MICROPROCESSORS, DIGITAL SIGNAL PROCESSORS AND MICROCONTROLLERS

M. D. Santambrogio, D. Sciuto, and **F. Sironi**
Politecnico di Milano, Dipartimento di Elettronica e Informazione, Italy

Keywords: Microprocessors, RISC, CISC, VLIW microprocessor, Superscalar microprocessors, Multi-processor, Multi-core microprocessor, Multi-threaded microprocessor, Threading, DSPs, Microcontrollers.

Contents

Summary

This chapter focuses on the main characteristics of today microprocessor architectural features by following an historical perspective. Main goal is to show how the architectural innovations in microprocessor design have been born and how they have impacted the diffusion of microprocessors, DSP, and microcontrollers in our everyday life and in many of the objects we use (from cars to washing machines to cellular phones, digital cameras, MP3 players, video games, and so on). After a brief introduction on the general functionalities of a microprocessor, the chapter will provide an historical view of the early days of microprocessors and developments until the 1970s. Then, the introduction of the Reduced Instruction Set Computers (RISC) and their main architectural features brought computer architecture a leap forward, providing more performance in less space. New architectural advances built upon such concepts allowed the creation of microprocessors exploiting the available instruction level parallelism inside inherently sequential code execution bringing computer architecture into the era of the more complex superscalar microprocessors. When the performance improvements obtained in such a way proved to be insufficient and offset by power consumption, in particular with the development of battery-powered devices, then the new advances were mainly related to the increase of the level at which parallelism is sought in the executed code, i.e., the thread level, and the replication of

microprocessors in the same chip. Main drawback is that, in order to fully exploit such capabilities, the software must be parallelized at a coarse grain level, i.e., tasks and/or threads. Following the same line of thought, the second part of the Chapter deals with Digital Signal Processors and their developments since their first developments in the 1970s to today's integration with microcontrollers in a number of devices that require more sophisticated mathematical operations for digital signal processing than those available in general in a microcontroller that can be considered a self-contained system with a simple microprocessor, memory, and peripherals.

1. Microprocessors

1.1. Basic Definitions of Microprocessors

A microprocessor combines most or all of the functionalities of a computer's Central Processing Unit (CPU) on a single integrated circuit.

A CPU represents the core of a computer providing the basic computational and control functionalities allowing the execution of machine instructions to perform basic arithmetic and logical operations on data. The tasks that usually a CPU must be able to perform are:

- Fetch instructions from memory.
- Decode instructions to identify which operations and data are requested.
- Fetch data from memory or any temporary storage inside the CPU, i.e., registers.
- Process data through arithmetic or logical operations.
- Write data to memory or any temporary storage inside the CPU, i.e., registers.

The simplified internal structure of a CPU able to perform such tasks is shown in Figure 1 alongside with the simplified internal structure of the computational part (or data path), which is shown in Figure 1 too. The data path is connected to the control part, i.e., Control Unit (CU), which manages its execution fetching instructions according to the Program Counter (PC) value from the memory and storing them, one after the other, in the Instruction Register (IR) where they are decoded. Based on the decoding the CU drives the data path to perform computations. The data path is constituted by the Arithmetic Logic Unit (ALU) that performs basic arithmetic and logic operations. The figure shows an example in which the ALU performs the addition of two operands stored in registers.

This simplified internal structure shows only the main components necessary to perform the main tasks of a CPU.

The increased availability of computers into society has requested more and more performance from them, which has been achieved mainly by increasing the speed at which the microprocessor is able to execute instructions. The driving factors allowing the increase of execution speed have been both technological, with the continuing reduction in size of the transistors integrated, and architectural, with clever solutions to employ this increasing number of transistors to speed up execution. These architectural advances will be shown in the rest of this Chapter, starting from a short history of the evolution of microprocessors in the market.

