Information Society Technologies

SEVENTH FRAMEWORK PROGRAMME

HiPEAC info 16

COMPILATION ARCHITECTURE

appears quarterly | October 2008

Network of Excellence on High Performance and Embedded Architecture and Compilation

Welcome to THALES – HiPEAC Fall Computing Systems Week

www.HiPEAC.net

HiPEAC 2009 Conference:
Paphos, Cyprus, January 25-28

Message from the HiPEAC coordinator

Dear friends,

This summer we could all enjoy the Olympic Games in Beijing. Unmistakably, there is a parallel between the world of sports and networks of excellence. The athletes meet each other regularly at specialized events, they are competitors for the same scarce resources, but at the same time many of them are friends. They belong to the very best in their field, they try to promote their field, and they are constantly looking for new talent.

Koen De Bosschere

Then there is the Olympic motto: citius, altius, fortius – faster, higher, stronger. Don't we have similar goals – to be ever faster, more power efficient, and more reliable? The upcoming FP7 call is exactly about this. We hope that our community will submit many excellent research proposals in this call.

This summer, 200 of us enjoyed the yearly ACACES summer school in L'Aquila. From the evaluation forms, we learnt that the participants were very pleased with the program. We are currently working on the 2009 event, which will be announced in the next HiPEACinfo.

There are three more important upcoming events for our community: there is the launch of the next FP7 call on November 25, there is the fall HiPEAC Computing Systems Week in the same week at Thales Paris, and there is the HiPEAC Conference in January 2009, in Paphos, Cyprus. All events are ideal occasions for finding the perfect project partners for the upcoming FP7 call on computing systems.

The end of this year also means the end of the HiPEAC1 project. In retrospect, we believe that HiPEAC did much groundbreaking work for our community in Europe. We started with a small network and we worked very hard to create a strong HiPEAC brand. Once we got all the processes in place, we were then able to open the network up to a wider community. Nobody could ever dream then, when we started, that we would have what we have today. But we do still realize today that this is only the first step.

We believe that the next step will be about consolidation - how to attract the necessary resources to undertake cutting-edge research, how to create new innovative products, how to convince policy makers that computing systems research is at least as important for the future of mankind as research in particle physics.

As a coordinator of HiPEAC2, I am very committed to making this happen. However, as we all know, in our community the credit should go to the inventor and not to the implementer! Therefore I would like to sincerely thank Mateo Valero who initially had a dream about a vibrant HiPEAC community Europe, and steered HiPEAC through its early years. The current generation can only be thankful for such an inheritance which we will protect and further develop in the coming years. Unfortunately, we will never be able to make more than incremental contributions to such a great concept.

Take care,

Koen De Bosschere ■

HPCA Group: NVIDIA Professor Partner Award

NVIDIA, the world's largest manufacturer of graphics processors, has distinguished the HPCA (High Performance Computing and Architectures) group (http://www.hpca.uji.es) from the Universidad Jaime I de Castellon with one of the 2008 NVIDIA Professor Partner Awards.

The research conducted by the members of the HPCA group under the leadership of Profs. Enrique S. Quintana-Orti, Rafael Mayo and Gregorio Quintana-Orti, pursues the development of faster and more reliable solvers for dense linear systems of equations using NVIDIA hardware. A second major goal of this research is to improve the programmability of parallel solutions to dense linear algebraic problems by approaching the design of the libraries with a high level of abstraction. Important applications that lead to large-scale problems of this class arise, e.g., in boundary elements methods for Aeronautics, Statistics,

Panos Tsarchopoulos
Panagiotis.Tsarchopoulos@ec.europa.eu

Message from the project officer

ICT 2008 - "I"s to the future: Invention, Innovation, Impact

With more than 4000 visitors expected, the biennial ICT Event is the most important forum for discussing research and public policy in Information and Communication Technologies at European level. The ICT Event brings together researchers and innovators, policy and business decision makers working in the field of digital technologies.

The ICT 2008 event takes place from 25th to 27th November in Lyon, France. It is organised by the European Commission's Directorate General for the Information Society and Media and is hosted by the French Presidency of the European Union.

Over €2 Billion for ICT research, 2009–2010. This year's ICT event - the largest research event in Europe in 2008 – will examine:

- European Union priorities in ICT research for over €2 billion of funding available in 2009–2010
- The major current technological trends which impact upon strategic research planning
- Public research policies to stimulate research and innovation

Setting the ICT research agenda for the next decade. The next ten years will see major transformations in the technological, industrial and business landscapes surrounding information and communication technologies. ICT 2008 will set the agenda for ICT research and innovation in Europe during this crucial decade. The event hosts leading visionaries from academia and industry and addresses topics as diverse as Europe's role in shaping the future internet, ICT's contribution to advancing the sustainability agenda and alternative research paths for future ICT components and systems. These and many other cutting-edge themes will be explored in depth at ICT 2008.

Inventing the Future: ICT technologies for the future. The event's "Inventing the Future" theme covers major research trends in information and communication technologies (ICT) such as new computing paradigms, ICT-bio and nano, photonics, cognition robotics and the use of ICT for science. The 2009–2010 Work Programme and Calls for Proposals for ICT research in the EU's Seventh Framework Programme (FP7) will be presented in detail - one dedicated session will present the next Call in Computing Systems. Other sources of EU research funding for ICT will also be examined, including the new Joint Technology Initiatives (JTIs) and the Competitiveness and Innovation Programme (CIP). These initiatives together represent over €2 billion in EU support for ICT research over the next two years.

Innovative Europe: new markets, new sectors, new players. This year's ICT event will seek to involve new players in the development of next generation ICT and to ensure the expansion of existing businesses. In the spotlight: SMEs, the future internet, new business sectors and creative industries. The relationship between research and successful innovation will be critically examined. Mobilising and inspiring today's young people as potential researchers and innovators of the future will also receive special attention at ICT 2008.

Impact through Policy. Focus on effective public policies to stimulate ICT research and innovation for growth and sustainable development. This will include issues such as community and public research spending, the creation of conditions favourable to innovation and better coordination of the European research effort in ICT.

Exhibition and Networking. The event is accompanied by an exhibition with more than 200 stands and numerous networking sessions designed to facilitate contacts between researchers, innovators and engineers from all ICT fields.
More information at:
http://ec.europa.eu/information_society/events/ict/2008/index_en.htm

Panos Tsarchopoulos
Project Officer

■

Computational Chemistry, and Earth Sciences. While the solution to these problems would require months using thousands of more traditional (general-purpose) multi-core processors, this research may yield a reduction in the number of resources or the solution time by one order of magnitude, and thus it is expected to have a significant impact in these application domains. ■

HPCA Group, Universidad Jaime I de Castellon

HiPEAC 2009 Conference

The three previous editions of the HiPEAC conference in Barcelona, Ghent, and Goteborg have continued to contribute greatly towards establishing the HiPEAC conference as a premier forum for dissemination and networking in the area of high performance and embedded architecture and compilers. A clear illustration of this is HiPEAC 2008's record participation of more than 250 attendees. It is a great honour and privilege to be the general co-chairs of the fourth HiPEAC conference that takes place in Paphos, Cyprus on January 25-28, 2009. The local organization has been arranged by Yanos Sazeides, University of Cyprus.

HiPEAC 2009 is shaping up to be an exciting event. Stefanos Kaxiras, the workshop and tutorial chair, is organising a set of seven workshops on topics central to HiPEAC: multi-cores, compiler optimizations, compilation framework, reconfigurable computing, interconnection networks, reliability and simulation/evaluation. Also, a tutorial on reliability will be offered by experts from Intel.

The conference technical program is also rich and diverse. The program chairs, Mike O'Boyle and Margaret Martonosi, together with their program committee, have selected 27 out of 97 submitted papers for presentation at the conference. The conference will also feature two keynotes from Tilak Agerwala (IBM Research VP, IBM Systems) and from Francois Bodin (CTO, CAPS Entreprise). There will be plenty of opportunities for interaction during the conference, which takes place at Amathus, a 5-star seaside hotel. The conference will include an excursion to various archeological sites. After the main conference, HiPEAC research cluster meetings will take place on 29th and 30th January.

In conclusion, the last week of January 2009 will be full of various events that are central to HiPEAC activities, and therefore provide many reasons for you to attend. Please check the conference web page (http://www.hipeac.net/conference) for conference information and news, and to register for these great events in time.

