

# HiPEAC<sup>info</sup>21

COMPILATION ARCHITECTURE

appears quarterly | January 2010

Network of Excellence on High Performance and Embedded Architecture and Compilation

- 2 Message from the HiPEAC Coordinator
- 3 Message from the Project Officer
- 4 HiPEAC Activity:
  - 4 - HiPEAC Computing Systems Week in Wrocław
  - 5 - ProRISC 2009 in Veldhoven the Netherlands
- 6 In the Spotlight:
  - 6 - EC FP7 2PARMA Project
  - 7 - EC FP7 MERASA Project
- 8 Announcement:
  - ACACES 2010
- 2 Community News:
  - 2 - Mateo Valero, New Member of the Academia Europaea
  - 3 - Release of RSlib and SIRAlib
  - 12 - HiPEAC Technical Reports Instrument Expanded
  - 16 - Christophe Dubach Received the BCS Distinguished Dissertation Award 2009
  - 16 - HiPEAC Member Named ACM Distinguished Member
- 9 Member Profile
  - 9 - Professor Andreas Herkersdorf, Technische Universität München, Germany
  - 10 - Professor Stefano Crespi Reghizzi, Politecnico di Milano, Italy
- 8 HiPEAC Start-ups
- 10 HiPEAC Students and Trip Reports
- 18 PhD News
- 20 Upcoming Events



**HiPEAC 2010  
Conference,  
Pisa, Italy,  
January 25-27**

[www.HiPEAC.net](http://www.HiPEAC.net)

HiPEAC Computing Systems Week in Edinburgh, UK, May 3-6 2010

# Message from the HiPEAC Coordinator

Koen De Bosschere

Dear friends,

A couple of weeks ago, we celebrated the transition to 2010. This reminds me that we are already living for a decade in the third millennium. When looking back, it is hard to believe how many things have changed over the last 10 years. To give a few examples: there was no Euro currency in 2000, there were only 16 EU member states instead of 27, we did not have a unified bachelor-master structure in European higher education. In 2000, 9/11 was still equal to 0.8181...



Today there are an estimated 4.6 billion mobile cellular subscriptions, compared to an estimated 0.82 billion in 2000. Broadband access just exploded over the last decade, turning the world into a global village, a process that is still continuing and even accelerating. A quarter of the world's population has access to Internet now. It is hard to picture how the world will look like ten years from now. One thing we know is that the mobile phone penetration in the developing countries today is already at the level it was in Sweden ten years ago!

This mailing contains the HiPEAC vision for the next decennium. The HiPEAC network has been working on this consolidated vision for almost two years now, and we hope that this long-awaited document will inspire you. At the same time, our community has

been actively involved in the preparation of the FP7 computing systems call that will be launched later in 2010, with a deadline in spring 2011. The selected projects will tackle some of the computing system challenges of the next decennium.

For HiPEAC, 2009 was a good year. We organized several well-attended networking events, we noted a 50% increase in HiPEAC awards, and our community has been growing steadily. We have also been very active in submitting high-quality proposals to the computing systems call in April, and several of these proposals were selected, and are now starting as FP7 projects.

In this newsletter, you will also find the first announcement of ACACES 2010, one of the flagship events for our network. This year, we succeeded

again in hiring world-class instructors for the summer school. ACACES is attracting an ever-growing number of applicants from all over the world. If you have never attended ACACES, this is your opportunity to attend this exciting event.

Our biggest spring event will be the computing systems week in the first week of May 2010 hosted by our HiPEAC members at Edinburgh. The main theme of this event will be academic-industrial collaboration with numerous events by and for start-up companies. We expect a record number of attendees for this meeting. As you can read, we are very committed to making 2010 as exciting as 2009.

Happy 2010!

Koen De Bosschere

## Community News

# Mateo Valero, New Member of the Academia Europaea

The ceremony for the introduction of new members took place on Friday, 24th September, in Naples (Italy). Valero, who was the coordinator of HiPEAC 1, said, "this honour encourages me to continue working hard in Research and Science, both fields I have always been devoted to".

The Academia Europaea was founded in 1988. It's an organisation of eminent, individual scholars from across

the continent of Europe. Its members cover the full range of academic disciplines that comprises the humanities, social, physics and life sciences as well as mathematics, engineering and medicine. Currently there are over 2100 members.

Further information:  
<http://www.acadeuro.org/>



# Message from the Project Officer



In June 2009, the European Commission has organised a brainstorming workshop to identify success stories in EU projects funded in the Computing Systems research area and to analyse ways of creating and nurturing future success in this research area.

One of the success stories is clearly the HiPEAC NoE. The best way of measuring the success of HiPEAC has been described in the workshop: "before 2004, when you wanted to meet your European colleagues, you had to go to a conference somewhere in the US. Today, there is a vibrant community of hundreds of researchers in computing systems. Some years ago, when the EU tried to attract leading talent from abroad, the researcher had to land into a university department somewhere and start to build a research group. Now the researcher can land in the middle of an existing ecosystem."

Another success story has been the recently finished Milepost project that has performed exceptionally well. Success, in this case, can be measured by the fact that the Milepost project has created the world-first open source

machine learning compiler able to achieve 20% improvements compared to proprietary compilers.

Workshop participants highlighted the boundary between industry and academia as one of the biggest challenges in creating success in computing projects. If the overall success of the EU Computing Systems research is measured through its socio-economic impact, natural indicators of success could be, for example, the number of start-up firms and revenues from commercial licensing of project outputs. As the EU Framework projects typically have academic partners from countries where the boundaries between academia and industry are different, the problem of integration spreads across the consortium, and limits the possibilities of generating project success that could be measured in such socio-economic terms. This is in contrast with computing research in US universities, which thrive on deep integration of business and research. From this point of view, there is an important disconnect between the objectives of EU industrial and academic partners; in computing systems, where research

priorities are often determined through industry relevance, this disconnect is particularly challenging.

In many rapidly moving areas of computing, the locus of knowledge creation has over the years moved to industry research centres and start-up firms, and more recently, to networked innovation ecosystems. Academic research, therefore, needs access to industry to keep up to date with state-of-the-art ideas and technologies. Academic researchers today have strong incentives to collaborate with firms, especially in the area of computing systems. Improvements in managing the boundary between industry and academic partners in future EU projects in Computing Systems would have great potential in producing successful projects.

Workshop presentations and report at [http://cordis.europa.eu/fp7/ict/computing/events\\_en.html](http://cordis.europa.eu/fp7/ict/computing/events_en.html)

Panos Tsarchopoulos

## Community News

Guaranteeing the absence of spilling before instruction scheduling is a crucial issue in many situations of embedded systems and high performance computing. But the absence of spilling must not hurt the exploitation of instruction level parallelism. The two antagonistic constraints are studied since many years from the theoretical and the practical points of view (see our contributions on register saturation and schedule independent register allocation). As far as we know, we are delivering here the first free independent libraries on this

topic that handle multiple register types and delayed access times to registers. This general processor model allows the integration of our libraries in many compilers and tools. Our libraries must be plugged before instruction scheduling in order to analyze and, if needed, extend the data dependence graphs with additional arcs. We have demonstrated that this process does not hurt the quality of instruction scheduling because our register pressure analysis and optimization are sophisticated enough to model the constraints of

software pipelining and acyclic scheduling very well. The free software is available here:

- For register saturation computation (RSlib): <http://hal.archives-ouvertes.fr/inria-00431103>
- For bounding the register requirement in data dependence graphs (SIRAlib): <http://hal.archives-ouvertes.fr/inria-00436348>

Sebastien Brais  
and Sid-Ahmed-Ali Touati  
INRIA, France



## Release of RSlib and SIRAlib

# HiPEAC Computing Systems Week in Wrocław



Cluster meeting audience.