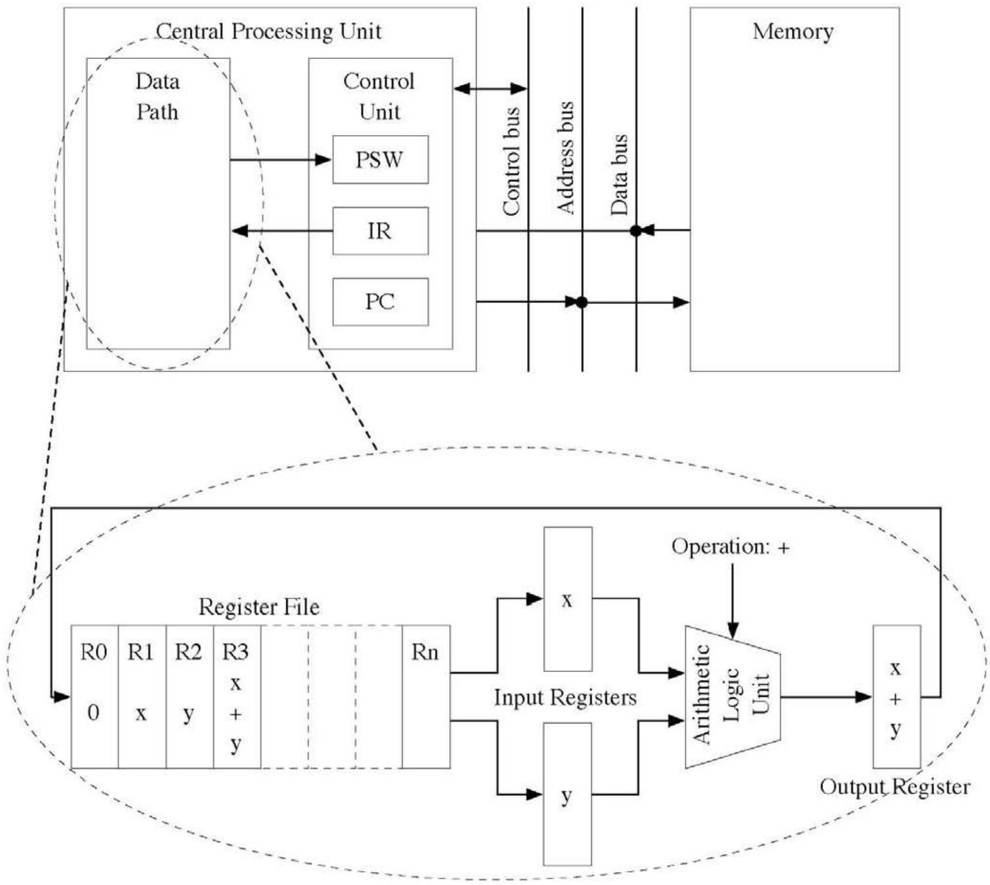

Figure 1. Simplified internal structures of a CPU (with memory access) and CPU's data path.

1.2. The Early Days of Microprocessors

The advent of microprocessors follows the advent of the integrated circuits. The idea of a computer on a single chip had been described in the literature as far back as 1952, and more articles on this topic appeared until the beginning of the 1970s. In 1968, Complementary Metal Oxide Semiconductor (CMOS) became available on the market and the number of transistors integrated on a chip started to increase constantly with the advances of integration processes. Observing this trend, Fairchild Semiconductor's director of Research & Development, Gordon Moore, who became afterwards the cofounder of Intel with Robert Noyce who was one of his colleague at Fairchild, observed in 1965 that the density of elements in integrated circuits was doubling annually, and predicted that the trend would continue for the next ten years. This observation has been known since then as Moore's Law, and was corrected in 1970 with the doubling of the number of transistors occurring every 18 months.

The first integrated circuits contained just a few transistors per wafer; by the beginning of the 1970s, integration processes allowed for thousands of transistors per wafer. It was only a matter of time before someone would use this capacity to put an entire computer on a chip; in fact, three projects delivered a microprocessor around the same time:

Intel's 4004, Texas Instruments (TI) TMS 1000, and Garrett AiResearch's Central Air Data Computer (CADC), all of them allowing 4 bits data management.

In 1968, Garrett AiResearch was invited to produce a digital system to compete with the electromechanical systems under development for the main flight control computer in the US Navy's new F-14 Tomcat fighter. The design was completed by 1970, and used a MOS-based chip as the core CPU. The resulting digital system was smaller and much more reliable than the electromechanical systems it competed against, and was used in all of the early F-14 Tomcat models. However, the system was considered so advanced that the Navy refused to allow the publication of the design, and continued to refuse until 1997. For this reason the CADC is fairly unknown even today.

