



Information Society  
Technologies



# HiPEAC *info* <sup>15</sup>

COMPILATION ARCHITECTURE

appears quarterly | July 2008

Network of Excellence on High Performance and Embedded Architecture and Compilation

- 2 Message from the HiPEAC coordinators
- 2 Message from the project officer
- 3 Community News
- 4 In the Spotlight
  - EC FP7 STREP Project VELOX
  - BSC and Microsoft Create Joint Research Centre in Barcelona for Parallel Computing
- 7 Guest Column
  - SHAPES and its DIOPSIS Tile
  - Tsuyoshi Isshiki, Tokyo Institute of Technology
- 10 HiPEAC Activity
  - HiPEAC Research Collaborations
- 11 HiPEAC Start-ups
- 12 PHD News
- 15 HiPEAC Students
- 16 Upcoming Events



Welcome to L'Aquila, Italy  
ACACES Summer School 2008

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

Second HiPEAC Computing Systems Week  
Paris, France, November 26-28, 2008

## Message from the HiPEAC coordinators

Dear friends,

Several of you will find this issue of HiPEACinfo when entering your room at the ACACES summer school. This is the fourth time that HiPEAC has organized its yearly summer school. Over these four years, more than 750 participants have enjoyed the classes and informal interactions with about 50 teachers. The number of applicants keeps increasing year upon year: this year, we received more than 370 applications for just 200 places. We know from past evaluations of the summer school by students and teachers that the level of appreciation among participants is very high. Organizing the summer school consumes a lot of resources from the HiPEAC network, but it is certainly worth the effort!



Mateo Valero



Koen De Bosschere

Recently, in June, we enjoyed the Spring Computing Systems Week in beautiful Barcelona, hosted by the BSC-Microsoft Research Center, UPC and HP Labs. It was a week full of exciting research meetings, high-quality events, and lots of opportunities to network. The event attracted almost 240 participants from all over Europe, and beyond. We would like to thank everybody involved in the organization of the Computing Systems Week for all the energy they devoted to making this a successful event.

The next Computing Systems Week will be hosted by Thales Paris in the week of November 24, 2008. You are

all cordially invited to join us there.

Six months since the start of HiPEAC2, the clusters and task forces are showing a healthy level of activity, and have begun to define their own research agenda. The first collaborations have now been funded, and several consortia are forming in preparation for the November 2008 FP7 call.

Another piece of positive information: the number of registered people on our database has been steadily growing. This summer, we expect to hit the 500-mark. We very much value that our members identify their PhD students and affiliated members, and

add them to the database. Since all registered users bring their own expertise, the HiPEAC database is gradually becoming an up-to-date database of European expertise in computing systems. Companies and universities looking for new European talent should definitely start their search at the HiPEAC website where they will find 250 PhD students.

Let us end by wishing all of you a pleasant summer holiday and a good start after the summer.

Koen De Bosschere and Mateo Valero

## Message from the project officer

Eight collaborative research projects have started recently as a result of the 2007 Call in Computing Systems. These eight projects, together with the HiPEAC2 Network of Excellence, represent 25 million euros of funding from the European Commission over the next few years. At the last Computing Frontiers conference, a special session was dedicated to these new projects. The project presentations will be available from the Events section of the Cordis Computing web page. <http://cordis.europa.eu/fp7/ict/computing/>

**Apple-CORE:** This project will develop compilers, operating systems and execu-

tion platforms for many-core computing systems. The project has a disruptive approach adopting a systematic model of concurrency implemented in the processor's ISA.

[www.apple-core.info](http://www.apple-core.info)

**CRISP:** This project will build a general-purpose, low-power, stream-processor chip that can be reconfigured dynamically. The chip includes radio-frequency interface modules, and it will be tested in streaming applications (beamforming and Galileo/GPS navigation).

[www.crisp-project.eu](http://www.crisp-project.eu)

**GENESYS:** This project will develop a

reference architecture for embedded systems, applicable across a wide range of domains with

a uniform architectural style. Four prototypes from various domains will help to evaluate the feasibility of the architectural concepts.

[www.genesys-platform.eu](http://www.genesys-platform.eu)

**ICT-eMuCo:** This project focuses on the platform architecture of future mobile devices. The project addresses the relevant controller elements as well as the



Panos Tsarchopoulos  
[Panagiotis.Tsarchopoulos@ec.europa.eu](mailto:Panagiotis.Tsarchopoulos@ec.europa.eu)

## HiPEAC Computing Systems Week

In the first week of June (from 2nd to 6th) experts from all over the world in multi-core architecture and programming attended the annual HiPEAC Computing Systems Week at the Technical University of Catalonia (UPC). During the entire week, approximately 240 delegates enjoyed the high-quality sessions, hosted by UPC and the Barcelona Supercomputing Center (BSC), as well as the chance to engage in extensive networking.

The program for the whole week included the following:

- Cluster meetings on Monday and Tuesday
- the Industrial Workshop at HP Labs in Sant Cugat del Vallès, bringing together researchers from the academia and industry to present their communications and discuss on methodologies, tools and compiler optimizations for multicores and distributed memory systems
- the Barcelona Multi-core Workshop (BMW'08) on Thursday and Friday, sponsored by the BSC-Microsoft Research Center. The workshop consisted of 14 invited talks and 2 panel sessions, discussing the challenges raised by the multi/many-core architectures of the future.



Furthermore, on Monday and Thursday all delegates had the opportunity to visit the supercomputer MareNostrum installed in the chapel of Torre Girona at the University Campus of Catalonia.

"This week has been a good opportunity to gather together the most prestigious experts in Barcelona to exchange future ideas and knowledge, and to discuss the state-of-the-art in computer architecture," said Jesús Labarta, Computer Sciences Director of BSC.

"Barcelona has been a preeminent intellectual center in computer architecture for some time, which is why so many companies invest in centers here. This week demonstrated once again the strength and breadth of the research going on here," said Doug Burger, Principal Researcher at Microsoft and Manager of the Computer Architecture Research Group.

operating system and application layers. A multi-core architecture will be the basis of a trial system in order to achieve the requirements of performance and low power-consumption with flexibility and scalability.