Zbigniew Chamski presenting activities at Infrasoft IT Solutions/Proximity Poland.

the event relied on a network of contacts of individual HiPEAC members across Europe, especially Prof. Rainer Leupers from RWTH Aachen. The preparation for the event and local arrangements were handled by the Wrocław Research Centre EIT+, led by Prof. Mirosław Miller, while the program of the Industrial Workshop was coordinated by Dr. Zbigniew Chamski (Infrasoft IT Solutions, Płock, Poland) with help from Prof. Wojciech Kabaciński (Poznań University of Technology).

The Industrial Workshop consisted of nine presentations, split into four main themes: infrastructure and support for HiPEAC-related R&D in Poland, ongoing industrial R&D in Poland, new multicore architectures and academic research in Poland.

The presentations from the Polish side illustrated the scope of ongoing and planned HiPEAC-related R&D activities. They also captured further key characteristics of the HiPEAC-related research and development community in Poland, where research is primarily driven by industrial applications and often funded directly by industry, implying direct and close co-operation between academic and industrial players. Secondly, leading-edge R&D activities in Poland are frequently driven by individuals who built up their expertise and networks while working in the industrial R&D centres in "old" member states of the EU or in the United

States, then returned to Poland to leverage the availability of well-educated and highly qualified graduates and set up projects with direct industrial exploitation paths. In the case of EIT+, this observation was made into an official strategy, determining the staffing policy and making it possible to handle a large portfolio of European projects.

The panel discussion that followed the Industrial Workshop focussed on the ways of strengthening the links between HiPEAC members and the Polish R&D community. The discussion identified three key enablers of better contacts: mobility, information flow and co-ordination, and contacts between HiPEAC NoE and the communities that in Poland are already well-structured - especially in the microelectronics area. The key instruments for supporting mobility are HiPEAC mobility funds (initially, through the associated member mechanism), Marie Curie fellowships and co-advised PhD theses. In the latter case, past experiences of panel participants were very mixed, but further investigations showed that the obstacles identified during the discussion could be successfully overcome.

Work on improving information flow and co-ordination between the HiPEAC NoE and the HiPEAC-related community in Poland has already started. Starting from direct contacts between Polish researchers and HiPEAC NoE members, it will focus on centralis-

The Autumn 2009 edition of the HiPEAC Computing Systems Week took place in Wrocław, Poland, between 26 and 28 October 2009. For the first time, the Computing Systems Week was held in a new member state of the EU, and attracted nearly 70 representatives of R&D centres from Europe, Israel and Turkey. The Polish academic and industrial community was well represented with 11 participants, affiliated with five academic centres and four industrial R&D centres.

The Wrocław edition of the Computing Systems Week was strongly influenced by the characteristics of Poland's high-performance and embedded computing community, which remains very dispersed - primarily because architecture and compilation tasks have not yet achieved the status of major R&D areas by themselves. In the absence of structuring mechanisms (formal networks, mailing lists, etc) and big, established industrial players, the organisation of



ing relevant information (HiPEAC announcements), making an "inventory" of teams and individuals directly involved in HiPEAC-related activities in Poland, and aim at creating "HiPEAC.pl", a virtual network of excellence across Poland.

The contacts between HiPEAC NoE and the Polish microelectronics community - which is very active, well structured and strongly integrated with the international community in this

domain - were already initiated during the HiPEAC Computing Systems Week. As the first tangible outcome, a dedicated HiPEAC session is planned at the 17th International Conference on Mixed Design of Integrated Circuits and Systems (MIXDES 2010), that will take place in Wrocław on 24-26 June 2010 (more at <http://www.mixdes.org>).

The venue of the Computing Systems Week (Hotel Sofitel Wrocław), located in Wrocław's historical Old Town,

provided very good infrastructure for the industrial workshop and cluster meetings (including flawless wireless networking), and a convenient environment for discussions and social contacts.

Last but not least, Polish gastronomy - both at the conference venue and in Old Town's restaurants - was greatly appreciated by all participants.

Zbigniew Chamski,  
Industrial workshop organizer

## HiPEAC Activity


# ProRISC 2009 in Veldhoven the Netherlands

The 20th edition of the annual two-day workshop on micro-electronic systems design has been successfully held in Veldhoven, the Netherlands on November 26 and 27.

Founded in 1989 by Philips Research Laboratories and the Dutch Technology Foundation-STW the workshop originally started with a focus on design methodology and currently covers research and development in the entire area between microelectronic devices and complex microelectronics based systems. The main goal of the ProRISC workshop is to bring together researchers from universities, research organizations and industry, and to serve as a platform for discussion of current trends in the fields of IC-technology and manufacturing. The workshop is held at the quite Koningshof hotel premises and its informal ambience makes it the ideal place for starting PhD students to take their first steps in presenting and discussing their research plans. The majority of the submitting authors present their work by means of poster sessions and flash presentations. This year ProRISC attracted 132 visitors from Germany, Belgium, Luxembourg, Bulgaria and the Netherlands. In addition to the traditional par-

ticipants from Philips Research, NXP Semiconductors and the three Dutch Technical Universities (Delft, Eindhoven and Twente), this year the workshop welcomed researchers from Saarland University, University of Technology Dortmund, RWTH Aachen University, Hamburg University, Ghent University, Leibniz University Hannover, University of Liege, KU Leuven, University of Lyon, Bulgarian Academy of Science, ASTRON, Holst Centre IMEC and IMS CHIPS Stuttgart. Keynote and invited speakers from Technische Universität Ilmenau, TUD, IBM Systems and Technology Group, Recore Systems BV, TUE, NXP Semiconductors, Eidgenössische Technische Hochschule Zurich and Avancis GmbH & Co. KG presented their recent developments and vision for the future of the different sub-fields covered by ProRISC. The talks of Dr. Ir. P.M. Heysters from Recore Systems BV "Research Reconsider, Resolve ... Recore" and Dr. H.P. Hofstee from IBM Systems and Technology group "From ASICs to Multicore and back again" were on important HiPEAC research agenda topics.

Since 1998 the ProRISC workshop has been co-located with the SAFE (Semiconductor Advances for Future



Poster session at the Workshop

Electronics) workshop. SAFE focuses on recent developments in semiconductor research, materials, fabrication technology and device technology. In 2010 the organizers of ProRISC and SAFE intend to expand the workshops research agenda with a third workshop aimed at Embedded Systems. The IEEE Benelux chapter on Embedded Systems showed interest in contributing to the organization of this event. From 2010 onwards, the combined workshops ProRISC, SAFE and Embedded Systems will be announced as STW.ICT congress.

