of Intel corporation, is an informal claim about the number of transistors on microchips doubling approximately every two years, leading to exponential growth in computing power. In recent years, challenges in maintaining this pace have emerged due to physical and economic limitations.

1. Evolution of modern computers

In general, five major generations of computers marked, basically, by new technological discoveries that have made it possible to build ever faster computers. Next, we briefly review what these generations have been and discover what the key moments are in the evolution of modern computers.

1.1.   First generation (1946 ~ 1954)

The first major generation of computers was born with the first all-electronic general-purpose calculating machine called ENIAC (Electronic Numerical Integrator And Computer), developed by professors John Mauckly, Presper Eckert, and Herman Goldstine of the University of Pennsylvania. This machine was the result of work commissioned by the United States government to resolve issues that had arisen following the Second World War. The computer structure occupied the entire first floor of the Moore School, weighed 30 tons and had more than 18 000 vacuum tubes.

Although it had already been more than ten years that scientists from all over the world had been working on the construction of machines that could make numerical calculations and/or solve specific problems, ENIAC is considered the first computer in history since it was completely electronics (other machines worked with electro-mechanical relays) and could be programmed to do different tasks. In its time it represented a great revolution in scientific environments, since it could perform more than 5 000 additions or subtractions in just one second compared to the two operations that its predecessor, the MARK-I, could calculate. The work that took MARK-I a week or two months at the US Army Ballistics Laboratory could be done with just one hour of ENIAC computing! With this machine it was possible to calculate, for the first time, the number n with more than 2 000 digits and it was the one that carried out the feasibility calculations to build the first H-bomb in history.

General purpose machines.

We refer to designs that are not expected to perform a single specific task, but can be programmed to perform any type of general purpose tasks.

The MARK-I computer.

It was developed at Harvard University by Howard Aiken between 1937 and 1944.

In 1948, a congress was held where scientists from all over the world were invited to discuss and lay the foundations of computing. John Von Neuman, unwittingly, was the main propagator of the ideas that were raised there and that have lasted until the fourth generation of computers.


Figure 1.1. The ENIAC is considered the first computer of our recent history.
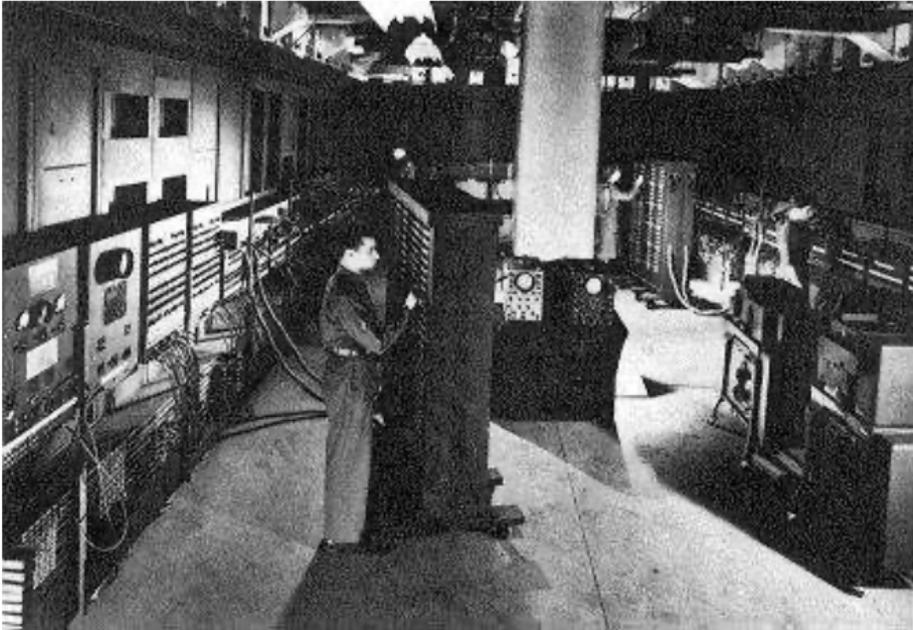
Von Neuman architecture.

