ARM Gets Serious About IP

Second in a Two-Part Series About the History of ARM

In just under 30 years, our use of technology and our expectations have changed dramatically, and semiconductor IP leader ARM has evolved to meet them.

And It Came to Pass . . .

Turning the clock back to 1979, technology was already shrinking in size and significant steps were being made in technology development.

Sony released the Walkman, which at US$200 cost a significant amount of money. Texas Instruments made its debut in the computer arena, and the Motorola 68K and Intel 8088 were released. Hermann Hauser and Chris Curry, together with a team of students and researchers at Cambridge University in England, were conceiving a personal computer company, dubbed "Acorn Computers," which was to change the face of computing as we know it.

Acorn Computers Creates the First PC

The Acorn Atom was the first home computer off the production line, complete with a 1-MHz processor and 12 kB of random operating memory (ROM) and random access memory (RAM). The Atom was the next generation of the metal oxide semiconductor (MOS) technology 6502-based machines that the company had been creating since 1979. It was a cut-down Acorn System 3 with an integral keyboard and cassette tape interface but without a disk drive. In 1980, it sold for the princely sum of £170.

In late 1982, Acorn released an upgrade ROM chip for the Atom which enabled users to switch between Atom BASIC and the more advanced BASIC used by the BBC Micro.

Digital Object Identifier 10.1109/MSSC.2009.933674

Apple Introduces the Lisa

Apple, however, was nipping at Acorn's heals and had its team work on a personal computer (PC) with a graphical user interface (GUI) targeted at business users. The result was the Lisa. The Lisa was first introduced in January 1983 at a cost of US$9,995. It was one of the first commercial PCs to have a GUI and a mouse.

Acorn fast realized that increased performance was the way forward. Having investigated several alternatives, it approached Intel to sample the 80286 processor. Intel said no, and this refusal—unknown to Acorn at the time—would revolutionize computing.

ARM—The Conception

As a direct result of Intel's decision, a development team was set up at Acorn Computers to build a compact RISC central processing unit (CPU). The team was led by Sophi Wilson, who had built all the versions of BASIC for the BBC Micro, and Steve Furber. Their goal was to create a low-latency input/output system, similar to the MOS 6502 technology used in Acorn's designs. The result was the Acorn RISC machine.

Acorn's ARM Microprocessors Debut in 1984

Furber defined the architecture, while Wilson developed the instruction set. The team produced development samples of the first ARM in the spring of 1985, yielding working silicon the first time it was fabricated using a 3 μm two-layer metal process at VLSI. It was the arrival of a microprocessor that was to create a whole new generation of computing power. The processor did the same amount of work as 16-b microprocessors, but with one tenth of the transistors, and hence a lower power supply requirement. The ARM1 was a prototype and was never actually released.

"At 1 p.m. on 13 April 1984, the first ARM microprocessors arrived back from the manufacturer—Plessey," Furber is quoted as saying. "They were put straight into the development system which was fired up with a tweak or two and, at 3 p.m., the screen displayed: 'Hello World, I am ARM.'"

Silicon Strategy Results in the First RISC Processor

The ARM processor was part of a strategy in which it was decided that a computer should be designed on silicon, rather than made up of third-party components. Hauser later said this made Acorn one of a small, select group of computer companies in that created its own technology from the ground up.

The small-die-size ARM2 quickly saw the light of day and was a very simple RISC processor with only 25,114 transistors—the same number as the Z80 or 6502 that Acorn utilized in its BBC Micros, but 20 times faster! This was a milestone—the birth of the world's first RISC processor.

Much of its simplicity came from not having to use microcode. It was used in the original Archimedes, the successor to the BBC Micro. It provided an average performance of 4–4.7 MIPS and was clocked at 8 MHz, which is quite unbelievable given today's speeds. It featured a true 32-b data bus and a 26-b (64-MB) address space, with 16 32-b registers and no on-chip cache. Program code had to lie within the first 64 MB of memory, as the program counter was limited to 26 b. This was primarily because the top 6 b of the 32-b register served as status flags. Although the ARM2 had its limitations, it heralded sound. A boon to home and educational computers, thanks to multiply and multiply accumulate instructions.