Looking forward to seeing you in Paphos in January next year!

André Seznec
IRISA/INRIA

Joel Emer
INTEL/MIT ■

HiPEAC Activity

HiPEAC Fall Computing Systems Week

THALES

November 24-28 2008, Thales Research and Technology, France

We are glad to announce that the "HiPEAC Computing Systems Week" will be hosted by Thales Research and Technology, Palaiseau, France from November 24th to November 28th, 2008. The facility is located in the suburbs of Paris and close to INRIA-Saclay and Paris XI University.
The week will include the following co-located events:
The 6th HiPEAC Industrial Workshop on November 26th. The goal of this workshop is to bring together researchers from academia and industry to investigate the challenges, techniques, tools and requirements of multi-core architectures in embedded systems, especially in the automotive, aerospace, and security industry.
The HiPEAC Cluster Meetings will take place on November 27th and 28th.
Extra meeting rooms will also be made available on Monday 23rd and Tuesday 24th (before the industrial workshop) for extra HiPEAC-related meetings, such as research-project meetings.
Besides these events, we also plan an exciting social event at one of the ancient 17th century typical domains to allow further opportunities for discussions and networking.
We are looking forward to seeing you and your collaborators during the HiPEAC Computing Systems Week!
Please take note of the following schedule; all but the social event are taking place at Thales Facilities:

Date/Time	Event
Monday 24th November, all day	Extra meetings
Tuesday 25th November, all day	Extra meetings
Wednesday 26th November, all day	6th HiPEAC Industrial Workshop
Thursday 27th November, all day	HiPEAC Cluster Meetings
Thursday 27th November, PM	Social Event
Friday 28th November, all day	HiPEAC Cluster Meetings

6th HiPEAC Industrial Workshop –Domain specific, real time and safety critical embedded systems in the multi-core era

November 26th, 2008

Organized by Embedded System Lab, Thales Research & Technology

http://www.hipeac.net/industry_workshop6

The ubiquity of embedded electronics in all aspects of human life as well as the emerging applications in automotive, aerospace and security industry suggest multi-core architectures as a natural path to scalable performance. Still, such applications have challenging requirements beyond those existing in high performance computing and high volume embedded systems.
As well as the increasing need for performance, low power consumption, and the challenges the industry is facing to effectively program parallel architectures, such specific systems are usually subject to conflicting requirements such as flexibility (or "genericity") to compensate for low volume production needs, and customizability to address each application domain needs.

Furthermore, a lot of these systems, especially in the automotive and the aerospace domains, are subject to hard real time and safety critical requirements that reduce the efficiency and potential of current multi-cores and high performance architectures.

The goal of this workshop is to bring together researchers from academia and industry to investigate the challenges, techniques, tools and requirements of multi-core architectures in embedded systems, especially in the automotive, aerospace, and security industry.

The topics of interest include, but are not limited to:

- Processor customization and application-specific architectures
- Parallel and multi-core systems for real time and safety critical systems (including performance predictability, reliability constraints, …).
- Low power techniques for multi-cores
- Reconfigurable and adaptive architectures
- Heterogeneous multi-core architectures and virtualization techniques
- Compilers, tools and methods for embedded, high performance and heterogeneous parallel architectures
- Mapping parallel applications to heterogeneous multi-cores
- Simulation tools, design space exploration and performance modeling of multi-core and/or application specific architectures.
- Iterative compilation techniques for embedded and high performance systems
- Domain/application specific accelerators ■

Good bye and Thanks Sylvie! Hello and Welcome Klaas!

As of September 1st, Sylvie Detournay, who was the HiPEAC technical staff member, left the network. Over the last year, Sylvie enhanced the HiPEAC website with many new services, and she helped with the logistics for several HiPEAC tasks, such as the cluster meetings, the conference, and the summer school. She will start a PhD in Mathematics at INRIA (the network works in many ways). We wish her a very successful career.

Sylvie is replaced by Klaas Millet. As well as possessing a Master in History, Klaas also has a Master in Applied Informatics. He has built several websites using the Drupal-system (e.g. www.halestra.be). We expect the transition to go smoothly. You may contact Klaas at webmaster@hipeac.net. ■

EC FP7 STREP Apple-Core Project: Architecture Paradigms and Programming Languages for Efficient Programming of Multiple CORES

Apple-CORE will develop compilers, operating systems and execution platforms to support and evaluate a novel architecture paradigm that can exploit many-core computer systems to the end of silicon.

Coordinator & Technical Manager
Name: Chris Jesshope
Institution: University of Amsterdam
Email: C.R.Jesshope@uva.nl
Project website:
http://www.apple-core.info
Partners: University of Amsterdam (Netherlands), University of Hertfordshire (UK), Institute of Information Theory and Automation (Czech Republic), University of Ionnina (Greece), Associated Compiler Experts (Netherlands), Gaisler Research (Sweden)

Objectives

The Apple-CORE project will evaluate a programming model developed by the University of Amsterdam in the FP6 AETHER project. This model, SVP (the Self-adaptive Virtual Processor) provides for concurrency creation, control and reflection. Its actions can be considered as an operating system kernel. The SVP model provides the following benefits:

- Safe composition of concurrent programs without inducing deadlock;
- It captures program concurrency abstractly without any mapping or scheduling - even at the ISA level;
- It captures resources as first-class objects and supports the dynamic allocation of resources.

In Apple-CORE this model is implemented in the ISA of a conventional RISC processor, which allows many-core chips to be designed with the following advantages:

- Provides a high degree of latency tolerance to asynchronous operations, such as memory references;
- Given the abstract nature of the binary code, it has position independence and is independent of the number of cores used to execute it;
- SVP cores can execute legacy binary code unchanged.

The project has produced a full software emulation of a cluster of SVP cores. This emulation is being used to evaluate memory systems, to evaluate code-generation schemes for high-level, parallelising compilers and to investigate the model's generality. Within the project we will also be implementing a hardware prototype of an SVP core based on the LEON III and we will port a Linux operating system on top of the SVP kernel.

The core compiler for this project is based on GCC. We have defined a language, µTC, that adds the SVP kernel actions to the C language and we have modified GCC to compile this language to SVP extended ISA. Currently the compiler produces Alphas code but will also be modified to support the SVP version of Leon III.

The project will promote the µTC language as a standard front-end to the GCC compiler and will use it as a target for all user-level compiler development. These include:

- a parallelising C compiler based on the ACE CoSy framework
- a data-parallel functional language – Single-assignment C (SaC)

Some Results

Some results from the Alpha multi-core emulator using different memory systems are given below and show the expected performance of a single cluster of SVP cores executing a 256-point FFT. This code generates 8 stages of 128 threads and is a relatively small problem. The performance assumes that the FFT was delegated to a set of bare cores using an SVP create instruction and also assumes all local caches were cold. A number of on-chip memories are evaluated. The dotted line shows the maximum theoretical performance. All results show very good scaling. Note that at 32 cores there are only 4 threads per core and very little opportunity to tolerate memory latency. It is not surprising therefore that performance begins to saturate at this point. As expected, the ideal memory saturates latest and the COMA earliest, as there is significant latency induced by the coherence protocol.

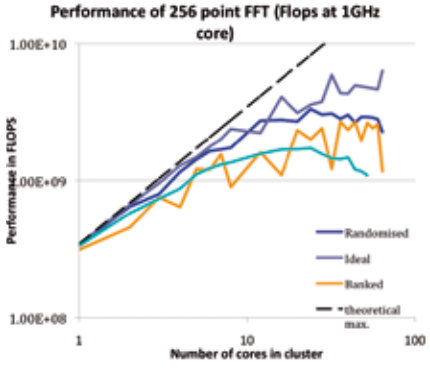

Expected Impact

With the infrastructure developed in this project, it will be possible to evaluate silicon products with many cores for a variety of applications, including embedded, commodity and high-performance systems. The project will also impact programmability aspects of this migration and support the software engineering of highly concurrent systems. This results from the natural separation of concerns between application and concurrency issues in the SVP model. The languages supported are: µTC for low-level system code, C for legacy applications and SaC for the rapid development of efficient code for high performance applications. Overall, this will allow the concept of massively threaded architectures to be explored by the European computer industry and the outputs of the project could have a large impact on all the above markets.