The Intel invention of the microprocessor is strictly connected to its potential customer, Nippon Calculating Machine Corporation, which in 1969 requested Intel to design 12 custom chips for its new Busicom printing calculator. Intel did not have enough manpower for building the 12 custom chips and therefore, instead of creating many custom chips, proposed a new design: a family of just four chips, including one that could be programmed and therefore be used in a set of products. The set of four chips included a microprocessor, the Intel 4004, a supporting Read-Only Memory (ROM) to store custom programs, a Random-Access Memory (RAM) to store data, and a shift-register chip for the Input/Output (I/O) port. Intel offered lower prices in return for securing the rights to the microprocessor design and the rights to market it for non-calculator applications, allowing the Intel 4004 microprocessor to be advertised on the 15 November 1971 issue of Electronic News. The Intel 4004 became the first general-purpose microprocessor on the market, a "building block" that engineers could purchase and to execute custom programs to perform different functions in a wide variety of electronic devices. Its simplified architecture is shown in Figure 2.

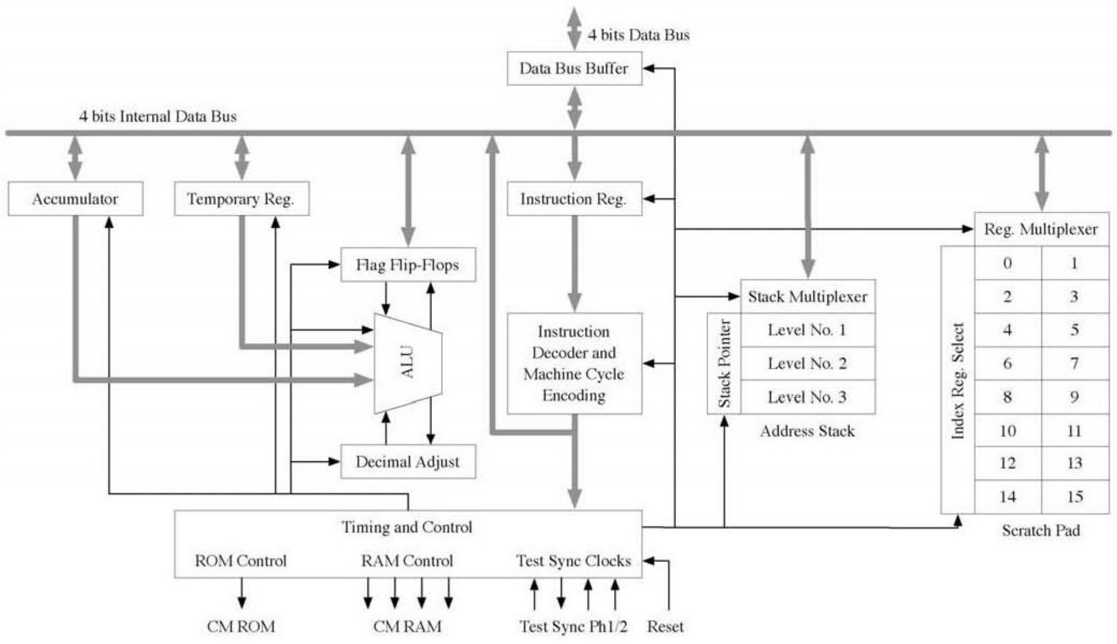

Figure 2. Simplified architecture of the Intel 4004.

A further step in the evolution of microprocessors was the birth of the 8 bits microprocessors. The first one was the Intel 8008 in 1972, followed by the more successful 8080 in 1974 and then Zilog Z80. The 8080 and Z80 were general-purpose microprocessors, which supported an extended set of instructions. The competing Motorola 6800 was released in August 1974. Its architecture was cloned and improved in the MOS Technology 6502 in 1975, rivaling the Z80 in popularity during the 1980s.

Both the Z80 and 6502 concentrated on low overall cost, through a combination of small packaging, simple bus requirements, and the inclusion of circuitry that would normally have to be provided in a separate chip (for instance, the Z80 included a memory controller). These features allowed the personal computer (PC) "revolution" to take off in the early 1980s, eventually delivering machines that sold for US $99.

