

# HiPEAC

COMPILATION ARCHITECTURE

info 33

APPEARS QUARTERLY  
JANUARY 2013

NETWORK OF EXCELLENCE ON  
HIGH PERFORMANCE AND EMBEDDED  
ARCHITECTURE AND COMPIRATION

WELCOME  
TO THE HIPEAC  
2013 CONFERENCE,  
BERLIN, GERMANY,  
21-23 JANUARY  
2013

AN AVALANCHE OF  
RECOGNITION GOES  
TO MEMBERS OF THE  
HIPEAC COMMUNITY



[WWW.HIPEAC.NET](http://WWW.HIPEAC.NET)

SPRING COMPUTING SYSTEMS WEEK, 2-3 MAY 2013, PARIS, FRANCE

## INTRO

# MESSAGE FROM THE HIPEAC COORDINATOR



First of all, I would like to wish you a happy 2013 in good health. Globally, the year 2012 turned out not to be so positive. The economy did not revive, all countries and many companies had to cut spending, and on top of that, there were some natural disasters and bloody conflicts too. I hope that all this bad news was compensated by successes in your personal life. I also hope that 2013 will bring us better news.

HiPEAC<sub>3</sub> has now been running for one year, and it is good to look back. All tasks are on speed, and they are starting to produce their first results. There are a few highlights I would like to share with you. Many of you will be reading this newsletter at the HiPEAC conference in Berlin. The conference has become the flagship networking event of HiPEAC. This year, we used the innovative journal-first publication model for the second time, leading to 43 thoroughly reviewed and revised high quality papers. We are pleased to see that other big conferences are also starting to discuss their publication model.

HiPEAC keeps growing. With over 1200 researchers, we are the biggest research

network in computing systems in Europe. We could, however, considerably increase our membership from the new member states and we have therefore started to organize activities there. We are planning a computing systems week during 2013 to take place in Tallinn, Estonia.

The industrial membership of HiPEAC has gone up to 70, and we are continuing to look for new industrial members. Last October in Ghent, the HiPEAC industry-partner program organized a successful event.

The jobs website is up and running, and we are seeing regular submissions of new jobs and CVs. Hopefully, it helps in attracting talent to European companies and universities.

The HiPEAC community has worked hard on a new version of the HiPEAC Roadmap. This document will serve as one of the inputs for the early calls in complex systems & advanced computing in Horizon 2020.

The year 2012 was a good year for the HiPEAC publication awards. The number of awards went up from a yearly average of 22 in HiPEAC<sub>2</sub> to a record breaking 50 in the

first year of HiPEAC<sub>3</sub> (visit [www.hipeac.net/](http://www.hipeac.net/) award for more details). We hope to further increase the number of high-quality publications coming out of the HiPEAC community.

In 2012 we awarded technology transfer awards for the first time. We received 8 nominations, and 6 of them got a technology transfer award.

In 2013, we will continue working hard on the technology needed for the innovations of tomorrow – innovation that should help our society to deal with its grand challenges. I was happy to see that, despite the harsh times, Europe is trying to further increase research funding in Horizon 2020. Investing in research is investing in the future.

Take care!

Koen De Bosschere

## CONTENT

### INTRO

#### 2 MESSAGE FROM THE HIPEAC COORDINATOR

#### 3 MESSAGE FROM THE PROJECT OFFICER

#### HIPEAC ACTIVITY

#### 4 HIPEAC AUTUMN COMPUTING SYSTEMS WEEK, GHENT

#### 6 HIPEAC WELCOMES TECHNICAL UNIVERSITY OF SOFIA

#### 7 REFLECT AND 2PARMA PROJECTS' FALL SCHOOL

#### 8 HIPEAC MINI-SABBATICAL - SASCHA UHRIG

#### HIPEAC ANNOUNCE

#### 9 ACACES 2013: 14TH - 20TH JULY, FIUGGI, ITALY

#### 10 BOOK ON UTLEON3: EXPLORING FINE-GRAIN MULTI-THREADING IN FPGAS

#### HIPEAC NEWS

#### 10 FOUR HIPEAC MEMBERS GETTING LEADING ROLES IN ACADEMIES, ASSOCIATIONS, AND COMMITTEES

#### 11 HIPEAC MEMBER CHRISTOPHE DUBACH RECEIVES INTEL EARLY CAREER FACULTY HONOR PROGRAM AWARD

#### 11 HIPEAC STUDENT PEJMAN LOTFI-KAMRAN RECEIVES INTEL DOCTORAL STUDENT HONOR PROGRAMME FELLOWSHIP

#### 12 EUPHORIA: A NEW EUROPEAN LAB IN PROGRAMMING AND DESIGN OF HETEROGENEOUS MANY-CORE SYSTEMS

#### 12 UGENT-CSL-HES GROUP WINS STAMATIS VASSILIADIS MEMORIAL AWARD

#### 13 AMAZON WEB SERVICES AWARD FOR CLOUD-BASED COMPILER TUNING

#### 13 MCC'12: 5TH SWEDISH WORKSHOP ON MULTI-CORE COMPUTING

#### 13 HIPEAC STUDENTS FROM UNIVERSITY OF CANTABRIA RECEIVE BEST PAPER AWARD AT ICPP 2012

#### 14 THE PAPAGENO PARALLEL PARSER GENERATOR

#### 14 COMPILER, ARCHITECTURE AND TOOLS DAY AT INTEL HAIFA

#### 15 NEW HIPEAC MEMBER: LUCIAN VINTAN, UNIVERSITY OF SIBIU, ROMANIA

#### HIPEAC STUDENTS

#### 16 INTERNSHIP REPORT - ANTONIO GARCÍA GUIRADO

#### 17 INTERNSHIP REPORT - PAOLO BURGIO

#### 18 INTERNSHIP REPORT - ALEXANDER JORDAN

#### 18 RAPID PROTOTYPING OF 3-D HYBRID VIRTUAL SYSTEM-ON-CHIPS

#### IN THE SPOTLIGHT

#### 20 T-CREST PROJECT: TIME-PREDICTABLE MULTI-CORE ARCHITECTURE FOR EMBEDDED SYSTEMS

#### 21 MONT-BLANC PROJECT SELECTS SAMSUNG EXYNOS 5 PROCESSOR

#### 21 PHD NEWS

#### 24 UPCOMING EVENTS

## MESSAGE FROM THE PROJECT OFFICER



During 2012, we ran a study at the European Commission as one of the inputs to develop the Horizon 2020 strategy in Computing, through an assessment of the changes to the hardware and software market driven by the current use and projected growth of multicore/manycore and heterogeneous computing systems.

The main goal of the study was to analyse in detail the actual and potential multicore/manycore Computing market in Europe for the period 2011-2020, focusing on the following issues:

- Analysis of the areas of strength and weakness for Europe, compared with the US and Asia in the computing ecosystem.
- Identification of the areas of the software and hardware ecosystem where Europe has an opportunity to specialize and grow with longer-term benefits for the European economy and society.
- Identification of the areas of the software and hardware ecosystem where research work at European level is required.

The study analysed all Computing market segments (general-purpose computing,

consumer electronics, communication and wireless, automotive, energy, industrial automation, healthcare, etc.). One important conclusion is that the adoption of multicore/manycore is not simply an incremental innovation (by, for example, just improving speed and power); it enables a new level of integration and innovation, accompanied by sophisticated applications and services, with advanced control and management capabilities.

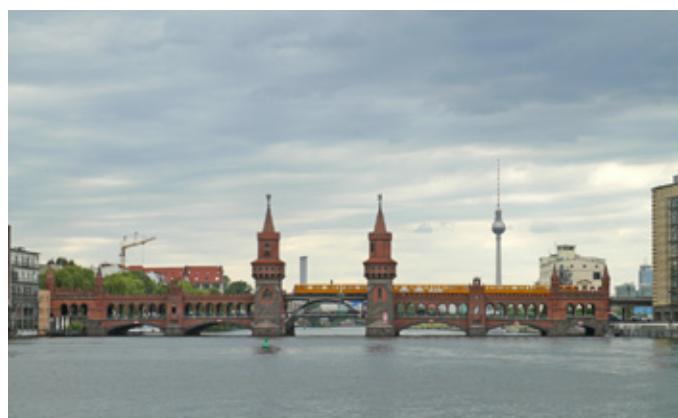
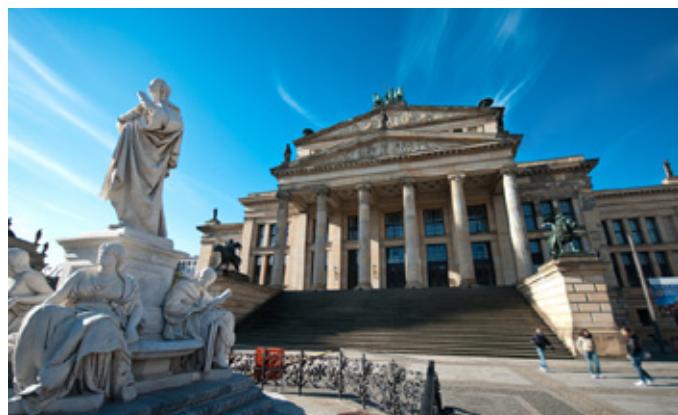
We expect to publish the final detailed findings of the study on our web site during the first months of 2013, but here are some preliminary points coming out:

- The emergence of a Low Power Computing Ecosystem is a clearly identified trend that will become much stronger in the future.
- How to remove barriers to growth for high-tech Computing start-ups and SMEs is a critical challenge that would require much more effort, building on the presence of strong ICT innovation clusters in Europe, and encouraging the development of smart coaching

networks, learning communities, as well as innovative forms of funding (e.g. crowd-funding).

- Main research challenges ahead include the development of cost-effective, advanced software tools and systems suited for parallelism, enabling the development of new software or re-engineering of legacy codes; the development of heterogeneous hardware and software architectures that are aligned with requirements in key industries such as automotive, healthcare, and energy; breakthroughs in exploiting efficiencies from parallelism in algorithms and codes.

*Panos Tsarchopoulos*



# HIPEAC AUTUMN COMPUTING SYSTEMS WEEK, GHENT

This CSW featured a program with 12 thematic sessions as well as the first edition of the HiPEAC Industry Partner Program (HIPP) event

During the Computing Systems Week in Ghent (October 15-17, 2012), we had the opportunity to exercise for the second time this new instrument in HiPEAC3. As a result of a call launched by mid June, a program with 12 thematic sessions was designed. Before presenting a short summary of each of the sessions, I just want to thank the organizers for their enthusiasm in preparing them and for their contributions to this summary for the HiPEAC Newsletter. Additional details about the presentations and slides for most of them are available through the HiPEAC website.