Computers up to the fourth generation are based on the Von Neuman architecture which, despite bearing his name, is the result of researchers from all over the world.

The creators of the ENIAC founded a company called UNIVAC that created the UNIVAC-I and later the UNIVAC-II, machines that were already commercialized. The IBM company also began to manufacture commercial machines of this type called 701 and 704. The big problem with these first computers, however, was their maintenance: very high both in electricity consumption and in that of empty vacuum tubes. Although vacuum tubes were the first technological innovation that made it possible to build electronic computers to perform calculations, their high error rates meant that they constantly had to be changed and were not at all reliable.

At the beginning of the first generation, the programming of these first computers was done from the internal wiring of the machine. Every time you wanted to change the program, a good part of the computer's circuits had to be redone (wired). The fact that this system was so manual and involved a great deal of effort made it necessary to change this way of proceeding for a more versatile one.
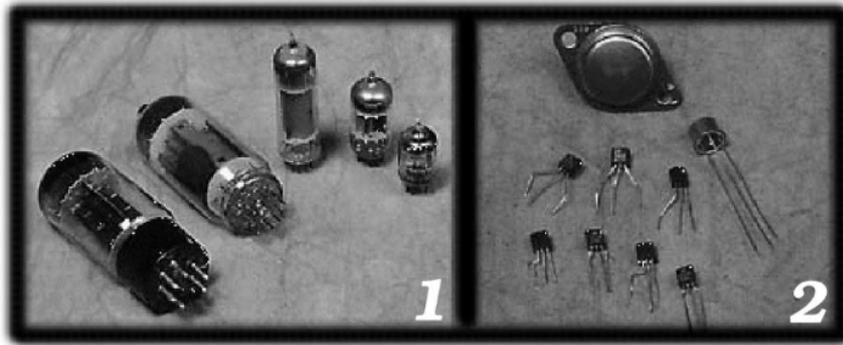
Figure 1.2. (1) Vacuum tubes. (2) Transistors.

From 1948 and after the congress, it began to be talked about storing the programs in an internal memory of the computer to be able to execute them later. Between 1947 and 1949, the EDSAC (Electronic Delay Storage Automatic Computer) was created, which as the main innovation incorporated a small operating system, standard routines to perform repetitive calculations and aids for locating errors in programs.

In 1947, the first transistor was patented, which performed the same function as vacuum tubes but with a much smaller size and much lower consumption. The industrial manufacture of transistors represented a great technological change, precursor to the next generation of computers. The main advantages of transistors compared to vacuum tubes can be seen in Table 1.1.

Operating system

Software component that manages the resources of the computer, cf. Section 3 of this chapter for further details

Table 1.1. Advantages of transistors over vacuum tubes

	Vacuum tubes	Transistors
Size	4 cm	$2 - 4$ mm
Cost	5 W	2 mW
Speed	—	500 times faster

1.2. Second generation (1955 ~ 1964)

The second generation appeared with the industrial development of transistors. Computer designers replaced the already obsolete vacuum tubes with the new, much more powerful, reliable and cheap transistors. The first machines made with transistors were the UNIVAC 1108, IBM 7040 and IBM 7090. At this time, technological innovations such as magnetic tapes or *ferrite* memories also appeared. However, these computers were still very expensive because in addition to having to be built entirely by hand, gold was used in *ferrite* core memories (which was the dominant technology from early 1950 to late 70s).

As transistors began to be used massively, new techniques appeared that represented a complete revolution in the world of computing: large-capacity memories began to exist, the power of computers began to be measured in milliseconds and not with seconds (as was done up to that point), systems were developed to perform floating-point calculations, and after all that, between the second and third generations, high-level programming languages began to appear. In addition, the popularization of these new machines advanced rapidly. From 1958, research began with new integrated circuit techniques. The manufacturing process involved putting more than one element into the same piece. This new way of proceeding began to be introduced and little by little the old elements were replaced by these new chips that made the manufacturing process of computers much easier.