Another important step in the development of the market of microprocessors was the introduction of the CMOS 65C02 in 1982 by the Western Design Center Inc. (WDC), which licensed the design to several companies. This microprocessor became the core of the Apple IIc and IIe personal computers, medical implantable grade pacemakers and defibrillators, automotive, industrial, and consumer devices. WDC pioneered the licensing of microprocessor technology that was later followed by ARM and other microprocessor Intellectual Property (IP) providers in the 1990s. Motorola introduced his successful MC6809 in 1978, arguably one of the most powerful, orthogonal, and clean 8 bits microprocessor designs ever designed, but also one of the most complex hardwired logic designs that ever made it into production for any microprocessor. The instruction set was orthogonal since the same instruction could access registers or memory through different addressing modes, making it more easily programmable, since it allowed a limited number of instructions to express the different operations. This also allowed keeping the control logic, responsible for decoding the instructions and the addressing modes and controlling the execution of the operations of the data path, to be implemented directly in hardware, given the relative simplicity. Microprogramming replaced hardwired logic at about this point in time for all designs more powerful than the MC6809, this decision was mandatory due to fact that the design requirements in terms of support of ever more complex instruction sets were getting too complex for hardwired logic. Microcode was developed as a simpler method of developing the control logic for a microprocessor, by substituting the combinatorial logic used by hardwired control with micro-program routines that allowed an easier implementation of complex machine instructions that needed different sequential operations to be executed. The microcode implementing each machine instruction is stored in a high-speed control store memory (ROM) in the microprocessor. Each micro-program is constituted by a sequence of microinstructions that control the components of the microprocessor.

In the mid seventies the first 16 bits microprocessors appeared in different types of computers from National Semiconductors, TI, Digital Equipment Corporation (DEC) with their PDP/11, a minicomputer capable of performing all functionalities of a mainframe (such as those dominant in the market from IBM) but at a much lower cost. WDC introduced the CMOS 65816 16 bits upgrade of the WDC CMOS 65C02 in 1984. The 65816 16 bits microprocessor was the core of the Apple IIgs and later the Super

Nintendo Entertainment System, making it one of the most popular 16 bits designs of all time.

Intel followed a different path, upgrading their 8080 architecture into the 16 bits 8086, the first member of the x86 family, which has powered most modern PCs. Intel introduced the 8086 as a cost effective way of porting software from the 8080 microprocessors, and succeeded with this strategy in acquiring more business. The 8088, a version of the 8086 that used an external 8 bits data bus, was the microprocessor that powered the first IBM PC, the 5150. Following up their 8086 and 8088, Intel released the 80186, 80286, and, in 1985, the 32 bits 80386, strengthening their PC market dominance with the microprocessor family's backwards compatibility.

A note should be made on Advanced Micro Devices (AMD). Jerry Sanders founded AMD in 1969. Like so many of the people who were influential in the early days of microprocessors (including the founders of Intel), Sanders came from Fairchild Semiconductor. AMD's business was not the creation of new products; it concentrated on making higher quality versions of existing products under license. For example, all of its products met the US military specifications (MILSPEC) requirements no matter what the end market was. In 1975, it began selling reverse-engineered clones of the Intel 8080 microprocessor.

The 16 bits designs dominated only briefly the market since they were shadowed quite quickly by the appearance of the 32 bits designs. While the first 32 bits microprocessor was the AT&T Bell Labs BELLMAC-32A, with first samples in 1980, and general production in 1982, the most famous of the 32 bits designs is the MC68000, introduced in 1979. The 68K, as it was widely known, had 32 bits registers; a 32 bits address space but used 16 bits internal data paths, and a 16 bits external data bus to reduce pin count. Motorola generally described it as a 16 bits microprocessor, though it clearly featured a 32 bits design. The combination of high speed, large (16 MB) memory space and fairly low costs made it the most popular microprocessor design of its class, used for instance in the Apple Macintosh personal computers. The first design was followed by different others that added new features, such as virtual memory support and floating point unit (FPU). The 68K family became the microprocessor of choice for any machine that was not running the DOS operating system. Motorola ceased production of the 68000 in 2000, but just before that, Motorola entered into the AIM - Apple/IBM/Motorola - alliance, which would eventually produce the first PowerPC microprocessor. Intel's first 32 bits microprocessor was the iAPX 432, which was introduced in 1981 but was not a commercial success. It had an advanced capability-based object-oriented architecture, but poor performance compared to other competing designs such as the Motorola 68000. In 1985 Intel launched the 80386, which was in line with the 8086 philosophy. Other designs included the interesting Zilog Z8000, which however arrived too late to the market to stand a chance and disappeared quickly.

These microprocessors were retrospectively named Complex Instruction Set Computers (CISC) since their instruction sets include instructions that can execute several low-level operations, e.g., load from memory, an arithmetic operation, and a memory store, and/or are capable of multi-step operations or addressing modes within single instructions.

-
-
-

Bibliography

Ars Technica (1999). RISC vs. CISC: the Post-RISC Era. http://tinyurl.com/75562. [This describes the evolution of microprocessors from RISC to CISC and then to hybrid solutions that cannot be called neither RISC nor CISC].