The ARM2 processor was manufactured by VLSI, which also had the rights to sell chips using the design. The first ARM2 product was the ARM Development system which incorporated an ARM processor and three support chips, 4 Mb of RAM, and a set of development tools coupled with BBC BASIC, with a price tag of around US$6,366. The three peripheral chips were designed by Tudor Brown, Mike Muller, and Steve Furber.

Acorn's Multimedia "Archimedes" PC Confronts the IBM Behemoth

Acorn then rolled out the Acorn Multimedia PC, the infamous Archimedes. The first models were released in June 1987, as the 300 and 400 series. Unfortunately, the behemoth IBM had cornered the PC market and was quickly seen as a standard. Acorn's system was seen as somewhat avant garde, offering a new operating system, a new processor, and no bank of software applications to support it. In May 1989, the 300 series was phased out in favor of the Acorn A3000. The A3000 used an 8-MHz ARM2 and was supplied with 1 MB of RAM.

Development continued and, in 1991, the A5000 was unveiled featuring a 25-MHz ARM3 processor with 4 KB on chip data and instruction cache, an option of 2 or 4 Mb of RAM, and either a 40-Mb or 80-Mb hard drive. Graphics also took a leap forward. The A5000 was capable of displaying VGA resolutions of up to 800×600 pixels.

Apple Bases its "Newton" PDAs on the ARM610 RISC Processor

At the same time, Hermann Hauser had turned his attention to the emerging personal digital assistant (PDA) market with the Active Book Company, which intended to use ARM in its line of personal communicators. Hauser was looking to exploit pen-based systems, which he thought would be much easier. At the same time, Apple was looking to make its mark in the PDA market and launched the Newton. Apple first checked out AT&T's low-power processor, dubbed the Hobbit. But John Stockton, research fellow of VLSI, and Larry Tesler, who headed the team, were disciples of ARM and believed it was the only way to go. The original Newtons were based on the ARM610 RISC processor and featured handwriting recognition software.

ARM Is Born

Apple was won over by ARM technology, but, for competitive reasons, Stockton and Tesler realized that Apple wanted a separate company from Acorn. In just six weeks, in 1990, a joint venture was negotiated between Acorn Computers, Apple Computer, and VLSI Technology that was intended to further develop the Acorn RISC chip. Acorn provided the people and retained a 40% equity stake, Apple Computers provided financial support and gained a 40% stake, and VLSI provided the design tool and received a 5% stake. The startup period was a tough one, and ARM nearly ran out of money. But, by 1993, it had silenced its critics and those who said such partnerships don't exist by turning in its first profitable year!

Hermann went on to do a deal with AT&T over the Active Book Company, which became known as EO Ltd. The EO design team switched from ARM to AT&T's Hobbit offering and the Apple Newton was delivered with an ARM processor! The rest, as they say, is history!

Hauser Recruits First ARM Team

A key twist in the tale involved the meeting of 12 engineers in an English pub over a pint. Hermann Hauser had created the Cambridge initial caps, and he was supplied chips by Robin Saxby, managing director of European Silicon Structures (ES2). Hauser, who had a good working relationship with Saxby, offered him the job of CEO after an initial interview. Saxby, however, said he wanted to meet the 12 engineers before accepting the job. Saxby later said that a key decision was going to be whether he could grow his team from the inside, or whether he'd have to go down the more expensive route of hiring outside engineers. It was no surprise that a number of Acorn staff moved across to ARM, under such inspirational management.

Saxby took the job in 1991, shortly after its launch date Advanced RISC Machines Ltd. on 27 November 1990 in Cambridge, England. The ARM development team moved out of Acorn's Cambridge headquarters to a converted 18th century barn in the chocolate box Fenland village of Swaffham Bulbeck, ten miles outside Cambridge.

The Apple Newton Uses ARM's 32-b RISC CPU

The ARM development team's mission was clear—to further develop the ARM processor and facilitate its use by system developers as a standalone processor, or alternatively as a macrocell with custom logic or other ARM components added to the mix to create a custom chip. This was outlined in the company's first press release, which stated that the new company had been created to "address and attack the growing market for low-cost, low-power, high-performance 32-b RISC chips."