Multicore days in Stockholm, Sweden

All speakers at Multicore days, 2008 in Stockholm

For two days (11th and 12th September, 2008) multi-core technology was the focus at the Multi-core days held in Stockholm, Sweden. The event followed up the success from last year with more than 300 registered visitors who listened to world renowned experts in the area, such as Anant Agarwal of MIT and Tilera corporation, Kunle Olukotun of Stanford Pervasive Parallelism laboratory, David Padua of University of Illinois and HiPEAC members Per Stenström of Chalmers and Erik Hagersten of Uppsala University, among others.

One of the more exciting sessions of the event was the panel discussion where the future direction of multi-core architectures and programming models were debated. It was clear that the panel members had strongly differing views in various areas, e.g. homogeneity vs heterogeneity; many thin cores or fewer fat cores; domain specific languages vs general purpose languages supporting parallelism. There is still a lot of interesting work to be done.

Multi-core days 2008, co-organized with Ericsson SW Research, is one of the activities in the newly established Swedish Multi-core Initiative. This is a concerted effort to address the tactical and strategic issues related to multi-core processor technology for the software intensive systems industry in Sweden. The Initiative brings together all parties interested in taking this technology forward. The founding members are from six academic institutions and three enterprises under the umbrella of Swedsoft (www.swedsoft.se).

The vision of the Initiative is to increase competitiveness of the Swedish software intensive industry, leading to an ability to leverage multi/manycore technology. Objectives of the initiative include:

- To ensure that graduates from Swedish universities have the highest competence in multi-core-related technology
- To help Swedish multi-core research remains competitive worldwide
- To help leading researchers seek collaborations with members of the network

The Initiative will help foster ties between Swedish and international industrial/academic partners, such as the ones at UC Berkeley/University of Illinois UC (Microsoft, Intel US $20m) and at Stanford University (AMD, HP, Intel, NVidia and Sun US $6m).

Activities organized by the Swedish Multi-core Initiative include:

- Multi-core day: A state-of-the-art annual seminar for industry and academia highlighting technology advances, research results and solutions to current problems.
- Swedish Multi-core Workshop: An annual workshop for academia and industry to present and discuss recent research results. The first workshop will be organized by Blekinge tekniska högskola in November 2008.
- Best practices workshop: An annual workshop for industry and academia to present and discuss best practices in multi-core software development. First BP workshop will take place in February 2009.
- Networking activities between research groups
- Working groups will be formed around a number of important topics:
  - A technology roadmap from a Swedish industrial perspective
  - Curriculum development
  - Coordination and marketing of Swedish multi-core competence

The Swedish multi-core initiative is coordinated by a steering committee consisting of representatives of the founding members (Blekinge Institute of Technology, Chalmers, Mälardalen University, Swedish Institute of Computer Science, Royal Institute of Technology and Uppsala University).

For more information about the Swedish Multi-core Initiative and its activities see www.sics.se/multicore. ■

Guest Column

Professor Norbert Wehn, University of Kaiserslautern

As a relatively new member to the HiPEAC network, I am very proud to be asked to write a guest column for the HiPEAC newsletter, and thus be given the opportunity to present some of the research conducted at my research group. My name is Prof. Norbert Wehn, and I currently head the Microelectronic Systems Design Research Group at the Department of Electrical Engineering and Information Technology at the University of Kaiserslautern, Germany. Today, I want to introduce to you our XScale Energy EMUlator XEEMU, a simulation tool which we developed in cooperation with the research group from Prof. Tibor Gyimóthy at the University of Szeged, Hungary.

Energy efficiency is key in embedded-system design. Understanding the complex issue of software power consumption in early design phases is of extreme importance for making the right design decisions. Power management schemes for CPUs have been thoroughly researched, dynamic voltage and frequency scaling (DVFS) perhaps being the most important one. A large amount of work is based on models, which assume a quadratic dependency of power consumption on supply voltage and a linear dependency on operational frequency. In reality, however, this assumption does not hold for today's complex CPUs and it also completely neglects the energy consumed by other system components, such as the memory subsystem, for example, which is known to contribute significantly to system energy consumption.

Thus, a meaningful evaluation of power management schemes and other energy related design decisions can only be done either by power measurements or by simulation. While power measurements often are complex, error prone, or simply infeasible in early design stages, simulators often lack capabilities of memory subsystem simulation, and do not model the time- and energy overheads introduced by voltage and frequency scaling.

In order to overcome the above mentioned problems we have developed XEEMU, a cycle accurate simulator for the wide spread XScale architecture. XEEMU was first presented to the research community at PATMOS 2007. In contrast to other well known simulators, such as the SimpleScalar tool set, it aims at the most accurate simulation possible for an existing platform. It combines cycle accurate behavioural simulation models for the CPU, the system busses, and a configurable SDRAM subsystem. All its power models are based on measurements on an evaluation board and include accurate simulation for the timing and energy overheads introduced by dynamic power management. XEEMU is still under active development and has recently been extended with accurate models for the simulation of SDRAM power management schemes and support for OS simulation. We feel that XEEMU is a very interesting tool for researchers in the field of power-efficient software design and power management.

XEEMU is freely available in source and binary form for Linux operating systems under GPL license from our webpage at www.inf.u-szeged.hu/xeemu. For further information please also visit www.eit.uni-kl.de/wehn or feel free to write an e-mail to wehn@eit.uni-kl.de. ■

Prof. Dr.-Ing. Norbert Wehn.
Microelectronic Systems Design Research Group. University of Kaiserslautern

Community News

The University of Belgrade Honorary Doctorate for 2008 awards to Prof. Mateo Valero

On September 12th, the University of Belgrade awarded Prof. Mateo Valero, Director of the National Supercomputing Center in Spain, an Honorary Doctorate of the University of Belgrade. The reasons for this important recognition are as follows:

- His research results have had an extremely strong impact on the development of world science in the field of computer architecture. Furthermore, Prof. Mateo Valero has once again placed the whole of Europe on the map for research in this field (the United States and Japan dominated this scientific area up to the mid 1950s).
- His contribution to the creation of the computer architecture department at the Technical University of Catalonia (UPC)
- His contributions since 1984 to the creation of several parallel computing

Professor Gerd Ascheid, RWTH Aachen University

As coordinator of the UMIC research cluster, I feel honoured to contribute a guest column for the HiPEAC Newsletter. About two years ago, Germany started an initiative that involved the funding of "clusters of excellence" at German universities, granted in a country-wide competition. The only research cluster in the mobile communication domain was awarded to RWTH Aachen University; its focus was on Ultra high-speed Mobile Information and Communication - abbreviated as "UMIC".

While the dominating application in mobile networks in the past and still today is voice communication, traffic figures show strong growths in data communication. It is expected that future mobile network traffic will be dominated by data services. A key driver for this traffic is the mobile internet. In addition, many new "mobile applications" are under development or envisaged. 3G technology and its evolution (LTE) - which will be deployed in the coming years - support increasing data rates. However, significant progress is still required to meet future needs. Key issues are not just higher data rates, but high data rates for many users in dense traffic areas (densely populated areas, business areas, etc.) and an improved country-wide coverage of high data rate services. Even more important is an improvement of the quality, as perceived by the user, which for mobile access should be as close as possible to wired access.

An example of a key research topics at UMIC is the interfaces between applications and the mobile network, enabling a more efficient use of the mobile network and providing a significantly improved quality of service. Moreover, suitable network architectures to improve ultra high-speed data rate coverage and an intelligent use of mobile network resources. Mobile devices must have much higher compute performance in the future to support such a strongly improved quality, much higher data rates and demanding applications, e.g. in the multi-media realm. Unfortunately, future silicon technologies will not allow much higher clock speeds. Therefore, massively parallel processing will be required. The processing for wireless communication includes hard real time constraints. Because mobile devices are battery operated, low power is another key requirement. As a result, wireless devices will be based on heterogeneous rather than homogeneous multi-processor systems on a single chip (MPSoC). Key UMIC research topics in this field, which has close links to the HiPEAC network of excellence, are power efficient processor architectures, ultra fast simulation of MPSoC to enable hardware/software co-design and programming of embedded MPSoC.