### PARALLEL AND CONCURRENT PROGRAMMING MODELS

This Thematic Session had two goals. The first goal was to promote mobility and research collaborations in the context of HiPEAC and to give visibility to HiPEAC tools, used in PhD work conducted by junior researchers on internships. In addition, senior HiPEAC members, from industry and universities, presented current trends in the field for stimulating mobility and collaboration. The three speakers in this first slot were Ana Varbanescu (TU Delft), Bilha Mendelson (IBM Haifa) and Judit Planas (BSC).

The second goal, addressed during the second slot, was to invite prominent European researchers in the fields of concurrency and parallelism to describe important research activities in Europe. The three speakers in the second slot were Luciano Lavagno (P.Torino), Yannis Smarag-



dakis (U. Athens) and Thomas Fahringer (U. Innsbruck). The keynote speech was given by Alejandro Duran, who presented the heterogeneous MIC architecture. He provided an interesting view of the MIC and its possibilities, from the programming model point of view. During the keynote, there was some discussion on the key points for programming future heterogeneous systems.

### MODELS AND ASSISTIVE TOOLS FOR PROGRAMMING EMERGING ARCHITECTURES

Emerging multi- and many-core architectures have led to a wealth of research in programming models and corresponding assistive tools. This thematic session compared and contrasted the approaches of a representative set of ongoing EU projects, to identify complementarities and similarities. The thematic session provided an excellent platform to encourage cooperation among projects, and an opportunity for igniting future collaborations and for disseminating results to one another and to the HiPEAC community.

Representatives from six EU projects that address the Computing Systems objective of FP7 contributed to this thematic session: Sven-Bodo Scholz (ADVANCE), Nikola Puzovic (ENCORE), Sabri Pllana (PEPPHER), Houssam Haitof (AUTOTUNE), Sven

Verdoolaege (CARP), and Kevin Hammond (PARAPHRASE). Each project presented its approach, illustrated with case studies and code excerpts. The programme was augmented by an invited talk of Jürgen Teich from FAU Erlangen on "Invasive Computing - or - How to Tame 1000 Cores on a Chip?" The thematic session concluded with a panel discussion on "ICT research in Horizon 2020," moderated by Panos Tsarchopoulos (EC), with the panel comprising the representatives of the six projects plus the invited speaker.

### PARALLELISM DISCOVERY AND MAPPING FOR MULTI AND MANY-CORE ARCHITECTURES

This thematic session focused on an outstanding challenge for systems researchers: how software techniques can make the best use of future many-core hardware. The session had seven speakers from industrial companies and leading universities in Europe. The exciting keynote speech was given by Dr. Jos van Eijndhoven from Vector Fabrics, concerning auto-parallelisation techniques for multi-cores. This talk was followed by six presentations covering a wide range of topics: from language abstractions, on to parallelism discovery and optimisation. This event provided an excellent platform for knowledge exchange between industry and academia, which is of significant



CSW participants during the reception at the City Hall.

benefit to both sides. It was a successful event with fruitful discussions both during and after the session.

## DATA CENTERS

The session presented the work of several datacenter research projects and discussed research challenges in the area. This is an exciting research area making a lot of headlines lately. The EUROCLOUD project presented its work in ARM-based servers and 3D server chips. The IOLANES project presented its work in advancing the scalability and performance of I/O subsystems in multicore platforms. The RELEASE project presented its work in adapting Erlang for clusters and datacenters. Two invited talks gave further insights in the emerging market of micro-servers and the use of optical interconnects in datacenters. It is clear that in Europe there is a nascent next-generation server/datacenter ecosystem with a lot of promising potential. Let's stay tuned for some more exciting developments in 2013.

## EMBEDDED AND MOBILE PLATFORMS AND TECHNOLOGIES FOR CLOUD SERVICES

This thematic session provided 1) an industrially-oriented perspective of the evolution of hardware platforms for modern mobile phone and similar devices to sustain the impact of the Cloud, 2) a quick glimpse of the diverse software/hardware infrastructure to support horizontal systems made of wireless sensors and actuators nodes that export their functionality to the Cloud, and 3) an interesting window into the potential of Cloud-aware applications to deliver complex real world interaction and functionality by way of Augmented Reality, in the context of the VENTURI FP7 project. The three speakers were Marco Cornero from ST-Ericsson, Fabien Castagner from STMicroelectronics and Michele Zanin from Fondazione Bruno Kessler FBK.



## MIXED CRITICALITIES / VIRTUALIZATION

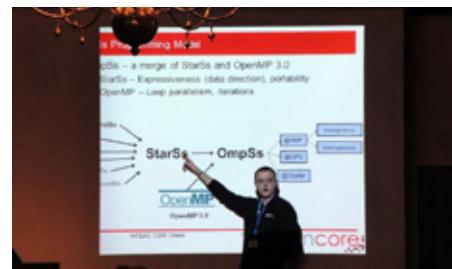
This thematic session consisted of three presentations. The first one, by Sergey Tverdyshev from SYSGO AG, explained the motivation and methodology behind the development of a Common Criteria level 7-certified separation kernel in the EURO-MILS FP7 project. A Multiple Independent Levels of Security (MILS), architecture certified based on verification using formal methods, should enable the shift towards the integration of applications with different criticality levels, reducing hardware cost and the amount of wiring required to network and power all those systems.

The second presentation by Marcello Coppola from STMicroelectronics presented the virtual FP7 project, which looks into software and hardware extensions for heterogeneous multicore platforms. By designing an embedded framework from top (programming models) to bottom (the hardware), the project aims to design a virtualization-ready and programmer-friendly embedded SoC platform, that at the same time fulfils other requirements in terms of performance, security, power consumption, security, QoS, reliability etc.

Finally, the presentation by Thomas Nolte from Mälardalen University gave an overview of the research at his university's real-time research centre (MRTC). One of the key challenges, when running virtualized execution partitions with differing real-time requirements all on a single chip, is how these environments should be scheduled. Researchers at the MRTC have designed and developed a hierarchical resource-based scheduling framework for dealing with these issues, and the audience were given a high level overview of how it works.

## RECONFIGURABLE COMPUTING

The Reconfigurable Computing thematic session addressed various topics dealing with the growing (design) complexity of Systems-on-Chip (SoCs). Ioannis Soudris (Chalmers University of Technology) presented research on the design of future reliable Systems-on-Chip conducted in the FP7 DeSyRe project: improving reliability, reducing power and performance overheads for fault-tolerance of future SoCs that will be prone to many faults. Christiaan Baaij (University of Twente) introduced "CλaSH", a functional language tailored to hardware specification for fine-grained



reconfigurable architectures. Kim Sunesen (Recore Systems) addressed the efficient programming of reconfigurable Multi-Processor Systems-on-Chip in the FP7 ALMA project. The ALMA project is developing a flexible architecture-independent tool chain that can be (re)configured to target different architectures as well as different (re)configurations of the same architecture. Dirk Stroobandt (Ghent University) addressed productivity improvements due to reconfigurable systems. The FP7 FASTER project provides a methodology that allows designers to easily implement and verify applications on reconfigurable platforms.

## THE INTERTWINING CHALLENGES OF RELIABILITY, TESTING AND VERIFICATION

The goal of this thematic session was to bring researchers from these areas to present their current project results as well as their views about upcoming challenges and opportunities. We were able to have twelve presentations, from a variety of sources (industry, FP7 projects and academia), and an open discussion related to the theme of the event. The session was an excellent opportunity to expose the HiPEAC community to recent development in these areas and to provide opportunity for networking and consortia formation for upcoming research calls.

## ULTRA-LOW (ZERO?) POWER SYSTEM ARCHITECTURE, DESIGN AND IMPLEMENTATION

This thematic session consisted of two sessions. The first session was a birds-of-a-feather meeting devoted to charting the power and - equally important - energy optimisation concerns of HiPEAC members, both academic and industrial. The key issues identified were linked to the absence of a common framework to reason about power and energy (unlike those already established for performance, code size and safety), which leads to often tackling the wrong issues. Researchers in processor architecture stressed the lack of common, reliable metrics and models for power and

## HiPEAC ACTIVITY



energy at all design levels (from gate to core to system). From the system level perspective, the main obstacle appears to be the absence of a common framework (concepts, methods, appropriate tools), making the process of designing power-and energy-efficient solutions slow and failure-prone. In the production and applications context, it was observed that power and energy considerations now impact the actual safety of systems (e.g., their ability to complete actions), thus going beyond the usual merchantability and time-to-market aspects.

During the second session, three technical presentations provided an insight into recent developments in industrial ultra-low power system design, power-centric micro-architecture research, and power management in datacenter applications. Peter J. Slikkerveer from MpicoSys introduced the challenges of designing and implementing ultra-portable, autonomous personal devices able to operate entirely on scavenged power and energy. Prof. Sorin Cotofana from Technical University Delft presented an in-depth study of using die

stacking to control power distribution in microprocessor designs. Finally, Prof. Rafael Mayo from Universitat Jaume I reported on research in power monitoring, management and optimisation in high-end compute servers.

### LOW POWER ECOSYSTEM - GEM5 PRACTICAL USER EXPERIENCE

This session focused on presenting experiences in using GEM5. Talks addressed fairly classical simulator issues as well as future and emerging technologies. Bernard Goossens (Perpignan U) started by describing his work on building a new CPU model, while Dmitry Knyaginin and Angelos Arelakis (Chalmers U) explained data collection for memory research. Erik Tomusk (Edinburgh U) presented his struggle to get GEM5 working with the McPAT power simulator and Anouk Van Laer presented his experience in integrating support for optical on-chip interconnects.

### INTERNATIONAL COLLABORATION

The session was a unique opportunity to take stock of what has been achieved up to now in international collaboration in Computing Systems research, and to discuss future perspectives. The two running EU-Russia projects, APOS and HOPSA, presented their work. APOS is optimising scientific and industrial codes that are scalable and portable across heterogeneous and homogeneous architectures. HOPSA is creating an integrated diagnostic infrastructure for combining

application and system tuning. Both projects agreed that, although initially there is a bigger overhead in running an international collaboration research project, the benefits are such that there should be more research projects with an international dimension. Three Support Actions investigating the feasibility of research collaboration with China (SCC-Computing Support Action), India (EU-INCOOP Support Action) and Latin America (RISC Support Action) presented their first results and preliminary ideas for common research priorities.

