Mainframe History Provides Lessons

Microprocessors Have Followed a Similar Path, But What Lies Ahead?

by Bernard L. Peuto, Consultant

For most of the past 25 years, microprocessors have followed a development path strikingly similar to that taken by mainframes in their first 50 years. This should not surprise us, as the principal application target for microprocessors has been personal data processing, closely related to the principal application of mainframes: data processing. Given this similarity of application target and development path, microprocessor designers, after some early hesitation, copied essentially all the architectural concepts invented for mainframes, reinforcing the similarities between them.

With the widespread popularity of graphical user interfaces (GUIs) and multimedia and with the emergence of the Web as a driving application, we may see an emerging architectural differentiation between mainframes and microprocessors. But this possible split is not yet certain.

This two-part article develops the theme of importing mainframe concepts into the development of microprocessors. First we compare the time lines for mainframes and microprocessors, then we explain how specific architectural features of microprocessors were influenced by mainframes.

Without undue distortion, it is convenient to simplify the history of mainframe development as a succession of six major phases, each 10 years in duration, starting in 1935.

1935–44: Pioneers

The early pioneers conceived and designed the first digital computers. Among them was the Mark I, a major relay computer designed at Harvard by Howard H. Aiken in collaboration with IBM. Started in 1937, the Mark I was completed in 1944. Its independent memories for data and programs gave the name of Harvard architecture to this design. It executed approximately one operation per 300–500 milliseconds.

The development of ENIAC, widely considered the first digital electronic computer, began in 1943. Designed by J. Presper Eckert and John W. Mauchly to compute ballistic tables, it began executing in late 1945 and was officially inaugurated in 1946. It was a decimal machine, using electronic tubes and a patchboard-like mechanism to enter data and programs. It executed approximately one operation every 200 microseconds (or 5 Kops/s).

Other pioneering efforts should be mentioned. George R. Stibitz designed an early mechanical relay calculator around 1939. Konrad Zuse designed the Zuse Z3 relay computer in Germany around 1941. The Atanasoff Berry machine, probably the first prototype of a binary computer using tubes and a capacitive memory, was designed between 1937 and 1942, though never fully completed.

The Colossus project was one of the most intriguing early “computers.” This dedicated code-breaking machine, conceived by Alan Turing, was first built as a relay-based machine in 1940, then with electronic tubes in 1943. Since the British Intelligence Agency kept its existence secret until the mid-1970s, little has been published that can tell us to what extent this special-purpose computer was a forerunner of ENIAC. The experience the British acquired with this project explained in part their early lead, described in the next phase.

1945–54: First Computers

Learning from the ENIAC experience, its designers and John von Neumann drew up plans for a next-generation computer: EDVAC, which introduced the concept of binary stored programs. But EDVAC was delayed by the departure of its design team in a dispute with the management of the Moore School of Engineering at the University of Pennsylvania. The first computer to implement the von Neumann architecture, as it is now called, was EDSAC, built in England in 1949 by Maurice Wilkes.

ENIAC’s designers founded the Eckert and Mauchly Corporation, later renamed the UNIVAC Corporation. In late 1951, they introduced the best known of the first commercial computers: the UNIVAC I. It used a mercury delay line for memory but is especially noteworthy for having introduced the first magnetic tape drive. The first commercial computer was actually the Feranti Mark I, built in England in early 1951.

During this phase, many other computers advanced the development of the scientific and commercial mainframe industry. Core memory was invented in 1953 by Jay Forester for MIT’s Whirlwind project, the first real-time computer system. The IBM 650, a commercial drum computer introduced in 1954, enjoyed great success and signaled the dominance of IBM over its rivals. As Herbert Grosh, himself a pioneer, used to say, the history of UNIVAC is one of a company snatching defeat from the jaws of victory. Its early lead, better products, and bad management were no match for IBM’s superior management and the genius of the T.J. Watsons, Senior and Junior. It was a story that would be repeated many times.

1955–64: Broad Products

By this time, the first computers had established the two tracks, scientific and commercial, that would characterize mainframes. This new phase represents a reinforcement of this split and a massive increase in the number of computers shipped.