Further information about ProRISC can be found at:  
[www.stw.nl/programmas/prorisc](http://www.stw.nl/programmas/prorisc)

Georgi N. Gaydadjiev,  
TU Delft

# EC FP7 STREP 2PARMA Project: PARallel PARadigms and Run-time MAnagement techniques for Many-core Architectures



## Project Coordinator:

Prof. Cristina Silvano  
Politecnico di Milano  
silvano@elet.polimi.it

## Project Technical Manager:

Prof. William Fornaciari  
Politecnico di Milano  
fornacia@elet.polimi.it

## Project website:

www.2parma.eu

## Partners:

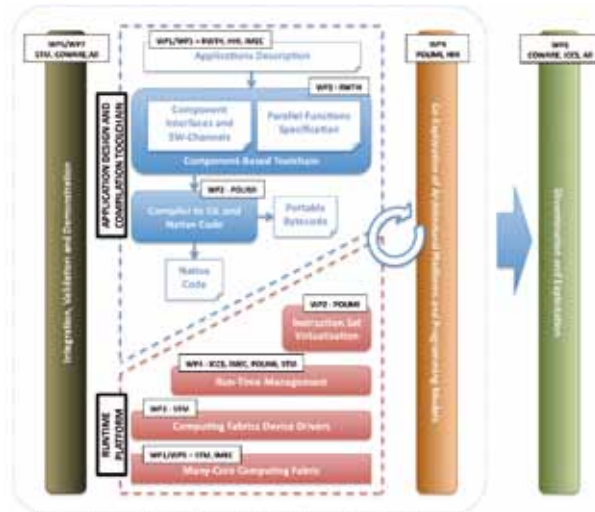
Politecnico di Milano (I)  
STMicroelectronics (I)  
Fraunhofer – HHI (D)  
IMEC (B), ICCS (G)  
RWTH Aachen Univ. (D)  
CoWare (B)

## Duration:

Jan. 2010 – Dec. 2012

## Main Objectives

The number of cores to be integrated in a single chip is expected to increase rapidly in the coming years, moving from multi-core to many-core architectures. This trend will require a global rethinking of software and hardware design approaches. This class of computing systems (Many-core Computing Fabric) promises to increase performance, scalability and flexibility if appropriate design and programming methodologies will be defined to exploit the high degree of parallelism exposed by the architecture. Other potential benefits of Many-core Computing Fabric include energy efficiency, improved silicon yield, and accounting for local process variations. To exploit these potential benefits,



Project structure and work flow

effective run-time power and resource management techniques are needed. With respect to conventional computing architectures, Many-core Computing Fabric offers some customisation capabilities to extend and/or configure at run-time the architectural template to address a variable workload.

The 2PARMA project aims at overcoming the lack of parallel programming models and run-time resource management techniques to exploit the features of many-core processor architectures. To this purpose, a proper Consortium has been set up to gather the required expertise in the areas of system/application software and computing architectures.

The 2PARMA project focuses on the definition of a parallel programming model combining component-based and single-instruction multiple-thread approaches, instruction set virtualisation based on portable bytecode, run-time resource management policies and mechanisms as well as design space exploration methodologies for Many-core Computing Fabrics.

## Technical Approach

The 2PARMA project will demonstrate methodologies, techniques and tools by using innovative hardware platforms provided and developed by partners, including the "Platform 2012", an early implementation of Many-core Computing Fabric provided by STMicroelectronics.

To ensure a wide range of application scenarios


comprising the typical computation-intensive workload of a general-purpose computing system, a set of industrial high performance demanding applications will be used and customized by using the techniques and methodologies developed in the 2PARMA project. Applications' architecture, development and integration will leverage in particular from the acknowledged experience of three partners from the Consortium: HHI for Scalable Video Coding application, RWTH for Cognitive Radio, and IMEC for Multi View.



## Project Coordinator

Prof. Cristina Silvano  
Politecnico di Milano, Italy  
silvano@elet.polimi.it

## EC FP7 MERASA Project: Multi-Core Execution of Hard Real-Time Applications Supporting Analyzability




The EC FP-7 MERASA project - now at the beginning of its third and last year - targets analyzable multi-core archi-

tectures, system software, and WCET (worst case execution time) analysis tools for execution of hard real-time applications. Most of the partners met in the HiPEAC NoE, namely partners of University of Augsburg, Barcelona Supercomputing Centre, University Paul Sabatier, of Toulouse, and Rapita Systems Ltd., expanded by Honeywell International of Czech Republic as an application company.

A higher performance than today's embedded processors can deliver will increase safety, comfort, services, and lower emissions of current and future automotive, aerospace, space and construction-machinery systems. However at the same time, in developing safety-related real-time embedded systems, there is a need to prove that the timing requirements are met. Multi-core processors achieve a high throughput by putting multiple cores on a single chip. However, mainstream multi-cores result in non-analyzable (or extremely pessimistic) worst-case timing behaviour that deems them unusable in the domain of safety-related real-time systems.

The MERASA project investigates analyzable embedded multi-core architectures with 2-16 cores in combination with hard real-time support for multi-cores in system software and WCET analysis tools targeting the combination of high performance features with time-predictable execution of single or multiple threads. The project addresses static WCET analysis tools (by the OTAWA toolset of Université Paul Sabatier, France) as well as hybrid measurement-based tools (by RapiTime of Rapita Systems Ltd., UK) and their interoperability.



Generic MERASA multi-core processor architecture with four cores

The general MERASA multi-core architecture (see Figure) is based on SMT cores and capable of mixed application execution of hard real-time and non real-time threads. The execution of one hard real-time thread per core is supported by isolation of threads or, where full isolation is not possible, by bounding of timing effects for the interferences of the threads. Inter-thread interferences appear when threads try to access the same shared resource at the same time, which concerns particularly the bus and memory system. The MERASA multi-core processor is designed in such a way that the inter-thread interferences are controlled, easing WCET analysis.

Each of the MERASA cores implements the Infineon TriCore ISA and consists of two different pipelines, an integer and an address pipeline, four thread slots (separate instruction windows and register sets per thread) able to accommodate one hard real-time and three non real-time threads, and an integrated real-time in-order issue/scheduling stage.

We defined an arbitrated real-time bus able to allow bounding of access times in case of interferences. The dynamically partitioned cache is an inter-core shared resource that avoids cache bank conflicts by assigning a private subset of cache banks to each hard real-time thread such that no other thread has access to it. Hard real-time threads access local instruction and data

scratchpads and a private cache partition, while non-hard real-time threads have access to the first-level instruction and data caches and a private cache partition that can be shared among all non-hard real-time threads.

The OTAWA tool set was adapted to the TriCore/MERASA core by a description of TriCore ISA in the modelling language nML that is part of the OTAWA tool, and by modelling the MERASA single-core based on parameterized execution graphs. The RapiTime tool instrumentation was integrated into the high-level and low-level MERASA simulators, which are able to catch timing traces of hard real-time thread execution on the MERASA multi-core. The generated traces are used by RapiTime to compute a WCET estimation. We also developed coding guidelines to support a WCET analysis and defined data formats used for the interoperability of the WCET tools based on a common object code reader.

The developed POSIX compliant system software supports isolation of hard and non real-time threads on a MERASA SMT core and of hard real-time threads running on different MERASA cores. All developed architectures and WCET tool adaptations were tested by running



benchmarks as well as by the Honeywell collision avoidance algorithm that runs on all simulators and WCET tools.

### Project Coordinator

Theo Ungerer, Department of Computer Science, University of Augsburg, 86159 Augsburg, Germany

### Project website

<http://www.merasa.org>

# ACACES 2010: Sixth International Summer School on Advanced Computer Architecture and Compilation for Embedded Systems

**11th–16th July, 2010,**

**La Mola, Barcelona, Spain**

We are proud to announce the sixth HiPEAC Summer School, which will again take place at La Mola, Barcelona during the second week of July. We start on a Sunday evening with an opening ceremony. The 12 courses start on Monday, spread over two morning and two afternoon slots. There are three parallel courses per slot, from which you can take one course. The courses have been allocated to slots in such a way that it will be possible to create a summer school program that matches your research interests. The following world-class experts will present the topics of this year's Summer School.