As a final note for this article, we hope you realize the importance of this HiPEAC3 instrument to promote your research areas in the HiPEAC community, to share your own research results, and to build a network of researchers from which you can form a consortium for a future project proposal. Volunteering to organize a thematic session is an opportunity to contribute to the HiPEAC community, helping the HiPEAC network to work on the challenges identified in the **HiPEAC roadmap** (which was discussed in another thematic session organized by Marc Duranton during the Computing Systems Week in Ghent). Thanks in advance for submitting a proposal in the forthcoming calls for thematic sessions.

*Eduard Ayguadé,  
BSC/UPC, Spain*

## HIPEAC WELCOMES TECHNICAL UNIVERSITY OF SOFIA

### A HiPEAC delegation visited TU Sofia to welcome new members from Bulgaria

On November 5, 2012 a HiPEAC delegation went to the Technical University of Sofia to welcome new members from Bulgaria's first class Technical University. The members of the delegation were HiPEAC coordinator Koen De Bosschere (Ghent University), Rainer Leupers (RWTH Aachen University), and Kostas Magoutis (FORTH), all three HiPEAC partners. Also present was Prof. dr. Ed Deprettere, Leiden University, who is actively involved in education and research programs at the Technical University of Sofia, and was asked to make contacts between HiPEAC and Sofia.

The day took off with a meeting with the rector, Prof. dr. Marin Hristov, and Prof. dr. Peter Yakimov, who was co-founder of the Leiden-Sofia Daedalus Laboratory at the Technical University of Sofia. After this formal, but useful, meeting a large audience gathered to learn more about HiPEAC, its activities and events, and the opportunities that HiPEAC can offer to the new members from Sofia. The presentation of the HiPEAC network of excellence was given by Koen De Bosschere, who also introduced his research programs in the Electronics and Information Systems (ELIS) department at

the Ghent University. Rainer Leupers and Kostas Magoutis also presented their work at The Institute for Communication Technologies and Embedded Systems (ICE), and the Institute of Computer Science, Foundation for Research and Technology -Hellas, respectively. Finally Ed Deprettere gave a brief history of the close cooperation between the Leiden Embedded Research center and the Department of Electronics at the Technical University of Sofia.

Discussions continued, during a delicious lunch at the faculty lunchroom, with the

rector Marin Hristof, vice rectors Valeri Mladenov and Georgi Mihov, and Prof. Angel Popov – who, together with Prof. Marin Marinov, has played a crucial role in the establishment of the Leiden-Sofia Daedalus Laboratory.

After lunch, the HiPEAC delegation visited the Department's Embedded Systems Laboratory, as well as the Leiden-Sofia Daedalus laboratory.

At a good-bye meeting with the rector, we all agreed that the visit had been extremely useful. HiPEAC can expect more faculty and students to become a member, and we are looking forward to the participation of the Technical University of Sofia in the various HiPEAC events.

*Ed Deprettere,  
Leiden University, the Netherlands*



*From left to right: Rector Marin Hristov, Prof. Peter Yakimov, Prof. Ed Deprettere, Prof. Kostas Magoutis, Prof. Rainer Leupers, and Prof. Koen De Bosschere.*

## REFLECT AND 2PARMA PROJECTS' FALL SCHOOL

Jointly organized by two EU-funded projects, the school on Programming Paradigms for Multi-core Embedded Systems was held in Freudenstadt, Germany, on October 2nd to 5th.



This event exposed the attendees to a wide range of programming and application mapping challenges raised by multi-core reconfigurable architectures. The meeting provided a friendly and informal environment that fostered in-depth technical discussions about key challenging topics related to multicore-based embedded computing systems. The core of the technical program combined two days of hands-on sessions where

attendees experimented with two programming tool chains developed in the context of the REFLECT and 2PARMA research projects. The hands-on sessions introduced attendees to an Aspect-Oriented Design Flow and to a domain-specific language (LARA) developed in the context of the REFLECT research project, for the transformations and mapping of application examples written in C to an FPGA-based computing platforms. These



*Group picture during social event of the fall school.*

## HiPEAC ACTIVITY

sessions included live demos of specific application case studies from avionics and multimedia codes with the use of customized compilation sequences. Specific sequences, specified in LARA, included the use of triple-modular redundancy (TMR) in hardware, word-length optimizations and a wide range of loop transformations and loop pipelining geared towards hardware synthesis. Attendees experienced the expressive power of LARA for design-space exploration, with both commercial and academic hardware synthesis tools. Also unveiled was a MATLAB-based weaving capability, thus demonstrating the flexibility of the approach for a wide range of

programming languages. The attendees were exposed to a system-level design methodology and tool flow developed in the context of the 2PARMA project, which encompasses a flexible DSE and several run-time management issues, whose scope cover vertically the flow. Designers can negotiate several design trade-offs navigating a rich design space in terms of computing and storage resources, when mapping applications to emerging architectures such as the STHORM many-core by STMicroelectronics. Specifically, emphasis was given to efficient memory utilization through dynamic memory (heap) management techniques for allo-

cating and de-allocating data at run-time, resulting in lower memory fragmentation. Moreover, a run-time resource manager called Barbeque (BBQ) is capable of taking full advantage of information from the design space exploration phase and the programming paradigms for many-core, to orchestrate system-wide resource allocation to meet QoS of mixed workloads, while optimizing a range of possible figures of merit.

Website:

[http://www.i-paco.com/fall\\_school\\_forest\\_2012](http://www.i-paco.com/fall_school_forest_2012)

## HIPEAC MINI-SABBATICAL - SASCHA UHRIG

### Report for mini-sabbatical at Barcelona Supercomputing Center, Spain

In September 2012, I had the great opportunity to visit the CAOS (Computer Architecture and Operating Systems) group at Barcelona Supercomputing Center for two weeks, sponsored by HiPEAC. I would like to thank the whole CAOS team for hosting me and especially Francisco J. Cazorla and Eduardo Quinones for making this visit happen.

I met Francisco and Eduardo several years ago in the context of HiPEAC, and we are now working together in our second common EC funded research project. We share an interest in processor architectures for hard real-time capable embedded systems, which also forms the background for the visit. Most work regarding real-time capable systems concerns the processor core, the memory hierarchy, and the predictability of these structures. However, since the tasks in embedded real-time systems mainly depend on interactions with their environment, it is also important to provide a real-time capable input/output system. Especially in the context of future many-core embedded systems, input/output has greater importance because of higher concurrency and complexity in the interconnection network. Hence, the main topics of the visit have been interrupt transportation and handling techniques, and direct memory access (DMA) from peripheral devices. We conducted an intensive literature study to investigate and discuss interrupt handling

and DMA techniques used in existing single- as well as multi- and many-core systems. The multi- and many-core systems taken into account provide powerful techniques to deal with input/output, but they are mainly optimized for average performance, which complicates or even impedes meaningful real-time timing analysis. During several brainstorming sessions, we formed two concepts for real-time capable DMA transfers, and one for interrupt transportation. All concepts have in common the ambition for the hardware requirements to be as low as possible, while providing the highest flexibility. Of course, they all provide analysability. Because of limited time, it was not possible to evaluate the proposals directly. Instead, they will be explored in more detail within the FP7 project parMERASA, which provides the required infrastructure.

As a side effect of the discussions about interrupt transportation, we developed a novel energy-aware and real-time capable locking mechanism for many-core systems. The drawback of common spin-lock or ticket-lock based techniques is the need for busy-waiting if the lock is not received immediately. The frequent memory accesses increase network traffic and energy consumption. Our energy-aware technique is based on the proposed interrupt technique and provides fair and predictable access to the lock. Moreover, it completely eliminates the repeated

memory accesses of the waiting thread, which reduces power consumption and concurrency at the memory.

In summary, the two weeks at BSC were very fruitful. Besides the two topics of DMA and interrupts, which were scheduled to be handled during the visit, the collaboration was enhanced by the additional topic of the locking mechanism. The proposed techniques will be evaluated in detail in the near future, and we expect that these new proposals will lead to the publication of several papers.

Sascha Uhrig

Technical University of Dortmund, Germany



## ACACES 2013: 14 - 20 JULY 2013, FIUGGI, ITALY

### 9th International Summer School on Advanced Computer Architecture and Compilation for High-Performance and Embedded Systems

We are proud to announce the ninth HiPEAC Summer School, which will take place in downtown Fiuggi, a historical hill town near Rome, during the third week of July. We start on Sunday evening with an opening keynote. The twelve courses start on Monday, spread over two morning and two afternoon slots. There are three parallel courses per slot, from which the participants 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.

Instructor	Course
Arvind - MIT, USA	Computer architecture: a constructive approach
Emery Berger - University of Massachusetts Amherst, USA	Software fault tolerance and correction
Hendrik Berglund - Chalmers University of Technology, Sweden	Creating new business based on innovative technology
Luis Ceze - University of Washington, USA	Approximate computing from programming language to hardware.
Samarjit Chakraborty - Technical University of Munich, Germany	Cyber-physical systems
Alex Garthwaite - CloudPhysics, USA	Memory hierarchies and their impact on virtualization
Jesus Labarta - Barcelona Supercomputing Center, Spain	Parallel programming with OmpSs
Hsien-Hsin Lee - Georgia Institute of Technology, USA	2.5D and 3D IC electronic systems
Sam Midkiff - Purdue University, USA	Automatic parallelization for computer architects
Onur Mutlu - Carnegie Mellon University, USA	Scalable many-core memory systems
Lesley Shannon - Simon Fraser University, Canada	Computing system design for reconfigurable platforms
Anand Sivasubramaniam - Pennsylvania State University, USA	Datacenter power management



On Wednesday afternoon, participants will be 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. Accommodation will be provided by a consortium of hotels in Fiuggi, all located closely together. There will be abundant Italian food, and the town of Fiuggi will provide plenty of opportunities to socialize in the evenings. At the end of the event, all participants will receive a certificate of attendance detailing the courses they took. If you are a student member of HiPEAC, you can apply for a grant that covers the registration fee. In this newsletter, you will find a summer school poster. Please post 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



### BOOK ON UTLEON3: EXPLORING FINE-GRAIN MULTI-THREADING IN FPGAS

M. Daněk, L. Kafka, L. Kohout, J. Sýkora, R. Bartosiński

The book presents the design, implementation and evaluation of instruction set extensions for fine-grain multithreading implemented in UTLEON3, a SPARC V8 processor with micro-threading derived from Aeroflex-Gaisler's LEON3, in the FP7 project Apple-CORE funded by the European Commission. The first part of the book defines new processor instructions for thread management in a way compatible with the existing SPARC V8 opcodes, and proposes the necessary hardware extensions to LEON3, describing them on a functional level.

