SEVENTH FRAMEWORK PROGRAMME

HiPEAC info 18

COMPILATION ARCHITECTURE

appears quarterly | April 2009

Network of Excellence on High Performance and Embedded Architecture and Compilation

HiPEAC Cluster Meetings and Industrial Workshop, Infineon Campeon

www.HiPEAC.net

ACACES 2009. July 12-18, 2009, L'Aquila, Italy

Intro

Message from the HiPEAC coordinator

Dear friends,

The last few months have been quite hectic for HiPEAC. After the HiPEAC conference, we started preparing the final review for HiPEAC1 and the first review of HiPEAC2; we worked hard on the HiPEAC roadmap, and many of us have been preparing one or more proposals for FP7 call 4.

Koen De Bosschere

HiPEAC1 started on September 1, 2004 and ended on December 31, 2008. The final review took place on March 20, 2009 in Barcelona and it was a good opportunity to look back on the whole project, and to see where we have been successful and where there is still room for improvement. The reviewers all agreed that HiPEAC1 has had a considerable impact on our community. We succeeded in creating a real computing systems community in Europe, and there is now much more collaboration between European research groups than before. There is also more interaction between universities and industry, and between computer architects and compiler builders. Furthermore, HiPEAC has also helped create a European computing systems identity, and we are moving in the direction of a common research agenda.

It is, however, clear that the job is not yet finished: we need to work harder to synchronize the research agendas of the 600+ researchers in the network, and we need to improve the relationship between HiPEAC and companies that innovate in the HiPEAC domain. We also need to increase the efforts to valorize the HiPEAC research results. These are therefore the main focal points of HiPEAC2.

I would hereby like to pay tribute to all the colleagues who helped to shape HiPEAC1. Above all, I would like to mention Mateo Valero - initiator, coordinator, negotiator, inspirator without whom there would never have been a HiPEAC1. I also want to mention Olivier Temam who worked day and night to get our community organized in the early days and Pilar Armas who helped us with all the administrative aspects of the project management. Finally, there was Stamatis Vassiliadis who left us too early and who would definitely have enjoyed the final review with all the good food and Catalan delicacies.

Koen De Bosschere ■

HiPEAC Activity

SDR Pioneers Visited SoC 2008

Photo: Pertti Kellomäki, TUT

General Chair Jari Nurmi introduces Joseph Mitola III before the keynote.

The special theme of the 2008 International Symposium on System-on-Chip organized in Tampere, Finland by TUT Department of Computer Systems was on Software-Defined and Cognitive Radio. The event attracted a number of high-calibre professionals in SDR. The tutorial day was instructed by Mark Cummings and Todor Cooklev. Cummings was the first chairman of the most important body in the field, SDR Forum, and he has patented the first SDR solution utilizing programmable logic. Cooklev, on the other hand, is leading a joint wireless technology center at Indiana University and Purdue University.

The symposium keynote speaker was the "godfather of software-defined radio" Joseph Mitola, III. He has had a long career in practical SDR projects in different military technology companies, and has recently started a new job as professor at Stevens Institute of Technology. Mitola also coined the term "Cognitive Radio". In his keynote, he approached the cognitive radio challenges from a cognitive linguistics perspective. In the "SoCnitive radio" the terminology and structures of linguistics and genetic programming can be applied to skills engineering for SoC implementations in radio systems.

Liesbet van der Perre from IMEC pre-

Panos Tsarchopoulos
Panagiotis.Tsarchopoulos@ec.europa.eu

Message from the project officer

Intellectual Property Architectural Blocks (IP cores) are reusable design components that are used to build advanced integrated circuits (IC) and systems-on-chip (SoC). IP cores are internally developed by semiconductor firms and electronic system manufacturers, and they can also be licensed from independent IP vendors. As IC designs are becoming highly complex it is typically impossible to create new IC designs without pre-designed IP blocks as a starting point. IP cores are perhaps the most knowledge-intensive link in the information economy value chain. They define the capabilities of billions of electronic devices produced every year.

I give below some of the preliminary findings of an ongoing European Commission study aimed at creating an overall view of the developments in the area (the study's findings do not necessarily represent the official EC view on the subject):

- There are over 150 European firms that license semiconductor IP. Most of these are small, with less than 10 employees, but several leading IP vendors are headquartered in the EU.
- Among the top 20 independent IP vendors worldwide, nine are headquartered in the EU or have substantial development activities in European countries. The industry operates on a global scale, so firms that have headquarters in Europe usually have research and development activities in several geographical regions, including India, China and the US.
- The IP market is only around 1% of the overall semiconductor market. However, about four or five times more reusable IP blocks are developed internally than are sold on the market. The volume of reusable IP design activities, therefore, may well be five times bigger than market studies estimate.
- Many IP vendors have difficulties with profitability and growth. The growth of small firms is limited partly because successful firms are often acquired by semiconductor design firms, electronic design automation (EDA) tool vendors, or larger IP vendors. Only some firms with highly competitive technologies, and firms that have been able to participate in successful IP ecosystems, seem to be able to grow profitably.
- For almost five decades, computing has been driven by continuous miniaturization. The size of transistors on a semiconductor die is now measured in nanometers. In the near future, this fundamental driving force will evaporate. Miniaturization is becoming increasingly expensive, its technical and economic benefits are declining, and new alternative sources of value are emerging in the knowledge economy.
- One of the possible future scenarios is that the next rapid expansion in the ICT industry will occur in low-cost computing architectures that are implemented using mature production processes. IP cores may play a crucial role in this scenario.

Panos Tsarchopoulos
Project Officer ■

sented IMEC's flexible air-interface baseband platform for SDR mobile terminal research. The platform is based on IMEC's ADRES reconfigurable architecture and software-programmable processors. In addition to the importance of software engineering for SDR, she also emphasized the power consumption on the actually implemented circuits.

Jukka Wallinheimo from Nokia research described the challenge for SDR was in computerizing the radio modem. For this, the standard interfaces become highly important to ensure interoperability between subsystems from different vendors. Jouni Isoaho from the University of Turku assessed future platform development towards autonomous computing architectures, which take into account not only the different radio air interfaces and standards but also hardware fault tolerance, as autonomously as possible.

Harri Saarnisaari from the University of Oulu's Center for Wireless Communications raised an interesting signal processing aspect in radio communication systems by describing them in a time-frequency plane, e.g., OFDM operates in frequency domain, WCDMA in time domain, and novel multicarrier-CDMA systems in both dimensions. The correspondence between the systems can be exploited by using time-frequency transformations to do mapping between different systems in SDR algorithms.

Eleni Patouni from the University of Athens concluded the invited talks in SoC 2008 by introducing self-management of radio communication networks. This closed the circle, bringing Joe Mitola's concepts to more autonomous network management. In addition to the invited talks on SDR and cognitive radio, various System-on-Chip issues were raised and solved

in 36 contributed papers and five industry forum presentations.