On Sunday evening there will be a keynote talk. On Wednesday afternoon, participants are given the opportunity to present their own work to other participants during a huge poster session; and finally, on Friday evening there will be a farewell dinner and party. Students and lecturers will be accommodated in rooms on campus, where they will stay for one week. This will

Instructor	Affiliation	Title
David Brooks Derek Chiou	Harvard University University of Texas at Austin	Variation-Aware Processor Design Fast and Accurate Computer System Simulators
Steven Hand Andreas Herkersdorf	Citrix TU Muenchen	System Virtualization Application-Specific (MP)SoC Architectures for Internet Networking
Mahmut Kandemir Scott Mahlke Vivek Sarkar	Pennsylvania State University University of Michigan Rice University	Embedded Systems: A Software Perspective Compilation for Multicore Processors Multicore Programming Models and their Compilation Challenges
Donatella Sciuto Michael Scott Dan Sorin Per Stenström and Andrzej Brud Theodore Ts'o	Politecnico di Milano University of Rochester Duke University Chalmers The Linux Foundation	FPGA-based reconfigurable computing Transactional Memory Fault Tolerant Computer Architecture How to transform research results into a business File Systems and Storage Technologies

provide plenty of opportunities to have discussions with teachers and other participants in the relaxing surroundings of La Mola. At the end of the event, all participants will receive a certificate of attendance detailing the courses they took.

Unfortunately, the number of participants will be limited. Therefore, we have an admission procedure in place to guarantee a fair distribution of available places among all qualified applicants from the various countries and institu-

tions. If you are a studentmember of a HiPEAC institution you can ask for a grant that covers the registration fee. In this newsletter, you will find a summer school poster. Please hang it at some visible place in your department. You can find more information about the summer school at <http://www.hipeac.net/summerschool>. We look forward to seeing you there!

Koen De Bosschere  
Summer school organizer

## Community News



# Recore Systems Receives US\$ 3 Million Funding

**The company aims to enter the market for digital radio/TV receiver chips and to expand global sales and marketing.**

Recore Systems, a fabless semiconductor company specialized in reconfigurable multi-core processors, received US\$ 3 million funding in September

2009. Point-One Innovation Fund and the East Netherlands Participation Company led the investment round. The new funds will be used to expand the sales, marketing and customer support organization. Moreover, the company will release a chip for receiving digital radio and TV. This receiver chip targets consumer electronics for

European and Asian markets, such as (car) radio systems, portable media players and navigation devices. Its features include reception and playback of DMB, DAB+ and DAB broadcasts. The underlying reconfigurable technology allows the receiver to be used in every region of the world and enables adapting to new or unfore-



# Professor Andreas Herkersdorf, Technische Universität München, Germany



Andreas Herkersdorf is Professor and Director of the Institute for Integrated Systems in the Department of Electrical Engineering and Information Technology at Technische Universität München (TUM), Germany. He received a Dipl.-Ing. degree in electrical engineering from TUM in 1987, and a Dr. degree, also in electrical engineering, from the Swiss Federal Institute of Technology (ETH), Zurich, in 1991.

Between 1988 and 2003 he has been with the IBM Zurich Research Laboratory, Rüschlikon, Switzerland. In technical and management positions he contributed to the design and development of advanced VLSI architectures for high-speed wire line data transmission and networking systems, such as SONET/SDH framers and add/drop multiplexers, ATM cell switches and network processors.

In September 2003 he became head of the Institute of Integrated Systems at TUM. Dr. Herkersdorf is inventor / co-inventor of 14 international patent applications and member of the editorial boards of Design Automation for Embedded Systems Journal

(DAEM), Springer, and Electronics and Communications International Journal, Elsevier. He serves as European representative in the executive committee of ICCAD. His research interests include reconfigurable multiprocessor VLSI architectures for IP networking and automotive applications, system level SoC modelling and design space exploration methods, and self-adaptive fault-tolerant computing.

## Institute for Integrated Systems

The Institute for Integrated Systems at Technische Universität München investigates and develops novel VLSI architectures, algorithms and innovative IP (intellectual property) building blocks for primarily targeting the following application domains:

- Embedded systems in automotive,
- Internet / ATM networking,
- Autonomic, self-organizing System on Chip.

Research focal points are the systematic investiga-



tion of homogeneous and heterogeneous multi-processor / many-core SoC (MPSoC) architectures consisting of standard RISC CPUs, application specific processors (ASIP) and hardware accelerators, multi-topology Network on Chip (NoC) interconnects, memory hierarchies and the efficient interaction with the system environment via standard I/O.

Established skills in the following areas form the methodological basis of his contributions:

- Modelling, design space exploration and simulation techniques at high levels of abstraction,
- Reconfigurable computing,
- Visual computing and scene analysis,
- TCP / IP networking and protocol processing,
- Bio-inspired principles of self-organization,
- Fault-tolerant, robust system design techniques, and
- Dynamic power and energy management.

## Contact information:

Email: [herkersdorf@tum.de](mailto:herkersdorf@tum.de)

Website: <http://www.lis.ei.tum.de/>

seen developments. Integrating the silicon tuner, baseband processing and media decoding in a single chip makes the solution cost competitive and easy to integrate in end user applications.

"Raising venture capital in the semiconductor industry in these challenging times is a great confirmation of the excellence of our technology and people", stated Recore Systems' CEO Paul Heysters. "During the past years, we strengthened our IP portfolio and

are now ready to enter the market with a chip of our own. We offer a future-proof digital radio/TV receiver chip that can be fine-tuned in software to the requirements of a specific region. Entering this market requires expansion of our business development team, which will also foster our IP sales. We are therefore very excited to have Point-One Innovation Fund and the East Netherlands Participation Company as part of our investor base."

Recore Systems' products enable highly efficient multi-core systems for applications such as broadcasting, multimedia, wireless (tele)communication and digital beamforming. These fully programmable systems reduce time-to-market and add flexibility through software upgrades.

## Contact:

Paul Heysters

[paul.heysters@recoresystems.com](mailto:paul.heysters@recoresystems.com)

Website: <http://www.recoresystems.com>

## Collaboration Grant Report - Marco D. Santambrogio



### Self-aware and autonomic systems

My name is Marco D. Santambrogio and I completed my PhD at Politecnico di Milano working with Prof. Sciuto in the DRESD research group in 2008. At the end of the same year I received a grant for a postdoc fellowship from the Progetto Roberto Rocca. My postdoc activities at MIT focus on self-aware systems and I'm closely working with Prof. Anant Agarwal and the member of the CARBON group, especially with Jonathan Eastep and Henry Hoffmann.

Self-aware computer systems will be

capable of adapting their behavior and resources thousands of times a second. This is done to automatically find the best way to accomplish a given goal despite changing environmental conditions and demands. This capability would benefit the full range of computer systems, from embedded devices to servers to supercomputers. Scenarios where self-awareness will be particularly useful include: mobile technologies, cloud computing systems, adaptive and dynamic compilation, multicore microarchitecture, and novel operating systems. We believe that semiconductor technology, computer architecture, and software systems have advanced to the point that the time is ripe to realize such a system. We are proposing not only a novel programming paradigm or architecture, but a completely new

way of thinking about and approaching computer systems that reflects 21st century constraints and opportunities.

Some of the main challenges in realizing such a vision are: to add adaptability capabilities to devices, in order to implement distributed self-training algorithms over such architectures, and to formulate application solutions using such a computing paradigm. The problem with existing approaches to adaptive systems is that they are largely ad hoc and often fail to incorporate the true goals (performance or otherwise) of the applications they are designed to support. The goal of this work is to present enabling technologies for adaptive computing that address these challenges. Examples of such technologies are the Application

## Member Profile

## Professor Stefano Crespi Reghizzi and Formal Language and Compiler Group of Politecnico di Milano, Italy



Stefano Crespi graduated at Politecnico di Milano, then obtained a Ph.D. from UCLA, and

served as professor at Università di Pisa, before joining the Dept. of Elettronica e Informazione (DEI) of Politecnico di Milano as full professor. He is the senior member of a research group (<http://compilergroup.elet.polimi.it>) that, as the title says, does research on compilation as well as formal languages and automata theory. The group performs theoretical research and R&D. Examples of theoretical research are picture grammars and tiling systems for 2D languages, and formal languages for model checking. Such activities

are part of the European Science Foundation Initiative AutoMathA, Automata Theory from Mathematics to Applications.