It is beyond the scope of this article to mention all the vendors of computers and their major products. Burroughs,

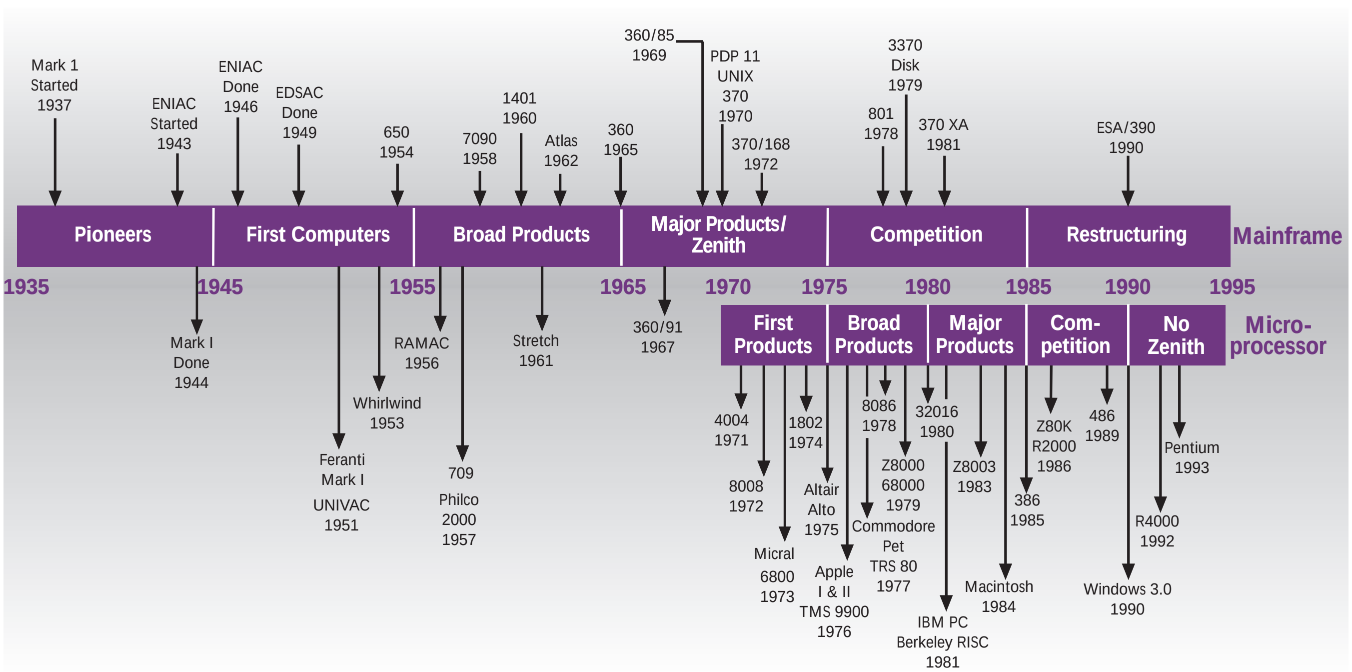

Figure 1. The history of microprocessors shows some parallels with that of mainframes, but progress has occurred roughly twice as quickly.

Control Data, General Electric, Honeywell, NCR, RCA, UNIVAC (the so-called “Seven Dwarves”), and many more, including Digital, were often the first to introduce commercial versions of new engineering concepts. But very early, IBM showed such an incredible dominance of the field, especially in technical innovation, that we will focus on IBM to characterize the next few phases of mainframe history.

Introduced in 1957, the IBM 709 was the last major scientific computer using tube electronics and CRT memory. The first transistor-based computer was probably the Philco 2000, introduced around 1957. With the 7090 in 1959, IBM itself introduced transistors and core memory.

In 1960, IBM introduced the 1401, a character-oriented commercial computer with removable disk drives. IBM had introduced the first disk drive in the 305 RAMAC in 1956. Equipped with a major new chain printer, the 1403, the 1401 became the first complete electronic system solution that was more economical to use than a punched-card system. The dominant vendor of punched-card systems had rendered its own product line obsolete!

1965–74: Breakthrough Products and Zenith

In part due to the Cold War, U.S. government funding for computer research dominated the field until 1965 when, for the first time, the rising percentage of private research and development funds invested in computers equaled the percentage of government outlays.

That year also saw IBM ship the first System 360. Its announcement in late 1964 stunned the competition by the magnitude of the product line and the incredible gamble that IBM was taking, betting the whole company on this project.

Today, more than 30 years later, the 360 instruction set and its system architecture are still represented in all IBM ESA/390 computers. There have been significant architectural changes in the I/O subsystem but very little in the instruction set, which I consider the very model of a CISC instruction set. Currently popular RISC architectures exhibit little or no advantage over the 360, in either performance or code density. In contrast, the x86 architecture performs quite poorly in comparison with popular RISCs. (Note that IBM invented RISC architecture with the 801 in 1978.)

In 1971, IBM introduced the 370, an upgrade for the 360 line that used semiconductor memory and integrated circuit logic. Announced in August 1972 and shipped in 1973, the 370/168 was the first widely available commercial IBM computer that supported virtual memory. It was not the first virtual memory computer from IBM or other vendors, but IBM’s mainstream blessing made virtual memory a feature required of all mainframe vendors.