SoC 2008 (4.-6.11.2008) was the 10th anniversary event in Tampere. The general chair of the event was, for the 10th time, too, Prof. Jari Nurmi. In the audience there was at least one faithful SoC-enthusiast: the managing director of Mentor Graphics Finland, Markku Tuomola. He has never missed the opportunity to participate in SoC. In 2009 SoC will take place in October co-located with the IEEE Workshop on Signal Processing Systems.

SoC: http://soc.cs.tut.fi/ (theme in 2009: SoC Technology for Signal Processing)
SiPS: http://www.sips09.org/ (theme in 2009: Signal Processing on System-on-Chip) ■

Call For Papers:
International Symposium on System-on-Chip
Tampere, Finland
October 6-7, 2009
Scientific paper deadline: May 15, 2009
Industry forum talk abstract deadline: August 7,2009
Special theme: SoC Technology for Signal Processing
Co-located with IEEE Workshop on Signal Processing Systems
http://soc.cs.tut.fi/

Technology Transfer: A New Methodology

Professor Jose Duato and Associate Professor Federico Silla from the Technical University of Valencia, Spain have established a new methodology to transfer technology to industry. Instead of developing a collaborative research project with a single company, or developing research with a small group of companies in the context of a European-funded project, the new methodology is based on simultaneously transferring technology to an entire consortium, composed of several companies. In this case, they have put this new mechanism into practice by transferring technology to the HyperTransport Consortium.

The HyperTransport Consortium consists of more than 50 worldwide technology-leading companies, such as AMD, Hewlett-Packard, Sun Microsystems, Cisco, Dell, IBM, Nvidia, and Cray.

The first result of this new way to transfer technology was presented by the consortium in February at the First International Workshop on HyperTransport Research and Applications, held in Mannheim, Germany. This first result is the definition of an extension to the HyperTransport 3.10 specification. The new extension is named High Node Count HyperTransport Specification 1.0, and it will be used as the basis for the implementation of high performance computing clusters. These clusters, used in data processing centres, internet servers, scientific database servers, etc. will be faster and easier to program because of the lower communication latency and the shared-memory architecture. Prior to the deployment of this new extension, nodes in a cluster had to use explicit message-passing communication in order to exchange information, thus involving the operating system in the communicating nodes and thereby noticeably increasing latency. By using the new extension, nodes in the cluster can share their memory space and therefore communication between processes can be achieved directly by hardware without requiring the involvement of the operating system and thereby reducing latency. Moreover, shared-memory communication is easier to program.

The new extension also allows a process running in one of the computers of the cluster to use all the memory available in the cluster. In this way, programs that could only be run in mainframes can be executed now in commodity computers, thus lowering cost. Additionally, costs are further reduced because memory resources in the computers in the cluster do not need to be oversized in order to accommodate large memory footprint processes. With the new extensions, when a process requires more memory than that available in the motherboard where it is running, it can borrow additional memory resources from other motherboards in the cluster.

Professor Silla's talk during the First International Workshop on HyperTransport Research and Applications in Mannheim.

The new extension has been developed in the context of the Advanced Technology Group (ATG), which is part of the HyperTransport Consortium. The ATG led by Professor Duato includes prestigious universities, such as the University of Heidelberg (Germany) and Georgia Institute of Technology (USA), and is additionally composed of leading companies from the HyperTransport Consortium.

Prof. Federico Silla
Universidad Politecnica de Valencia, Spain, fsilla@disca.upv.es ■

Joint Computer Architecture Seminar: RWTH Aachen University and University of Edinburgh

On February 11, 2009 a computer architecture seminar between two HiPEAC partners took place: the University of Edinburgh and RWTH Aachen University. For this event a research group from the Edinburgh Institute for Computing Systems Architecture (ICSA) group visited the Institute of Integrated Signal Processing Systems (ISS) of RWTH Aachen. Both groups have a long standing partnership, and the seminar has been the third of its kind since 2006. However, for the first time, it took place at the premises of ISS, or more precisely, in the newly opened headquarter building of the UMIC Excellence Cluster.

During the seminar, different subjects from the research areas of both groups were presented and intensively discussed. The major focus was on processor customization for embedded applications, fast simulation, advanced profiling, and multi-core programming. By the end of the seminar, both groups had found large overlaps in their research interests, indicating the relevance of those topics in the domain of HiPEAC and embedded systems design in general. The seminar day was rounded up with a joint dinner and a guided tour through Aachen's medieval town centre, where all attendees had a chance to get together in an informal setting. Both groups concluded that these seminars were a great opportunity for exchanging both technical and cultural knowledge. Therefore, planning for the next event started as soon as the seminar was over. ■

Prof. Bjoern Franke introducing ICSA activities

Release of DDG version 2.0

DDG is a generic C++ library that handles data dependence graphs for optimising compilation. It is built on top of the LEDA graph library. Our graph library is particularly for research purposes, where people are willing to make quick, robust and modular implementations of code-optimisation techniques for basic blocks and innermost loops (modeled by regular mono-dimensional data dependences). We manage directed acyclic graphs (DAGs) for basic blocks and cyclic graphs for innermost loops.

The user is able to take advantage of many advanced algorithms for graphs that are not present in boost. It is also possible to configure the library for different instruction set architectures and multiple register types.

We can also deliver, upon request, all the data dependence graphs for many standard benchmarks (mediabench, spec2000, spec2006, ffmpeg). For now, we can be contacted by email for such delivery. We are willing to provide these data dependence graphs to the community as a set of benchmark suites for instruction scheduling technology - either for acyclic or cyclic scheduling.

Version 2.0 of the DDG library exceeds the functionality of version 1.5 in a number of points:

- Now, DDG may handle complex ISA with multiple register types. Each instruction may write multiple results of distinct types.
- We have added new classes : ARCHITECTURE, REGISTER_TYPES
- Enhancements of precedent classes: adding check methods, adding multiple register types
- Register saturation computation can be done for any register type.
- Bug fix in Register Saturation Computation
- New enhanced XML formats for input/output DDG and architecture description.
- Deprecated DDG file formats: gl.
- Now, DDG is distributed under the LGPL software licence instead of GPL.

http://www.prism.uvsq.fr/~touati/sw/DDG

■

CUDA Seminar at the University of Patras

For two days (19th and 20th March, 2009) CUDA was the focus of attention at the University of Patras, the third largest university in Greece. During that time, a seminar was held by Richard Membarth, a PhD student of the Hardware/Software Co-Design Chair (Prof. Teich) at the University of Erlangen-Nuremberg, in cooperation with Professor Stefanos Kaxiras, University of Patras. On Thursday, a seminar on CUDA, a new programming paradigm to leverage the potential of current highly parallel graphics cards hardware, was presented at the crowded conference centre. The architecture of current graphic cards as well as the programming model was discussed during the first day.