[www.emuco.eu](http://www.emuco.eu)

**JEOPARD:** This project will develop a platform-independent, software-development interface for complex multi-core systems. The interface extends existing Java technologies that do not currently provide support for multi-core systems. A real-time Java implementation will maintain predictable execution on multiple parallel processors with real-time memory management including garbage collection.

[www.jeopard.org](http://www.jeopard.org)

**MERASA:** This project will develop and prototype multi-core designs from 2 to 16 cores for hard real-time embedded applications, together with timing analysis techniques and tools to guarantee the analysability and predictability of every component feature. The hardware and software techniques developed will be evaluated by application studies in aerospace, automotive and construction machinery.

[www.merasa.org](http://www.merasa.org)

**MOSART:** This project develops a flexible, multi-core, on-chip platform, design methods and tools with a NoC backbone to allow the scaling and optimisation of various applications in multimedia and wireless communications. The project addresses the challenges of program-

ming heterogeneous, multi-core architectures and overcoming the interconnect and memory bottlenecks.

[www.mosart-project.org](http://www.mosart-project.org)

**VELOX:** This project will provide coordinated hardware and software support for the transactional memory programming paradigm by developing an integrated, transactional memory stack spanning from the underlying hardware to the high-end application. The results will be demonstrated using software and FPGA-based simulators.

[www.velox-project.eu](http://www.velox-project.eu)

**Panos Tsarchopoulos**

Project Officer

# EC FP7 STREP Project VELOX: An Integrated Approach to Transactional Memory on Multi- and Many-core Computers



A huge challenge facing the computing community today is how to make programming multi-cores easier. With this goal in mind, the VELOX project, funded by the European Commission with € 4 million, launched its activities at the beginning of this year. The objective of VELOX is to deliver seamless transactional memory (TM) systems that integrate well at all levels of the system stack. This three-year project hopes to obtain important research breakthroughs that will enable Europe to lead in integrated TM systems.

The adoption of multi-core chips as the architecture-of-choice for mainstream computing will undoubtedly bring about profound changes in the way software is developed. In this brave new era, programs will need to be rewritten in a parallel way for computers that have multiple processing cores. One of the fundamental issues in developing parallel programs is a coordinated and orderly way of accessing shared data. The use of previous techniques, such as fine-grained locking as the multi-core

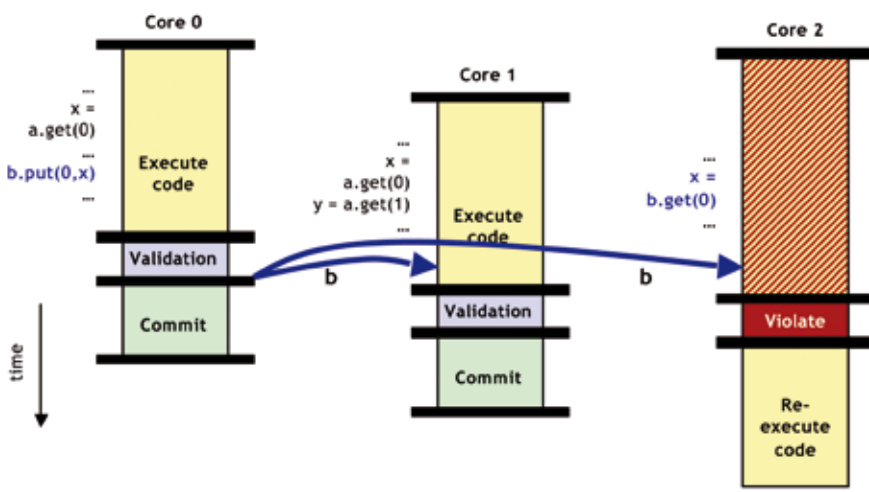
programmer's coordination methodology is viewed by most experts as a dead-end since locking is too complicated for the average programmer.

The TM programming paradigm is a strong contender to become the approach of choice for replacing those coordination techniques and implementing atomic operations in concurrent programming. Combining sequences of concurrent operations into atomic transactions promises a great reduction in the complexity of both programming and verification by making parts of the code appear to be sequential without the need to program fine-grained locks. Transactions remove from the programmer the burden of figuring out the interaction among concurrent operations that happen to conflict when accessing the same locations in memory.

To make TM an effective tool, TM systems will need the right hardware and software support to provide scalability not only in terms of number of cores, but also in terms of code size and complexity. The objective of the VELOX project is to understand how to provide such support by developing an integrated TM stack. Such a TM stack would span a system from the

underlying hardware to the high-end application and would consist of the following components: CPU, operating system, runtime, libraries, compilers, programming languages and application environments. The team includes internationally recognized TM experts in each of those components. These fully integrated TM systems will not only improve the understanding of TM designs, but will greatly help in the adoption of the TM paradigm by the European software industry, making it a tool-of-choice for concurrent programming on multi-core platforms.

The consortium includes top research and system integration organizations with non-overlapping skills in the areas of computer architecture, operating systems, compiler and runtime systems, programming language design, programming models and application and benchmark design. Coordinated by the Barcelona Supercomputing Center, the VELOX consortium gathers nine different partners that include top research and system integration organizations such as the University of Neuchâtel, the Technische Universität Dresden, Ecole Polytechnique Fédérale de Lausanne, Tel Aviv University, Chalmers University of Technology as well as leading integrators from the IT industry such as AMD, Red Hat and VirtualLogix SAS. Microsoft, Intel and SUN are the official observers to VELOX. The project has five technical workpackages: Multi-Core Architectures for TM (leader: Adrian Cristal from Barcelona Supercomputing Center), Operating Environments for TM (leader: Christof Fetzer from Technische Universität Dresden), TM Runtime (leader: Nir Shavit from Tel Aviv University), Language Extensions and Integration for TM Programming (leader: Pascal Felber from University of Neuchâtel) and TM Benchmark and Applications (leader: Rachid Guerraoui from Ecole Polytechnique Fédérale de Lausanne).