All these successful products and investments by IBM continually increased its dominance of the mainframe business. IBM seemed invincible then; the Seven Dwarves were getting smaller. GE abandoned the business in 1970; RCA followed. The remaining competitors became the BUNCH (Burroughs, UNIVAC, NCR, CDC, and Honeywell).

Knowing the eventual impact of microprocessors, we believe today that August 1972 represented the zenith of IBM and the mainframe industry. Despite a few more years of excellent revenue growth, IBM would never again be the exclusive controller of the computing industry.

One of the saddest episodes of early computer history was its inventors’ battle for recognition. In 1973, the courts declared Eckert and Mauchly's patents for ENIAC and UNIVAC invalid and named John Vincent Atanasoff the inventor of the computer. ENIAC was the first electronic computer that was fully built and executed actual programs. Therefore, Eckert and Mauchly should be recognized as the inventors of the first working electronic computer.

1975–84: Competition From Microprocessors

Today, the microprocessor has completely displaced the mainframe as the largest provider of aggregate data-processing MIPS sold. But when the Intel 4004 was introduced in 1971, such an outcome was not a realistic projection. Even the 8080, introduced in 1975, was still too primitive to foreshadow any significant competition with mainframes.

During this period, larger and faster disks, like the 512M 3370 in 1979, dominated the interest of data-processing users. In 1981, IBM introduced the 370/XA, which removed the 24-bit addressing limitation of early 360s and changed their I/O channel architecture. In a bold move, IBM recognized the importance of microprocessors by introducing the PC (in 1981) and investing in Intel.

At the end of this phase, marked by the introduction of the 80386 in 1985, the competitive nature of microprocessors in the delivery of data-processing cycles finally became obvious to most people. It had taken only 14 years from the introduction of the 4004!

1985–94: Mainframe Restructuring

In these 10 years, mainframe vendors underwent wrenching restructuring, caused in part by the different business model created by the microprocessor's delivery of data-processing cycles. IBM had its first layoffs, and more vendors disappeared or are still today in the middle of a very difficult phase, as Digital is.

With the ESA/390 introduced in 1990, IBM continued to dominate the mainframe business. Its ESCON I/O architecture provided even greater support for the large data repository of legacy applications.

For a while, the mainframe was distinguished from microprocessor-based systems by a multichip implementation of its central processing unit. This implementation provided increased performance through either faster technology (like ECL) or a larger number of gates dedicated to pipelining and caches.

Today, CMOS density and speed have essentially eliminated any other competing technology, and the CPU of an ESA/390 mainframe is a one-chip microprocessor implementing the current extensions of the IBM instruction-set architecture compatible with the 360. So the difference between mainframes and microprocessor-based systems now resides in the I/O architecture, support for multiprocessors, and the memory hierarchy.

Mainframes are not likely to be eliminated, but one thing is sure: their vendors do not grab the daily newspaper headlines; Intel, Microsoft, and the Web do!

Selecting Representative Events

Writing about the history of microprocessors is very satisfying because it celebrates inventiveness and the pioneering people responsible for many innovative solutions to challenging problems. But the tremendous difficulty of being accurate and understanding the true reasons for historical developments makes it also very humbling. In these few pages I can only hope to give a glimpse of what happened.

Writing a short historical article is also a difficult challenge because the number of illustrating examples that can be used is significantly restricted. For lack of space, I could not cover the impact of minicomputers and workstations on mainframes and microprocessors. Despite their critical importance to the societal impact of microprocessors, microcontrollers (or embedded microprocessors) are outside the scope of this article.

IBM and Intel have such dominance in the mainframe and microprocessor industries that it is difficult to mention other vendors like Univac, Honeywell, Control Data, Zilog, National, and Texas Instruments that today are not as important as they used to be. They were often first to introduce significant advances, later copied by their now-dominant competitors. It would take a much longer article to do justice to all the first inventors of mainframe and microprocessor features.

Microprocessors Follow the Mainframe Path

The most interesting conclusion to draw from a time line of microprocessor history is that its phases are similar to those of mainframes. But what took 50 years for mainframes to achieve seems to have taken the microprocessors only 25 years, with each phase lasting five years instead of ten. In truth, there is no pioneer phase, since the computer had already been invented.

1970–74: First Microprocessors

In late 1971, Intel announced the first microprocessor, the 4004, which it had started designing in 1970 and shipped in early 1971 to Busicom—a Japanese vendor of desk calculators. At 5 Kops/s, it executed at approximately the same speed as ENIAC. It was followed quickly in 1972 by the 8008, which executed approximately one operation every 35 microseconds (or 29 Kops/s).

During this phase, many new microprocessors were introduced, among them Texas Instruments' 1795 in 1971, TMS 100 in 1972, and TMS 1000 in 1974; the Rockwell PPS-4 in 1973; the Motorola 6800 in late 1973; and the RCA 1802 in 1974. The 1795 was designed at the same time as the 4004 but never marketed. The PPS-4 and 1802, although successful as first products, never had any significant follow-ons. Only Motorola would become a formidable competitor to Intel.