Prof. Dr.-Ing Gerd Ascheid
Institute for Integrated Signal Processing Systems (ISS) RWTH Aachen University

The UMIC research cluster involves more than 20 chairs of the Electrical Engineering and Information Technology Faculty and the Computer Science Department of RWTH Aachen University. UMIC operates a research centre, where the interdisciplinary research teams, four new UMIC professorships and a joint research lab are located. In December the centre will move from its temporary site to the UMIC research centre building which is under construction.

For further information please also visit www.umic.rwth-aachen.de or contact us via email: public.relations@umic.rwth-aachen.de ■

centers made the construction of the National Supercomputing Center in 2004 possible, where he is the director. This public consortium and the computer architecture department at UPC are strongly linked to the computer architecture departments of the most prestigious international universities, such as the University of Belgrade.

- Prof. Mateo Valero has built strong bridges with the IT industry, especially Microsoft, Intel and IBM.
- His activities have had a strong impact in Serbia.

The Rector of the University of Belgrade, Branko Kovačević presented the award to Prof. Mateo Valero and thanked him for his contributions and continuous dedication to science and especially to the relevant progress made in the field of computer architecture. ■

Promotion of HiPEAC Start-ups

Marco Cornero

In this issue of the start-up promotion section we focus on Nanochronous Logic and Nema Labs. Although very different in nature, they still share the same ambition to address some of the fundamental problems of the electronics industry today: at one extreme Nanochronous Logic offers innovative EDA tools for the design and fabrication challenges of deep submicron circuits, while at the other end NEMA Labs deals with the issue of effectively programming multi-cores, offering interactive automatic software threading tools. We are pleased to see that HiPEAC offers visibility and networking opportunities in such a wide range of fundamental topics. ■

Nanochronous Logic

Nanochronous Logic is a dynamic start-up company providing breakthrough EDA technology for ASIC and SoC developers to address design challenges at 65nm, 45nm and below. The company is a co-located entity with corporate activities in San Jose, CA and R&D facilities located in Heraklion, Crete, Greece. Nanochronous Logic's "De-Synchronization" technology enables for the first time in EDA the automatic transformation of any conventional ASIC or SoC circuit to a functionally equivalent circuit with embedded variation-sensitive and timing-adaptive structures. A de-synchronized device improves IC implementation in many ways:

- Reduces power by enabling effective, continuous DFVS (Dynamic Frequency Voltage Scaling), effective MSMV (Multi-supply Multi-voltage), and PSO (Power Shut-Off) Operation
- Enables effective Yield Correction for Power or Performance by allowing fabricated die to individually self-diagnose, report and even self-tune their speed based on operating conditions and applied voltage
- Increases achievable performance by (i) adapting a circuit's timing to the actual, not the worst-case predicted conditions, and (ii) exploiting a time-borrowing mechanism which can compensate for inter/intra-die variations
- Enables EME (Electromagnetic Emissions) constrained design through clock-control mechanisms

The Nanochronous team is composed of Brian Hamel, Co-founder/CEO, Christos P. Sotiriou, Co-founder/CTO, and Spyros Lyberis, VP Engineering/lead software developer.

The Technology

Nanochronous Logic provides a novel timing methodology creating designs that are variation-aware; addressing the hurdles presented by processes at 65nm, 45nm and below. Utilizing the novel concept of creating timing partitions for a design, an IC's global timing can be optimally tuned at the presence of post-fabrication variations.

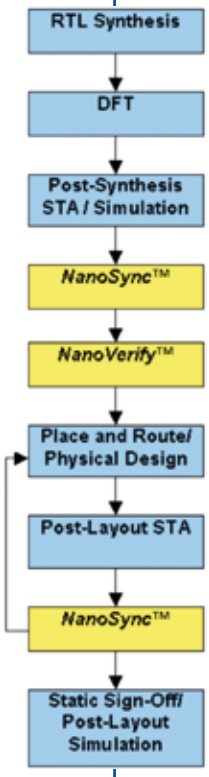

NanoSync™, Nanochronous' first product is an EDA tool for applying de-synchronization to any conventional ASIC or SoC design. It integrates smoothly into existing front-to-back-end EDA tool flows. Based on an analysis of the clocking, the datapath structure and the timing constraints of an IC, NanoSync is capable of dividing an IC into local regions of timing. Each region is driven by an adaptive, variation-sensitive clock and it can borrow time from neighboring timing regions. Therefore, global timing is not directly impacted from every local problem.
Power-saving techniques, such as MSMV and PSO, can be incorporated as timing region properties. DFVS is inherently supported by the methodology. The timing capabilities of the tool include rationally related clocks, asynchronous clocks and Globally Asynchronous Locally Synchronous (GALS) systems. Back-end integration is another key strength of the tool. Post P&R information, such as placement locations and back-annotated parasitics, can be fed back into NanoSync in order to fine tune and re-calibrate the desynchronized clocking network parameters prior to tapeout.

NanoVerify™ is a formal verification EDA tool that works in conjunction with industry standard formal verification engines. It is able to formally verify the correctness of the de-synchronization conversion process by comparing the synchronous and desynchronized netlists for equivalence. NanoVerify is a total solution for desynchronized designs. It is capable of verifying the logic, the timing partitions and the clocking structures of the desynchronized design with respect to the reference (original design) through its own equivalence and assurance checks.

Nanochronous Logic has filed a number of U.S.P.T.O. patent applications for its technology.

Company status

Nanochronous Logic has already released versions of its software tools, **NanoSync™** and **NanoVerify™** to tier one IDMs. The company is looking to expand its initial customer relationships throughout 2008 and plans to deliver additional tools to complement

NanoSync and **NanoVerify** throughout 2008 and beyond.
Nanochronous Logic has passed multiple designs through the NanoSync de-synchronization tool and has already demonstrated working silicon on an older, 250nm technology node. A significant number of industrial complexity designs have been de-synchronized in 90nm and 65nm processes, where the voltage scaling ability and the low electromagnetic emissions have been successfully demonstrated through simulation.
The company has closed on a seed funding round and is now targeting a Series A round in the following months. This will allow Nanochronous to ramp up business development and technical support resources, as well as support the continued development of additional products derived from the core technology. ■

Nema Labs – Fast and Reliable Threading on Multicores

Near the turn of the century, processor architects began to use the extra logic to place multiple cores per die because extra investment in single-core processors generated extra heat, not greater performance. All major silicon processor manufacturers were providing commodity processors with at least two cores per die, and system manufacturers rapidly brought to market new machines that were, in fact, not optimum for most software on the market and in development. This mismatch between the hardware capability of commodity systems and the software architecture of decades of products has been termed a 'software crisis,' an expression that truly understates the magnitude of the problem facing the trillion dollar information technology industry.

It will take many years of parallel programming education and innovation for software vendors to create products that will blaze when executed on multi-core systems. In the meantime, nearly 90% of today's C/C++ programmers can't program for multi-cores, and at least 70% of legacy C/C++ software won't run better on multi-cores. Multi-core systems, with the capability to execute multiple threads, are in fact mostly running single-threaded code.

Nema Labs' mission is to offer tools that allow software developers to concentrate on delivering functionality and to leave the task of reliably threading the code to leverage on the performance of multi-cores to intelligent tools. While tools and methods to parallelize C/C++ code are available, they either reduce software development productivity by offering too low a level of abstraction (e.g., by using OpenMP) or by only occasionally leveraging the performance potential of multi-cores (auto-parallelizers). By contrast, Nema Labs roadmap realizes tools that allow software developers to thread legacy C/C++ code fast and reliably by using a sequential abstraction with the help of automatic threading in an interactive fashion.

At Chalmers University, Dr. Per Stenstrom has contributed to architectural advances for parallel systems for well over two decades. With the emergence of multi-cores on the horizon, Dr. Stenstrom leveraged both his knowledge of parallel systems and his mission to educate new generations of computer scientists into a new company, Nema Labs, that is offering a roadmap of threading tools that innovates on top of automated C/C++ source code threading products to accelerate reliable threading of legacy code. The first product, called FASThread (Fast Automated Software Threader), will be released in early 2009 as a Linux plug-in for the Eclipse Integrated Development Environment (IDE), and will provide developers a familiar setting for threading their C programs with a few simple, automated steps. Shortly thereafter, plug-ins for other IDEs and operating systems will be released followed by a C++ version as well.