**Transactional memory synchronizes threads optimistically. Core 0 and core 1 read the same memory location but do not conflict; hence both can commit. Core 0 writes to a memory location read by core 2, which creates a conflict; therefore, one of them (core 2) must abort and restart its execution.**

Many of the project members are already leaders in Transactional Memory. The members had important contributions

## In the Spotlight

that are considered to be important TM milestones. Examples include the first Software Transactional Memory (Tel Aviv University), the first TM benchmark suite (Ecole Polytechnique Fédérale de Lausanne), the first time-based Software Transactional Memory (Technische Universität Dresden and University of Neuchâtel), and the first fully-adaptive Hardware Transactional Memory (Barcelona Supercomputing Center and Chalmers University of Technology). Individually, they cannot carry such an ambitious endeavor because it requires significant resources and, most importantly, multiple complementary competencies. Therefore, a pan-European

approach was necessary to carry out this project.

"Thanks to the complementary skills of its partners, it will pave the way for key European researchers to make significant contributions to the ongoing revolution to make parallel programming easier for the masses", says Osman Unsal, coordinator of the VELOX project. Mateo Valero, director of BSC, stressed that "the VELOX project is crucial to enable the supercomputing applications of today to run on the laptops of the near future."



Osman Unsal  
Barcelona Super Computing Center  
Osman.unsal@bsc.es

## Community News

### The 2008 Aragón Award



*Mateo Valero receiving the Aragón 2008 award by Marcelino Iglesias, President of the Autonomous Community Government of Aragón (Spain)*

On 23rd April (Saint George's day in the Autonomous Community of Aragón, Spain), Prof. Mateo Valero received the most important honour of the region of Aragón, the 2008 Aragón award - consisting of a diploma and a statue. The Regional Government of Aragón awarded the prize in recognition of Prof. Mateo Valero's and his PhD students' dedication to science over the

years, as well as to Prof. Valero's contribution to society thanks to his scientific findings.

The President of the Autonomous Community, Mr. Marcelino Iglesias, presented Prof. Valero with the prize and thanked him for his great and continuous dedication to science. Mr. Iglesias emphasized Prof. Mateo Valero's scien-

tific and research contribution to society and the relevant progress made in the field of computer architecture.

Prof. Mateo Valero was very honoured to receive this award from his colleagues and friends from Aragón and emphasized his desire to continue to work hard in both research and science, the two fields he has always been devoted to.

# Barcelona Supercomputing Center and Microsoft Create Joint Research Centre in Barcelona for Parallel Computing

**Centre seeks solutions to challenges and opportunities associated with multi-core processors.**

BARCELONA — 18 Jan 2008 — During a press conference at the Universitat Politècnica de Catalunya in Barcelona, the Barcelona Supercomputing Center – Centro Nacional de Supercomputación (BSC) and Microsoft Corp announced the creation of the BSC-Microsoft Research Centre, which will focus on the way microprocessors and software for the mobile and desktop market segments will be designed and interact over the next 10 years and beyond. The advent of many- and multi-core processor computing architectures will make it possible to deliver enormous computational power on a single chip with profound implications for the way software is developed. Optimizing the design and interaction of hardware and software architectures to take advantage of the new computing power will require tight integration across the industry.

Computer architecture experts at BSC have teamed up with computer scientists at Microsoft Research Cambridge (MSRC) in the United Kingdom to look for innovative solutions to the challenges and opportunities that parallel processing represents. The vision of the centre is of a top-down computer architecture in which software requirements drive the hardware innovation forward rather than letting the hardware design condition software development. In addition to fundamental and applied research in transactional memories, a promising technology that facilitates writing of parallel programs for multi-core processors, hardware support for managed runtimes will be conducted in the initial research projects.

The official inauguration of the centre was attended by Josep Huguet, Minister of Innovation, Universities and Enterprise of the Generalitat de Catalunya; Prof Antoni



Giró, Rector of the Technical University of Catalonia; Prof Mateo Valero, BSC's director; Dr Andrew Herbert, managing director of the Microsoft Research Laboratory in Cambridge; Tony Hey, corporate vice president of External Research in Microsoft Research; and Rosa Garcia, president of Microsoft Spain.

"Two years after the initial agreement, we are reaffirming our research commitment by establishing a research centre in Barcelona, building upon the successful collaboration of a group of researchers from the Computer Architecture team in BSC and Microsoft Research. To optimize the designs and interactions of multi-core processors and software, we need to start from parallel programming. The way to deal with this multi-core architecture challenge is to bring together computer architects and programming language experts," stressed Mateo Valero, director of BSC.

"We are pleased to partner with the Barcelona Supercomputing Center to create this new research centre," said Tony Hey, corporate vice president of External Research in Microsoft Research. "Partnerships like this help us to reach our goal of supporting the global research community and ultimately assisting researchers and scientists to address


some of the toughest, most urgent societal and technological challenges worldwide. Microsoft Research has a strong track record of collaboration with public research centres and academic institutes to advance the state of the art in computing. The BSC-Microsoft Research Centre is our most recent step, combining our respective areas of expertise to address the challenges of writing trustworthy software for multi-core processors."

"At Microsoft we are committed to joining and building partnerships between businesses, research centres, education institutions and governments to promote Knowledge Society across Europe and to seeing Spain and Catalonia's excellence in research and innovation achieve the worldwide recognition it deserves through increased investment and commercial collaboration," explained Rosa García, president of Microsoft Spain.

# SHAPES and its DIOPSIS Tile

I am honored to be asked to contribute a guest column to the HiPEAC newsletter about the SHAPES tiled architecture and its DIOPSIS building block. Modern MPSoCs must integrate hundreds of million-gates to satisfy high performance demands from embedded applications. Tiled architectures provide a scalable design style addressing two key issues: global wiring and design complexity on deep submicron technologies. The usage of a set of "small" mesochronous processing tiles (a few million gates each) with "short" intra-tile wires solves the clocking issue. The replication of stable and validated processors and tiles maintains a manageable complexity at system level.