At the end of the first day, the students had the possibility to register for a hands-on workshop on Friday. Due to the high level of interest, the number of places for the workshop was limited. During the workshop the students learned how CUDA differs from other parallel programming models like OpenMP. They also had the chance to present their own work on CUDA and to experience CUDA for themselves. ■

Load Balancing for Mobile Devices (LoBaMo)

A new paradigm for future mobile devices

With the recent advances in wireless networks and the exponential growth in the usage of multimedia applications, multi-core platforms have emerged as the key for future mobile devices. A new paradigm has simultaneously been created, which we refer to in the ICT-eMuCo project as Load Balancer for Mobile Devices.

In general, a load balancer can be seen as a component that spreads the work between two or more entities such as: CPUs, clustered computer systems, network links or other resources, in order to get optimal resource utilization. Load balancing for mobile devices enables the concurrent and parallel execution of applications and control signals by the efficient execution of multiple threads on processors with multiple cores.

The load balancer for mobile devices enables:

- effective exploitation of the multi-core hardware platform
- reduction of the power consumption
- scalability
- and most importantly from users' perspective: assurance of the expected user experience.

Simultaneously meeting deadlines of real-time transactions and shortening the response time of online transactions is critical in mobile communications as it directly influences the user experience. As is shown in Figure 1, in a mobile device, there are threads generated in the modem subsystem and in the application subsystem. All these threads have native priorities by manual programming assignation or by an operating system assignation. Nevertheless, all the threads have to be distributed to be executed on the multiple cores. The role of the load balancer is to monitor thread execution and distribute the threads on the available cores based on: a) the capacity limit of the unit processing; b) the capacity limit of the communication channels;c) the profiling information of the processing stages; and d) approximate knowledge of the processing demand of the expected workload, which is in part provided by the protocol stack. Therefore, the scheduling policy of the load balancer depends not only on the performance metrics obtained from the hardware platform and applications, but by the

ICT-eMuCo is a European project supported under the Seventh Framework Programme (FP7) for research and technological development (www.emuco.eu)

Project Coordination:
Attila Bilgic
Ruhr-University Bochum
Attila.bilgic@is.rub.de

rich amount of information provided by the protocol stack and the predictable interdependency of the protocol stack processes, which mark the difference between a general purpose load balancer and a load balancer for mobile devices.

The research group involved with the "load balancer for mobile devices" is part of the eMuCo project and its research provides a natural answer to the requirements of the system platform for future convergence mobile devices. It is expected that the output of this research group will have an impact on market trends for mobile devices. ■

Maria Elizabeth Gonzalez
Institute for Integrated Systems
Ruhr Universität Bochum
D-44780 Bochum, Germany
Elizabeth.gonzalez@is.rub.de

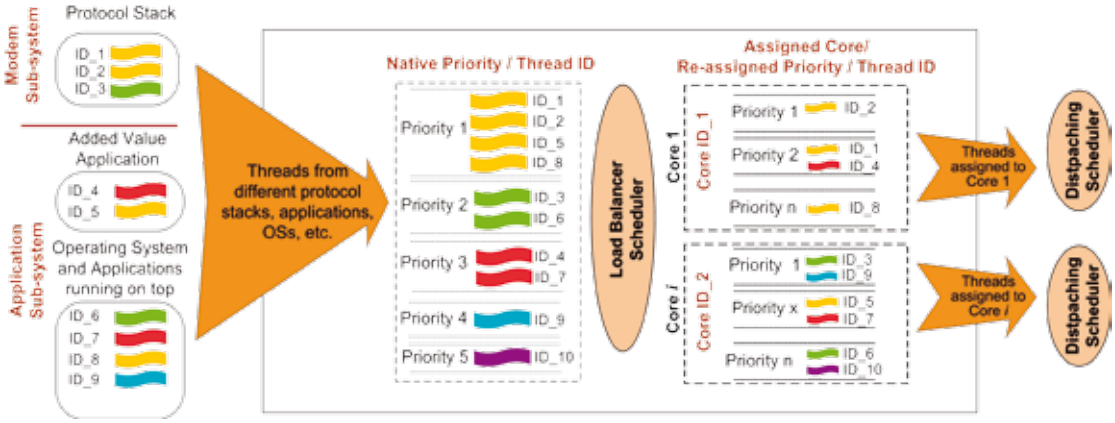

Functionality view of the Load Balancer

Access to the Collective Optimization Database

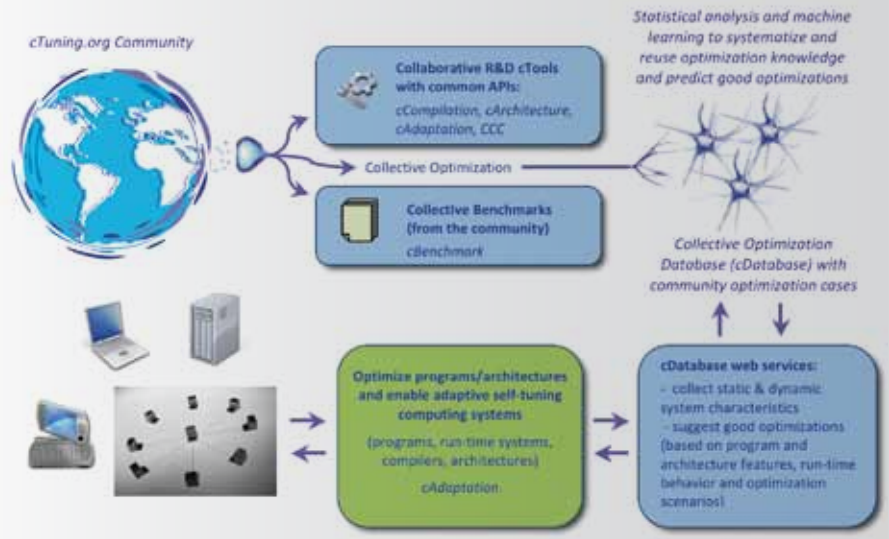

Public access to the Collective Optimization Database (COD)* has been arranged to provide a common repository for sharing knowledge about program and architecture optimizations. Optimization cases, including program transformations or architecture configurations to improve execution time, code size, power consumption, etc can be submitted either manually through the web form or automatically using Continuous Collective Compilation Framework** through COD web-services.

The COD is intended to improve the quality of academic research by avoiding costly duplicate experiments and providing replicable referable results. It can be combined with the Collective Benchmark/MiDataSets** (collection of publicly available benchmarks and datasets) and provide detailed performance analysis and comparison of different programs, datasets, compilers and architectures.