The following research topics are more central to HiPEAC: dynamic compilation, parallel programming paradigms, and automatic parallelization.

**Dynamic compilation:** the team lead by S. Campanoni, with partial support from STMicroelectronics, develops and exploits the increasingly popular ILDJIT compiler for ECMA 335 CIL byte-code. This is a parallel compiler running on multi-core or multi-processor machines (with X86 and ARM processors) and supporting C# and C

(and Java in the near future).

It comes from the past experience on dynamic Java compilation for VLIW machines.

Some current investigations are: code specialization (in collaboration with INRIA-IRISA for the GCC4NET frontend), reflection and generic classes, and dynamic look-ahead compilation for reducing JIT latencies.

ILDJIT is the dynamic compiler used in the European project OMP Open Multimedia Platform. It is also used in research on the prevention of voltage emergencies in CPUs by David Brooks of Harvard University.

**Parallel Programming Paradigms:** the team lead by G. Agosta is investigating current languages for paral-

Heartbeats (or just Heartbeats) and Smartlocks.

The Application Heartbeats framework provides a simple, standardized way for applications to monitor their performance and make that information available to external observers. The framework allows programmers to express their application's goals and the progress the application is making using a simple API. As shown in Figure 1, this progress can then be observed by either the application itself or an




Figure 1 - (a) Self-optimizing application using the Application Heartbeats framework. (b) Optimization of machine parameters by an external observer.

external system (such as the OS or another application) so that the application or system can be adapted to make sure the goals are met.

Within this context, a second enabling technology for adaptive computing systems has been implemented. We designed an open-source self-aware synchronization library for multicores and asymmetric multicores called Smartlocks. Smartlocks is a spin-lock library that adapts its internal implementation during execution using heuristics and machine learning. Smartlocks optimizes toward a user-defined goal, programmed using the Application Heartbeats framework (see Figure 2) which may relate to performance, power, problem-specific criteria, or combinations thereof.

Smartlocks takes a different approach to adaptation than its closest predecessor, the reactive lock. Reactive locks

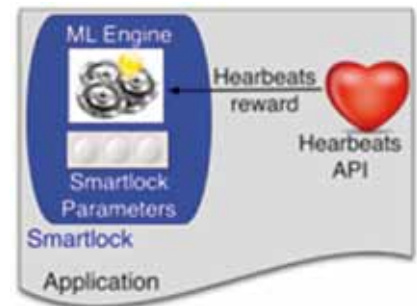


Figure 2 - Application-Smartlocks Interaction. Smartlock ML engine tunes Smartlock to maximize Heartbeat reward signal which encodes application's goals.

optimize performance by adapting to scale, i.e. selecting which lock algorithm to use based on how much lock contention there is. Smartlocks use this technique, but use an additional novel adaptation - designed explicitly for asymmetric multicores - that we term lock acquisition scheduling. When multiple threads or processes are spinning, lock acquisition scheduling is the procedure for determining who should get the lock next, to maximize long-term benefit.

lel data-intensive computation (such as Cuda and OpenCL) in order to increase compiler support for automatic deployment of data on a distributed and shared memory architecture. The aim is to increase the portability of parallel application code across different target platforms. The work is part of the European project 2PARMA Parallel Programming and Run-time Management Techniques for Many-core Architectures.

**Automatic Parallelization.** This is a case of synergy between automata theory with emphasis on the potential of the theory of traces (or partially commutative languages) for modeling control/data dependences and parallelizing code transformations.

Two research lines are active: the theoretical study of dependences and schedules in nested loops; and, more advanced towards experimentation, the Hydra project for auto-parallelizing control-intensive CIL programs, exploiting the notion of Control-components.

Another research problem is how to parallelize syntax and semantic analysis (parsing).

#### Teaching activities:

- The course Formal Languages and Compilation is based on the recent book (<http://www.springer.com/computer/foundations/book/978-1-84882-049-4>) offering a unified methodological presentation of classical methods used in front ends.

- At ALaRI (<http://www.alari.ch/>) master of Sc. in Embedded Systems Design Prof. Crespi and his group teach the course on Software Compilers.



#### Contact information:

Email: [stefano.crespireghizzi@polimi.it](mailto:stefano.crespireghizzi@polimi.it)

Website:

<http://www.elet.polimi.it/people/crespi>

## CAPS is speeding up its expansion on the international scene of hybrid programming



**CAPS enterprise takes a key step forward in its development by announcing major evolutions of its HMPP™ hybrid compiler (Heterogeneous Multicore Parallel Programming) that reinforce its position as a technology leader on the market of programming tools for hybrid systems and by developing its global reseller network.**

### **A leading product in directive-based compiling technology**

Alongside the high growth market of hybrid systems (i.e. associating GPUs to micro-processors) that have been deployed in 2009 on all hardware ranges, from the workstation to teraflop machines, CAPS announces the release of its two new AMD CAL/

IL and OpenCL back-ends as well as Windows support. CAPS thus broadens the potential of new HMPP users: besides Linux x86\_64, HMPP opens to a wider community of developers with its future Windows x86\_64 version under Visual Studio.

The CAL/IL backend is now available since early December and a release that supports OpenCL and Windows is planned for 2010 Q1.

This next HMPP release will also integrate new functionality improving its ease of use and the performance of the accelerated code it generates.

HMPP enables users to easily and rapidly port existing applications onto hybrid systems while maintaining a unique C or Fortran source code annotated with directives. From this annotated source code a hybrid application

for machines integrating NVIDIA or AMD/ATI GPUs is generated automatically. Its users can therefore take the most out of GPU accelerators performance without having to sacrifice their applications portability, thus protecting their software assets.

With this new CAL/IL back-end, HMPP now supports the most widely used hardware accelerating technologies on the market: AMD FireStream and NVIDIA Tesla. HMPP CAL/IL brings a real added value by allowing developers to quickly leverage the tremendous power of AMD GPU processors for accelerating their scientific applications; the integration of GPUs in clusters being indeed one of the major HPC trends confirmed during SuperComputing 2009 in Portland.

## Community News

### HiPEAC Technical Reports Instrument Expanded

The HiPEAC Technical Reports were recently enriched by a new time stamping service. This service offers HiPEAC members the possibility to safely record the timing of new ideas long time ahead of their publication even when only very preliminary results are available. In addition to certified timing, this service is intended to promote sharing of early ideas that will advance the HiPEAC research field. The HiPEAC steering committee made sure to create a very low overhead service easily available to all HiPEAC members and suitable for any

electronic document of interest. The submitted documents are not subject of any review that places the responsibility for the technical content, its originality and soundness right where it belongs: in the hands of the document author. Needless to say that the very same is true for the copyrights of the submitted manuscripts. Another way of seeing this new service can be as a platform for HiPEAC wide dissemination of technical reports.

Submission in "HiPEAC Technical Reports" does not preclude subsequent publication in conferences or

journals. In contrary we believe that the HiPEAC time stamped reports will eventually result in one or more high-quality scientific publications when all of the experimental work is performed and the analysis is completed. This confidence is due to the fact that the focus of this service is early time stamping of novel proposals. In cases when the work is already mature enough it can better be submitted as a normal paper at any venue from the HiPEAC research domain.

After the submission of each time stamped document the service will



Excellent performances generated by HMPP during demonstrations on this tradeshow have attracted great interest in the HPC community where the need for high level programming software has been re-asserted more than ever.

### **An international sales expansion strategy**

CAPS counts among its customers major European actors in the fields of energy, oil & gas, defence and research, using HMPP for key application deployments like its latest collaboration with GENCI, aiming at developing hybrid computing in France.

