

Topic 09

Parallel Computer Architecture - What Is Its Future?

Chris Jesshope

Global Chair

The Past

It gives me great pleasure to write this introduction to the Workshop 9 on Parallel Computer Architecture. This subject has a long and interesting history and is still of great interest to me even though I have been involved in it since the mid 1970s, having been one of the first users of the pioneering Illiac V Computer, which was installed at NASA Ames. Much has happened since those early pioneering times. I am sure that you are as aware as I am of the driver of the rapid changes in this field over its first two decades. This has primarily been the exponential growth of the underlying silicon technology. Ironically the braking factor, in innovation at least, over the last decade has been the market forces in the same industry and its economies of scale. It is now no longer viable to set up a silicon foundry and design a processor and compete in the market, the investment is too large. What has happened therefore is that we have seen the diversity in this field dissipate, due to economic factors. But not only this, I must add, because there is a maturity of the knowledge in this area. Indeed the impression gained of the subject over the last decade has been one of consolidation and the application of sound engineering. Parallel Computer Architecture has come of age but the outcome is not as the early pioneers may have imagined it. The major commercial thrust is in the business sector where the development of server architectures has provided much of the innovation, using commodity microprocessors and bus-based symmetric multi-processor architecture. There is also development of massively parallel "supercomputers" but the number of manufacturers now in this sector of the business has been severely curtailed and again the commodity processor rules the roost.

The Future

So, where are we going from here? For what it may be worth, I will give you my opinion of the longer term future, so that you can put the papers to be presented in this workshops in perspective. I believe that the continued growth and abundance of silicon real estate has allowed inefficiency to creep into the design of microprocessors. The current thrust has been to extract as much instruction-level parallelism as possible from existing code and to use out-of-order execution

and speculation in order to overcome the many dependencies that are to be found in compiled code. Just looking at the silicon real estate used in support of these techniques, one is led to wonder why a more explicitly parallel solution has not evolved. The silicon area used in these current, complex microprocessors would support many simpler pipelined datapaths (and their register sets). Why then has such an explicitly parallel approach not been followed, when at best, even with 6-way, superscalar issue it is seemingly impossible to get an IPC of more than 2. The problems in this approach are soluable; multi-threading in microprocessor design allows efficient use of high latency shared memories and caching can limit the parallel slackness required to exploit a number of parallel datapaths on a single chip. The answer is that such a scheme would require recompilation, or at least very good analysis tools that could extract explicit parallelism from compiled code. If we look at the history of microprocessor design however, it has made one major paradigm shift, from CISC to mainly RISC based design. A shift from implicit to explicit parallelism, which would allow the development of on-chip, parallel architectures, would not therefore be so far fetched and the knowledge required is already being published in today's literature, including these proceedings. Another area that the abundance of silicon has aided is the massive advance in custom computing. This field exploits the FPGAs that were initially developed for integrating custom logic onto a single die. Because these chips can be programmed from an array of uncustomised logic gates, at system design time, there are the same economies of scale that have made microprocessors so inexpensive. What started as a means to compress glue logic on motherboards was quickly taken up by a community of researchers looking for custom solutions to (usually) highly parallel problems. Today there is much research being undertaken in this area and it is also represented in this workshop.

The Workshop Papers

Let us now look at the papers in this workshop. They can be divided into three groups. In the first we have a collection of papers concerning conventional parallel computer architecture. These include instruction set architectures for vector computing, cache-coherence protocols, multi-threading and graphics engines:

- Vector ISA Extension for Sparse Matrix-Vector Multiplication - Stamatis Vassiliadis, Sorin Cotofana and Pyrrhos Stathis
- Implementing snoop coherence protocol for future SMP architectures - Wisam Hlayhel, Jacques Collet and Laurent Fesquet
- An Adaptive Limited Pointers Directory Scheme for Cache Coherence of Scalable Multiprocessors - Cheol Ho Park, Jong Hyuk Choi, Kyu Ho Park and Daeyeon Park
- Two Schemes to Improve the Performance of a Sort-last 3D Parallel Rendering Machine with Texture Caches - Alexis Vartanian, Jean-Luc Bechennec and Nathalie Drach-Temam

- A Study of a Simultaneous Multithreaded Processor Implementation - Dominik Madon, Eduardo Sanchez and Stefan Monnier

In the second group we have a number of papers looking at the general area of custom computing:

- The MorphoSys Parallel Reconfigurable System - Guangming Lu and Hartej Singh, Eliseu M. C. Filho and Nader Bagherzadeh
- Design and Analysis of Fixed-size Systolic Arrays for Modular Multiplication - Hyun-Sung Kim, Sung-Woo Lee, Jung-Joon Kim, Tae-Geun Kim and Kee-Young Yoo
- ManArray Processor Interconnection Network: An Introduction - G. G. Pechanek and S. Vassiliadis and N. P. Pitsianis

And finally we have two theoretical papers:

- A graph-oriented task manager for small multiprocessor systems - Xavier Verians, Jean-Didier Legat, Jean-Jacques s, Quisquater and Benoit Macq s
- The Algebraic Path Problem Revisited - Sanjay Rajopadhye, Claude Tadonki and Tanguy Risset

I commend these papers to you and welcome you to this workshop on Parallel Computer Architecture.