The second part of the book describes the implementation details of the architectural

extensions required to execute micro-threads, both in the processor pipeline and in specialized hardware accelerators. An effort has been taken to describe the structure of the new blocks in a way that would provide guidance to the actual VHDL sources of the blocks without going into unnecessary details.

The description is accompanied by an analysis of performance gains; these were evaluated by comparing the cycle count of microthreaded assembler programs, executed on UTLEON3 in the microthreaded mode, with reference legacy assembler programs, coded using the standard SPARC V8 instructions and executed on LEON3,

and UTLEON3 in the legacy mode. The main text is supplemented with appendices that provide additional information on LEON3, together with a scheduling example and resource requirements for UTLEON3. As the book is released in parallel with the UTLEON3 sources, a tutorial is also provided that gives instructions on the basic set-up of the VHDL package with the UTLEON3 sources.

M. Daněk, L. Kafka, L. Kohout, J. Sýkora, R. Bartosiński  
Institute of Information Theory and Automation, Czech Republic

### FOUR HiPEAC MEMBERS GETTING LEADING ROLES IN ACADEMIES, ASSOCIATIONS, AND COMMITTEES

More evidence of HiPEACs leadership in architecture of computing systems

Four senior HiPEAC members were recently acknowledged for important roles in leading academies, associations and committees. Manolis Katevenis of FORTH and professor of University of Crete, considered the "father" of computer architecture in Greece and known as a major contributor to scalable multicore technology and co-founder of HiPEAC, recently became a member of Academia

Europaea. André Seznec, INRIA, a long-term contributor in computer architecture and an ERC advanced grant holder was elevated to the highest grade in IEEE (the Institute of Electrical and Electronic Engineers) for contributions to design of branch predictors and cache memory for processor architectures. Per Stenström, professor of Chalmers University of Technology, a long-term contributor to

parallel architecture and a co-founder of HiPEAC, has been appointed as a member of the ACM Turing Award Committee, which annually hands out the most prestigious prize in computing. Finally, Mateo Valero, Director of Barcelona Supercomputing Center and professor of UPC, has been appointed Correspondent Academic of the Mexican Academy.



Manolis Katevenis.



André Seznec.



Per Stenstrom.



Mateo Valero.

## HiPEAC MEMBER CHRISTOPHE DUBACH RECEIVES INTEL EARLY CAREER FACULTY HONOR PROGRAM AWARD

This annual award is designed to help Intel connect with the best and brightest early career faculty members at the top universities around the world.

Christophe Dubach, lecturer at the University of Edinburgh, UK, has been selected as a recipient of the Intel 2012 Early Career Faculty Honor Program Award. This award provides financial and networking support to faculty members who show great promise as future academic leaders in disruptive computing technologies. This year, a total of 75 highly qualified candidates applied for the award. The selected winners were from 9 European and 11 US universities including Carnegie Mellon, UC Berkeley, Cornell and MIT.

Christophe's research interests lie in the area of software/hardware runtime adaptation for heterogeneous systems. His current focus is on dynamic adaptation for heterogeneous multicore processors. The support from this award will allow Christophe to establish a stronger link with Intel and work towards addressing some of the major challenges facing industry.



*Christophe (on the right) receiving his award at Intel Barcelona.*

## HiPEAC STUDENT PEJMAN LOTFI-KAMRAN RECEIVES INTEL DOCTORAL STUDENT HONOR PROGRAMME FELLOWSHIP

The fellowship is granted to Pejman Lotfi-Kamran from EPFL for his contributions on scale-out processors

Datacenters are the backbone of cloud computing. Emerging cloud workloads like web search and social networking are scale-out in nature. Scale-out workloads run on thousands of servers, each processing a shard of a massive memory-resident dataset, to serve independent client requests. Because the combination of processors and memory dominates server acquisition and operating costs, datacenters require processors that efficiently use the memory and maximize throughput per dollar of ownership expense.

The PhD thesis of Pejman Lotfi-Kamran, who is a PhD candidate at EPFL, targets the design of server processors optimized for scale-out workloads. These workloads have three main characteristics that can be leveraged to reach high processor efficiency: abundant request-level parallelism, large instruction footprint, and large memory-resident datasets. These characteristics favor server processors with many cores for high parallelism, a modestly sized last-level cache (LLC) and low-latency interconnects for rapid instruction delivery; and rich off-chip bandwidth for data transport. Pejman demonstrated that processors based

on many lean cores provide superior throughput to conventional "fat-core" server processors. His research relies on two critical observations: First, large LLCs waste precious silicon real estate that could have been better used to integrate more cores. Second, the organization of a many-core processor has a significant impact on its performance. Existing many-core chips, such as those offered by Tilera, sacrifice much of the on-die real estate to LLC and employ a tiled organization that incurs a high on-chip communication overhead. In contrast, Pejman proposes a many-core processor based on the notion of pods. A pod is a module that tightly couples many cores to a modestly sized LLC through a low-latency interconnect. The proposed processor integrates many pods where each pod is a self-contained server-on-a-chip running a full software stack. Pejman formulated a methodology to determine the optimal number of cores and LLC capacity to integrate in a pod for peak throughput. The proposed design, called the Scale-Out Processor, delivers peak throughput in today's process technology and affords near-ideal scalability as the technology scales.



*Pejman receiving his award at Intel Barcelona.*

Pejman is one of the recipients of the Intel doctoral student honor programme fellowship for the 2012-2013 academic year. This award is given to exceptional PhD candidates pursuing leading-edge innovation in fields related to Intel's business and research interests in the European Union, Switzerland and Russia. This is a prestigious and highly competitive program with a limited number of fellowships awarded annually. Recipients of this award are recognized as being amongst the best in their areas of research.

## EUPHORIA: A NEW EUROPEAN LAB IN PROGRAMMING AND DESIGN OF HETEROGENEOUS MANY-CORE SYSTEMS



On 10th December 2012 the EuPhoria lab had its kick-off event in Edinburgh, UK. There was an invited industrial keynote from Jem Davies, VP of Technology, Media Processing Division at ARM, followed by a panel session including speakers from Imagination Technology, Samsung, Freescale, Herta Security, Vector Fabrics, Wolfson, CodePlay, Agilent and Critical Blue.

Euphoria is a new laboratory in the programming and design of tomorrow's heterogeneous many-core systems. It brings together five international leading universities to work together with industry on the grand challenges in our area. The academic team consists of:

- University of Edinburgh - Coordinator Prof. Michael O'Boyle

- BSC/UPC - Prof. Xavier Martorell
- University Ghent - Prof. Koen De Bosschere
- INRIA - Prof. Albert Cohen
- RWTH Aachen - Prof. Rainer Leupers

### VISION

The future is heterogeneous many-cores, but there is no clear way to design and program them. As the number and diversity of cores increases, the software gap between theoretical and actual performance grows larger, and this will be the critical issue for computing systems in five to ten years' time. Our goal is to research and develop core technology for embedded to exascale many-core computing. It will take an end-to-end approach based on real applications and deliver the intervening stack. It brings together Europe's strongest teams in system software and programming languages.

### PROGRAMME OF WORK

EuPhoria will tackle the central challenge of how to design and use heterogeneous many-cores. It will tackle these with five coordinated activities:

- Research: This spans programming languages to hypervisors, from embedded controllers to exascale computing. We will exploit the rapid convergence of previously vertical domains to develop scalable innovation.
- Technology transfer: Research excellence must have a real impact. We will develop knowledge transfer into companies and exciting new start-ups from our research.
- Next generation innovators: A real two-way dialogue. PhD education linked with industrial internships and industrial engineers with academic sabbaticals.
- Next generation applications: Too often system research focuses on tomorrow's solutions to yesterday's problems. We will explore new mobile and crowd-sourced applications.
- Tools and prototypes: The way we do systems research is changing. Leveraging community development can lead to rapidly prototyped ideas and deep community penetration.

*Michael O'Boyle  
University of Edinburgh, Scotland*



*Fatma Aboueella receives the award from FPL program chairs Dirk Koch and Satnam Singh.*

## UGENT-CSL-HES GROUP WINS STAMATIS VASSILIADIS MEMORIAL AWARD

This award recognizes the best paper at the conference on Field Programmable Logic and Applications (FPL) 2012

At the 22nd International Conference on Field Programmable Logic and Applications (FPL), the HES (Hardware and Embedded Systems) group of Ghent University's Computer Systems Lab (CSL) received the Stamatis Vassiliadis Memorial Award for the best paper. During the conference dinner, the award was handed for the paper "Automatically Exploiting Regularity in Applications to Reduce Reconfiguration Memory Requirements" by Fatma Aboueella, Karel Bruneel and Dirk Stroobandt. The prize was an Altera Nios II Embedded Evaluation Kit featuring a touch screen. With this award the Hardware and Embedded

Systems (HES) group of Prof. Dirk Stroobandt received recognition for its groundbreaking work on parameterized run-time reconfiguration, an invention of post-doctoral researcher Karel Bruneel. More information on this work can be found at the group's webpage at <http://hes.elis.ugent.be/>. The award is named after Prof. Stamatis Vassiliadis (IEEE Fellow, ACM Fellow, Member of the Dutch Academy of Sciences, and Professor at Delft University of Technology). He was well known as an outstanding computer scientist and for establishing the SAMOS conference in 2001. When he passed away much too

early on 7 April 2007, the community lost one of its most skilled and inspiring actors. Stamatis Vassiliadis was involved in the organization of many scientific conferences and he was the general chair of FPL2007.

*Dirk Stroobandt  
Ghent University, Belgium*

## AMAZON WEB SERVICES AWARD FOR CLOUD-BASED COMPILER TUNING

A new project investigates JIT compiler tuning on various AWS instance types

Amazon Web Services has awarded \$2500 to HiPEAC member Jeremy Singer at Glasgow University School of Computing Science. Singer aims to investigate JIT compiler heuristic tuning using search-based techniques on the Jikes RVM. Such tuning methods are normally prohibitively

expensive for single-system specialization. However Singer hopes to use the massive computing resources in an AWS data center to tune the Jikes RVM system for particular AWS instance sizes. The expense of search-based tuning will be offset by the efficiency gain in running a

specialized Jikes RVM in each standardized AWS instance size.