ARM's first development was the next step forward from the ARM3 processor, which was dubbed the ARM6 and included full 32-b addressing, an improved video controller, and a floating-point processor. It also included endian-ness (byte sex) support, one of the features Apple requested. ARM's first big commission was to design a CPU for Apple to be used in a PDA. This device was christened ARM600, which later became the ARM610, used in the Apple Newton personal organizer.

ARM Introduces Cross-Platform Development and Hardware Evaluation Toolkits

At the same time, ARM recognized the necessity for development kits and developed the ARM Cross Development Toolkit, a suite of tools enabling designers utilizing a variety of platforms to use ARM

(Continued on page 68)

Dr. Erich F. Haratsch (right) receiving a certificate for the VLSI-DAT Best Presentation Award from Prof. Tzi-Dar Chiueh, General chair of 2009 VLSI-DAT.

Dr. Eng-Huat Toh (right) receiving a certificate for the 2008 VLSI-TSA Best Student Paper Award from Dr. Roger De Keersmaecker, General chair of 2009 VLSI-TSA.

2008 VLSI-TSA Best Student Paper Award

The 2008 Student Paper Award winner was "P-Channel I-MOS Transistor Featuring Silicon Nano-Wire with Multiple-Gates, Strained Si1-yCy I-Region, in situ Doped Si1-yCy Source, and Sub-5 mV/Decade Subthreshold Swing," by Eng-Huat Toh, Grace Huiqi Wang, Doran Weeks, Ming Zhu, Trevan Landin, Jennifer Spear, Lap Chan, Shawn G. Thomas, Ganesh Samudra, and Yee-Chia Yeo of the Department of Electrical and Computer Engineering, National University of Singapore. The VLSI-TSA Symposium first presented a Best Student Paper Award in 2005. A conference committee determines the one paper that is best written and presented by a full-time student.

2008 VLSI-DAT Best Presentation Award

The 2008 VLSI-DAT Best Presentation Award winner was "A Radix-4 Soft-Output Viterbi Architecture," by Dr. Erich F. Haratsch of LSI Corporation, USA, Erich F. Haratsch, and Kelly K. Fitzpatrick. Since 2006 the VLSI-DAT Symposium has presented a Best Presentation Award to the paper that is best organized each year. In every session, the audience is asked to complete a survey evaluating each paper for technical content, slide quality, presentation, and question and answer handling.

Dr. Pouria Bastani of IC Santa Barbara won the 2009 DAT Best Paper award.

2008 VLSI-DAT Best Student Paper Award Presented for the First Time at VLSI 2009

The VLSI-DAT symposium presented a Best Student Paper Award for the first time this year to the paper that was best written and presented at the 2008 conference by a full-time student.

The winner was "An Improved Feature Ranking Method for Diagnosis of Systematic Timing Uncertainty," by Dr. Pouria Bastani of the University of California Santa Barbara, USA, Nicholas Callegari, Li-C. Wang, and Magdy Abadi. Dr. Pouria Bastani was unable to attend the conference.

Submissions were evaluated by a technical committee that assessed the written quality of the paper as well as the oral presentation during the conference. The Best Student Paper Award for 2009 will be presented at the next conference. Students must indicate with their submission that they would like to be considered for this award.

For more information on both of these conferences, please contact Ms. Clara Wu, VLSI-TSA Symposium secretariat, vlsitsa@itri.org.tw and Elodie Ho, VLSI-DAT Symposium secretariat, vlsidat@itri.org.tw. SSC

ASSOCIATE EDITORS' VIEW *(Continued from page 9)*

development tools, assembler, compilers, and debugging and evaluation programs. ARM also rolled out hardware evaluation kits enabling designers to test the ARM6 processor. This allowed them to develop operating system and support software for use with their designs before availability of finished product.

Coming of Age

ARM had long seen the importance of a licensing model, prompting it to license its designs to chip foundries that could sell the chips, providing ARM with royalties, rather than having to establish its own fabrication facilities. Under this strategy, the intellectual property (IP) is designed into products by other companies who subsequently sell and market products utilizing the technology. VLSI Technology, which had built previous ARM chips, together with Apple and Acorn's custom logic devices, was the very first licensee. GEC Plessey and Sharp followed in licensing ARM technology.

Sharp Corporation, NIF, and TI Sign Licensing Deals

Establishing a relationship with a major Japanese manufacturer was a key component of ARM's strategy. In March 1993, Sharp Corporation of Japan signed a deal to manufacture and market ARM processors and associated products.

