Chips with everything

February 24 2009

XMOS

(PhysOrg.com) -- While the technology to make computer chips smaller and cheaper progresses each year, the fundamental structure of the chip - the computer architecture - has remained the same for decades. This led Professor David May in the Department of Computer Science to think about what a computer chip should look like for the twenty-first century. Today, his technology has unified the hardware and software worlds into one environment, such that hardware is software.

Take a look around and you will soon realise that you are surrounded by small computers - in your mobile phone, your iPod or music player, your TV and wireless network and, of course, in the computer on your desk.

But there are many others, embedded and invisible, in lighting systems, security systems and your car. The number of embedded computers is growing exponentially each year, but they are often based on ideas that are decades old. A new architecture only comes along occasionally, but when it does, it heralds tremendous changes.

Back in 2000, it was becoming clear to David May that although the cost of new electronic designs was increasing - along with the time needed to complete designs - the new market opportunities would require a wider range of products that could be rapidly adapted as fashions changed. This led him to think about how electronic systems could be built in the twenty-first century and whether standard electronic components could be adapted to specific needs simply by programming, thereby providing customisation before, at, or even after the point of sale.

His new idea in 2000 was to find a way to design electronic systems as collections of small programmable computers

May's early ideas - over 30 years ago - were inspired by research on robotics and on looking at how collections of small computers could be used to build intelligent systems; how they could communicate with each other and, just as importantly, with the outside world. His new idea in 2000 was to find a way to design electronic systems as collections of small programmable computers, at a cost and power consumption low enough for them to be embedded in toys or even clothes.

"Research in computer design is an interesting concept, as there are many different ways of acquiring knowledge and pushing back the boundaries of our understanding," says May. "Computer science, and computer architecture in particular, really is a different kind of activity as it is more focused on creating and building things than on scientific

investigation. The difficulty is that many of the things we design are so complex that we have to employ scientific methods to understand how the design will work, and even how to understand them after we have built them. It is sometimes said that we are engaged in the science of the artificial."

The first draft of the new computer architecture used a large array, a 'sea', of low-cost processors on a chip that would be configured in the software, rather than in the hardware. But as there were no potential investors around in the aftermath of the Dot.com crisis, that was as far as it went. Four years later, along came a promising fourth-year undergraduate student looking for a project. Ali Dixon took the draft, which included some ideas for the instructions that control the computer, as well as the communications technology and links, and used this to create a model of the processor that could be simulated on a PC.

XMOS has set out to make chips itself and build a significant company around the new ideas

Dixon came back with fundamental questions about how to write the software that would run on the processors, and so a software tool called a compiler had to be built to take the instructions and convert them into the 1s and 0s that control the processor.

This was no simple task as the key to the new architecture was to have lots of small processors all working independently, but running concurrently, which creates a significant problem for traditional software that is essentially designed to be sequential. Moving from this sequential world to one where processors run concurrently required tremendous amounts of research.

A lot of attention was focused on the ability of the array of processors to

handle the input and output of data. "This has been a neglected area over the past 30 years of computer technology, but providing a flexible way for the concurrent software to control data passing in and out of the chip is fundamental to the new architecture," explains May, "and so much of the development was breaking completely new ground". This led to 17 patent applications, covering new ways of handling input and output, memory access and instruction scheduling.

In July 2005, the company XMOS was formed in order to commercially develop the new architecture, now called Xcore, and, in September 2007, XMOS secured $16 million of venture capital funding. The new company has taken the ideas from research, and the investment from the globally renowned venture capital firms, and built a team of experienced and enthusiastic designers and engineers.

But instead of just providing the design to other chip makers, XMOS has set out to make the chips itself and build a significant company around the new ideas. "Building the devices is an expensive business and outside the remit and capabilities of the world of research. It has always required new companies to bring together the engineering, business expertise and funding to turn a research idea into reality," says May.

XMOS has taken the concurrent architecture and created the idea of Software Defined Silicon, where the chips provide the performance and low power that usually comes with a hardware implementation, but which can be programmed like any computer, with the input-output and communication between the multiple processors controlled simply through software instructions. The company is now producing a chip with four processor cores that will cost just $10, while a single core version can be even cheaper. There are now two development kits that engineers are using to write all kinds of applications and May wants to see these used in universities to teach computer science and engineering.

The speed at which the company's technology has moved from university science to a commercial proposition has been remarkable. What started as an undergraduate project just five years ago is now a major semiconductor business with plans to revolutionise the consumer electronics market.

Provided by University of Bristol

Citation: Chips with everything (2009, February 24) retrieved 9 April 2024 from
https://phys.org/news/2009-02-chips.html