*Jeremy Singer  
University of Glasgow, UK*

## MCC'12: 5TH SWEDISH WORKSHOP ON MULTI-CORE COMPUTING

The event took place at KTH Royal Institute of Technology in Kista, Sweden, on 22-23 November 2012

The program consisted of 13 submitted presentations divided in five sessions: Energy efficiency, task scheduling, architecture and performance analysis, applications and composable software, and automatic parallelization. Prof. Martti Forsell from VTT Finland opened the workshop on Thursday with a keynote talk on parallelism, architecture-independent abstraction of computation and architectural support of them on CMPs. The first day concluded with a poster session with 10 poster presentations

while the workshop participants enjoyed some drinks and snacks. On Friday, Prof. Stefanos Kaxiras from Uppsala University gave a keynote talk on complexity-effective multicore cache coherence. This two-day event was successfully attended by almost hundred participants from academic and research institutions all over Scandinavia.

*Mats Brorsson  
KTH, Sweden*



*Stefanos Kaxiras during his keynote talk at MCC'12.*

## HiPEAC STUDENTS FROM UNIVERSITY OF CANTABRIA RECEIVE BEST PAPER AWARD AT ICPP 2012

Their paper explores the effects of local link saturation in interconnection networks for exascale supercomputers

Interconnection networks constitute a key subsystem in the architecture of modern high-performance computers. Technology trends suggest that the available pin bandwidth can be best exploited using high-degree routers. To this aim, it has been suggested that the interconnection networks of forthcoming exascale supercomputers should be built as a two-layered hierarchical network known as a Dragonfly. These topologies are divided in groups of routers denoted as "supernodes," which are interconnected by optical wires denoted as "global links". The routers within a

supernode are interconnected by electrical wires denoted as "local links". Although minimal routing provides good performance for uniform traffic, under certain adversarial patterns non-minimal routing is necessary to avoid saturated global links. Adaptive policies have been designed to handle variable traffic patterns, by selecting between minimal and non-minimal routing. However, those policies do not take into account the effect of saturation of the local links. In "On-the-Fly Adaptive Routing in High-Radix Hierarchical Networks", the authors

study how local link saturation can be common in these networks and show that it can significantly reduce their performance. This work received the best paper award at the International Conference on Parallel Processing 2012. This conference is a forum for engineers and scientists in academia, industry and government to present their latest research findings in any aspects of parallel and distributed computing, and was held on 10-13 September 2012 in Pittsburgh, Pennsylvania. This work has been indirectly fostered by the HiPEAC Network of Excellence, since

## HiPEAC NEWS

most of its authors are members of HiPEAC. In addition to the PhD students Marina García, Cristobal Camarero and Miguel Odriozola, professors Enrique Vallejo and Ramón Beivide from the University of Cantabria, Mateo Valero and Jesús Labarta, professors at Universidad Politécnica de Catalunya, and Germán Rodríguez and Cyriel Minkenberg, researchers at the System Fabrics Group in IBM Zurich Research Laboratory were co-authors of this work.

In such work, the authors propose using non-minimal paths for avoiding saturated local links. However, this solution extends the maximum path length. Previous routing proposals prevent deadlock by relying on

an ascending order of virtual channels, which would imply higher cost and complexity in the network routers. This paper presents a novel routing/flow-control scheme that decouples the routing and deadlock avoidance mechanisms. Deadlock is prevented by employing a deadlock-free escape sub-network based on a Hamiltonian ring with Bubble flow-control. The results show that the model obtains low latency, high throughput, and quick adaptation to transient traffic, consuming packet bursts 43% faster than previous proposals

Ramon Beivide  
University of Cantabria, Spain



Marina García receives the award from ICPP program chair Manish Parashar.

## THE PAPAGENO PARALLEL PARSER GENERATOR

A tool for producing optimized C parallel parsers exploiting POSIX threads to efficiently split the load on multiple processors.

One field where exploitation of machine parallelism has been slow to materialize is language parsing. The obvious idea is to split a long text into chunks to be independently parsed by workers. However, this idea is incompatible with classical parsing algorithms because they have been designed to obtain efficient execution on serial computation models. In particular, the well known LR(1) technique imposes sequential execution of the parsing actions, because the choices to be made by the algorithm may depend on past parser states.

However the classical deterministic parsing algorithm known as operator-precedence parser, due to R. Floyd, does not suffer from these limitations, as parsing decisions can be made locally, an ideal property for producing a data-parallel parser.

Exploiting this locality property, we have devised a parallel parsing algorithm, which has proven effective in parallelizing the parsing process on modern symmetric multicore processors, exhibiting a parallel code

portion in the 80-85% range, reaching 95% on Simultaneous MultiThreading (SMT) enabled CPUs. We tested our parallel parsing algorithm on different languages, and, more extensively, on documents written in the JavaScript Object Notation language (JSON), a data description language, which has seen widespread use, especially in modern web applications.

The product of this research effort is PAPAGENO, a tool for producing optimized C parallel parsers exploiting POSIX threads to efficiently split the load on multiple processors.

The generated C parser can be run on a parametric number of parallel workers, to be chosen at runtime by the developer or user, and it only depends on the standard C library, to retain maximal portability.

PAPAGENO accepts as input a grammar specification written with the same syntax of Bison, for the sake of easy portability, and it supports the use of Flex generated scanners. Notice that our generated parsers, when

run on serial computers, are at least as fast as Bison parsers; however not all source grammars that are acceptable for Bison are good for PAPAGENO.

PAPAGENO is written in Python to allow for easy extensibility, and the source code is available freely at <https://code.google.com/p/parallelparsing/>

Currently, we are experimenting with using the tool on other languages, in particular Javascript, where the operator form of Floyd's grammars requires some adaptation; and we are evaluating the possibility of also parallelizing the scanning phase. We acknowledge the financial support of Google Research Award and we thank Brad Chen of Google Research.

Alessandro Barenghi, Stefano Crespi-Reghizzi,

Dino Mandrioli, and Matteo Pradella

Politecnico di Milano, Italy

## COMPILER, ARCHITECTURE AND TOOLS DAY AT INTEL HAIFA

On Monday, November 26th, about 150 attendees from Israel, Europe and the US gathered at Intel's Israel Development Center in Haifa to attend a Compiler, Architecture and Tools Day. Prof. David

August from Princeton University opened the event with an intriguing keynote talk, reflecting on what techniques worked well in the past and suggesting how to restore the trend of computing performance in

the future. A full program followed, with four sessions on Programming GPUs and OpenCL, Compilation Techniques, Tools, and Parallelization Techniques, together with a second keynote talk by Intel Fellow

Geoff Lowney on Data Parallel Programming - from its history up to current platforms - providing a couple of "good ideas" to help programming converge.

Half of the fourteen talks were given by European guests from the Netherlands, UK, Poland, Russia and Germany. Local presenters from academic and industrial institutions in Israel completed the diverse agenda. Additional guests attended the event, from France, Ireland, Cyprus and the US. Lively discussions took place throughout the full day, during the breaks and over lunch, ending after 6pm. Positive and appreciative feedback relating to the technical content, participation, and organization was received from many participants.

The event marks the revival of previous events held in the past, the most recent taking place at IBM Haifa in 2007, coinciding with the 3rd HiPEAC Industrial Workshop. Similar to the past workshops, the main theme covered at the event involved various aspects of compilation, ranging from a variety of programming

language design issues and data-parallel programming, through dynamic, adaptive and machine-learning techniques, to the efficient use of contemporary GPU and VLIW architectures. A new session on Tools was introduced this year, revealing new tools, features and infrastructure related to simulation, performance analysis and design for heterogeneous systems, all using and supporting efficient compilation processes.

The event was organized jointly by Ayal Zaks and Gadi Haber from Intel Haifa, Dorit Nuzman from IBM Haifa, and Erez Petrank from Technion. It took place at the "Green" IDC9 building (LEED® Gold certified), was hosted by Intel's Software and Services Group (SSG), and endorsed by HiPEAC.

Please visit  
<http://www.intel.co.il/compilerConf2012>  
 for more information.

Ayal Zaks, Gadi Haber, Dorit Nuzman  
 and Erez Petrank,  
 IBM, Intel and Technion, Israel



*The event at Intel IDC9 in Haifa.*

## NEW HIPEAC MEMBER: LUCIAN VINTAN, 'LUCIAN BLAGA' UNIVERSITY OF SIBIU, ROMANIA

Lucian N. Vintan is currently a Professor in Computer Engineering at "Lucian Blaga" University of Sibiu, Romania. He leads the Computer Architecture Research Group from this university – see <http://acaps.ulbsibiu.ro>. He received the MSc degree in Computer Engineering from "Politehnica" University of Timisoara, Romania, and the PhD degree in Computer Engineering from the same university.

Professor Vintan has received various awards and recognitions for his work. For his professional merits he received The Romanian Academy "Tudor Tanasescu" Prize in 2005. In 2012 he was elected full-member of The Academy of Technical Sciences from Romania (correspondent member in 2005). He is actively involved in EU-funded projects from the evaluation of proposals to the managing and reviewing of projects. Professor Vintan is a European Commission DG Information Society Expert, Visiting Research Fellow at University of Hertfordshire, UK, etc. From October 2012, he was

accepted as a HiPEAC member. Also, four members of his group have became affiliate members and two HiPEAC PhD students.

Professor Lucian Vintan is an expert in the areas of instruction/thread level parallelism, SMT, multi-core and many-core systems, automatic design space exploration, prediction techniques in UbiCom systems and text mining. He has published over 100 scientific papers in many prestigious journals and international conferences. He introduced some well-known original concepts in Computer Architecture domain: Dynamic Neural Branch Prediction, Pre-Computed Branches, Unbiased Branches, Value Prediction focused on CPU's Context, Selective Dynamic Value Prediction and Dynamic Instruction Reuse, Automatic Design Space Exploration augmented with Computer Architecture Fuzzy Micro-ontologies,etc. Furthermore, he has authored six books on these subjects. Two of them, written in English, are used as textbooks on the subject at leading universities

around the world. He finalized as a Project Manager ten research grants obtained through competition. Until now, five PhD students graduated under his supervision. He has served on the technical program committee of over 50 international computer systems conferences and has peer-reviewed hundreds of research papers for numerous international journals and conferences.

Website: <http://webspace.ulbsibiu.ro/lucian.vintan/html/>



## INTERNSHIP REPORT - ANTONIO GARCÍA GUIRADO

### Cache Organizations and Management for On-Chip Photonic Networks in CMPs



I am a PhD student at the University of Murcia, Spain, working on novel cache coherence protocols and cache organizations for chip multiprocessors (CMPs), under the supervision of Ricardo Fernández Pascual and José M. García. Thanks to a HiPEAC collaboration grant, I had the chance to work together with Sandro Bartolini at the University of Siena, Italy, where his research group is working on the application of silicon photonics to CMPs within the National Project FIRB "Photonica".