FASThread is straightforward to use and does not require any prior experience with parallelization. The user loads a project to be parallelized into Eclipse, sets the number of threads desired, builds, and runs the program. A push of a button then profiles the code and produces a call graph from which the user can see where the program is spending most of its time. The user can then zoom into the function of interest and right-click on the loop to be threaded, selecting it for threading analysis. A push of another button threads the loop and informs the user of actions to take to extract parallelism from what was otherwise intractable code. A press of the next button determines whether the threading process produced code that is free of deadlock and race conditions, the two most notorious problems faced by programmers attempting to parallelize code. Pushing the final button will report the effective performance speedup obtained by the threading process.

This seemingly simple process conceals a wealth of innovations underlying the entire method. With unique and powerful algorithms, Nema Labs is able to analyze and safely thread loops in seconds that might otherwise take many weeks or months by a programmer trained in parallel programming. And this is accomplished for both novice and experienced programmers, regardless of their parallel programming experience. The combination of these two features—threading difficult code, and doing it fast and automatically—makes the Nema Labs' product release of high interest to customers and users.

Nema Labs has released an alpha product for targeted work with customers in Sweden to help them discover how they might significantly improve their products with minimal new investment. The company has also launched a new US operation based in Silicon Valley, and has initiated projects with several equipment manufacturers and software development groups. The company plans to release its first product in the spring of 2009.

The innovations and discoveries in the Nema Labs' products are made possible by an exceptionally talented and focused group of software engineers and developers. Nema Labs is currently expanding rapidly and seeking highly creative and skilled people to join the team. ■

EC FP7 STREP eMuCo Project: Embedded Multi-Core Processing for Mobile Communication Systems

ICT-eMuCo project addresses the system platform of future mobile devices based on multi-core architecture

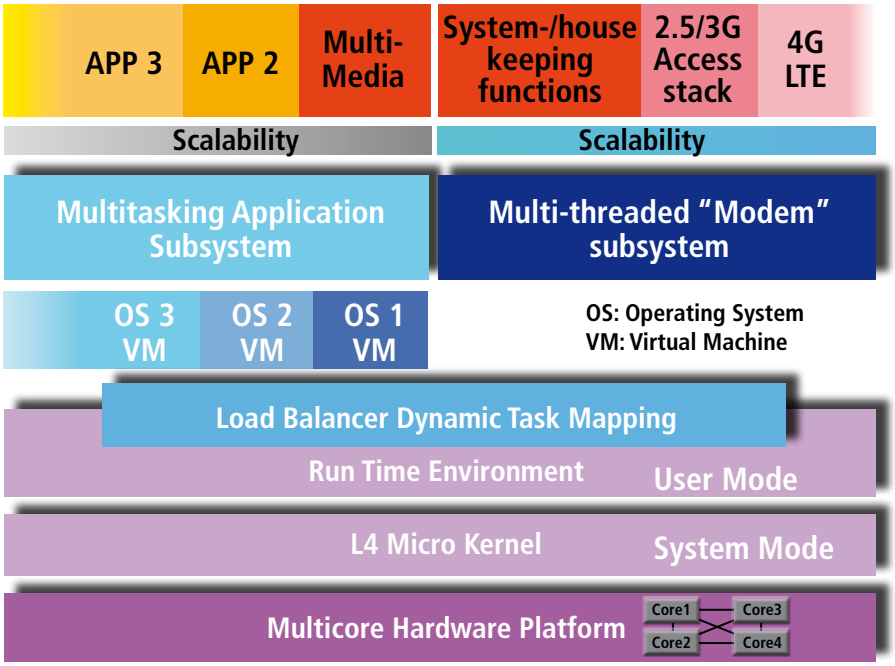

Flexible System Architecture for Future Mobile Devices

ICT-eMuCo (www.emuco.eu) is a European project with a total cost of 4.6m€ which is supported by the European Union under the Seventh Framework Programme (FP7) for research and technological development with 2.9m€. This project is coordinated by Ruhr-Universität Bochum, which is known as one of the biggest universities in Germany. There are also many strong academic, as well as industrial, partners participating, such as Technische Universität Dresden (Germany), University of York (United Kingdom), "Politechnica" University of Timisoara (Romania), Infineon (Germany), Telelogic (Sweden), ARM (United Kingdom) and GWT-TUD (Germany).

Mobile communication has become the dominant branch in the telecommunication's sector over the last decade and is still rapidly growing in the market. Mobile devices for standards like the Universal Mobile Telecommunication System (UMTS) and the future Long Term Evolution (LTE) incorporate multiple radio access technologies to enable the best quality of service in the current environment of the user. In addition, on one hand, it is expected that there will be exponential growth in the usage of multimedia applications, such as video streaming, video conferencing, complex graphics etc., thereby increasing the demand for computational power, which can not be pursued further by accelerating the processor clock. On the other hand, the coexistence of multiple software environments will be necessary for delivering all demanded services to the user.

Today's mobile communication systems usually contain multiple processing units, but individual units are typically dedicated to specific tasks rather than being general purpose, e.g. specifically for protocol stack handling. Other design approaches allocate special tasks to dedicated systems resulting in unnecessary hardware for situations other than high load.

The aim of eMuCo is to address the platform for future mobile convergent devices based on multi-core architecture offering:

- Flexible system platform by a modular architecture.
- Seamless system platform allowing the coexistence of difference radio access technologies and software environments on the same platform.
- Improvement of computational processing capacity of applications
- Reduction of power consumption
- Scalability
- Real time handling of the system control signals and real time applications.

It is expected that eMuCo will have an impact on the future of mobile devices due to the revolutionary approaches in the system architecture. Two major European companies are active in this market: ARM as the world's leading supplier of cores for mobile devices and Infineon as the world market leader for communication RF integrated circuits. Both will participate in this project and will ensure the technology transfer. Many other suppliers in Europe depend on this market either directly or as a supplier for the major market players. ■

Attila Bilgic
Institute for Integrated Systems
Ruhr Universität Bochum
D-44780 Bochum, Germany
Attila.Bilgic@is.rub.de

Trip Report: ACACES 2008

This summer, I was able to attend ACACES-2008, the fourth edition of the annual HiPEAC summer school held in L'Aquila, Italy. What follows is a brief trip report of a week full of learning new stuff from field experts, meeting interesting people and crappy wireless internet.

After getting up early in the morning on Sunday July 13th to catch my 10am flight to Rome Fiumicino airport and catching the bus in Tiburtina station near downtown Rome, we arrived at the campus in L'Aquila just in time for the small opening reception.
Although most students were a bit apprehensive at first, I soon noticed people having interesting discussions over a glass of bubbles. I was able to have a nice informal talk with Mary Lou Soffa, the lecturer of the compilers course I was going to take the rest of the week.
In the evening, Koen De Bosschere, the main organizer, held his opening speech, which was immediately followed by a keynote talk by Yale Patt, professor at UT Austin. He explained why we should not give up on uniprocessor performance in the multi-core era we are in.

The first classes got started early Monday morning. Each day, students took classes in four time slots, choosing one of three courses in each slot. Thus, 12 field experts gave lectures on a variety of interesting topics, including (but not limited to) compilers, networks on chip, transactional memory, etc.

It was clear the first day that the lecturers were not going to limit themselves to a mere introduction to the topic they were lecturing on. Most made it clear that they would also touch on hot ongoing research and invite the audience to actively participate in the classes. Several people grasped the opportunity given to them by answering questions, giving insightful comments and questioning the lecturers themselves.

Another invited talk was given on Tuesday evening by Peter S. Magnusson, founder of VirtuTech. He entertained the audience by talking about entrepreneurship, going over his top 10 reasons to do a start-up in IT (most of which were money-related). I was thoroughly amused by the people asking him questions like "Are you a millionaire?" and "Once you have all that money, what comes next?"

On Wednesday, the afternoon course slots were replaced by a student poster session. Each participating student was invited to present his/her poster on ongoing or future work. I actively participated and was glad to find multiple people interested in my work, including fellow PhD students, people from industry and lecturers.
After classes, the annual soccer tournament was held, traditionally against the campus personnel, but this year also against a team of young Italians. Although the ACACES crowd did the best they could to cheer for their team(s), we still lost the trophy to the young Italian team.