1.3. Third generation (1965 ~ 1980)

The change between the second and third generation of computers was not as abrupt between the first and second generation, caused by the appearance of the transistor. Yet, early architectures were starting to evolve with the spreading of integrated circuits. Computers like the IBM 360, the IBM 370 and the UNIVAC-3600 are examples of computers fully produced with integrated circuits. The production cost of these computers started to lower considerably and market computing boom began to rise. Many companies partnered to found computing centers where computers were bought to perform massive computing operations more quickly. The number of computers installed all over the planet began to grow exponentially. The so-called semiconductor memory created from transistors, which greatly reduced its cost, began to be installed in all computers. The fact that this type of memory could not save information without electrical energy prompted the appearance of the first magnetic disks to store information.

Computers began to be manufactured all over the world; in particular, France was one of the states that did the most research in micro-electronics. It is also at this time when we start to hear about telecommunication, e.g., communication via telephone between two computers. The main problem, however, was that computing was still not within the reach of small businesses due to the high cost of computers. Aware of this drawback, Ken Olsen and Harlan Anderson, engineers who had been working at MIT (*Massachusetts Institute of Technology*) brought to the market in early 70s a new computer project called PDP-11, which served to satisfy reduced computing needs and where small applications could be executed. Generically, these new devices were later referred to as minicomputers. Seeing the success of their idea, they created a company called DEC (*Digital Equipment Corporation*) and began marketing the resulting product. Before the big pioneering companies in the world of computing realized that this product had a very large potential market, DEC had already sold more machines in two years than IBM in its entire history.

During 1972 the US Army called a competition to develop a single-piece computer (even if it was not as powerful as the big machines — mainframes or super-computers — from IBM or UNIVAC) since it needed to control somehow their missiles. The only company that was able to carry out this project was the house called Intel that presented a micro-processor (the 4004) of 4 bits. Although the military did not find it useful, they marketed it through civilian means. Later, Intel presented the second microprocessor it had developed, the 8008, just like the previous one but based on a 8-bit design. This was the mother of the subsequent evolutions of the Intel house that has continued with its philosophy and compatibility with previous versions until reaching the microprocessors that are still manufactured today. Other companies, such as Motorola, also began to manufacture chips, although the internal architecture has always been very different.

In 1974, those previous designs materialized into a new family of computers referred to as microcomputers. A first microcomputer, the Altair 8800 by MITS (Micro Instrumentation and Telemetry Systems), started reaching out the general market with a cost of about 5 000 dollars, making it affordable for the general public.[1] Later similar microcomputers like the Apple II and the Apple III, became pioneering machines as home desktop computers. IBM, to avoid ending left behind, started a new project that soon reconquered the market: the popular IBM Personal Computer (known as the IBM PC). Faced by the new competitor (IBM), Apple started new innovative strategies, embedding new microprocessors and accessories to their products, soon starting the development of other successful families, like the Apple Lisa and the Apple Macintosh.

> IBM did not patent the IBM PC, since most of the functional components were not directly designed by them. Alternative designs reached soon the market, offering microcomputers known as PC clones, and characterized by their low price (but also by their poor performance). Even though PC clones could have helped IBM to expand their product in the home market (since the original products from IBM were superior in terms of quality and performance), IBM created and patented a new family of microcomputers known as IBM Personal System second generation (or IBM PS/2). This new family of microcomputers was in the strange situation of using components that were incompatible with common hardware standards for the home PC industry. The IBM PS/2 did not sell as expected, and soon IBM moved back to compatible designs with the PS/1 model (also known as PS/1000 to avoid numbering ambiguity) to return back to the home market in 90s.

[......]

---

[1] Until that time, computers sold by, e.g., IBM, could cost more than a million dollars.