The Micral, from R2E in France in 1973, has the distinction of being the first commercial microprocessor-based personal computer. It used the 8008.

1975–79: Broad Products

Introduced in 1974, the Intel 8080 is the symbol of this new phase, characterized by better hardware-development tools, better peripheral chips, better software support, and increased competition from the Motorola 6800—and from Zilog's Z80, introduced in 1976. The 8080 executed approximately one operation every 3.5 microseconds (or 290 Kops/s).

One of the last major microprocessors designed by TI, the TMS9900 was introduced in 1976. It had a fairly elegant architecture and the unusual characteristic of keeping its registers in main memory instead of on the chip itself.

The Altair, introduced in 1975, was the first personal computer to use the 8080. Microprocessors were still used as logic replacements, but their potential as minicomputer replacements became obvious. It is amusing to note that early personal computers shared with ENIAC a very primitive way to load data and programs: through switches!

The Apple I and II (based on the MOS Technology 6502, a variant of the Motorola 6800) in 1976, the Commodore Pet (based on the 6502) in 1977, and the Radio Shack TRS 80 (based on the Z80) established the base for personal computers, whose numbers would explode during the next phase.

In 1979, Visicalc, the first spreadsheet, and Wordstar, the first major word processor, established themselves as "killer applications," jump-starting the personal-computer software industry.

1980–84: Breakthrough Products

If there had been any doubts about the nature of microprocessors, they were eliminated by the Intel 8086, introduced in 1978, and the Zilog Z8000 and Motorola 68000, both introduced in 1979. The designers of these chips clearly acknowledged the influence of mainframes and minicomputers on their designs. IBM then gave a tremendous boost to the personal computer and the microprocessor with the introduction of its PC (based on the Intel 8088) in 1981.

Intel, threatened by its competitors, turned the inferior 8086 into the cornerstone of its dominating x86 architecture. The 432, an overly ambitious product introduced after much delay in 1981, was quickly forgotten in favor of the 80286, another flawed architecture that still succeeded by providing support for the PC/AT.

Minicomputer vendors introduced single-chip variants of their architectures, including the LSI-11 and the Micro Nova. Researchers at U.C. Berkeley designed the first single-chip RISC microprocessor in 1981.

1985–89: Competition From the x86 Architecture

By the time the Intel 386 shipped in 1985, microprocessors that supported many applications other than the PC were finally perceived as the main engines of this personal-data-processing revolution. The combination of IBM, Microsoft, and Intel proved to be too much for many semiconductor vendors that had introduced products in the previous five years but that now saw themselves squeezed out of the dominant market.

Zilog followed its Z8001 with a paged virtual memory version, the Z8003, in 1983. It announced and fully documented the Z80000 in 1983. But the version of the chip that supported segmented paged virtual memory, a cache, and TLB became available only in 1986, killing any competitive hopes for Zilog in the personal-data-processing market.

National designed an elegant architecture for the 32016, introduced in 1980, but despite several follow-ons like the 32032, the 32332, and the 32532 (in 1982, 1985, and 1987, respectively), it too had no significant presence beyond 1989.

Only Motorola would survive the consolidation, with design wins in workstations at Apollo and Sun (started in 1982) and in the personal computer market for the Apple Macintosh, introduced in 1984.

1990–94: Are We at the Zenith? Not Yet!

The comparison with mainframes probably breaks down with this phase. It is obvious that Pentium, introduced in 1993, was a major milestone in the evolution of microprocessors, but we are too close to this phase to predict its zenith. The only predictable decline is in the relative importance of personal-data processing compared with other applications, like games, multimedia, and networking.

With Merced coming in 1999, well outside this phase, we might wish for a decline of the x86 architecture. Unlike the IBM 360, the constraints of compatibility have loaded down the x86 architecture with many questionable features whose abandonment would benefit everyone by removing performance bottlenecks. For the IA-64 architecture to replace the x86 architecture while abandoning the warts of the past, it must provide full migration for all users and a significant improvement in performance. Nothing is certain at this point.

Thus, in the future, we may see the zenith and decline of one architecture (the x86), just as we saw the decline of one form of computing delivery (the mainframe). But as microprocessors are currently the symbol of programmable computing, we certainly cannot project their decline. We first need a new symbol of programmable computing to emerge; none has yet. M

In the next issue, we will cover in more detail how mainframe architectural concepts were adopted by microprocessor designers.

Dr. Bernard L. Peuto was the architect for the Z8000, the Z8, and several other Zilog products. Now the president of Concord Consulting, he develops business strategies for microprocessor-based products. He can be reached at 415.851.0194 or by e-mail at blpeuto@peuto.com.