SHAPES (Scalable Software and Hardware Architecture Platform for Embedded Systems) engines are based on the replication of multi-processor heterogeneous tiles, derived from the DIOPSIS MPSoC, the RISC + floating-point GigaFlops DSP System-on-Chip designed by Atmel Roma. European researchers can contact me (pier.paolucci@atmel.com) to get access to a DIOPSIS Board and its development environment, one of the first public deliverables of the SHAPES project. The DIOPSIS board is a superb environment to experiment advanced compilation techniques, like those explored by the HiPEAC community. Using DIOPSIS, you can apply those techniques to real-world applications, especially in the multi-microphone, multi-loudspeaker sound, acoustic, and speech-processing domains, and execute it on the heterogeneous RISC + floating point DSP DIOPSIS MPSoC. Starting from DIOPSIS, the SHAPES elementary tile is created adding the DNP, a Network Processor designed by INFN (Istituto Nazionale di Fisica Nucleare). A team of the "Sapienza" University of Roma is taking care of issues on deep submicron technology (45 nm). A typical tile in a SHAPES system typically includes: a VLIW floating-point DSP (mAgicV GigaFlops Digital Signal Processor from ATMEL Roma), a RISC controller (ARM926), and a DNP (Distributed Network Processor from INFN). Each tile can be individually connected with a DXM (Distributed External Memory) and a set of peripherals (e.g. ADC/DAC) and always includes sev-



eral banks of intra-tile memory. A routing fabric connects on-chip and off-chip tiles, weaving a distributed-packet, switching network. The design of the Network-on-Chip is done by STM, Cagliari and Pisa Universities, while 3D next-neighbours toroidal connections are adopted for off-chip networking and maximum system density, leveraging on the INFN experience gained thanks to the previous designs of several generations of APE Massive Parallel Machines for theoretical physics computations. The SW challenge is to provide an efficient programming and simulation environment managing four levels of parallelism:

- **Level 1:** Inside each processor: e.g. exploiting the 15 operations/clock cycle of the multiple-issue mAgicV floating-point VLIW gigaflops DSP, and its software-managed, multi-bank memory system;
- **Level 2:** Inside each tile, e.g. activating the parallel operation of the RISC, DSP and DNP processors inside the tile, supported by a SW managed multi-layer bus matrix, which permits multiple simultaneous intra-tile data transfers between the intra-tile processors, the peripherals and the Distributed External Memory attached to each tile.
- **Level 3:** Inside each multi-tile chip, managing the inter-tile DNP-NoC-DNP packet transfers, for a number of

simultaneous data transfers proportional to the number of tiles;

- **Level 4:** At multi-chip system level, using the 3D toroidal off chip Network.

In short, RWTH-AACHEN provides the multi-tile simulator. The application is described as a "network of processes". The best mapping of the application on the multi-tile architecture is automatically discovered and the generation of the executable, including supporting, OS is fully automated. ETH Zurich provides the optimizer for automatic mapping (i.e. binding and scheduling). TIMA Lab Grenoble provides a system for automatic generation of OS support. TARGET Compiler technologies the VLIW C compiler. SHAPES is benchmarked on Acoustic Wave Field Synthesis by Fraunhofer IDMT, on Ultrasound Scanners by ESAOTE, MedCom and Fraunhofer IGD and Lattice QCD by INFN. SHAPES is a Future Emerging Technologies – Advanced Computer Architectures FP6 Integrated Project (2006-2009). See [www.shapes-p.org](http://www.shapes-p.org) for details. ■

Pier Stanislaio Paolucci  
Technology Director,  
Atmel Roma  
(part-time) researcher,  
INFN



## Tsuyoshi Isshiki, Associate Professor, Tokyo Institute of Technology

It is a great pleasure to have the opportunity to contribute to this guest column and introduce myself, as well as our laboratories at Tokyo Institute of Technology (TITech), to the HiPEAC community. I work very closely with Professor Hiroaki Kunieda and we jointly run our two laboratories at TITech in the Department of Communications and Integrated Systems. Our research activities cover various aspects of SoC and embedded system solutions, focusing on several key application domains, such as video codec systems and fingerprint identification systems. Providing solutions to any application-specific SoCs and embedded systems requires a concerted effort by all design levels, starting from algorithm designs, system architecture designs (SW/HW partitioning, memory architecture, interconnect architecture), down to detailed implementations of SW and HW components. Our efforts have concentrated mainly on algorithm designs on these applications that enable efficient HW implementation (large parallelism, small gate counts, low power) and/or efficient SW implementation (small memory footprint, small code size, fast response), and also on the system architecture designs that enhances testability, design modularization, and ease-of-design integration.

During the course of our research activities, we came to realize the strong


need for a solid design methodology that bridges the gap between algorithm design and system architecture. We therefore started a new project on this issue a few years ago, which has led to the development of our new design framework, which we call the "Tightly-Coupled Thread" model (TCT model).

Our TCT model contains several aspects that we believe could eventually become one of the key solutions to the enormously complex MPSoC designs, including a very simple programming model on C, where the programmer simply inserts "thread-scopes" directly onto the C programs to describe the program partitioning for the MPSoC elements - compiler technology that extracts all dependences between "threads" and automatically inserts communication instructions - and also its MPSoC platform that allows efficient message passing implemented by a dedicated communication module embedded inside each processing element.

We were fortunate enough to start a collaboration with the MAPS project in RWTH Aachen, headed by Professor Rainer Leupers, which was highlighted in the January issue of the HiPEAC newsletter. Despite the MAPS-TCT collaboration still being in its early stages, it has nonetheless already produced promising results, thereby demonstrat-

ing the potential of a new MPSoC design methodology that utilizes the standard C programming on top of a powerful set of tools to aid the programmers to intelligently explore the vast MPSoC design space.