In 2009 CAPS initiated its international deployment by signing several reseller agreements in the United States

(with ParaTools, Inc.), in Japan (with JCC-Gimmick) and in Taiwan (with ARAvision).

With 30 employees this year, CAPS shows an increase by two of its revenue in 2009 and reveals great ambitions for 2010, namely to become the global leading provider of development tools for hybrid systems.

#### **CAPS Contact:**

Estelle Dulsou

[estelle.dulsou@caps-entreprise.com](mailto:estelle.dulsou@caps-entreprise.com)

Website: [www.caps-entreprise.com](http://www.caps-entreprise.com)



CAPS booth at SC09 in Portland

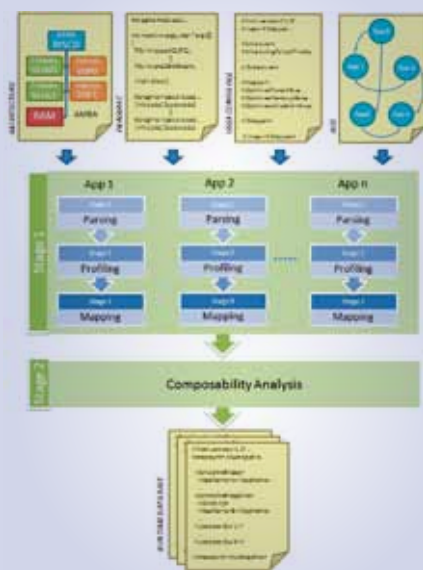
generate a unique identification number, e.g., TR-HiPEAC-00123 that can be used for future references. During the submission process, the author is asked to indicate his/hers preferences on making the content publically available or not. The following two options are available: 1) providing everyone with read-only access or 2) granting only the document author a read-only access. In the latter case, the author is asked to provide a short abstract summarizing the document content. In both cases, the time stamped documents cannot be modi-

fied by anybody including the original author and the service administrator. Withdrawing of already time stamped documents can be requested via the HiPEAC web master or the HiPEAC Technical Reports contact person. Even the best initiative cannot create impact without an active user community. This is why we would like to emphasize on the advantages this new HiPEAC automated service offers. Do you have the brightest idea you would like to time stamp? Have you thought of provoking an interesting discussion by proposing

something out of the box? Would you like to expand the outreach of your technical reports to the whole HiPEAC research community? If the answer to at least one of the above questions is "YES" the HiPEAC Technical reports is the right tool for you. Let us all make this unique service one of the many other successes the HiPEAC community created during the past years. More information can be found at: [http://www.hipeac.net/tech\\_reports](http://www.hipeac.net/tech_reports)

Georgi N. Gaydadjiev,  
TU Delft

# Collaboration Grant Report - Ricardo Velasquez



Multiple KPN Applications Flow

I am a student at Università della Svizzera Italiana (USI) in the Advanced Learning and Research Institute (ALaRI). Last year I received a HiPEAC Collaboration Grant that gave me the opportunity to work with Jeronimo Castrillon and Prof. Rainer Leupers at SSS group at RWTH University, on extensions for the MPSoC Application Programming Studio (MAPS) framework.

Parallel compilers are one of the main trends in parallel programming. Addressing this approach, the MAPS framework has proven to be excellent at extracting parallelism from sequential C programs. The performance results obtained are encouraging, but still lay slightly behind the ones obtained when the parallelism is specified explicitly. To overcome the difficulty of (semi-) automatic parallelism extraction, parallel programming models need to be supported as well. In this work support for parallel programming models is enabled by using the Khan Process Network (KPN) Model of Computation (MoC).

Multiple applications with different performance requirements and

resource utilization might run in parallel. This leads to a large number of possible use-cases that the system must be able to run. The number of use-cases grows exponentially with the number of applications, and verifying the correct operation of all variants becomes impractical. Moreover, applications might have soft/hard real time constraints, and thus further increase the complexity of the verification process. Proposed solutions to this problem include composability, virtualization, and some exact approaches. This work follows the composability approach.

A Pragma C language is used as input description for the applications. This language adds extensions to the C language in form of pragmas in order to support explicit parallelism. The programming model used within this work is based on the KPN MoC, and therefore the provided pragmas describe processes and channels. The pragmas also capture real-time constraints and mapping preferences. Furthermore, custom data type extensions have been introduced in order to simplify channel operations.


In this context, the MAPS framework has been extended with two design flows. The first denominated Single KPN Application Flow (SKAF) is used to analyze and synthesize applications in isolation. The flow is composed of three stages: Parsing, Profiling and Mapping. The parsing stage takes the pragma C description, and turns it into an intermediate representation graph (KPNGraph). The KPNGraph stores all specifications provided by the pragma C description. During the profiling stage, time traces of events in every channel are generated. The mapping stage uses these traces to reconstruct the KPN behaviour for different mappings and therefore obtain performance measures to feed mapping algorithms.



The second flow is called Multiple KPN Applications Flow (MKAF). MKAF is composed of two stages. During the first stage MKAF uses the single flow (SKAF) to obtain mappings with different properties, for each one of the applications in the system. In this way every application has a set of mappings fulfilling application requirements but with diverse resource utilization patterns. Then in the second stage, MKAF analyzes every feasible use-case. During this process prospect mappings per application are selected, then a composition function determines if the combination of mappings can run simultaneously on the platform fulfilling application constraints. If that is the case, the process continues with the next use-case, otherwise a new selection of mappings is made and the composition function is executed again.

Finally, in order to test the flows, a case study with three applications was designed: JPEG, GSM and MPEG-2. The case study shows how the tool can ease the design effort of parallel applications, produce good scheduling and mapping configurations and perform composability analysis. The results of this work will appear in DATE 2010 in the area of Compilers and Code Generation for Embedded Systems in the paper entitled "Trace-based KPN Composability Analysis for Mapping Simultaneous Applications to MPSoC Platforms".

# Collaboration Grant Report - Francesco Paterna



I am a PhD student in Electronics, Computer Science and Telecommunications at the University of Bologna. My advisor is Prof. Luca Benini.

My research activity focuses on Variability and Reliability for MPSoC, and in particular on adaptive software techniques for performance improvements, reduction of energy consumption, and lifetime requirements for MPSoC using sub-50nm CMOS technology.

Thanks to HiPEAC I have spent three months at the STMicroelectronics site of Cornaredo (Milan) from September to December 2009, working with Giuseppe Desoli and Francesco Papariello on the xStream multi-processor platform, composed of a ST231 host core, and a regular array of xPE processors. Each processor has its own distributed but uniform memory address space.

I started my employment at STMicroelectronics by working on a MPEG2 decoder originally written for the ST231 multi-threaded processor. This program was designed to run on 1, 2, or 4 threads. The aim of the effort was to transform the code into a realistic benchmark for a class of parallel multimedia codec's suitable of being deployed on massively parallel-embedded multiprocessor arrays.

The task graph is composed of four parts: a scan of the current frame, a slice decoding, an inverse discrete cosine transform (IDCT), and the commit of results.

I restructured the application so that the host core performs the scan of the current frame and the commit of the previous results, while the slice decoding and the IDCT, which are executed on the xPEs, have been divided in independent tasks whose number can be equal or greater than the number of xPEs. Regarding the latter case, a dispatcher has been implemented on the host core to schedule the different tasks on the xPEs.

To increase performance I further modified the code to execute the commit of the previous frame during the execution of the current frame on the xPEs.

Once the parallel benchmark was ready for prime time I focused the rest of my effort to demonstrate the effectiveness of heuristic task mapping policies in the presence of process variability.