At the same time, ARM came of age on the world stage by receiving a significant investment from the Japanese investment house, NIF. ARM's investors now included European companies, Acorn, and via Olivetti, U.S. companies Apple and VLSI Technology, and NIF in Japan. Shortly after these agreements were signed, Texas Instruments was added to the list of ARM partners. The company looked to use ARM macrocells as a basis for its custom embedded controllers.

ARM now boasted offices in Silicon Valley and Tokyo to maintain a close relationship with its licensees and customers and to promote its existing ARM devices and roadmap.

ARM's QuickDesign Revolutionizes Low-Cost, Rapid Development Partnerships

In 1992, ARM chose the technology expo Comdex to roll out QuickDesign. ARM had always been able to develop custom processors and controllers rapidly from its library of standard macrocells. With QuickDesign, it could design a custom part from standard parts and interface it with custom technology developed by ARM. Alternatively, it could work in partnership with a customer to quickly produce a low-cost product with a fast time to market.

ARM's objectives were now clearly stated—developing processors with RISC design principles to provide high performance, low price points, easy customization, and short design time.

IPO Joint Listed on LSE and NASDAQ in 1998

By the end of 1997, ARM had grown to be a £26.6m business with £2.9m net income. In 1998, when the company changed its name from Advanced RISC Machines to ARM, it was shipping 50 million cores a year. It was time to take out an IPO listing on both LSE and NASDAQ. The decision to take out a joint listing covered two bases, the two major shareholders in ARM were English and American and the company wanted existing Acorn shareholders in the United Kingdom to remain involved in the company. Second, ARM was confident that NASDAQ would provide it with the valuation for which it had worked so hard. The move proved a success—the listings were oversubscribed and reached a 20–30% premium on the first day of trading!

ARM Matures and So Does IP

Way back at ARM's conception, it had seen the importance of IP and licensing within its business model. The company has never departed from this route.

In December 2004, stockholders from both ARM Holdings PLC and Artisan Components Inc. voted their approval of ARM's acquisition of Artisan. There was a great synergy between both companies providing silicon intellectual property (IP). While ARM focused on processor cores, Artisan's products included embedded memory, standard cell, input/output, analog, and mixed-signal components. Under the ARM umbrella, the Artisan operation became known as the PIPD (Phyical Intellectual Property Division). The deal was sealed in 2005.

In 2006, ARM took the wraps off a high-performance implementation of the ARM1176JZF-S processor, enhanced with the ARM Artisan Advantage cell library and memories. This achieved a frequency of more than 750 MHz in a high-performance 90-nm foundry process while occupying only 2.4 mm^2 of silicon area. This entailed a significant performance increase through a combination of collaborative design, advanced physical IP, and process technology.

The combination of the ARM processor business and the physical IP business acquired from Artisan had delivered technically, as well as commercially, creating a best-in-class 90-nm design for semiconductor partners serving the high-volume and rapidly expanding digital consumer market.

ARM has continually looked to top-tier licensees of ARM processors to add the physical IP product to the processor product. The reason for doing so, ARM knows, is compelling. If physical IP and processors are developed together, it makes the whole process much faster and more reliable.

In 2006, ARM acquired Soisic, a developer of physical IP based on silicon-on-insulator (SOI) technology. In addition, ARM and Taiwanese foundry provider United Microelectronics Corp. (UMC) jointly developed an SOI-based processor core at the 65-nm node.

At the time, ARM CEO Warren East accepted that increasing design costs would prompt more collaboration in the industry. East said that ARM intended to offset the risk to companies by helping chipmakers get a jump on the physical IP market at the 32-nm node.

IP and a Flexible Business Model Propel ARM's 32-b CPU Product Family

While ARM had been known around the globe for its processor cores, it was fully aware that physical IP—memory cores, phase-locked loops, standard cells, and other IC building blocks for system-on-chip design—were key to its growth. Moreover, the company understood the importance of a flexible business model with a variety of entry points for potential customers.

***

Read how the ARM product family captured 75% of the world market for 32-b RISC CPUs after 2006 in the Winter 2009 issue of *IEEE Solid-State Circuits Magazine* (volume 1, number 1).

—Tony Harker
tonyharker1@btinternet.com