I am very excited to become a new associate member in the Design Methodology and Tools cluster, and I look forward very much to attending many HiPEAC activities to exchange new ideas on MPSoC design issues and in seeking new collaboration opportunities among the HiPEAC members in industry and academia. ■



Tsuyoshi Isshiki,  
Associate Professor, Tokyo Institute of  
Technology

## Community News

### Chalmers University Honorary Doctorates 2008

On 10th May, 2008, Prof. Mateo Valero received the award of Honorary Doctorate of Engineering from the Faculty Board at Chalmers University at the Doctorate Conferment Ceremony. The award was in recognition of his eminent and great

professional achievement in the field of computer architecture - the engineering discipline of designing computers to carry out computations as fast as possible within financial and technological constraints. ■



## 25 Anniversary in UPV Computer Science Awards 2008



Professor Dr. Jose Duato is the recipient of the “25th Anniversary in UPV Computer Science Awards 2008” for his outstanding contribution to the studies of computer science at the Universidad Politecnica de Valencia. Professor Duato has demonstrated his leadership through the creation and consolidation of a strong investigatory group, as well as through the impact and recognition his scientific works have had at national and international level. ■

## Best Paper Award: International Workshop on Applied Reconfigurable Computing

Carlo Galuzzi from the Computer Engineering Lab at Delft University of Technology, the Netherlands received the best paper award at the International Workshop on Applied Reconfigurable Computing, which was held at Imperial College, London in March of this year. The paper entitled “Instruction-Set Extension Problem: A Survey” investigates the issues concerning the customization of an Instruction-Set as a function of the specific requirements of

an application. The paper provides an overview of all the relevant aspects of the problem and compensates for the lack of a general view of the problem in the existing literature.

Carlo Galuzzi and Koen Bertels, The Instruction-Set Extension Problem: A Survey Conference: International Workshop on Applied Reconfigurable Computing (ARC) pp. 209-220, March 2008, Imperial College, London (UK) LNCS Series ■



## Prestigious Royal Academy of Engineering / EPSRC Research Fellowship Grant




Timothy Jones has been awarded five years of funding from the Royal Academy of Engineering (RAEng) and EPSRC in the UK to investigate energy efficiency in the next generation of microprocessors. The RAEng / EPSRC Research Fellowships are awarded to outstanding researchers from all branches of engineering who have up to three years of post-doctoral research experience.

The goal of Timothy’s work is to use the compiler to help make energy-saving decisions in multi-core processors. He will

tackle this problem at design, at compilation and during program execution. The overall aim is to create a dialogue between the software and hardware, using knowledge gained at both levels to influence energy-saving decisions.

Timothy is a post-doctoral research fellow within the Compiler and Architecture Design (CArD) Group at Edinburgh University. He is also technical leader of the Adaptive Compilation Cluster in HiPEAC. ■

## HiPEAC Research Collaborations



During 2007, HiPEAC focused a lot of its efforts on interconnects. The main avenue for this development has been the creation of the interconnect's cluster in the former HiPEAC NoE. For this, several members agreed to build a general cluster, focusing on networks (off-chip and on-chip). This initial effort was initiated by UPV, FORTH-ICS and by new institutions, which at that time were non-HiPEAC members: Jönköping University (Sweden) and University of Catania (Italy).

The main outcome of this development has been an increase in the numbers of new members, and currently, most of the researchers at European level are linked through the HiPEAC-2 NoE, including the Technical University of Valencia (UPV, Spain), FORTH-ICS (Crete, Greece), Jönköping University (Sweden), University of Catania (Italy), Simula Research Labs (Norway), University of Ferrara (Italy), University of Bologna (Italy), Universidad de Valencia Estudi General (Spain), Universidad de Castilla-La Mancha (UCLM, Spain), and Poznan University of Technology (PUT, Poland).

Since the initial meetings, several research activities between different members have been started. What follows summarizes the main research collaborations among the cluster members.

### Research on new topologies and routing requirements for NoCs (UNIFE-UNIBO-UPV):

The tight interaction between UNIFE, UNIBO and UPV within the activities of the HiPEAC interconnect cluster for the very first time allowed researchers from the on-chip and the off-chip interconnect domains to cooperate and to bring a significant contribution to the advancement of knowledge in the field of NoCs. The first contribution was an in-depth assessment of multi-dimensional topologies for general purpose, tile-based MPSoC platforms, relying on the validation of the results of abstract, system-level analysis with the feedback from the physical synthesis. Francisco Gilabert (UPV) did a 3-month internship at UNIFE. Ongoing activities are extending the scope of the analysis to more promising, yet more challenging (silicon-wise) topologies, such as fat-trees. Moreover, following current research trends, new activities on fault-tolerance have been recently initiated focused on routing mechanisms for NoCs. The implementation trade-offs of a logic-based, distributed-routing scheme (from UPV) is currently being explored as an efficient way of tackling the problem. Samuel Rodrigo (UPV) is doing a 3-month internship in UNIFE.

### Research on routing algorithms for NoCs (Jönköping-UNICAT-UPV):

The cluster project helped to extend collaborations between Jönköping University

and UNICAT in the area of Application-Specific, Routing Algorithms. Earlier work has been extended and a methodology has been designed for the development of application-specific, routing algorithms that distribute load uniformly over the network. Fault-tolerance properties and Hierarchical Routing (a new methodology) are now being explored for the application-specific routing. As a result, 3 different internships have taken place between UNICAT, Jönköping and UPV.

### Research on congestion management and switch architectures (FORTH-UPV):

Nikos Chrysos (FORTH-ICS) visited UPV to discuss preliminary ideas on possibly combining his congestion-management scheme (proactive Request-Grant) with the UPV reactive scheme (RECN). Georgios Passas (FORTH) visited UPV and worked on buffer-memory organizations and crossbar scheduling in high-radix crossbar switches, in connection with the "proximity communication" packaging and signalling technology of Sun. A 4-month internship at Sun Research Labs in California followed.

### Research on task-mapping strategies (UV-UNICAT-UPV):

Joint applications of task-mapping techniques (developed at UV) with a topology-agnostic routing algorithm (developed at UPV) are being explored. In addition, UV are collaborating with UNICAT to integrate the mapping technique with the APSRA design methodology. In this regard, Rafael Tornero (UV) did a one-month research fellowship in Catania.

Other ongoing research activities are power-efficient NoCs (FORTH-UNIFE-UPV), efficient routing implementations for NoCs (Simula-UPV), quality of service in NoCs (UCLM-UPV), network-processor interfaces (FORTH), packet dispatching algorithms and arbiter implementations for switching fabric architectures (PUT).