Photonics can potentially solve the scalability problems of networks-on-chip (NoCs) when moving towards CMPs containing hundreds of cores. An appropriately designed photonic network enables any part of the chip to be reached by means of a single light pulse, preventing costly retransmissions that are commonplace in electric networks, especially as increasing core counts introduce too many intermediate retransmission points. A smart use of photonic technology, at the chip-architecture level, is needed to efficiently use its outstanding latency and energy features.

To that end, during this collaboration we have designed novel dynamic network management policies to exploit hybrid photonic/electronic NoCs in a NUCA-based CMP. We tested our policies on a CMP containing an ordinary electrical mesh and a ring-based photonic network that cooperate to minimize the overall network latency and energy consumption. Our policies are based on available real-time message-grain information such as message size, distance between endpoints and photonics queuing time. By selectively using the photonic network to send those messages most likely to benefit (distant,

short, and keeping low waiting times), we reduce the number of retransmissions in the electrical mesh, and at the same time we prevent the excessive and harmful use of the photonic network (e.g., preventing message serialization). We found that the fine granularity of our policies allow them to adapt to changing traffic conditions and make effective use of any amount of photonic resources, enabling a more efficient use of photonic networks than previous proposals. Also, we have developed line-grain adaptive LLC organizations for stand-alone photonic networks, based on novel cache coherence protocols. These organizations dynamically adapt to changing execution characteristics, trading off LLC resources for faster LLC accesses. By adaptively allowing the local replacement of frequently accessed blocks (at the cost of adding pressure to the LLC), we can modulate network traffic to prevent peaks of long queuing times on the photonic network that would notably harm performance and energy efficiency. Moreover, these LLC organizations enable a more efficient use of modest amounts of photonic resources that can leave extra room on chip for other critical resources. As a result of this collaboration, strong bonds between the universities of Murcia and Siena have been created. I want to thank HiPEAC for giving me the opportunity to carry out this collaboration. I would also like to thank Sandro and other colleagues at the University of Siena for their support and for making my stay such a pleasant experience.

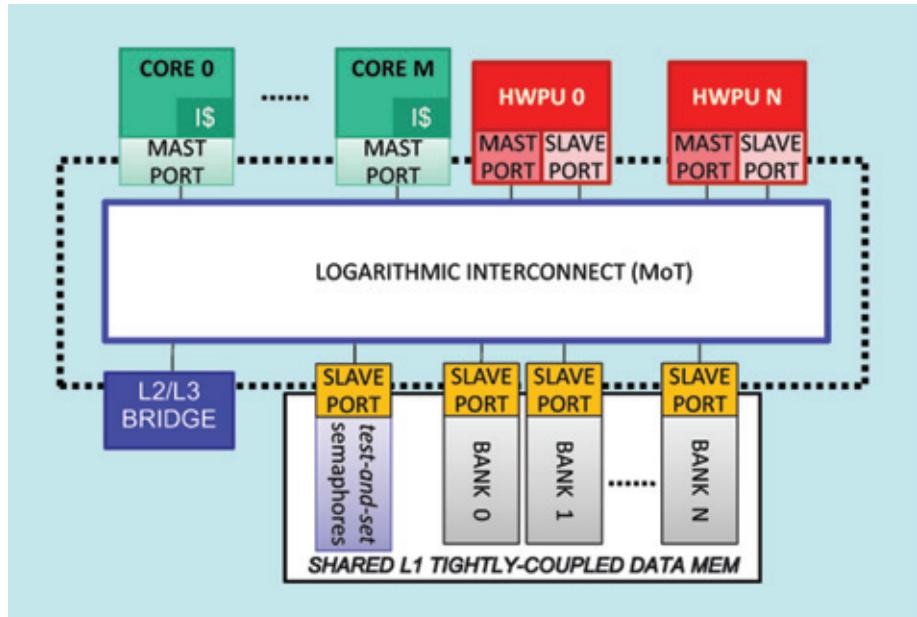
*Antonio García Guirado,  
University of Murcia, Spain*

# INTERNSHIP REPORT - PAOLO BURGIO

## OpenMP-based Synergistic Parallelization and Hardware Acceleration for On-Chip Shared-Memory Clusters

I am a student in a joint PhD program between Università degli Studi di Bologna, Italy, and Université de Bretagne-Sud, France, which started in 2011. My advisors are Prof. Luca Benini and Prof. Philippe Coussy, respectively. Thanks to the HiPEAC mobility grant, I spent four months in UBS, working on a framework to design accelerator-based MPSoCs. This collaboration brings together design expertise from UniBo and UBS to streamline the design of new generation many-core SoCs platforms featuring hardware accelerators. Hardware accelerators are increasingly being coupled to cores in modern many-core SoCs, to trade energy efficiency for application (domain) specificity of platforms. The design process has been traditionally carried out “by hand”: starting from the high-level specifications of an algorithm (e.g., C code), the system architect manually extracts key computations (kernels) which benefit most from a hardware implementation, replacing them with the call to a runtime which interacts with the actual hardware modules. The latter are (semi-manually) implemented and optimized, then the platform is composed and the application deployed. Although this process can be partially automated using HLS tools, either academic (UBS GAUT) or commercial (Calypto CatapultC), its inherent multidisciplinarity requires knowledge on embedded software and hardware design, algorithmics, as well as an extensive phase of validation of the overall platform and software using simulators. We tackle these issues by proposing the template of a generic accelerated platform and implementing an automated design toolflow for it. In more detail, our contributions are:

1. the template for a generic accelerator-based, multi-cluster many-core SoC, inspired by state-of-the-art architectures such as STM’s Platform 2012/STHORM, Plurality HW Architecture Line and GPGPUs as NVIDIA Fermi;
2. the template for a generic HW accelerator (i.e., whose “business” functionality is not defined in advance) named HardWare



*Scheme for a single cluster of the proposed architecture.*

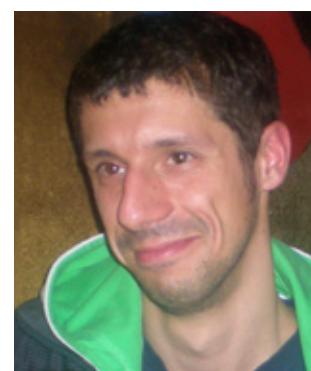
Processing Unit (HWPUs), designed to share the same SW-managed data memory that the cores also work on (zero-copy scheme); 3. we proposed extensions to the OpenMP programming model to enable fast and streamlined exploitation of HW acceleration opportunities directly from within the application code. We introduced a custom *accelerate* pragma to identify key regions of a C program that are suitable for acceleration in Hardware.

Our integrated toolflow processes the application source files, replacing the annotated code with the calls to a highly optimized library to interact with the target HWPUs. This is done by our customized GCC compiler, which also extracts the code of the kernels and gives it as input to an HLS tool (UBS GAUT). GAUT provides the RTL description of the HWPUs, as well as a behavioral SystemC model that is automatically integrated in a cycle-accurate simulator (UniBo MPARM) for fast prototyping and exploration of the target many-core platform. By extending OpenMP, we also enable programmers to mix parallelization and HW acceleration in their software, using a simple pragma-based programming interface. Using our toolflow we wrote a paper titled

“OpenMP-based Synergistic Parallelization and HW Acceleration for On-Chip Shared-Memory Clusters”, which I presented at Euromicro DSD 2012, and a poster with the same name at DAC 2012.

During the last month of my internship we also identified future research directions. Until now, our approach targeted a single-cluster system. Moving to a multi-cluster scenario is the next step of the project. We also plan to identify programming model extensions to address platforms with pre-existing HWPUs, instead of designing them from scratch.

*Paolo Burgio,  
University of Bologna, Italy, and Université  
de Bretagne-Sud, France*



## INTERNSHIP REPORT - ALEXANDER JORDAN

### Can you afford instruction scheduling in a lightweight JavaScript JIT compiler?

Last summer I had the opportunity to do a four-month internship with Christophe Guillon at STMicroelectronics in Grenoble. While my PhD research at the Vienna University of Technology is focused on integrated and optimal models of problems in compilation, this internship gave me the chance to do something different: optimizing a JavaScript Just-In-Time (JIT) compiler. At STMicroelectronics (STM) the V8 JavaScript virtual machine (VM), originally developed by Google, has previously been ported to SH4 (SuperH). This architecture can be found in STM's consumer electronics systems. Thus, improving JavaScript performance in this context, means improving the responsiveness and speed of applications and user interfaces that might be running on your TV or set-top box.

SH4 can execute up to two instructions in parallel, but it relies on the compiler to order instructions appropriately in advance. The latter is a problem in a JIT compiler, where the cost of scheduling may result in an unwanted increase of a program's total runtime. As a solution, striking a balance between runtime improvement and compile time overhead, Christophe proposed to use a lightweight scoreboard scheduling algorithm.

Given the particular properties of the hardware, JavaScript VM and the scheduling algorithm, my task was to find out to what extent scheduling improves the code

generated by V8 and whether it can achieve an overall speedup in applications.

To understand how the code generated by V8 executes on SH4, I started with the analysis of JavaScript programs from the V8 benchmark suite. Instrumentation added to a QEMU emulator allowed profiling of JIT-compiled code, which at the same time had been disassembled to MinIR, a compiler-independent intermediate representation in YAML. Together, this enabled off-line analysis using a simplified architectural model: reordering instructions and measuring schedule lengths was possible without any changes to the VM. We could see that the basic blocks produced for JavaScript are exceptionally small (only six instructions on average), but they can still benefit from scheduling. Encouraged by these results, my next step was integrating scheduling into V8. While the complexity of the scoreboard algorithm itself stayed within the expected bounds, it became clear that the overhead due to additional bookkeeping would play a key role. Essentially, all data structures related to scheduling had to be made efficient for the most frequent instructions. Ultimately my implementation, though still prototypic in places, closely matched the earlier estimates and showed that JIT-compiled JavaScript code can run up to six percent faster through scheduling alone.

What I take away from my internship, first and foremost, is how a methodical approach paired with solid engineering can lead to success in a tricky optimization project. Second is the importance of tool support. In my day-to-day work, I relied on many tools that would aid development, analysis, testing and integration of embedded code. Creating and continuously maintaining these tools comes at a cost, but not having them is something an enterprise like STMicroelectronics simply cannot afford.

I want to take this opportunity to thank the internship organizers at HiPEAC, as well as everyone at the Compilation Expertise Center who supported my activities during my time there.