The Collective Optimization Database is also intended to help end-users and companies optimize their computing systems. The Collective Optimization Database is currently used in the EU FP6 MILEPOST project (http://www.milepost.eu) to automate compiler and architecture design and program optimization based on statistical and machine learning techniques. ■

Project leader: Grigori Fursin,
INRIA Saclay, France
grigori.fursin@inria.fr

(*) Website: http://ctuning.org/cdatabase/
(**) Tools: http://ctuning.org/wiki/index.php/CTools

Best Paper Award: 1st Data Prefetching Championship, HPCA-15

The team formed by the HiPEAC researchers Luis M. Ramos, J. L. Briz, P. Ibáñez and V. Viñals (Computer Architecture Group, Univ. Zaragoza) has been awarded the Best Paper Award at the 1st Data Prefetching Championship (DPC-1) that was held during the HPCA-15 at Raleigh, North Carolina, on February 14-18, for their contribution entitled "Multi-level Adaptive Prefetching based on Performance Gradient Tracking".

Sponsored by Intel, JILP and IEEE TC-uARCH, the DCP is a competition for data prefetching algorithms, where contestants are given a fixed storage budget to implement their best prefetching algorithms on a common evaluation framework provided by the DPC organizing committee. In this first DCP-1 contributors were able to submit up to three different prefetching codes that were tested by the Program Committee under the common simulation framework. Results are publicly available at http://www.jilp.org/dpc/online/DPC-1 Program.html

The contribution, awarded as the Best Paper, introduces a multi-level prefetching framework with three set-ups, respectively aimed to minimize cost, minimize losses in individual applications or maximize a with moderate cost. Performance is boosted in all cases by a sequential tagged prefetcher in L1 with an effective static degree policy. Prefetch filters are applied in both L1 and L2. In L2 a brand-new adaptive policy is used that selects the best prefetching degree within a fixed set of values by tracking the performance gradient. An accurate, home-made, delta-correlating prefetcher is used in two out of the three proposals. ■

Sabbatical: Enrique F. Torres

I received an MS degree in Computer Science from the Polytechnic University of Catalunya in 1993, and a PhD degree in Computing Science from the University of Zaragoza in 2005. I am Assistant Professor in the Computer Science and Systems Engineering Department at the University of Zaragoza, Spain. I am a member of the Computer Architecture Group of Zaragoza (gaZ) led by Victor Viñals. My research interests include processor microarchitecture, memory hierarchy, and parallel computer architecture.

Memory hierarchy research is one of the oldest and most important fields in computer architecture. Processor and memory speed keep on growing at different rates. The Cache Coherence Protocol ensures that all processors in a system observe a consistent view of the shared memory. The coherence shared memory programming model is used in a wide range of applications. As the number of processors and threads integrated into the same chip increases, the cache coherence protocol can become a bottleneck. Many works have addressed this scalability problem since long ago and commercial systems have been built with many processors, but the lack of proper tools has limited the scope of research. Running an extensive number of threads in a large number of processors for long enough is impossible or at least impractical with software-only simulators.

Thanks in part to HiPEAC (Sep.-Nov. 2008) and the Spanish Science and Innovation Ministry, I am on a kind of sabbatical leave for study and research at the University of California at Berkeley, where I am working at the International Computer Science Institute (ICSI), collaborating with Krste Asanovic. I am participating in the Parlab and RAMP projects working on scalable cache coherence protocols and its related hardware structures.

ICSI is a research center in computer science and brings together researchers from around the world in areas, such as computer networking, speech and language processing, bioinformatics, computer architecture, and artificial intelligence. This is accomplished through active international visitor programs in which researchers from abroad visit ICSI to work with staff scientists and their networks of academic, government, and industrial partners. ICSI maintains affiliation with the University of California at Berkeley and some international sponsors like Germany, Spain, Switzerland, and

Finland. The Director of the Institute is Professor Nelson Morgan of the UC Berkeley Electrical Engineering faculty. There are approximately 130 scientists in residence at ICSI, including principal investigators, postdoctoral fellows, visiting researchers and students.

With around 70 researchers, the UC Berkeley Parallel Computing Laboratory, or Par Lab, is a five-year, $10m, multi-disciplinary research project, exploring the future of parallel processing. Directed by David A. Patterson and funded by Intel and Microsoft, the aim is to productively create efficient and correct software that scales smoothly as the number of cores per chip doubles biennially. To that end, the task is divided into two layers, an efficiency layer that aims at low overhead for 10 percent of the best programmers, and a productivity layer for domain experts that use software created by the other layer. The Berkeley lab is exploring everything from novel chip architectures to new parallel applications as well as parallel languages, autotuners, runtime environments and the test and emulation tools needed to create them.

The Research Accelerator for Multiple Processors (RAMP) is a collection of projects from different universities with the common idea of building a infrastructure based on FPGAs to help the development of hardware and software for manycores. Advances in FPGA capacity and speed now enable a 1024-processor system to be simulated with just a few FPGAs and fast enough to develop and debug new parallel applications and operating systems before the actual chip comes out. Within RAMP, the RAMP Gold project is developing a model of a manycore processor to be used by the UC Berkeley Par Lab in support of experiments in architecture, operating systems, languages, and parallel application development. To attain high emulator capacity and performance, the RAMP Gold emulation uses split timing and functional models, both of which are highly multithreaded. The goal is to obtain more than one GIPS with 8 BEE3 boards in cycle-level timing emulation.

Contact Information:
Enrique.torres@unizar.es

Microtech International S.A. New Member of HiPEAC Network

In January 2009, the HiPEAC Network steering committee accepted Microtech International S.A. as a new member of the HiPEAC community.

Microtech International is an SME that was established in 1990 by research workers from Wroclaw University of Technology with the purpose of running business activities in the area of the application of modern electronic and information technologies in both industry and in research. Microtech is located in Wroclaw – the capital of the Lower Silesia region in Poland – and employs some 35 people. The core staff of the company acquired their original experience at the Wroclaw University of Technology through extensive cooperation with industry. The company employs scientists from Wroclaw University of Technology on a regular basis.

For over 15 years Microtech has designed, developed and implemented digital equipment for data processing. The most frequently used technologies are FPGA circuits and DSP processors. We specialize in research, development & prototyping of advanced electronics for digital signal processing purposes. Microtech provides military as well as industrial solutions, especially for those requiring real-time signal processing, advanced signal analysis (for example spectral analysis) and large processing power at very high input/output signal speed. We offer improvement of processing algorithms to boost efficiency of designed solutions. Microtech has also expertise in business software development. This expertise encompasses: application programming; large databases for industry; communication protocols; as well as software integration, including integration of software applications delivered by different providers. We use various system platforms.

Among the many projects accomplished by Microtech in the field of electronics solutions, there are dedicated embedded systems for military purposes, as real-time signal processing modules for radars, sonars or telecommunication. The other areas of our applications are: industrial controllers, video processing modules and measurement equipment. Microtech International is equipped with advanced technical infrastructure necessary for successful implementation of both research and application activities.