Ars Technica (2002). *Introduction to Multithreading, Superthreading and Hyperthreading*. http://tinyurl.com/3gqxc7j. [This highlights various Simultaneous Multi-Threading technologies].

Atmel (2011). *Microcontrollers.* http://tinyurl.com/3dpr3qz. [This presents Atmel's microcontrollers].

Computer History (2011), *History of Microprocessors.* http://tinyurl.com/3uvay53. [This discusses microprocessors evolution].

CPU-World (2011), AMD 2901 Bit-Slice Processor Family. http://tinyurl.com/3wd43hb. [This details the Bit Slice microprocessor family].

Dyer D. (1998), *TRW: Pioneering Technology and Innovation Since1900*. Harvard Business Press. [This tells TRW's history].

Giovino W. A. (2002), Microcontrollers and DSPs – Will the Two Worlds Intersect? http://tinyurl.com/3ch38ua. [This outlines possible evolutions of microcontrollers and DSPs towards a single entity].

Hennessy J. L., Patterson D. A. (2006), *Computer Architecture: A Quantitative Approach*. Morgan Kaufmann, forth edition. [This discusses computer architectures from the ground to the top].

Intel (2011). Intel Museum: Corporate Archives. http://tinyurl.com/3pxf7v2. [This tells the history that shaped Intel].

Junior D. (2010), *DSP Multimedia Handset* – Android Based Home Phone. http://tinyurl.com/3ofcu2t. [This represents a futuristic example of DSP-powered device].

Patterson D. A., Hennessy J. L. (2008), *Computer Organization and Design: The Hardware /Software Inter-face*. Computer Architecture and Design. Morgan Kaufmann, forth edition. [This discusses computer organization looking at the intersection of the hardware and the software].

Silicon-Far-East (2011), *Transistor-Transistor Logic (TTL)*. http://tinyurl.com/3dfo9vk. [This describes the inner working of TTL].

Strauss W. (2003), *DSP Technology: Pivotal To Consumer Electronics.* http://tinyurl.com/3oal2k3. [This highlights viable way to use DSP in consumer electronics].

Strauss W. (2003). Early history of the DSP chip. http://tinyurl.com/3dvm87v. [This discusses DSPs evolution].

Valenzuela A. (2005). FIR Filters. http://tinyurl.com/3sajudc. [This talks about FIR filters].

Warner W. (2004). Great moments in microprocessor history. http://tinyurl.com/684dkt. [This discusses really important steps in microprocessors evolution].

Xilinx (2011), *PicoBlaze 8-bit Embedded Microcontroller User Guide*. http://tinyurl.com/3todpc9. [This describes the Xilinx Pico Blaze 8bit microcontrollers].

YourDictionary (2011), *Bit Slice Processor.* http://tinyurl.com/44ljgg3. [This defines the Bit Slice microprocessor].

Biographical Sketches

Marco D. Santambrogio received his Laurea (M.Sc. equivalent) degree in Computer Engineering from the Politecnico di Milano in 2004, his second M. Sc. degree in Computer Science from the University of Illinois at Chicago (UIC) in 2005 and his PhD degree in Computer Engineering from the Politecnico di Milano in 2008. Dr Santambrogio was at the Computer Science and Artificial Intelligence Laboratory (CSAIL) at Massachusetts Institute of Technology (MIT) as postdoc fellow and he is now assistant professor at Politecnico di Milano and research affiliate at MIT. He conducts research and teaches in the areas of reconfigurable computing, self-aware and autonomic systems, hardware/software codesign, embedded systems, and high performance processors and systems.

Donatella Sciuto received her Laurea in Electronic Engineering from the Politecnico di Milano in 1984. She received her Ph. D. in Electrical and Computer Engineering in 1988 from the University of Colorado, Boulder. She received her MBA from the Scuola di Direzione Aziendale, Bocconi University, in 1991. Since 2000 she is Full Professor at the Politecnico di Milano and she is now Vice Rector of the University. Her research interests include embedded systems design methodologies and architectures.

Filippo Sironi received his Laurea (M.Sc. equivalent) in Computer Engineering from the Politecnico di Milano and his Master of Science in Computer Science from the University of Illinois at Chicago in 2010. He is currently pursuing his Ph.D. in Computer Engineering at the Politecnico di Milano, Dipartimento di Elettronica e Informazione. His research interests include self-adaptive computing, operating systems design and implementation, and computer architectures.