José Flich  
Associate professor  
Universidad Politécnica de Valencia, UPV

## Promotion of HiPEAC Start-ups

We continue the HiPEAC start-ups promotion action by focusing on two recently created start-ups, beginning in this edition's HiPEAC info and continuing in the next two editions. In this issue, we learn about Acumem and CAPS, both of which target the most urgent challenge in computing today - to improve the effectiveness of programming multi-processors. They both rely on extensive experience and academic achievement for offering innovative tools in this domain. There is no doubt that the need is there, and the challenge too!

### Acumem: Analyzing and Removing Multicore Bottlenecks

the lunch may not be free, but the low-hanging fruit is



The move to multi-core architectures often results in lower-than-expected performance. Limited per-thread cache capacity, memory bandwidth bottlenecks and inefficient thread interactions ruins the potential performance, but only a handful of experts are capable of understanding and fixing such problems.

This led to the creation of the start-up company Acumem, partly based on the work by Professor Erik Hagersten and his research team at Uppsala University. Many years ago, they saw the need for tools that enable non-experts to work efficiently with multi-cores, but it can be equally important to provide tools to make the performance experts more productive, since they can be expected to be faced with a multitude of performance problems over the next decade.

One part of Acumem's core technology is a sampler that efficiently captures sparse architecturally-independent runtime information from native execution -- the application's fingerprint. A second part is a modelling technology that models a given multi-core's behaviour based on a fingerprint. These two parts together provide with high accuracy and high efficiency an unprecedented



Marco Cornero  
Director Compilers,  
Operating Systems  
and Applications  
STMicroelectronics  
HiPEAC partner

insight into what really goes on inside a multi-core.

Acumem has developed its first product, SlowSpotter™, on top of the core technology. It finds application SlowSpots™ - places in the code that can be modified to improve the application's performance. About 20 different types of SlowSpots™, related to multi-threaded execution and cache usage, are identified and fixes suggested at a level of detail allowing for non-experts to perform the optimization. A second product, SlowSpotter-Pro™ targets the performance experts with an aim to improve their productivity. The freely available SpotLite™ can tell if an application has SlowSpots™ and classifies them according to type.

The Acumem tools are included in HP's multi-core toolkit and AMD's recommended 3rd party performance tools. Sun Microsystems is another close partner. Acumem's tools currently support x86 architectures under Linux or Solaris 10.

Through HiPEAC collaboration, Stefanos Kaxiras and his group at the University of Patras have been heavily involved in the Acumem development.

Acumem is described in the feature article in [www.hpcwire.com](http://www.hpcwire.com).

Acumem is in an expansive phase and is looking for highly skilled people from HiPEAC's expertise area. More info about Acumem: [www.acumem.com](http://www.acumem.com).

**Contact: Erik Hagersten**



**Built on over five years of advanced research and development and founded in 2002 by members of an INRIA research team, CAPS is today a leading provider for innovative, standard, programming tools that leverage the computing power of evolutive, many-core, hybrid platforms.**

### How CAPS answers the new HPC technology mutation ...

Everybody agrees that multi-core is a huge revolution in the semi-conductor market. Processor performance that was essentially driven by the increase of their clock frequency is now fastened to concurrent cores parallelism. Today, to benefit from the computing power offered by multi-core architectures, software industry needs to adapt the way they design applications. From now on, applications performance depends on how efficiently software applications take advantage of the different specialized cores (GPUs, Cell, SSE...).

While hardware technologies are very promising in terms of performance, they still, however, require developers to have a detailed knowledge of the underlying architecture to achieve performance. As a result, this time-consuming learning effort inhibits their productivity.

CAPS answers this issue with HMPPT™, a new standard way of easily and rapidly programming many-core applications that automatically adapts to the configuration of hybrid multi-core systems.

# Promotion of HiPEAC Start-ups



Innovative tools for a new paradigm



HMPPT™ workbench consists of a set of compiler directives inserted into the application source to define hardware-specific implementation of functions, their execution, the data transfers and synchronization barriers between the host CPUs and the accelerators, as well as a dynamic codelet manager.

Thanks to the HMPPT™ workbench, software developers can gain the most out of multi-core processors, while preserving the legacy source of their applications. Portability and interoperability are both key benefits provided by HMPPT™.

Today, CAPS addresses industries requiring high-performance computing, such as oil & gas (seismic modeling), defence (facial recognition, communication analysis), but also finance (risk management) and life sciences (genes and cells process modeling).

CAPS joined the Multi-core Association as a member of its developmental working group on multi-core, virtual-machine standardization and multi-core resource management.

CAPS is also actively involved in many

French and European Research and Development projects concerning the development of multi-core compiling technologies and optimization methods:

## FRENCH COMPETITIVENESS NETWORKS and FRENCH RESEARCH AGENCIES

SYSTEM@TIC PARIS-REGION (FAME2, POPS)

MINALOGIC (SCEPTRE)

VIDEO & NETWORKS (SCALIMAGE, SERENITEC)

ANR-PARA

ANR-PARMAT

## EUROPEAN PROJECTS

MILEPOST

PARMA

More info on [www.caps-entreprise.com](http://www.caps-entreprise.com)

Contact: **François Bodin**

## PhD News

### A Reconfigurable Application-specific Instruction-set Processor for Trellis-based Channel Decoding

By **Timo Vogt** ([vogt@eit.uni-kl.de](mailto:vogt@eit.uni-kl.de))  
**Prof. Dr. Norbert Wehn**  
TU Kaiserslautern, Germany  
January 2008

Future mobile and wireless communication networks require flexible modem architectures to support seamless services between different network standards. Hence, a common hardware platform that can support multiple protocols implemented or controlled by software, generally referred to as software-defined radio (SDR), is essential, providing high performance with minimum power budget.

This thesis presents a family of dynamically reconfigurable application-specific instruction-set processors (ASIP) for channel coding in wireless communication systems. As a weakly programmable IP core, it can implement trellis-based, channel decoding in a software-defined, radio environment. It features Viterbi and Max-Log-MAP decoding for binary convolutional codes, and turbo decoding for binary as well as duobinary turbo codes. Due to its high flexibility, it allows for adaptation to future systems.