The company introduced standards related to the safety of the design and production of electronic systems and equipment applied in the military and aeronautic industry (DO-160E, DO-178B, DO-254, MIL-217F, MIL-STD-1629A, MIL-STD-202G and others).

Since 2004 Microtech has cooperated with European aeronautic and

Guest Column

space industry representatives (EADS, Airbus, Dassault Aviation, Thales and others) in the frame of research projects.
In 2008, after the certification process and formal audit at the company premises, Microtech was approved as a supplier to Thales in the area of research and development, prototyping and production of modern electronic systems and equipment for the aeronautic and space industry.

Agnieszka Skotarczyk
International Cooperation
Plenipotentiary

Microtech International S.A.
ul. Wolowska 20, 51-116 Wroclaw

tel. +48 71 346 35 00,
fax +48 71 346 35 99
tel. kom. +48 609 666 811
a.skotarczyk@microtech.com.pl
www.microtech.com.pl

HiPEAC Students

Collaboration Grant Report: Vincenzo Rana

I am a PhD. student from Politecnico di Milano under the supervision of Donatella Sciuto. Thanks to a HiPEAC collaboration grant, I had the opportunity to spent some time at the Integrated Systems Lab at EPFL - Ecole Polytechnique Fédérale de Lausanne, Switzerland. The main activity performed at the host's site has been he definition of a novel design flow for the mapping of multiple applications on reconfigurable FPGA-based systems based on 2D reconfigurable Networks-on-Chip, starting from the high-level specification of communication requirements and constraints.
The main objective of the flow is the generation, at design-time, of a set of different mappings of the cores, each one optimized for a particular application that has to be run on the target device. The information about these mappings can be used at runtime in order to configure on the FPGA device the best mapping of the cores for the application that is currently running on the system. The differences among the mappings have to be minimized in order to reduce the reconfiguration time overhead that is needed in order to pass from one application to the next. The reconfiguration overhead is, in fact, directly proportional to the number of reconfigurable regions (each one holding a single island of cores) that have to be reconfigured.

The design flow exploits a multi-stage design optimization algorithm that can achieve reductions in reconfiguration latency up to 43% with respect to state-of-the-art application mapping methods for FPGA. Taken into account are the reconfiguration costs and SoC block reuse between the different applications that need to be executed dynamically on the FPGA.

Furthermore, the design flow can be tuned with an α parameter in order to considerably reduce the average number of hops of the final system, thus reducing the average latency of each communication. The main drawback is the increment of the average number of reconfigurable regions that have to be reconfigured. For this reason, a value around 70% is usually used in the proposed design flow in order to obtain a good trade-off, even if the designer can manually tune this parameter to achieve the desired solution.

The presented design flow can reduce the reconfiguration overhead from 29.7% (when the percentage of the cores shared by the input applications is set to 0%) to more than 71% (when the percentage of the cores shared by the input applications is set to 70%) w.r.t. a scenario in which the underlying communication infrastructure can be reconfigured only by means of complete reconfigurations of the target device (16 reconfigurable regions over 16 have to be reconfigured in order to change the underlying communication infrastructure). Thus, the developed design flow is able to take great advantage of both the intrinsic and extrinsic similarities among the input applications, and scales very well when the number of input application increases, if they are characterized by enough similarities.

EC FP7 CRISP Project: Cutting edge Reconfigurable ICs for Stream Processing

Project Management:
Recore Systems
Coordinator: Dr. P. M. Heysters
Project Manager: Dr. Kim Sunesen
Website: www.crisp-project.eu
Email: info@crisp-project.eu
Participants:
Recore Systems (NL)
University of Twente (NL)
Atmel (D)
Thales Netherlands (NL)
Tampere University of Technology (FIN)
NXP Semiconductors (NL)

Main Objectives

The CRISP project develops a single highly scalable and efficient reconfigurable platform for a wide range of streaming applications.

Streaming applications have very high market potentials and will drive demand for reconfigurable platform chips. They range from low-end, consumer electronics and automotive applications to demanding, high-end, medical and defense applications.

In particular in the low-end markets, it becomes increasingly difficult to predict which applications are going to be successful in the future. For this reason, system designers look to programmable platform chips, which enable expected market trends to be anticipated, but which allow flexibility when the market develops differently.

The ambition of the CRISP project is to devise a single processing platform solution that will drive future technological innovations in the consumer, automotive, medical and defense markets. These innovations are enabled by smaller, cheaper, more reliable and efficient electronic components and systems. The envisioned platform solution includes a massive multi-core processing architecture in combination with innovative design-time and run-time tools.

Technical Approach

Targeting streaming applications, CRISP addresses optimal utilization, efficient programming, and dependability of reconfigurable many-cores. The project is organized around four central themes:

Streaming applications

The target is the wide range of streaming-based applications. As proof-of-concept, stream processing for beam forming and global satellite navigation is implemented.

General Stream Processor (GSP)

A GSP is a dynamically reconfigurable, many-core platform for streaming applications. Its scalable architecture combines flexibility and high performance with low power and a small footprint. The project develops a prototype GSP.

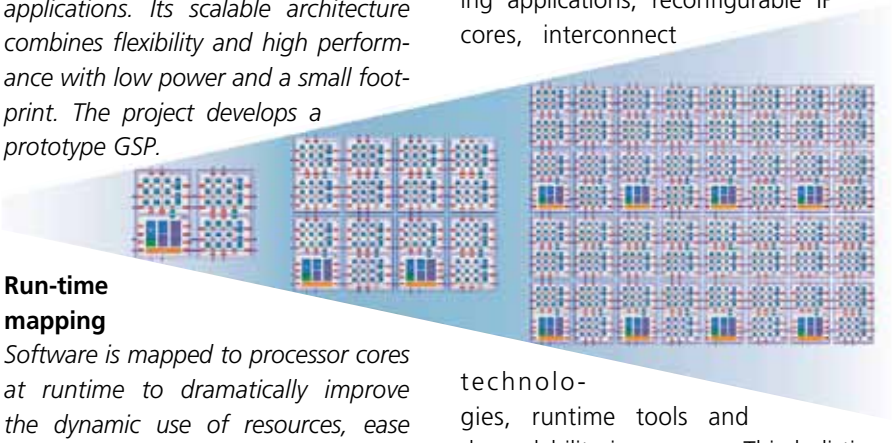

Run-time mapping

Software is mapped to processor cores at runtime to dramatically improve the dynamic use of resources, ease programmability, and enable dynamic hardware fault diagnosis and repair.

Dependability

Dependability and yield of deep-submicron chips are improved using new techniques for static and dynamic detection and localization of faults and (dynamically) circumventing faulty hardware.

Key Challenges

CRISP seeks to answer the following questions which are fundamental to the effective development of multi-core systems and for which no solutions are readily available.

