Mateo Valero, Director of the BSC

BARCELONA SUPERCOMPUTING

CENTER AND MICROSOFT CREATE JOINT RESEARCH CENTRE FOR PARALLEL COMPUTING

On 18 January 2008, Microsoft and the Barcelona Supercomputing Center (BSC) announced the creation of the BSC – Microsoft Research Centre, which focuses on the way microprocessors and software for the mobile and desktop market segments will be designed and interact over the next 10 years and beyond.

The advent of many- and multi-core processor computing architectures will make it possible to deliver enormous computational power on a single chip, with profound implications for the way software is developed. Researchers from BSC and Microsoft are addressing the challenges and opportunities that massively parallel processing presents, and are specifically focusing on optimising the design and interaction of hardware and software architectures to take advantage of the new computing power.

Futures recently asked Mateo Valero, Director of the BSC, about the work of the BSC-Microsoft Research Centre.

QUESTION: It is becoming increasingly impossible to build faster processors; and at the same time, more miniaturisation of processors brings heating problems that possibly we are not able to handle. Is the innovation of computing hardware and with it, computing itself, slowing down?

MATEO VALERO: We narrowly averted crashing onto a wall head-on: the power-density wall. It was a narrow miss, the thousands of engineers that worked on cancelled processors can testify on that. We can talk about other walls as well: the memory wall, the Instruction Level Parallelism (ILP) wall, etc. The way around those walls is multi-core architectures with many processors running in parallel.

So now, rather than having one monolithic fast processor, you have 128 processors. Each of those processors is definitely much smaller in area, dissipates much less power and is somewhat slower when compared to the large monolithic way. However, it is the aggregate performance that matters. You break down a large problem into 128 smaller problems that can then be executed much faster in parallel – and then you have a winner. We will also see more heterogeneous architectures; and high performance processors will adopt a systems-on-chip approach with many specialised cores; those architectures will likely reorganise themselves based on application needs.

Multi-core architectures open up new horizons. You can do neat tricks such as over-clocking the frequency of some cores even beyond the single core Thermal Design Power (TDP) limit; whereas if most cores are idle, you will still be substantially below the chip TDP. In this era, it becomes possible to run today's supercomputing applications on those hundreds of cores on chip; but we need further research into how the interconnection mechanisms influence this mapping. This is one of the topics we are working on in the Barcelona Supercomputing Center (BSC).

QUESTION: If getting a large number of processors to run in parallel is a way out, isn't it limited by how many we can get to run in parallel? And doesn't it become increasingly complicated for engineers to develop software that matches the requirements of those systems with many processors?

MATEO VALERO: The aim is to apply latest concepts from research in computer science, such as transactional memory. It also means that theoretical concepts from computer science and programming concepts will guide the design of processors. For a single application, we are limited by Amdahl's Law, i.e. by the inherently serial section of the application. More research is necessary at both the algorithm and application levels, since we always come up with penalisations of sub-problems that we thought were inherently serial before.

There is increased pressure on the software engineers to write code that efficiently utilises those new many-cores. We can say that if we leave things as they are, we may be facing a new wall: a software efficiency or performance wall. The key to scaling this wall is to overhaul the way we design processors. Processor design in the many-core era needs to be driven by software needs and requirements. Theoretical concepts from computer science and programming concepts will be the stars in this new era - for example, transactional memory. Hardware support for transactional memory enables software developers to write many-core code more efficiently.

Sophisticated run-time environments will provide an abstraction layer between the computer hardware and infrastructure on the one hand and the application layer on the other hand. However, it is important to shield the complexities of the hardware from the programmers. Carefully thought-out programming models could lead the way to making the life of the programmer easier. At BSC, we are working on programming models such as the Cell Superscalar which makes it easier to parallelise applications for one of the current multi-cores, the Cell Processor.

QUESTION: It seems that, by combining software concepts and hardware architecture with the software concepts guiding hardware design, software and computer science innovation is at the edge of an exciting era. Do you agree?

MATEO VALERO: Absolutely, the old Chinese curse 'May you live in interesting times' is our motto. I believe that we are witnessing another revolution in computing, one that will fuse software and hardware researchers firmly together in holy multidisciplinary matrimony forever! We are certainly proud to be one of the contributors to this revolution.

WHEN COMPUTERS COME TO THE AID OF MATHEMATICS

"Mathematics must make a great leap forward": this is the challenge that Georges Gonthier has set himself.

Georges Gonthier, Senior Researcher, Microsoft Research Cambridge

Since the autumn of 2006, the Microsoft Research Cambridge researcher has been leading the 'mathematical components' project at the joint research centre set up by Microsoft and the French National Research Institute for Computer Science and Applied Mathematics (INRIA). The 13 researchers working on the project are tackling the thorny problem of computer proof, which proves that a mathematical demonstration is thoroughly accurate. A problem that, according to Gonthier, "becomes increasingly difficult as the theorems become more complex."

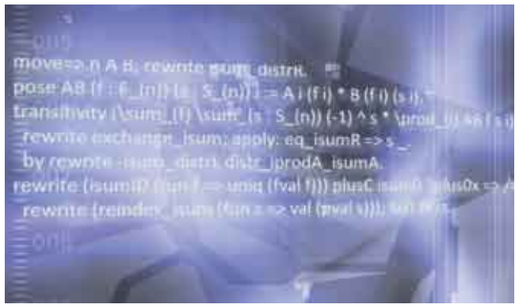

To achieve their ends, the research team is developing a computerised representation of mathematical theories. In other words, a detailed description, not only of this theory, but also of the way it is used, that will enable a computer to automatically check the accuracy of the theorems. It's a field where much remains to be done.

"In most cases," says Gonthier, "when mathematicians describe a theory in an article or a paper, they present a formal description, but they do not explain how to use the theorem." With this work, Gonthier, a former INRIA researcher, is continuing research that started at the beginning of the new millennium: in 2005, he was the first to demonstrate by computer the Four Colour theorem, which states that regions on a map may be coloured using no more than four colours in such a way that no two adjacent regions are of the same colour. Gonthier's present work takes some inspiration from the unprecedented development of software components, which was made possible by modern computer languages.

The joint Microsoft/INRIA project aims to foster the use of increasingly sophisticated formal mathematics. Ultimately, it will improve the certification of software - with possible applications including air traffic control, radiotherapy and banking. "Today, the construction of proof is done in a rather 'artisanal' way," observes Georges Gonthier. "Our aim is to provide proof that the theories used are reliable, rather than assuming this is the case based on the authors' fame or the claims made for the product."