After two more days of increasingly interesting classes, coffee breaks during which people got to know each other better, dinner with people from countries all around the world and evenings filled with card games, beer and tons of fun, people got ready for the gala dinner and closing party on Friday night.
Everyone enjoyed the free ice cream, wine, beer and sangria, and the live band playing funky music; some brave souls even got up to dance (after a few drinks!) We also exchanged contact info to keep in touch.

All and all, the HiPEAC summer school is an experience every PhD student should try and enjoy at least once in his/her career.

More information, including an extensive set of pictures can be found at http://www.hipeac.net/summerschool/.
And don't forget: next year is an anniversary edition, i.e. the fifth annual HiPEAC Summer School. So, don't forget to apply in time, because only a limited number of attendees are allowed!

Kenneth Hoste
(kenneth.hoste@elis.ugent.be)
Ghent University
http://www.elis.ugent.be/~kehoste

Trip Report: MPSoC 2008 – an exciting week in St. Gerlach

MPSoC 2008 attendees and Speakers (slides available: mpsoc-forum.org)

Maybe I would never have come across Château St. Gerlach in my life if it were not for the venue of the MPSoC 2008 event. Fortunately enough, in a charming June week, I not only enjoyed the irresistible beauty of this 800-year-old luxury Château near Maastricht/Aachen, but also had great fun in learning about state-of-the-art MPSoC technologies from presentations and face-to-face meetings with many top experts around the globe. It has been an unbelievably rewarding experience for a PhD student like me.

MPSoC 2008 is the eighth event of the forum series, which focuses on the fundamental and strategic issues to master multi-processor SoC design. This year the event featured more than 50 world-class speakers from both industry and academia, giving an excellent coverage of the challenges and opportunities in the MPSoC. The 4-day, info-packed presentations delivered a reliable and clear message on where we are now in the MP-infant age and where we are heading to on the road to maturity. I particularly like the breadth of the topics – to give a few examples, the new technology fronts like 3D chips, efficient hardware implementations, e.g. in the streaming/SDR areas, the exploding MPSoC software programming challenges, and many proven good practices from the leading industrial companies, which are tremendously helpful and valuable for the university researchers. Being "confined" in the glamorous Château also allowed me to approach and talk to many senior executives and famous professors for in-depth details about their presentations and other interesting topics face-to-face, a privilege that you probably won't find in many other conferences.

Besides the dense technical activities, MPSoC 2008 also provided plenty social situations for the attendees. The coffee breaks in the sunny garden and the wonderful dinners in the Château were perfect for networking and bringing people together easily. A city tour of the town of Aachen was also arranged during the week – the nice European summer weather and a fun meet-up in the guest house of Aachen University simply made this great event even greater.

It has been a fruitful and memorable week for me in the marvelous St. Gerlach. I hope the MPSoC 2009 edition finds its next superior spot well. It'd be great if I can be one of the crowd again next year at the MPSoC, and hopefully meet you there.

Weihua Sheng
(weihua.sheng@iss.rwth-aachen.de)
ISS/SSS, RWTH Aachen University

HiPEAC Students

Polychronis Xekalakis, Institute for Computing Systems Architecture, University of Edinburgh

I started my research while still being an undergraduate at the University of Patras in Greece. There, under the supervision of Assistant Professor Stefanos Kaxiras, I was working on thermally-aware leakage control policies and way-selection techniques for cache memories. Finishing my undergraduate studies I wanted to do my PhD in a well established research group, while still living in a very nice city so I picked the University of Edinburgh!

My name is Polychronis Xekalakis and I am currently a fourth-year PhD candidate at the School of Informatics in the University of Edinburgh, under the supervision of Assistant Professor Marcelo Cintra. Being a member of the Compiler and Architecture Design (CArD) group, I have had the opportunity to work and develop my ideas in a really nice enviroment where people do a lot of interesting world-class research.

With the advent of multicore systems, one of the problems our group is trying to solve is how to utilize them as much as possible so as to be able to gain some performance benefits even from sequential applications. In an effort to relieve

New Book

João M. P. Cardoso and Pedro C. Diniz: Compilation Techniques for Reconfigurable Architectures, Springer, 2008, ISBN 13: 978-0-387-09670-4.

The extreme flexibility of reconfigurable architectures and their high-performance potential have made them a vehicle of choice in a wide range of computing domains, from rapid circuit prototyping to high-performance computing. The increasing availability of transistors on a die has allowed the emergence of powerful reconfigurable architectures with a large number of hardware resources, flexible interconnection topologies, and storage resources, making thus possible to accommodate complex computing structures.

To exploit the potential of these reconfigurable architectures with efficient reconfigurable computing implementations, programmers must currently navigate a maze of program transformations, mapping, and synthesis steps. The richness and sophistication of any of these steps make the mapping of computations to reconfigurable architectures an increasingly daunting process. It is thus widely believed that automatic compilation from high-level programming languages is the key to the success of reconfigurable computing.

This book describes a wide range of code transformations and mapping techniques for programs described in high-level programming languages, most notably imperative languages, to reconfigurable architectures. The book presents a brief but broad description of reconfigurable architectures and contemporary synthesis flows, followed by a detailed description of important code transformations and mapping techniques when targeting reconfigurable architecture. In this description, the authors distinguish between the applications of mapping techniques to fine- and coarse-grained reconfigurable architectures. The book also presents a comprehensive survey of significant research prototype compilation efforts and concludes with the authors' perspective on possible avenues of research for reconfigurable computing to truly become a mainstream computing paradigm.

This book is primarily intended for practitioners, researchers, and graduate students in the areas of study of hardware compilation and advanced computing architectures in the fields of Electrical and Computer Engineering and Computer Science. Since it focuses on the specific topic of compilation from high-level program descriptions to reconfigurable architectures, this book can easily support advanced compiler and computer architecture courses related to reconfigurable computing. Through this book, the authors hope to motivate further discussions on the challenging topic of compilation for reconfigurable architectures. It is also hoped that this book can be a reference for researchers, educators and students to learn more about the most prominent efforts on this subject in recent years.

Despite the tremendous progress in recent years, compilation techniques for reconfigurable computing platforms still offer many exciting research and development opportunities. The authors hope this book will motivate further research efforts in this domain and serve as a base for a deeper understanding of the overall compilation and synthesis problems, current solutions, and open issues. ■

the programmer from the parallelization burden, we are trying to design systems able perform automatic hardware parallelization, with some minimal compilation support. Under this effort, we are performing research in architecture and compilation using both conventional and machine learning based techniques.

My current work aims at designing micro-architectural support for efficient Thread Level Speculative (TLS) systems. Increasing the efficiency of these systems involves identifying which are the critical micro-architectural mechanisms for them, gaining a deep understanding of how the new execution paradigm affects and is affected by them and finally employ techniques that will minimize their power consumption while increasing their performance. One of the things I am currently looking at, is whether we can distinguish cases where parallelism does not exist and try to utilize the underlying threading mechanism to increase the ILP of the application by converting the TLS threads to helper threads.

During the course of these years, I have been lucky enough to travel a lot and meet and exchange ideas with many researchers across Europe. I was also able to work in a company in a Sweden for three months under a HiPEAC grant, which for me was an invaluable experience. Of course this was all due to the financial support and to the organization of events like the summer school by the HiPEAC community.

contact:
p.xekalakis@ed.ac.uk
http://homepages.inf.ed.ac.uk/s0565349/
■

New Book

Stefanos Kaxiras, Margaret Martonosi: Computer Architecture Techniques for Power-Efficiency, Morgan & Claypool Publishers, 2008, ISBN: 13 978 1 598 292084

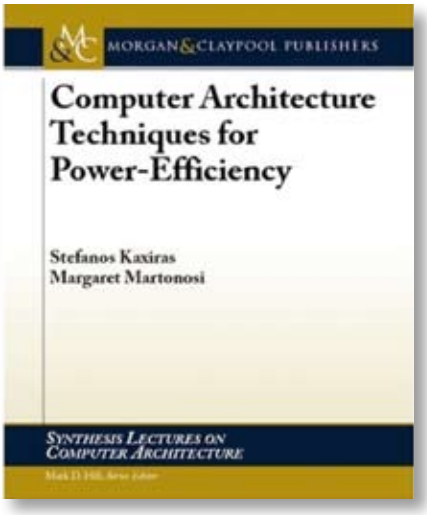

A new book in the *Synthesis Lectures on Computer Architecture Series* (Morgan & Claypool Publishers, Mark Hill Ed.) was presented this June at the ISCA 2008 (Beijing) conference. The book entitled **"Computer Architecture Techniques for Power Efficiency"** is co-authored by Stefanos Kaxiras (University of Patras) and Margaret Martonosi (Princeton University). Both are affiliated with the HiPEAC Network of Excellence: Kaxiras has been an active member in HiPEAC-1 and is currently the leader of the Task Force on Low Power in HiPEAC-2, while Martonosi, a HiPEAC-2 associate member, has been participating in many HiPEAC activities from across the Atlantic: collaboration in many HiPEAC funded European groups, ACACES 2006 teacher and HiPEAC2009 PC Chair among others.