To meet the deadline mandated by the frame rate requirements and to minimize energy consumption, the host processor, during frame decoding, executes the LP+BP policy to allocate tasks for the next frame. The LP+BP policy, developed in my research activity, is low overhead and solves a Linear Programming problem followed by a Bin Packing problem.

I used the xStream Instruction Set Simulator (ISS) and the energy consumption figures estimated from RTL synthesis of xPEs. For variability emulation I used the heuristic characterization provided by IMEC that is applied to a similar datapath within the context of the REALITY FP7 project.



In order to drive the task allocation policies and to profile the application, I implemented a plug-in for the xStream ISS that estimates the energy consumption and clock frequencies of each core (based on the longest path delay predicted by a variability emulation model); the plug-in also monitors the cycle counter and models temperature and aging effects.

Experimental results of LP+BP policy have been presented at ESTImedia 2009 in a work titled "Variability-tolerant Workload Allocation for MPSoC Energy Minimization under Real-time Constraints" (Francesco Paterna, Andrea Acquaviva, Francesco Papariello, Giuseppe Desoli and Luca Benini).

For a subset of real-time image processing workloads, another piece of work that proposes an on-line implementation of LP+BP has been submitted to Computing Frontiers 2010.

The results obtained by the application of methodologies and policies described in the parallel MPEG2 decoder application above will be submitted soon to an International Journal.

# Christophe Dubach Received the BCS Distinguished Dissertation Award 2009

Christophe Dubach, an RAEng/EPSC Research Fellow at ICSA, University of Edinburgh, has received the British Computer Society (BCS) Distinguished Dissertation Award 2009 for his dissertation entitled: "Using Machine-Learning to Efficiently Explore the Architecture/Compiler Co-Design Space". This annual award selects the best British PhD dissertations in computer science and has recognized this year the application of machine learning (AI) in the area of computing systems as important and significant.

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 optimizations are also considered. Once the architecture is selected, a new compiler must be developed and tuned.

Christophe's thesis proposes the use of machine-learning to address architecture/compiler co-design. The techniques developed in his thesis represent a new methodology that has the potential to speed up the design of new processors and automate the generation of the corresponding opti-

mizing compilers, resulting in higher system efficiency and shorter time-to-market.

The thesis is available on Christophe's homepage: <http://homepages.inf.ed.ac.uk/s0567037/>



Christophe Dubach at Award Ceremony, Royal Society, London (BCS, the Chartered Institute for IT)

# HiPEAC Member Named ACM Distinguished Member


The Association for Computing Machinery "recognizes those ACM members with at least 15 years of professional experience and 5 years of continuous Professional Membership who have achieved significant accomplishments or have made a significant

impact on the computing field" with the Distinguished Member Grade. For 2009, Prof. Stefanos Kaxiras, a HiPEAC member, was one of the recipients of the Distinguished Scientist award. Prof. Kaxiras was recognized for his contributions in computer architec-

ture, specifically for contributions in power-efficient computer architecture and in the design of memory systems. This recognition also reflects on the world-class research community and work of the HiPEAC network of excellence.



# Collaboration Grant Report – Darío Suárez Gracia



The cache tile layout

My name is Darío Suárez Gracia. I am a PhD student at the University of Zaragoza, Spain under the supervision of Víctor Viñals Yúfera and Teresa Monreal Arnal from the University of Zaragoza and the Universitat Politècnica de Catalunya (UPC), respectively.

As technology scales, processors include more cache levels opening a new latency gap between the small and fast first-level caches and the large and slow last-level caches. Our work tries to close this gap by merging the first cache levels (L1 and L2) into a single tiled structure; trading latency and size at a fine granularity (1 cycle). In order to provide such granularity, our proposed Light NUCA (L-NUCA) integrates one cache access (4-32 KBytes) and one hop routing in a single cycle.

Performing a cache access and one-hop routing in a single processor cycle is the critical hypothesis of L-NUCAs. A detailed analysis of all involved circuits within the tiles is required in order to verify this timing assumption. One of the best approaches for carrying on this analysis is the development of a VLSI implementation in a semi-custom ASIC library. Since we had neither the expertise nor the tools for performing the implementation, we applied for a HiPEAC collaboration grant for

visiting Giorgos Dimitrakopoulos and Manolis Katevenis at FORTH (Greece), because they are experts on both VLSI implementation and interconnection networks as used in L-NUCAs.

The collaboration spanned from July to October 2009. We started modeling the cache and then the router in Verilog. Later on, in order to validate the Verilog model, we compared its outputs with the outputs of an existing microarchitectural simulator running real memory traces from SPEC CPU2006. Once the model was validated, we continued with the synthesis of code with Synopsys tools using a 90 nm library. This step gave us positive feedback about the feasibility of integrating single cycle cache access and one-hop routing. The fast channel allocator, the avoidance of virtual channels, and the time required to generate signals for updating the valid, dirty, and LRU bits allow performing the routing in parallel with cache updating tasks almost without affecting the cycle time.

After synthesis, we continued with placement and routing using Cadence SOC Encounter. With the placed-and-routed model it is possible to obtain very accurate estimations of the delay, area, and energy of L-NUCA tiles in order to validate our hypothesis.

Apart from the implementation of the L-NUCA tiles, we discussed how to improve the performance of L-NUCAs, finding several interesting unexplored paths. We hope to publish the implementation and the rest of the results of this internship soon, so please stay tuned.

Summarizing, it has been a pleasure having this opportunity to learn about digital design at FORTH with the help of HiPEAC funds. I encourage all HiPEAC PhD students to apply for



these grants because they are a great experience from both the personal and professional point of view. Finally, I would like to acknowledge all the people at the CARV group for making my stay there so enjoyable.

# Collaboration Grant Report – Konrad Trifunovic



My name is Konrad Trifunovic. I am a PhD student at INRIA Saclay -Ile-de-France. I started my PhD at the end of 2007, in the ALCHEMY team, under supervision of Dr. Albert Cohen. My work is on automatic parallelization and compiler optimizations using a polyhedral model.

More specifically, I am trying to contribute to the automatic parallelization of scientific and media-processing kernels, as well as other compute intensive applications (found in SPEC2006 for example). I am mainly interested in the parallelization of loop nests, since the majority of runtime is spent inside loops.

For the purpose of parallelization we are using the polyhedral model. The polyhedral model enables us to reason

about a well-defined class of programs, to perform loop transformations and eventually parallelize the loops. So far, polyhedral techniques were only confined to research compilers. Together with our partners, we are actively developing GRAPHITE, a polyhedral model framework inside GCC (<http://gcc.gnu.org/wiki/Graphite>). It is now part of GCC 4.5. Using GRAPHITE we are able to perform powerful loop transformations to optimize for memory hierarchies, enabling loop parallelization and vectorization, just to name a few.

Thanks to a HiPEAC internship grant, I was able to join IBM Research Haifa in Israel for three months in 2008. I worked together with Dr. Ayal Zaks, Dorit Nuzman and Razya Ladelsky. I worked on several topics, one of them being cost-modelling and transformations for automatic loop vectorization. This work resulted in a published paper

at the PACT'09 conference, named "Polyhedral-Model Guided Loop-Nest Auto-Vectorization". I also had the chance to work on automatic parallelization inside the GCC compiler. The idea is to find loops where independent iterations could be split among different threads. This work has not been completed during my stay at IBM, but is now ongoing collaboration between our labs. Currently we are considering coupling automatic parallelization with streaming, in order to enable the parallelization in the case where synchronization between threads is needed.

I would like to thank to HiPEAC for providing this great opportunity. It allowed me to see how research is conducted in industry, to meet great people and to establish long-term collaborations. I can only recommend to take advantage of HiPEAC sponsored internships and collaborations to all HiPEAC students.