*Alexander Jordan  
Vienna University of Technology, Austria*



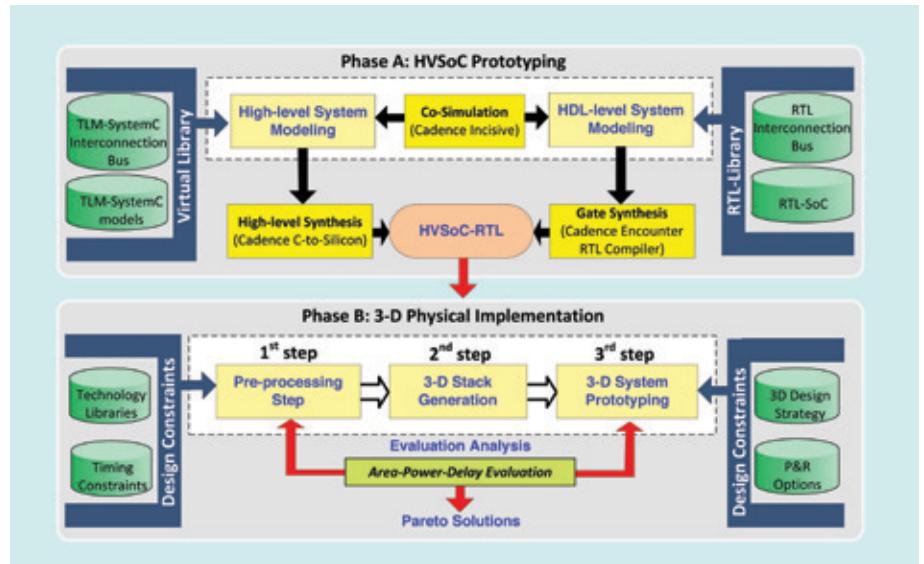
## RAPID PROTOTYPING OF 3-D HYBRID VIRTUAL SYSTEM-ON-CHIPS

HiPEAC student Dionyssios Diamantopoulos from NTU Athens visits Cadence's office in Munich as a finalist of their sponsored thesis contest for automotive embedded systems

I am an MSc student at Microlab, School of Electrical and Computer Engineering at National Technical University of Athens, Greece, where I work under the supervision of HiPEAC member Prof. Dimitrios Soudris. Over the past years, Cadence Design Systems has successfully sponsored several student design contests. This article describes my submission to the thesis contest announced for this year, in

the topic of automotive embedded systems. My submission was part of my master thesis.

In the automotive domain, there is a continuing increase in developing safety critical solutions. Such a requirement imposes that the development of distinct components cannot be thought affordable for the automotive industry, whereas a more holistic approach is necessary for



*The methodology used for rapid prototyping.*

deriving optimal solutions. At the same time, the requirement for integrating more functionality in a smaller form factor pushes traditional semiconductor technology scaling to its limits. Three-dimensional (3-D) chip stacking is touted as the silver bullet technology that can keep Moore's momentum and fuel the next wave of consumer electronic products. This work aims to introduce a new design paradigm, named Hybrid Virtual System-on-Chip (HVSoC), in order to support rapid evaluation of different technologies for IC product development. Our framework initiates from SystemC, whereas the target architecture consists of a 3-D chip. Rather than similar approaches, which mainly are based on academic tools, the 3-D HVSoCs is evaluated with the Cadence tools.

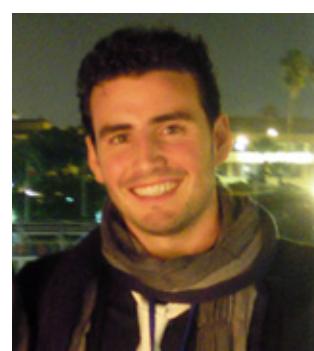
The proposed framework consists of two distinct design phases: a) HVSoC prototyping and b) 3-D physical implementation. The former deals with the modeling of a System-on-Chip using transaction-accurate SystemC modules, while the latter supports the implementation of the derived system onto silicon using 3-D technology. The supporting tool-flow for HVSoC platforms is currently based on the Cadence tools, C-to-Silicon HLS, which deals with the optimized high-level synthesis process of SystemC-described modules, and Incisive Simulator, which simultaneously offers simulation of transaction-level models and RTL descriptions. The second phase affects the architecture exploration and implementation of RTL SoC, produced by the first

step, to 3-D design, using TSV technology. By creating virtual metal layers and extra metal layers for TSV's into the physical libraries, our methodology rapidly explores the implementation cost of the 3-D IC, using the Cadence SoC Encounter tool suite and new academic tools.

After my thesis was accepted among the five top ranked submissions, I was invited by Cadence to perform a 60 minutes presentation of my work. During my visit at the Cadence center in Munich, I had the opportunity to meet outstanding professionals and receive constructive comments and recommendations about my ongoing and future work. The evaluation of the whole event was excellent. I want to thank Cadence for offering me the opportunity of stimulating discussion of research ideas with industry. The final contest results will be announced at CDNLive! EMEA 2013.

I strongly encourage HiPEAC students to participate in similar competitive events.

*Dionyssios Diamantopoulos, ECE, National Technical University of Athens, Greece*



## T-CREST PROJECT: TIME-PREDICTABLE MULTI-CORE ARCHITECTURE FOR EMBEDDED SYSTEMS

It is Time for a Computing Platform that is Time-predictable

**Project title:**

T-CREST (Time-predictable Multi-Core Architecture for Embedded Systems)

**Project coordinator:**

Scott Hansen, The Open Group - X/Open Company

**Technical Lead:**

Martin Schoeberl, Technical University of Denmark

**Partners:**

- Technical University of Denmark
- AbsInt Angewandte Informatik GmbH
- GMV
- Intecs SpA
- Eindhoven University of Technology
- Vienna University of Technology
- The Open Group - X/Open Company
- University of York

**Project website:**

<http://www.t-crest.org/>

**Project start date:**

September 2011

**Duration:**

36 months

Safety-critical systems are important parts of our daily life. Those systems are also called dependable systems, as our lives can depend on them. Examples are controllers in an airplane, braking controllers in a car, or a train control system. Those safety-critical systems need to be certified, and the maximum execution time needs to be bounded and known, so that response times can be assured when critical actions are needed. Note that just using a faster processor is not a solution for time predictability. Even with high-performance processors in our desktop PCs, we notice once in a while that the PC is "frozen" for a few seconds. For word processing we accept this minor inconvenience, but for a safety-critical system such a "pause" can result in a catastrophic failure.

The mission of T-CREST is to develop tools and build a system that prevents pauses by identifying and addressing the causes for possible pauses. The T-CREST time-predictable system will simplify the safety argument with respect to maximum execution time.

Thus the T-CREST system will result in lower costs for safety-relevant applications reducing system complexity and at the same time faster time-predictable execution.

architectures with a new paradigm: make the worst-case fast and the whole system easy to analyse.

Within T-CREST we will propose novel solutions for time-predictable multi-core and many-core system architectures. The resulting time-predictable resources (processor, interconnect, memories, etc.) will be a good target for WCET analysis, and the WCET performance will be outstanding compared to current processors. Time-predictable caching and time-predictable chip-multiprocessing (CMP) will provide a solution for the need of more processing power in the real-time domain. In addition to the hardware (processor, interconnect, memories), the project will develop a compiler infrastructure. WCET-aware optimization methods will be developed, along with detailed timing models, such that the compiler benefits from the known behaviour of the hardware. The WCET analysis tool aiT will be adapted to support the developed hardware and guide the compilation.

T-CREST will be evaluated by use cases from two industrial partners: GMV and INTECS. The first use case is in avionics, with several applications developed for the aeronautical domain. The second use case is from the railway domain that is based on a monitoring system for GSM-R (GSM railway) radio link.

The T-CREST hardware and compiler will be open-source under the simplified BSD license, and they are hosted on GitHub: <https://github.com/t-crest>



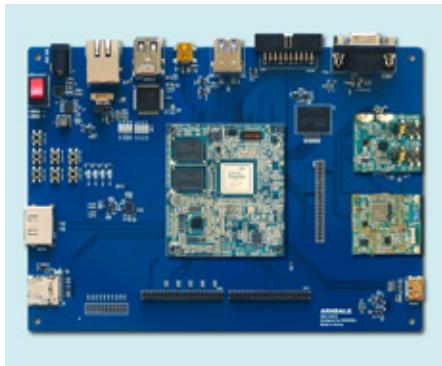
### TIME-PREDICTABLE COMPUTING

Standard computer architecture is driven by the following paradigm: make the common case fast and the uncommon case correct. This design approach leads to architectures where the average-case execution time is optimized at the expense of the worst-case execution time (WCET). Modelling the dynamic features of current processors, memories, and interconnects for WCET analysis often results in computationally infeasible problems. The bounds calculated by the analysis are thus overly conservative.

We need a sea change and we shall take the constructive approach by designing computer architectures, where predictable timing is a first-order design factor. For real-time systems we propose to design

## MONT-BLANC PROJECT SELECTS SAMSUNG EXYNOS 5 PROCESSOR

The project continues its research effort towards an energy efficient HPC prototype using low-power embedded technology



The Mont-Blanc European project has selected the Samsung Exynos platform as the building block for powering its first integrated low power- High Performance Computing (HPC) prototype. The aim of the Mont-Blanc project is to design a new type of computer architecture capable of setting future global HPC standards, built from today's energy efficient solutions used in embedded and mobile devices. Recently, the project has selected the Samsung Exynos platform as the building block for powering its first integrated low power- High Performance Computing (HPC) prototype. The Samsung Exynos 5 Dual is built on 32nm low-power HKMG (High-K Metal Gate), and features a dual-core 1.7GHz mobile CPU built on ARM®

Cortex™-A15 architecture plus an integrated ARM Mali™-T604 GPU for increased performance density and energy efficiency. It has been featured and market proven in consumer and mobile devices such as Samsung Chromebook and Google's Nexus 10. This will be the first use of an embedded mobile SoC in HPC, which enables the Mont-Blanc project to explore the challenges and benefits of deeply integrated energy-efficient processors and GPU accelerators, compared to traditional homogeneous multicore systems, and heterogeneous CPU + external GPU architectures.

"The Exynos 5 Dual packs the most powerful ARM processors with a programmable GPU in a low-power mobile device that would normally be in someone's pocket and running on a battery. Its performance density, energy efficiency, and low market price make it an extraordinary building block for prototyping a new generation of HPC systems." says Alex Ramirez, coordinator of the Mont-Blanc project.