The book is the first that deals with power efficiency (and low-power) techniques at the *architectural level.* While there are many great books at the circuit or VLSI level discussing power efficiency, this book intends to cover an emerging need: the synthesis of a significant body of recent architecture work on power efficiency. The abstract below, quoted from the book, gives a flavour of the book's style:

Abstract from the book (copyright Morgan & Claypool Publishers)
"In the last few years, power dissipation has become an important design constraint, on par with performance, in the design of new computer systems. Whereas in the past, the primary job of the computer architect was to translate improvements in operating frequency and transistor count into performance, now power efficiency must be taken into account at every step of the design process. While for some time, architects have been successful in delivering 40–50% annual improvement in processor performance, costs that were previously brushed aside eventually caught up. The most critical of these costs is the inexorable increase in power dissipation and power density in processors. Power dissipation issues have catalyzed new topic areas in computer architecture, resulting in a substantial body of work on more power-efficient architectures.
Power dissipation coupled with diminishing performance gains, was also the main cause for the switch from single-core to multi-core architectures and a slowdown in frequency increase. This book aims to document some of the most important architectural techniques that were invented, proposed, and applied to reduce both dynamic power and static power dissipation in processors and memory hierarchies. A significant number of techniques have been proposed for a wide range of situations and this book synthesizes those techniques by focusing on their common characteristics."

The book is divided into three parts: the first part (comprising chapters 1 and 2) presents introductory material, the second part (comprising chapters 3 and 4) deals with dynamic power consumption and the third part (chapter 5) deals with static (leakage) power. The book ends with a short summary chapter (chapter 6). Material in the first part includes introduction to modeling, measuring, and simulating power in processors and memory hierarchies. The second part (dynamic power) is structured after the terms *a, C, V*, and *f* in the dynamic power equation: $P = a\ C\ V^2\ f$. Chapter 3 discusses the techniques that affect *V* and *f*, namely dynamic frequency/voltage scaling from the system level to the flip-flop level. Chapter 4 presents techniques that affect a and *C* aiming to reduce (at run-time) the effective switched capacity. The third part discuses the increasingly important problem of leakage power. It presents techniques invented to reduce various forms of leakage power by manipulating (at run-time) physical values such as the supply voltage or the threshold voltage. The chapter concentrates on the architectural-level policies for applying such techniques.

In a fast-changing field the book aims to synthesize some of the most important recent work. The target readers of this book are engineers, researchers, computer architecture graduate students, or advanced undergraduates who are fairly fluent in computer architecture concepts, but who want to build their understanding of how power-aware design influences architectures without the need to delve deep into transistor or circuits details, beyond the basics of CMOS gate structures.

Trip Report: DAC 2008 – A sunny week in the city of Anaheim

Anaheim, the sunny south Californian city, is a famous place thanks to Disneyland. After years of dreaming, I found a way to go there for a good reason -- the 45th annual Design Automation Conference (DAC), which is something amazing that I'd really love to recommend to you.

As one of the premier EDA events, DAC consists of a conference, an exhibition, and several other activities, such as a number of collocated workshops and tutorials, among others. The technical program featured 138 novel papers, not only on electronic design and verification, but also covering architectures, tools, applications and a broad spectrum of topics. Apart from that, more than 200 exhibitors showed their state-of-the-art techniques and products. New techniques were unveiled, new products were showcased, and new ideas were discussed. The only problem for me was how to survive such a huge amount of information. If you are a technology nerd, you'll never want to miss the close contact with a lot of cool stuff, e.g., iPhone 3G or Android. There were several pavilions and panel sections to hear some crazy thoughts, and the discussions were really enjoyable.

DAC 46th will be held in San Francisco in July of next year. But before you plan the trip, allow me a minute to give you some extra recommendations. DAC is a perfect place for getting feedback on your work (both praise and critique) and, who knows, even good offers. However, since it's large, and in case you don't know where to go, maybe try two events which are extremely attractive to students: the University Booth and the PhD Student Forum. The first one, organized by ACM SIGDA and located in the expo, is exclusively reserved for students and is a perfect platform to show you work. You should apply in advance to get a place and travel funding. The second one, scheduled after the conference talks have finished, is the place where posters meet reviews, and food meets people. Anyway, it's a lovely place to go and I would really like to go back and meet some HiPEAC folks over there.

The author presenting his work at ACM SIGDA University Booth

Lei Gao (lei.gao@iss.rwth-aachen.de),
ISS/SSS, RWTH-Aachen University

Best Paper Award: International Conference on High Performance Computing and Simulation HPCS'08

Samir Ammenouche (PhD student), Sid Touati (Assistant Professor) and William Jalby (Professor) from the University of Versailles have been awarded the Best Paper Award for their contribution entitled "Practical Precise Evaluation of Cache Effects on Low Level Embedded VLIW Computing". The article was presented and published in the Proceedings of the International Conference of High Performance Computing and Simulation (HPCS'08, Nicosia, Cyprus).

Abstract:
The introduction of caches inside high performance processors provides technical ways to reduce the memory gap by tolerating long memory access delays. While such intermediate fast caches accelerate program execution in general, they have a negative impact on the predictability of program performances. This lack of performance stability is a non-desirable characteristic for embedded computing. We present the progress of our experimental study about the influence of cache effects on embedded VLIW processors (*ST2xx* processors). We are trying to understand qualitatively and quantitatively the interactions between cache effects (Data cache) and instruction level parallelism at different granularities: applications and functions (coarse grain), program regions (medium grain) and instructions (fine grain). Our aim is to come up with experimental arguments to help decide whether non-blocking caches would be a reasonable architectural design choice for embedded VLIW processors. By reasonable, we mean looking at opportunities at two levels: 1) program execution acceleration with tolerable performance predictability, and 2) active interactions with compiler optimization techniques. Our study is based on many months of full-time simulations on numerous workstations producing many terabytes of data to analyze.

Mechanisms for out of order memory accesses management

By Fernando Castro (fcastror@fis.ucm.es)
Prof. Daniel Chaver, Prof. Luis Pinuel and Prof. Manuel Prieto
Complutense of Madrid , Spain
June 2008

This work is included in the research area concerning new designs for microprocessors. Three new implementations for the hardware responsible for load and store instructions management, i.e., the load-store queue (LSQ) are made. As a result of the proposed designs, we obtain significant reductions in hardware complexity and energy consumption.

The first implementation, named split LQ, splits conventional and associative LQ into two new structures: an associative one and another banked and non-associative one. This way, a back-up mechanism is needed in order to detect potential memory ordering violations in the banked queue. To do this, we employ a simple and quick mechanism: a Bloom filter.

The two other designs (IMDC and DMDC) are included in the global definition of age-based mechanisms. Both remove the associative LQ and replace it with much simpler structures. IMDC (Issue-time Memory Dependence Checking) uses a hash table instead of LQ to record explicitly the age of load instructions. Therefore, stores only need to check this structure and compare their own age with the value recorded in the table to detect memory dependence violations.

The third design, DMDC (Delayed Memory Dependence Checking) replaces the LQ with an address-based table, which only writes a very small fraction of all stores. The scheme works in a decoupled way, so that in the first step an age-based filtering is performed. As a result of this process, store instructions potentially dependent with some in-flight load, are marked. Furthermore, using a straightforward age register, an instruction window containing potential offending loads is delimited. In the second phase (previous to commit time), simple address comparisons are established between both kind of instructions.

All proposed designs achieve the common goal of reducing LSQ hardware complexity, managing to very significant energy savings in this structure and the whole processor, with negligible performance losses.