1. *How can the intrinsic processing potential of a massive multi-core architecture be exploited optimally for a wide range of applications?*
2. *How can multi-core systems be programmed efficiently?*
3. *How can large, deep-submicron, multi-core integrated circuits be made reliable and self repairing?*

Expected Impact

Due to the increasing costs and complexity of ASIC design and production, and the demand for shorter time-to-market, it is expected that future streaming applications will favour generic, reconfigurable, multi-core computing platforms over conventional architectures.

CRISP anticipates this expected architectural shift by researching streaming applications, reconfigurable IP cores, interconnect technologies, runtime tools and dependability issues now. This holistic approach is necessary for European companies to achieve global leadership in computing solutions and products for streaming applications.

The market for fine-grained (FPGA-based) reconfigurable logic devices is largely US dominated. The new market for coarse-grained reconfigurable computing will become an important market segment and provides an excellent opportunity for Europe to take world-leadership.

The CRISP project aims to build the European foundations for worldwide leadership in embedded computing systems by addressing the inevitable paradigm shift that embedded systems for streaming applications will go through in the near future.

■

Collaboration Grant Report: Alberto Ros

My name is Alberto Ros and I am a PhD student at the University of Murcia, in Spain. I am currently writing up my dissertation which mainly focuses on cache coherence protocols for multiprocessor systems. The main focus of my thesis is on a family of cache coherence protocols called Direct Coherence Protocols, which avoid the indirection suffered in directory protocols by accessing the home node for most cache misses. By avoiding this indirection, cache misses are accelerated and, therefore, the execution time for most parallel applications is significantly reduced.

Last summer HiPEAC gave me the chance to spend four months at the University of Edinburgh collaborating with Assistant Professor Marcelo Cintra. The time spent at Edinburgh was a good opportunity for me to learn more about another research field that I am also very interested in: the mapping of memory blocks in a distributed shared cache, i.e., a NUCA (Non Uniform Cache Access) cache. This is a very hot research area with important recent publications and I would like to continue working in this area for some time.

The problem that we faced during the collaboration was the trade-off between cache access latency and cache miss rate in a NUCA cache for many-core chip multiprocessors (CMPs). In particular, there are two main policies for performing the mapping in a distributed shared cache: physical mapping policy and first-touch policy.

In a physical mapping policy, a set of bits in the block address defines the home cache bank (i.e., the bank where each block maps) for every memory block. In this way, blocks are assigned to banks in a round-robin fashion, which does not take into account the distance between requesting cores and home banks. This distribution of blocks leads to large L2 cache access latencies.

On the other hand, if the mapping is managed by the OS, more efficient policies can be easily implemented, such as a first-touch policy. This policy tries to map memory blocks to the cache bank closer to the core that more frequently access them. Although this mapping enables a reduction of the average access latency to a L2 cache bank, it can increase L2 cache miss rate when the working set of the application is not well-balanced among cores.

We proposed a distance-aware, round-robin mapping policy, an OS-managed scheme which tries to map memory pages to the local cache bank of the first requesting core, like a first-touch policy, but also introduces an upper bound on the deviation of the distribution of memory pages among cache banks, which lessens the cache miss rate and, therefore, the number of expensive off-chip accesses. This upper bound can be controlled by the OS through a threshold value. If the number of pages mapped to a cache bank reaches that value, all banks at one-hop distance are checked. If all of them have also reached the threshold value, banks at two hops are then checked. This algorithm iterates until a bank, whose value is under the threshold, is found. One of the main advantages of the proposed policy is that a good trade-off can be achieved between cache access latency and hit rate without requiring extra hardware structures.

We also observed that distance-aware mapping policies managed by the OS can, in some cases, hurt the hit rate of private caches. This happens when the same bits that define the home bank are used for indexing the private caches. In these cases, some sets in the cache can be overloaded while others remain almost unused. This imbalance increases conflict misses. This observation led us to explore the significance of the bits used to index caches and we found some interesting future work for collaboration between the University of Edinburgh and the University of Murcia, which we will try to carry out in the near future. ■

HiPEAC Student at ICT Lyon 2008

From 25th to 27th November 2008 Lyon was crowded with more than 4500 delegates for the biennial ICT Event organized by the European Commission. The goal of the event is to set the research agenda for the coming years and to tackle the problems we can expect in the next decade. People from both academia and industry, and policy makers came to the most important forum for discussing research and public policy in information and communication technologies at European level. Among the people were also 80 "young reporters", who were PhD students from various European Universities and networks. Sean Rul was the delegate for the HiPEAC NoE. ■

HiPEAC Students

The opening of the event was attended by many notable and influential people. First of all, there was Viviane Reding, European Commissioner for Information Society and Media. We were also addressed by Denis Ranque, CEO of Thales, and Gérard Collomb, the Mayor of Lyon. A peculiar fact was that the opening talks were given partly in English and partly in French. In the afternoon Albert Fert, Laureate of the 2007 Nobel Prize in Physics, gave a keynote in which he presented his view on the technology that could succeed CMOS once it is no longer possible to miniaturize transistors any further (the end of Moore's Law).

Besides opening speeches and keynotes, there were also a lot of parallel sessions during those three days in which all the different aspects of ICT-related research in Europe were discussed. Also our beloved subject of embedded systems and compilers was on the agenda. Topics such as robotics, digital libraries, technology-enhanced learning, 3D-Internet and many more (sometimes very exotic) subjects were also supported by large research communities. Throughout the conference there were also several exhibitions and networking sessions from the commission, SMEs and other interests groups.

Since there were so many parallel sessions, each young reporter was assigned to report on several sessions, but he/she was also free to attend any other session of interest to them. In order to get to know one another better, the young reporters spent time together visiting Lyon during the evening. One of the local students, Romain Vuillemot, took us to see some of the city's highlights as well as a few bars to warm up in. So

Albert Fert, Laureate of the 2007 Nobel Prize

thanks to the young reporter initiative of the European Commission we came into contact with a lot of interesting people from different domains and learned about the large research spectrum within Europe.