## PhD News

### Dynamic Memory Optimizations for Embedded Systems Using Software Metadata

**By Alexandros Bartzas**  
([ampartza@ee.duth.gr](mailto:ampartza@ee.duth.gr))

**Advisor: Prof. Dimitrios Soudris**  
**(currently at National Technical University of Athens, Greece)**  
**Democritus University of Thrace, Greece**  
**November 2009**

Emerging embedded systems offer rich services. Applications are composed of multiple threads and rely on the usage of dynamic data to adapt to user and environment constraints. Embedded platforms will use a multitude of processor and memory modules. Data storage, dynamic allocation, de-allocation and accesses greatly affect the performance and energy consumption of the system. Furthermore, designers have to efficiently design the on-chip interconnection among the components (taking

into consideration the advances from emerging 3-D process and integration technologies). The developed methodologies and tools target the efficient utilization of available memory resources, when the system executes dynamic multi-threaded applications.

The first step is to characterize the behaviour of software applications with special emphasis on dynamic data usage. Towards this, we developed a software metadata structure, representing such behaviour. The developed tools for profiling and analysis allow for easy construction of software metadata. Additionally, we developed a new methodology for dynamic data block transfers extraction. The most important advantage of software metadata is that it enables the unification of separate data management methodologies into a single design flow.

Another aspect of their usage is that when combined with vertical interconnection patterns it allows one to explore 3-D NoC architectures. To this end, we developed a high-level NoC simulator, triggered by synthetic and real-application inputs. Additionally, we have linked the simulator to a physical prototyping EDA tool, being able to get accurate estimations on wire-lengths and energy consumption of the on-chip interconnects from early design steps. The results presented in this thesis offer solutions to: a) characterizing the behaviour of dynamic multi-threaded applications; b) dynamic data management and c) exploration of alternative 3-D NoC architectures. Results of my PhD work were published in: 1 book, 1 book chapter, 5 journal papers, 12 conference papers and 2 PhD Forums.

## Efficient and Scalable Cache Coherence for Many-Core Chip Multiprocessors

**By Alberto Ros**

**(a.ros@ditec.um.es)**

**Advisor: Manuel Eugenio Acacio**

**and José Manuel García**

**Universidad de Murcia, Spain**

**September 2009**

The increasing number of cores integrated on a single chip prevents the popular snooping-based protocols from being used to keep cache coherence in future many-core CMPs. On the other hand, although directory-based protocols constitute the best alternative in these large-scale systems, they have two important issues that restrict their scalability: the directory memory overhead and the long cache miss latencies caused by the access to directory information (indirection problem). Additionally, since both data and direc-

tory caches are commonly distributed across the chip, the access latency to these structures can become another performance issue for large-scale systems. Our efforts in this thesis have focused on these key issues.

First, we present a scalable distributed directory organization that stores coherence information as duplicate tags and also uses fine-grained interleaving for distributing directory banks. This organization requires less area than a traditional directory to keep the same information, and its memory overhead does not increase with the number of cores.

Second, we propose direct coherence protocols. These protocols are aimed at avoiding the indirection problem of

directory-based protocols, but without relying on broadcasting requests. The key property of these protocols is the assignment of the task of keeping cache coherence to the cache that provides the data block in case a cache miss occurs. Indirection is avoided by directly sending requests to that cache.

Finally, we develop a novel mapping policy managed by the OS that reduces the long access latency to a distributed cache. It tries to map memory pages to the local cache bank of the first core that requests them, but it also introduces an upper bound on the deviation of the distribution of memory pages among cache banks. In this way, we reduce the average cache access latency and the number of off-chip accesses.

## Improvement of Variable Latency Algorithms for the Calculation of Division, Square Root and their Reciprocals

**By Daniel Piso Fernández**

**(daniel.piso@gmail.com)**

**Advisor: Javier Diaz Bruguera**

**University of Santiago de**

**Compostela**

**November 2009**

The improvement of computer applications throughout the past years has further increased the demand for high-performance floating-point computation. Especially division and square root operations are present in several problems. It has been proved that improvements in these functions have a noticeable impact on processor performance.

Multiplicative algorithms are one of the preferred solutions to calculate division and square root. However, these algorithms produce neither a correctly rounded result nor the remainder. The remainder has to be calculated once the result is obtained. Hence remain-

der calculation overheads the total calculation time of the algorithm. This work proposes a set of variable latency methods devoted to the improvement of the rounding of multiplicative algorithm results.

The first method determines the optimal size for intermediate operations of multiplicative algorithms when a certain result precision is required. It is based on an accurate error analysis that takes into account the errors introduced by hardware computation of arithmetic operations. This new capability is useful for designing the rest of the proposed method.

The rest of the proposed techniques are variable latency rounding methods. The first of them is a modification of the classical rounding method for multiplicative algorithms. This method reduces the number of cases where the

remainder has to be calculated to do the rounding by introducing additional information in the rounding table.

The other two methods are based on an alternative way of obtaining the remainder. One of them calculates a remainder estimation in parallel with the algorithm execution using an additional multiplication. The other obtains this remainder estimation by doing a very simple operation of the Goldschmidt algorithm magnitudes consuming less hardware. Both methods achieve noticeable reductions in the number of cases where the remainder calculation is necessary with respect to previous methods.

## Upcoming Events

**The 9th IASTED International Conference on Parallel and Distributed Computing and Networks (PDCN 2010),**  
16-18 February 2010, Innsbruck, Austria, <http://www.iasted.org/conferences/home-676.html>



**The Design, Automation and Test in Europe conference (DATE 2010),**  
8-12 March 2010, Dresden, Germany, <http://www.date-conference.com/>



**The 6th International Symposium on Applied Reconfigurable Computing (ARC 2010),**  
17-19 March 2010, Bangkok, Thailand, <http://www.arc2010.org>



**International Conference on Compiler Construction (CC 2010),**  
20-28 March 2010, Paphos, Cyprus, <http://www.cs.ucr.edu/~gupta/CC%202010.htm>



**The 25th ACM Symposium on Applied Computing (SAC 2010),**  
22-26 March 2010, Sierre, Switzerland, <http://www.acm.org/conferences/sac/sac2010>  
(Special Track on Embedded Systems <http://www2.ing.unipi.it/sac10/>)



**2010 IEEE International Symposium on Performance Analysis of Systems and Software (ISPASS 2010),**  
28-30 March 2010, White Plains, NY, USA, <http://ispass.org/ispass2010/>



**The 3rd IEEE International Conference on Software Testing, Verification and Validation (ICST 2010),**  
06-09 April 2010, Paris, France, <http://vps.it-sudparis.eu/icst2010/>



**The 24th IEEE International Parallel and Distributed Processing Symposium (IPDPS 2010),**  
19-23 April 2010, Downtown Sheraton, Atlanta (Georgia), USA, <http://www.ipdps.org/>



**The 4th ACM/IEEE International Symposium on Networks-on-Chip (NOCS 2010),**  
3-6 May 2010, Grenoble, France, <http://www.minatec.org/nocs2010/>



**The 47th Design Automation Conference (DAC 2010),**  
14-18 June 2010, Anaheim, CA, USA, <http://www.dac.com/>



**The 10th IEEE International Conference on Systems, Architectures, Modeling, and Simulation (SAMOS X),**  
19-22 July 2010, Samos, Greece, <http://samos.et.tudelft.nl/>



**The 19th International Conference on Parallel Architectures and Compilation Techniques (PACT 2010),**  
11-15 September 2010, Vienna, Austria, <http://www.pactconf.org/>



**The 28th IEEE International Conference on Computer Design (ICCD 2010),**  
3-6 October 2010, Amsterdam, the Netherlands, <http://www.iccd-conference.com>



### 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](mailto:leupers@iss.rwth-aachen.de)