The ASIP consists of a specialized pipe-

line with 15 stages and a dedicated communication and memory infrastructure. Logic synthesis revealed a maximum clock frequency of 400 MHz and an area of 0.11 mm<sup>2</sup> for the processor's logic using a low power, low leakage, 65 nm standard cell technology. Memories require another 0.31 mm<sup>2</sup>. Simulation results obtained for Viterbi and turbo decoding demonstrate maximum throughput of 196 and 34 Mbps, respectively. The ASIP outperforms state-of-the-art decoder architectures, targeting software-defined radio by at least a factor of three, while consuming only 60% or less of the logic area.

## Self-Adaptive Hardware/Software Reconfigurable Networks

**By Thilo Streichert**  
([streichert@cs.fau.de](mailto:streichert@cs.fau.de))  
**Prof. Dr.-Ing. Jürgen Teich**  
**University of Erlangen-Nuremberg,**  
**Germany**  
**February 2008**

Embedded networks are systems that consist of communicating nodes specialized for certain purposes. Typically,

these systems are subject to constraints such as fault-tolerance, availability, but also flexibility. This thesis presents a novel framework for increasing fault-tolerance and flexibility by separating functionality from hardware. Based on Field-Programmable Gate Arrays (FPGAs) in combination with a CPU, the presented methodology enables a task implemented in hardware or software

to migrate from one node to another, and in the case of a node defect, tasks are automatically replicated to obtain redundancy. As proof of concept, these methods have been integrated in a distributed operating system that hosts an application for lane-departure warning on a network of FPGA-based nodes.

## Loop Transformations for the Optimized Generation of Reconfigurable Hardware

**By Harald Devos**  
([Harald.Devos@elis.UGent.be](mailto:Harald.Devos@elis.UGent.be))  
**Prof. Dirk Stroobandt**  
**Prof. Jan Van Campenhout**  
**Ghent University, Belgium**  
**February 2008**

Current high-level design environments offer little support for implementing data-intensive applications on heterogeneous-memory systems; they instead focus on parallelism. This thesis

addresses the memory-hierarchy problem in high-level transformations of loop structures. The composition of long transformation sequences by combining shorter subsequences is studied together with the influence of the order of applying transformation steps. Several methods are presented to estimate bounds on Ehrhart quasi-polynomials, which can be used to statically evaluate program properties, such as memory usage. Since loop transforma-

tions not only influence the data-access pattern, but also the control complexity we present a hardware loop controller architecture which supports hardware generation from the polyhedral representation used for loop transformations. The techniques are demonstrated by the semi-automatic generation of an FPGA implementation of an inverse discrete wavelet transform.

## Run-Time Management for Future MPSoC Platforms

**By Nollet Vincent** ([nollet@imec.be](mailto:nollet@imec.be))  
**Prof. Henk Corporaal.**  
**TU Eindhoven, The Netherlands**  
**Prof. Diederik Verkest.**  
**KU Leuven, Belgium**  
**April 2008**

By using a flexible, easily programmable (heterogeneous) MPSoC platform in our everyday embedded devices and by using scalable applications with a predictable behaviour, manufacturers are able to keep the total cost of such devices under control. The platform run-

time manager plays an important role in matching the requirements of the application with the platform capabilities and in enabling this flexibility, programmability, scalability and predictability.

This thesis first provides a comprehensive survey of the contemporary MPSoC run-time management functionality and its implementation space. Consequently, it introduces several novel algorithms, mechanisms and components for performing run-time management of multiple scalable applications running on

a heterogeneous multiprocessor SoC. This includes run-time FPGA fabric management combined with efficient and fast resource allocation, task migration policies and mechanisms, a NoC communication management policy and mechanism, and a quality management component for handling adaptive applications. In addition, we describe a real-life MPSoC proof-of-concept implementation and demonstrator. The thesis concludes by providing a peek into the future of MPSoC run-time management.

## Transform Domain Transcoding Systems for Static and Dynamic Video Composition

**By Nuno Roma**  
([Nuno.Roma@inesc-id.pt](mailto:Nuno.Roma@inesc-id.pt))  
**Prof. Leonel Sousa**  
**TU Lisbon, Portugal**  
**May 2008**

Video delivery systems and service providers often face the need to manipulate compressed video-streams. As a consequence, video transcoding has emerged as a new research area. It concerns a broad set of processing, manipulation and adaptation techniques to convert one video bit stream into another, with a more convenient set of parameters

targeted to a given application. The prime focus of this thesis is the development of efficient and flexible transcoding algorithms and architectures for video composition in the DCT-domain. Two distinct applications were targeted: static video composition, consisting of the insertion of stationary visible data over the received video sequence; and dynamic video composition, consisting of the composition of one or more foreground video sequences over the display area, corresponding to the background video sequence. To support the implementation of the pro-

posed transcoding structures, several common video-processing operations had to be adapted and transposed into the DCT-domain, such as video space-scaling by an arbitrary integer scaling factor, and motion estimation using the blocks of DCT-coefficients obtained from the received video-streams. When compared with the corresponding pixel-domain counterparts, the proposed DCT-domain transcoders have shown to provide significant advantages, both in terms of the video quality, coding efficiency and computational cost.

## Hardware/Software codesign methodologies for dynamically reconfigurable systems

**by Marco D. Santambrogio**  
([marco.santambrogio@polimi.it](mailto:marco.santambrogio@polimi.it))  
**Prof. Donatella Sciuto**  
**Politecnico di Milano, Italy**  
**May 2008**

This dissertation deals with the definition of a specification-to-bitstream and autonomous design flow for dynamically reconfigurable platforms based on standard tools, where possible. The main goal is the definition and imple-

mentation of a flow that can output a set of configuration bitstreams used to configure and, if necessary, partially reconfigure a standard FPGA to realize the desired system. The main contributions introduced by the proposed framework can be found both in the definition of a complete methodology to implement self-reconfigurable embedded systems, taking into consideration both the HW and the SW side of the final architecture that can also

be easily adapted to work in different scenarios, and in the definition of a set of techniques to efficiently manage the reconfiguration at runtime. An example of the latter is the design of an Operating System solution that is able to support and manage the reconfiguration and the runtime bitstreams relocation. This approach has been validated using different technologies, and in different application scenarios.