Reconfigurable Optical Interconnection Networks for Shared-Memory Multiprocessor Architectures

By Wim Heirman (Wim.Heirman@elis.UGent.be)
Prof. Jan Van Campenhout and Prof. Dirk Stroobandt
Ghent University, Belgium
July 2008

Current multiprocessor communication technologies, using electrical signaling, are reaching the end of their capabilities. Moreover, the traffic on such communication networks is very irregular, making some parts of the network highly saturated while other parts are barely used. The precise patterns also depend on the application, and can even change during the run time of a single application. This makes it very difficult to build an efficient communication network.

This thesis explores the possibilities of optical, reconfigurable networks. Optical connections are one part of the solution to the communication problem, since they allow for much higher data rates. Moreover, optics allows the network to be reconfigured, during the course of a program, in a data-transparent way, to the current traffic patterns. This way, all network connections are utilized more efficiently, which can both increase the network's performance and reduce power usage.

Efficient Mechanisms to Provide Fault Tolerance in Interconnection Networks for PC Clusters

By Jose Miguel Montañana (jmontana@gap.upv.es)
Prof. Jose Flich and Prof. Antonio Robles
Technical University of Valencia
Valencia, Spain
July 2008

Fault tolerance in interconnection networks for clusters is becoming a key issue as far as the number of end nodes increases. However, most of the fault-tolerant routing strategies proposed for massively parallel computers cannot be applied to PC clusters. The thesis proposes two fault tolerant routing approaches to select alternative paths at the source end nodes when a fault appears. The first one provides in advance several disjoint paths between every source-destination pair of nodes, whereas the second one is able to dynamically provide new alternative paths by joining fragments of existing paths in a time-efficient manner. Also, a simple and fast method for dynamic network reconfiguration in the presence of faults is proposed. These mechanisms dynamically tolerate a reasonable number of faults and can be implemented on current commercial network technologies for PC clusters.

Avoiding Conversion and Rearrangement Overhead in SIMD Architectures

By: Asadollah Shahbahrami,
First promoter: Prof. Stamatis Vassiliadis*
Second promoter: Prof. Kees Goossens
Copromotor: Dr. Ben Juurlink
TU Delft, The Netherlands
September 2008

In this dissertation, a novel SIMD extension called Modified MMX (MMMX) for multimedia computing is presented. Specifically, the MMX architecture is enhanced with the extended subwords and the matrix register file techniques. The extended subwords technique uses SIMD registers that are wider than the packed format used to store the data. It uses 32 bits extra for each 64-bit register. The extended subwords technique avoids data type conversion overhead and increases parallelism in SIMD architectures. This is because promoting the subwords of the source SIMD registers to larger subwords before they can be processed and demoting the results again before they can be written back to memory incurs conversion overhead.

The matrix register file technique allows data that is stored consecutively in memory to be loaded into a column of the register file (where a column corresponds to the corresponding subwords of different registers). In other words, this technique provides both row-wise as well as column-wise access to the media register file. It is a useful approach for matrix operations that are common in multimedia processing. In addition, in this work, new and general SIMD instructions addressing the multimedia application domain are investigated. It does not consider an ISA that is application-specific. For example, special-purpose instructions are synthesized using a few general-purpose SIMD instructions. The performance of the MMMX architecture is compared to the performance of the MMX/SSE architecture for different multimedia applications and kernels using the sim-outorder simulator of the SimpleScalar toolset. Additionally, three issues related to the efficient implementation of the 2D Discrete Wavelet Transform (DWT) on general-purpose processors, in particular the Pentium 4, are discussed. These are 64K aliasing, cache conflict misses, and SIMD vectorization. 64K aliasing is a phenomenon that happens on the Pentium 4, which can degrade performance by an order of magnitude. It occurs if two or more data items whose addresses differ by a multiple of 64K need to be cached simultaneously. There are also many cache conflict misses in the implementation of vertical filtering of the DWT, if the filter length exceeds the number of cache ways. In this dissertation, techniques are proposed to avoid 64K aliasing and to mitigate cache conflict misses. Furthermore, the performance of the 2D DWT is improved by exploiting the data-level parallelism using the SIMD instructions supported by most general-purpose processors.

A Cache-based Hardware Accelerator for Memory Data Movements

By Filipa Duarte
Prof. Stephan Wong
TU Delft, The Netherlands
October 2008

This dissertation presents a hardware accelerator that is able to accelerate large (including non-parallel) memory data movements, in particular memory copies, performed traditionally by the processors. As today's processors are tied with or have integrated caches with varying sizes (from several kilobytes in hand-held devices to many megabytes in desktop devices or large servers), it is only logical to assume that data to-be-copied by a memory copy is already present within the cache. This is especially true when considering that such data often must be processed first. This means that the presence of the caches can be utilized to significantly reduce the latencies associated with memory copies, when a "smarter" way to perform the memory copy operation is used.

Therefore, the proposed accelerator for memory copies takes advantage of the presence of these caches and introduces a redirection mechanism that links the original data (in the cache) to the copied addresses (in a newly added indexing table). The proposed solutions avoid cache pollution and duplication of data, and efficiently schedule the access to the main memory, thus effectively reducing the latency associated with memory copies. Moreover, the proposed accelerator supports copies of cache line and word granularity, and can be connected to a direct-mapped or a set-associative cache, and can efficiently reduce the memory copy bottleneck in both single core or multi-core processors, executing a message passing communication model.

The proposed solutions have been implemented in a FPGA as a proof of concept and incorporated in a simulator running several benchmarks to determine the performance gains of the proposal. In particular, for the receiver side of the TCP/IP stack, the proposed solutions can reach speed-ups from 2.96 to 4.61 times and reduce the number of instructions executed by 26% to 44%.

Upcoming Events

International Symposium on System-on-Chip 2008. Call for Participation
Tampere, Finland, November 5-6, 2008. http://soc.cs.tut.fi/

41st Annual IEEE/ACM International Symposium on Microarchitecture (MICRO 2008)
Lake Como, Italy, November 8-12, 2008. http://www.microarch.org/micro41/

International Conference on Computer Aided Design (ICCAD 2008)
San Jose, USA, November 10-13, 2008

International Conference for High Performance Computing, Networking, Storage and Analysis (SC08)
Austin, Texas, November 15-21, 2008. http://sc08.supercomputing.org/

HiPEAC Fall Computing Systems Week
Paris, France, November 24-28, 2008, http://www.hipeac.net/computing_systems_week_paris

HiPEAC 6th Industrial Workshop
Paris, France, November 26, 2008

Information and Communication Technologies (ICT 2008)
Lyon, France, November 25-27, http://ec.europa.eu/information_society/events/ict/2008/index_en.htm

20th International Conference on Parallel and Distributed Computing and Systems (PDCS 2008)
Orlando, USA, November 25-27, 2008. http://www.iasted.org/conferences/home-631.html

8th USENIX Symposium on Operating Systems Design and Implementation (OSDI 08)
San Diego, USA, December 8-10, 2008. http://www.usenix.org/event/osdi08/

International Symposium on Parallel and Distributed Processing with Applications (ISPA-2008)
Sydney, Australia, December 10-12, 2008. http://www.cs.usyd.edu.au/~ispa2008/workshops.html

36th Symposium on Principles of Programming Languages (POPL 2009)
Savannah, USA, January 21-23, 2008. http://www.cs.ucsd.edu/popl/09/

International Conference on Parallel and Distributed Computing, Applications and Technologies (PDCAT 2008)
Dunedin, New Zealand, December 1-4, 2008, http://www.cs.otago.ac.nz/pdcat08/

HiPEAC 2009 Conference, Paphos, Cyprus, January 25-28, 2009
3rd Workshop on Statistical and Machine learning approaches applied to ARchitectures and compilation (SMART'09). Co-located with HiPEAC'09. Paphos, Cyprus, January 25,2009, http://www.hipeac.net/smart-workshop.html

Contributions
If you are a HiPEAC member and would like to contribute to future HiPEAC newsletters, please contact Rainer Leupers at **leupers@iss.rwth-aachen.de**

HiPEAC Info is a quarterly newsletter published by the HiPEAC Network of Excellence, funded by the 7th European Framework Programme (FP7) under contract no. IST-217068.
Website: http://www.HiPEAC.net
Subscriptions: http://www.HiPEAC.net/newsletter

HiPEAC² Cluster Meetings
27th of November, Paris