During the first year of activities, Mont-Blanc has focused on deploying successfully an HPC system software stack and full-scale scientific applications on ARM platforms, proving that ARM-based archi-

tures are feasible alternatives for HPC. Now the efforts gear towards integration of the Exynos platform in a HPC solution, and software exploitation of the embedded GPU.

The Mont-Blanc project brings together a purely European consortium federating industrial technology providers and research supercomputing centres: Bull, as the major HPC system vendor, ARM, as the world leader in embedded high-performance processors, and Gnodal, as interconnect partner that focuses its new product on scalability and power efficiency. Besides the technology providers, Mont-Blanc brings together the supercomputing centres from the four Tier-0 hosting partners in PRACE with leading roles in system software and Exascale application development: Germany (Forschungszentrum Jülich, BADW-LRZ), France (GENCI, CNRS), Italy (CINECA), and Spain (BSC).

For further information:  
[www.montblanc-project.eu](http://www.montblanc-project.eu)

## THERMAL CHARACTERIZATION AND OPTIMIZATION OF SYSTEMS-ON-CHIP THROUGH FPGA-BASED EMULATION

**By Pablo García del Valle**  
Universidad Complutense de Madrid, Spain  
**Advisors:** Prof. David Atienza Alonso and Prof. José Manuel Mendías Cuadros  
June, 2012

Tablets and smartphones dominate the consumer electronics market. Internally, they rely on Systems on Chip (SoCs) to meet the tight design constraints: performance, size, power consumption, etc. In a bad design, the high logic density may generate hot-spots that compromise the chip reliability. Pablo's thesis introduces his FPGA-based emulation framework for easy exploration of SoC design alternatives. It provides fast and accurate estimations of performance,

power, temperature, and reliability in one unified flow, to help designers tune their system architecture before going to silicon. Pablo's expertise is focused on embedded devices: Microcontrollers, FPGAs, Real-time OSs and drivers.



### ADAPTING THE POLYTOPE MODEL FOR DYNAMIC AND SPECULATIVE PARALLELIZATION

By Alexandra Jimborean,  
Université de Strasbourg, France  
Advisors: Prof. Philippe Clauss, and  
Assist. Prof. Vincent Loechner  
September, 2012



We present a Thread-Level Speculation (TLS) framework designed to speculatively parallelize a sequential loop nest, by re-scheduling its iterations at runtime. We perform code transformations by applying the polyhedral model that we adapted for speculative and dynamic code parallelization. Hence, we build a parallel code pattern at compile time, which is patched by our runtime system, guided by online profiling. The system discovers new parallelization opportunities,

enables partial parallelism and performs optimizations prior to parallelization. We show on several benchmarks that our framework yields good performance on codes that could not be handled efficiently by previous TLS systems.

### OPTIMIZING SIGNATURES IN HARDWARE TRANSACTIONAL MEMORY SYSTEMS

By Ricardo Quislant  
Universidad de Málaga, Spain  
Advisors: Prof. Oscar Plata, Prof.  
Eladio Gutierrez and Prof. Emilio L.  
Zapata  
October, 2012



This thesis proposes signature optimizations in the context of hardware conflict detection for Transactional Memory systems. We focus on exploiting both spatial locality of reference and asymmetry of transactional data sets. Our Interval Filter and Locality-Sensitive Signatures harness locality to efficiently maintain transactional data sets in a more concise way. Multiset and Reconfigurable Asymmetric Signatures deal with the fact that read and write sets of

transactions show uneven cardinalities, so the amount of storage devoted to each set is not known in advance. Results show a significant performance improvement and an increased scalability over conventional signatures.

### FLOORPLAN-AWARE HIGH PERFORMANCE NOC DESIGN

By Antoni Roca Pérez  
Universitat Politècnica de  
València, Spain  
Advisors: Prof. José Flich Cardo  
and Prof. Federico Silla  
October, 2012



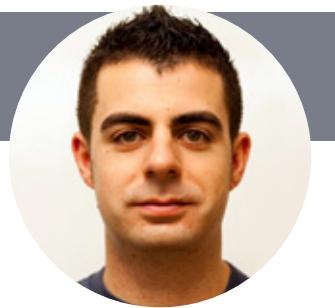
We present a floorplan-aware NoC design methodology that minimizes the implementation drawbacks of complex NoC designs, as high-radix switches and long links. The network is made of blocks that have their own buffering, arbitration, and crossing capabilities. Then, we make their placement flexible, allowing a distribution of the blocks all over the chip. By using our NoC design methodology, we implement a distributed crossbar that clearly outperforms in terms of latency and throughput

the rest of configurations. For very large networks, our crossbar design is unaffordable. To reduce resource requirements, a hierarchical approach is required, which keeps almost all the crossbar benefits.

## DESIGN TECHNIQUES FOR SMART AND ENERGY-EFFICIENT WIRELESS BODY SENSOR NETWORKS

By Francisco Rincon  
Universidad Complutense de Madrid, Spain  
Advisors: Nadia Khaled, David Atienza, Marcos Sanchez-Elez  
October, 2012

Wireless Body Sensor Networks (WBSNs) for health monitoring and diagnosis are gaining popularity and will deeply change healthcare delivery in the next years. The use of these networks enables continuous biomedical monitoring and care, helping the prevention and early diagnosis of diseases, while enhancing the patient's autonomy. This thesis proposes the development of a WBSN-based long-lived wireless ECG monitoring system that features on-board real-time analysis and



diagnosis, instantly providing information about the operation of the heart. Moreover, this work seeks to improve WBSN energy efficiency, proposing a set of design techniques to reduce energy consumption.

## EFFICIENT SYNCHRONIZATION AND COMMUNICATION IN MANY-CORE CHIP MULTIPROCESSORS

By José L. Abellán  
Universidad de Murcia, Spain  
Advisors: Prof. Juan Fernández and Prof. Manuel E. Acacio  
December, 2012

Most of today's chip-multiprocessors (CMPs) rely on an intuitive shared memory model. As the core count increases in these so-called many-core CMPs, architectural innovations are essential to keep up with high performance. In this thesis, we propose distinct and complementary HW-based solutions to overcome three fundamental performance bottlenecks in many-core CMPs: highly contended synchronization in the context of barriers and locks, and the maintenance of cache coherence. Through



both a novel full-custom technology and a current standard cell design methodology, we implement and evaluate our proposals. Our experimentation reveals important performance advantages to overcome such obstacles.

## FROM PARALLEL PROGRAMS TO CUSTOMIZED PARALLEL PROCESSORS

By Pekka Jääskeläinen  
Tampere University of Technology, Finland  
Advisor: Prof. Jarmo Takala  
December, 2012

This thesis proposes a novel Multicore Application-Specific Instruction Set Processor (MCASIP) co-design methodology that exploits parallel programming languages as the application input format. In the methodology, the designer can explicitly capture the parallelism of the algorithm and exploit specialized instructions using a parallel programming language, in contrast to relying on the compiler or the hardware to extract the parallelism from a sequential input. The



thesis proposes a multicore processor template based on the Transport Triggered Architecture, compiler techniques involved in static parallelization of computation kernels with barriers and a datapath integrated lock unit.

## UPCOMING EVENTS

### **THE 26TH INTERNATIONAL CONFERENCE ON ARCHITECTURE OF COMPUTING SYSTEMS (ARCS 2013)**

19-22 February 2013, Prague, Czech Republic, <http://arcs2013.fit.cvut.cz/>

### **THE 19TH IEEE INTERNATIONAL SYMPOSIUM ON HIGH-PERFORMANCE COMPUTER ARCHITECTURE (HPCA-19)**

23-27 February 2013, Shenzhen, China, <http://www.hpcaconf.org>

### **THE 18TH ACM SIGPLAN SYMPOSIUM ON PRINCIPLES AND PRACTICE OF PARALLEL PROGRAMMING (PPoPP 2013)**

23-27 February 2013, Shenzhen, China, <http://ppopp2013.ics.uci.edu/>

### **THE INTERNATIONAL SYMPOSIUM ON CODE GENERATION AND OPTIMIZATION (CGO 2013)**

23-27 February 2013, Shenzhen, China, <http://www.cgo.org/cgo2013/>

### **THE 16TH DESIGN, AUTOMATION AND TEST IN EUROPE CONFERENCE (DATE 2013),**

18-22 March 2013, Grenoble, France, <http://www.date-conference.com/>

### **THE 18TH INTERNATIONAL CONFERENCE ON ARCHITECTURAL SUPPORT FOR PROGRAMMING LANGUAGES AND OPERATING SYSTEMS (ASPLOS 2013)**

16-20 March 2013, Houston, USA, <http://asplos13.rice.edu/>

### **THE INTERNATIONAL CONFERENCE ON COMPILER CONSTRUCTION (CC 2013)**

16-24 March 2013, Rome, Italy, <http://www.etaps.org/2013/cc13>

### **THE IEEE INTERNATIONAL SYMPOSIUM ON PERFORMANCE ANALYSIS OF SYSTEMS AND SOFTWARE (ISPASS 2013)**

21-23 April 2013, Austin, USA <http://ispass.org/ispass2013/>

### **THE EUROPEAN CONFERENCE ON COMPUTER SYSTEMS (EUROSys 2013)**

15-17 April 2013, Prague, Czech Republic, <http://eurosys2013.tudos.org/>

### **18TH INTERNATIONAL CONFERENCE ON RELIABLE SOFTWARE TECHNOLOGIES (ADA-EUROPE 2013)**

10-14 June 2013, Berlin, Germany, <http://www.ada-europe2013.org/>

**contributions** If you are a HiPEAC member and would like to contribute to future HiPEAC newsletters, please visit <http://www.hipeac.net/hipeacinfo>



HiPEAC INFO IS A QUARTERLY NEWSLETTER PUBLISHED BY THE HiPEAC NETWORK OF EXCELLENCE, FUNDED BY THE 7TH EUROPEAN FRAMEWORK PROGRAMME (FP7) UNDER CONTRACT NO. FP7/ICT 287759  
WEBSITE: [HTTP://WWW.HIPEAC.NET](http://www.hipeac.net).  
SUBSCRIPTIONS: [HTTP://WWW.HIPEAC.NET/NEWSLETTER](http://www.hipeac.net/newsletter)

**9TH INTERNATIONAL SUMMER SCHOOL ON ADVANCED COMPUTER ARCHITECTURE AND COMPILEMENT FOR HIGH-PERFORMANCE AND EMBEDDED SYSTEMS (ACACES 2013), 14-20 JULY 2013, FIUGGI, ITALY**