If you want more information about ICT 2008 you can visit http://ec.europa.eu/information_society/events/ict/

Sean Rul
Ghent University - Department ELIS
Sean.Rul@elis.UGent.be

PhD News

Design of Hardware Accelerators for Embedded Multimedia Applications

By Claudio Brunelli (claudio.brunelli@tut.fi)
Prof. Jari Nurmi
Tampere University of Technology, Finland
December 2008

A RISC core alone cannot achieve the desired computational capability needed to meet the requirements of modern applications, especially demanding ones like audio/video compression, image processing and 3D graphics. Thus some additional dedicated resources are needed. The research presented in this thesis consists mainly of the design-space exploration of two coprocessors, MILK and BUTTER. MILK is a floating-point unit (FPU), while BUTTER is a coarse-grain reconfigurable machine. Their design space was explored using a parametric, synthesizable VHDL model, which was implemented on ASIC standard-cell technologies and on FPGA devices. The implementation on FPGA was useful as a prototype but also as a demonstration platform for running real-world applications like Mp3 and H.264 decoders. A GUI-based tool for the simulation and debugging of the execution within the FPU was also developed. The second part of this work describes the architecture of the coarse-grain reconfigurable machine named BUTTER. It was initially meant to be an accelerator to enable running real-time audio/video processing applications, but the range was broadened to image processing, GPS signal acquisition and tracking, and 3D graphics, enabled by the introduction of special architectural features like support for subword and floating-point operations, which represent an absolute novelty in the field of coarse-grain reconfigurable machines. A Graphical User Interface (GUI) based tool was developed for the automatic generation of the configware used to configure BUTTER. An entire System on Chip (SoC) featuring a 32-bit RISC core, MILK FPU and BUTTER takes 65173 ALUTs on a Stratix II EP2S180 FPGA device and runs at 34 MHz. Using BUTTER and MILK, some algorithms can achieve a speed-up from one up to two orders of magnitude compared to plain software implementation.

An OS-Based Alternative to Full Hardware Coherence on Tiled Chip-Multiprocessors

By Christian Fensch (cf309@cam.ac.uk)
Prof. Marcelo Cintra and Prof. Bjoern Franke
University of Edinburgh, UK
December 2008

The interconnect mechanisms (shared bus or crossbar) used in current chip-multiprocessors (CMPs) are expected to become a bottleneck that prevents these architectures from scaling to a larger number of cores. Tiled CMPs offer better scalability by integrating relatively simple cores with a lightweight point-to-point interconnect. However, such interconnects make snooping impractical and, thus, require alternative solutions to cache coherence.

This thesis proposes a novel, cost-effective hardware mechanism to support shared-memory parallel applications that forgoes hardware maintained cache coherence. The proposed mechanism is based on the key ideas that mapping of lines to physical caches is done at the page level with OS support and that hardware supports remote cache accesses. It allows only some controlled migration and replication of data and provides a sufficient degree of flexibility in the mapping through an extra level of indirection between virtual pages and physical tiles.

The proposed tiled CMP architecture is evaluated on the SPLASH-2 scientific benchmarks and ALPBench multimedia benchmarks against one with private caches and aa distributed directory cache coherence mechanism. Experimental results show that the performance degradation is as little as 0%, and 16% on average, compared to the cache coherent architecture across all benchmarks for 16 and 32 processors.

Part of this thesis has been published as a paper at the 14th IEEE International Symposium on High-Performance Computer Architecture in 2008. The paper got credited with a HiPEAC Paper Award.

System-Level Design Methods and Tools for Multimedia Heterogeneous Multiprocessor Platforms

By Victor Reyes
Tomas Bautista
Institute for Applied Microelectronics, IUMA
University of Las Palmas de Gran Canaria, Spain
December 2008

Electronic System Level design methods and tools are being proposed as a complement to conventional solutions to improve the productivity of the designers. ESL is rapidly settling down thanks to standards such as SystemC and TLM2.0 and the acceptance of Virtual Prototyping in the Semiconductor industry. VP can be applied to several use-cases during the design cycle. However, different use-cases have different requirements (in terms of accuracy, simulation speed, etc), which makes it very difficult to come up with a single modeling style that fits all these requirements. In addressing this problem, the main contributions of this thesis are an innovative modeling method, tools, and extensive libraries, which enable model reuse and refinement along different ESL abstraction levels relying on the SystemC and TLM 2.0 standards. Advanced modeling concepts, such as hierarchical modeling, dynamic layout, structural composition, advanced clock objects and TLM2.0 CA interfaces are introduced as an efficient way to create speed optimal refineable architectural models with limited effort. The efficiency obtained is quantified and demonstrated for several validation examples in complex multimedia codecs and MPSoC targets.

High-Performance Architectures for High-Radix Switches

By Gaspar Mora Porta (gaspar.mora.porta@intel.com)
Jose Francisco Duato Marin, Jose Flich Cardo
Technical University of Valencia, Spain
February 2009

As the optimal radix for switches increases due to the benefits in lower latencies, there is an overall reduction in cost and power consumption; the traditional switch architectures are no longer valid because of either low-performance or non-scalability with the number of ports.

This dissertation proposes a new switch architecture suitable for high-radix switches called Partitioned Crossbar Input Queued (PCIQ) that deals with one of the main constraints in high-radix switch design, the excessive memory requirements.

PCIQ relies on a smart partition of the crossbar into sub-crossbars, thus requiring less memory resources than other proposals for high-radix, yet obtaining high-performance and also increasing the arbiter efficiency. PCIQ uses two round-robin packet-based arbiters (one for each crossbar) that exhibit a linear cost and a logarithmic response time as the radix of the switch increases.

The other big issue on high-radix switches is the HOL blocking problem, which dramatically reduces the switch performance. Traditional solutions for removing the HOL blocking problem were based on VOQ schemes, but having a high number of ports on a high-radix switch prevents the use of any of them. In this dissertation, a new congestion management technique has been proposed. This solution is called RECN-IQ and is specific for IQ switches. It differs from the original RECN idea (suitable only for CIOQ switches) in being highly efficient and easy to implement, reducing the memory requirements to the maximum. RECN-IQ introduced a novel statistical approach for detecting congestion using just a single queue per input port.

The idea behind RECN-IQ is to add some extra queues dynamically allocated for congested packets, starting with a simple IQ switch with a single queue per input port. Congestion is detected as soon as HOL blocking begins to act, setting aside (in those extra queues known as SAQs) the congested packets in an efficient manner. Therefore, HOL blocking is completely eliminated (as proven by the simulation results). The hardware requirements for RECN-IQ are reduced, making feasible its implementation on any IQ-based switch architecture like PCIQ.

By combining the PCIQ microarchitecture with RECN-IQ, a new switch architecture (called here PCIQ-enhanced) is derived and evaluated in this dissertation. The PCIQ switch architecture inherits the benefits of the Partitioned Crossbar microarchitecture in reducing the memory requirements for high-radix designs with the power of a congestion management technique that dynamically removes the HOL, thus achieving maximum switch performance under all types of traffic.

Results have proven that by using switches with the RECN-IQ mechanism, the network will benefit from low cost switches and high-efficiency under any type of traffic pattern or network circumstances. All this makes the network predictable and stable in performance, no more drops in throughput because of congestion.

Using Machine-Learning to Efficiently Explore the Architecture/Compiler Co-Design Space

By Christophe Dubach (christophe.dubach@ed.ac.uk)
Prof. Michael O'Boyle
University of Edinburgh, UK
February 2009

Designing new microprocessors is a time-consuming task. Architects rely on slow simulators to evaluate performance and a significant proportion of the design space has to be explored before an implementation is chosen. This process becomes more time consuming when compiler optimisations are also considered. Once the architecture is selected, a new compiler must be developed and tuned. What is needed are techniques that can speedup this whole process and develop a new optimising compiler automatically.

This thesis proposes the use of machine-learning techniques to address architecture/compiler co-design. First, two performance models are developed and are used to efficiently search the design space of a microarchitecture. These models accurately predict performance metrics such as cycles or energy, or a tradeoff of the two. The first model uses just 32 simulations to model the entire design space of new applications, an order of magnitude fewer than state-of-the-art techniques. The second model addresses offline training costs and predicts the average behaviour of a complete benchmark suite. Compared to state-of-the-art, it needs five times fewer training simulations when applied to the SPEC CPU 2000 and MiBench benchmark suites.

Next, the impact of compiler optimisations on the design process is considered. This has the potential to change the shape of the design space and improve performance significantly. A new model is proposed that predicts the performance obtainable by an optimising compiler for any design point, without having to build the compiler. Compared to the state-of-the-art, this model achieves a significantly lower error rate.

Finally, a new machine-learning optimising compiler is presented that predicts the best compiler optimisation setting for any new program on any new microarchitecture. It achieves an average speedup of 1.14x over the default best gcc optimisation level. This represents 61% of the maximum speedup available, using just one profile run of the application.

Analysis, Design and Management of Multimedia Multi-processor Systems

By Akash Kumar (a.kumar@tue.nl)
Prof.dr. Henk Corporaal
TU Eindhoven, The Netherlands
April 2009

Modern multimedia systems are becoming increasingly multiprocessor and heterogeneous to match the high performance and low power demands placed on them by the large number of applications. The concurrent execution of these applications causes interference and unpredictability in the performance of these systems. In this thesis, an analysis mechanism is presented to accurately predict the performance of multiple applications executing concurrently. With high consumer demand, the time-to-market has become significantly lower. To cope with the complexity in designing such systems, an automated design-flow is needed that can generate systems from a high-level architectural description such that they are not error-prone and consume less time. Such a design methodology is presented for multiple use-cases -- combinations of active applications. A resource manager is also presented to manage the various resources in the system, and to achieve the goals of performance prediction, admission control and budget enforcement.

Design Methodologies for Embedded Multiprocessor Architectures on FPGA

By Antonino Tumeo (tumeo@elet.polimi.it)
Prof. Fabrizio Ferrandi,
Prof. Donatella Sciuto
Politecnico di Milano, Dipartimento di Elettronica e Informazione, Italy
April 2009

Multiprocessor Systems-on-Chip (MPSoCs) are the de facto standard for embedded systems design. Reconfigurable devices, like Field Programmable Gate Arrays (FPGAs) are today starting to be integrated in the design flow. They can be used in multiple ways: as one of the available resources for the system, as the target device for all the system, or as a prototyping device. This thesis proposes a flow for embedded multiprocessor systems design, exploiting reconfigurable devices from all these points of view.

In the first part of the dissertation, we introduce CerberO, a prototyping framework for MPSoCs on FPGA. We detail a series of proposals to assemble shared memory and pipelined, homogeneous and heterogeneous multiprocessors on FPGA, presenting hardware and software elements to manage synchronization, communication, and interrupt management. We validate the framework with realistic applications from the multimedia and communication fields, and discuss how our solutions can be exploited in the context of real-time multiprocessor solutions. We also use the framework to study the challenges and the advantages introduced by the support of partial dynamic reconfiguration.

The CerberO framework is available at the site: http://trac.elet.polimi.it/cerbero.

In the second part of this work, we move on to higher layers of the design flow for embedded systems, discussing some methodologies to reduce the efforts of embedded system developers when dealing with the growing complexity of these architectures. We firstly describe an automatic parallelization technique for homogeneous MPSoCs, and validate it with our prototyping platform. We then present a flow for automatic generation of hardware accelerators for embedded applications, targeting FPGA technology that uses evolutionary algorithms to find the best area/performance trade-offs. We test the produced accelerators on

PhD News

FPGA-based systems. Finally, we introduce a novel algorithm for simultaneous mapping and scheduling of task-partitioned applications on MPSoC architectures with configurable resources. We describe an Ant Colony heuristic algorithm that gradually constructs feasible solutions and searches the best ones. This algorithm is also extended to support partial dynamic reconfiguration, managing one-dimensional placing and configuration prefetch alongside mapping and scheduling.
These methodologies are integrated in the PandA framework available at: http://trac.elet.polimi.it/panda.

Upcoming Events

2009 IEEE International Symposium on Performance Analysis of Systems and Software (ISPASS 2009)
Boston, USA, April 26–28, 2009, http://ispass.org/ispass2009/

The 3rd ACM/IEEE International Symposium on Networks-on-Chip (NOCS 2009)
San Diego, USA, May 10–13, 2009, http://circuit.ucsd.edu/~nocs2009/

23rd International Conference on Supercomputing (ICS 09).
IBM T.J. Watson Research Center, Metro New York City Area, USA, June 08–12, 2009, http://www.ics-conference.org/

International Symposium on High-Performance Distributed Computing (HPDC).
Munich, Germany, June 11–13, 2009, http://www.lrz-muenchen.de/hpdc2009/

ACM SIGPLAN Conference on Programming Language Design and Implementation (PLDI 2009).
Dublin, Ireland, June 15–20, 2009, http://www-plan.cs.colorado.edu/~pldi09/

8th International Symposium on Parallel and Distributed Computing (ISPDC 2009).
Lisbon, Portugal, June 30–July 04, 2009, http://www.ispdc.org/

Acaces 2009, Fifth International Summer School on Advanced Computer Architecture and Compilation for Embedded Systems
L'Aquila, Italy, July 12 to July 18, 2009, http://www.hipeac.net/summerschool/

International Conference on Parallel and Distributed Processing Techniques and Applications (WORLDCOMP 2009).
Las Vegas, USA, July 13–16, 2009, http://www.world-academy-of-science.org/worldcomp09/ws

International Symposium on Systems, Architectures, Modeling, and Simulation (SAMOS IX).
Samos, Greece, July 20–23, 2009, http://samos.et.tudelft.nl/samos_ix/

Design Automation Conference 2009 (DAC).
San Francisco, CA, USA, July 26–31, 2009, http://www.dac.com/46th/index.aspx

Euromicro Conference on Digital System Design (DSD 2009).
Patras, Greece, August 27–29, 2009, http://www.dsdconf.org/

The Eighteenth International Conference on Parallel Architectures and Compilation Techniques (PACT).
Raleigh, North Carolina, USA, September 12–16, 2009, http://pact09.renci.org/

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 18
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

HiPEAC Cluster Meetings and Industrial Workshop.
Munich, June 2-4, 2009