## Designing Globally-Asynchronous Locally-Synchronous On-Chip Communication Networks

**Communication Networks**  
**By Xin Wang** ([xin.wang@tut.fi](mailto:xin.wang@tut.fi))  
**Prof. Jari Nurmi**  
**Tampere University of Technology, Finland**  
**June 2008**

In System-on-Chip (SoC) implementations of embedded systems, network-based technology is increasingly needed to replace conventional buses. In order to avoid excess power consumption and

problems in routing the synchronization clock signals, multiple clock domains are used. This calls for techniques to asynchronously transmit data across the chip and resynchronize again at the receiving clock domain. The design of logic modules for asynchronous operation is a challenge in the contemporary design tools targeted for synchronous design. In this work, methodology for such implementations is studied. Also a system-level scheme to build CDMA

(Code-Division Multiple Access) type of networks on chip has been developed and compared to more conventional crossbar and ring network implementations. The results show that the CDMA Network-on-Chip has its niche in complex SoC implementation due to its constant latency and reduced wiring needs compared to crossbar-based communication.

# Sean Rul, Electronics and Information Systems Department, Ghent University

I started my PhD in the summer of 2005. I had just finished my Master's thesis on selective value prediction at Ghent University when my supervisor, Koen De Bosschere alerted to me the possibility of doing a PhD in his research group. Now almost three years later, I'm glad I seized this opportunity with both hands.

Since I have a keen interest in computer architecture, the research group Paris (Parallel Information Systems group) was the right environment for me. For computer architects these are interesting times, since the rise of the multi-core poses a lot of problems to our community. I therefore decided to do research on the parallelization of complex sequential programs. My research involves building a tool to extract coarse-grain parallelism from programs based on dynamic dependences attained by profiling.

A nice research subject is one thing, but if you lead a secluded life separated from

the outside world, you will have a difficult time getting your work noticed. Like a missionary, you have to spread your research by networking and presenting your work at conferences. Since I like to travel, this is another aspect of PhD life which attracts me a lot. Moreover, thanks to the HiPEAC network and its numerous activities, such as the summer school and cluster meetings, I have been able to get in contact with a lot of fellow researchers from all over Europe.

A unique experience for me was my internship at the Haifa Research Lab of IBM in Israel from November 2007 until January 2008. This collaboration was made possible by the HiPEAC network and I advise every young PhD student to try it out. At IBM I worked on a framework that should ease the programmability of multi-cores (in particular the Cell BE). Besides the interesting research topic, it was also enriching to see the differences between research in academia and industry.



I think you can compare a PhD with a big rollercoaster ride. It's fun and has a lot of interesting twists and turns, allowing you to look at things from a different perspective. Of course, there are also the ups and downs, where you don't always go as fast as you want to, but nothing is insurmountable. And as with all fun things, it's all over before you know it.

**Contact:**

[www.elis.ugent.be/~srul](http://www.elis.ugent.be/~srul)

[sean.rul@elis.ugent.be](mailto:sean.rul@elis.ugent.be)

# Christos Strydis, Computer Engineering Laboratory, TU Delft



"Take one gallon of flour, two pounds of currants and one pound of butter, or better still, a quarter pound of sugar, a quarter pint of rose-water, half an ounce of nutmeg and half an ounce of cinnamon, two eggs. Then warm the cream, break the butter into the flour, temper all these with the cream, and mix in a quantity of yeast - a pint to three gallons. Cover your cake with two layers of the mixture when it comes hot out of the oven. Let it stand one-and-a-half hours in the oven. The oven should not be kept too hot. When the cake is two days old, slice it, then sugar it, and keep it in such a form all year round."

Replace a few keywords and the above medieval cake recipe could be used for cooking up a new biomedical implant. In many ways, implant design has, for a

large part, been based on secret, old and/or rigid recipes over the years. Designs have traditionally been highly inflexible, customized ones and, nowadays, increasingly based on off-the-shelf components – the lucky pick of the day. Meanwhile, in this field of highly mission-critical applications, design for fault-tolerance and reliability has been seriously absent or overlooked, resulting in devices with much to be desired. Can we do better than that? Here in Delft, we believe we can.

My name is Christos Strydis and I am currently entering my fourth year of a PhD at the Computer Engineering Laboratory of Delft University of Technology, under the supervision of Prof. Georgi N. Gaydadjiev. Our project is called "Smart implantable Medical Systems" (or "SiMS", for short) and, to abuse the recipe metaphor further, our proposal is of a new cake mold that does not break, does not turn hot in the oven and does not burn the dough.

In plain engineering jargon, SiMS is a fresh look at implant design, i.e. a framework supporting the educated, structured design of microelectronic implants explicitly provisioned for fault-tolerant and ultra-low-power operation over a large range of biomedical applications. Serious effort has already been spent on drawing the specifications of a novel, biomedical processor lying at the heart of SiMS and we are currently at the phase of exploring various design-space alternatives. We expect our version of the recipe to rise further up through the know-how yeast of the two freshly kick-started HiPEAC task forces on Reliability and Availability and on Low Power.

**Contact:**

[christos@ce.et.tudelft.nl](mailto:christos@ce.et.tudelft.nl)

<http://ce.et.tudelft.nl/~christos>

## Upcoming events

### **SAMOS VII: International Symposium on Systems, Architectures, Modeling and Simulation**

Samos, Greece, July 21-24, 2008, [http://samos.et.tudelft.nl/samos\\_viii/](http://samos.et.tudelft.nl/samos_viii/)



### **11th Euromicro Conference on Digital System Design – Architecture, Method and Tools**

Parma, Italy, September 3-5, 2008, <http://dsd08.iet.unipi.it/index.htm>



### **Design Automation & Test in Europe (DATE) 2009**

Nice, France, 20-24 April 2009, <http://www.date-conference.com/index.php>, Call for papers deadline: September 7, 2008

### **High-Performance Computing (HPC) Symposium 2008 – 37th JAIIO**

Santa Fe, Argentina, September 8-12, 2008, <http://www.37jaiio.org.ar/hpc/index.php>



### **The 37th International Conference on Parallel Processing (ICPP-08)**

Portland, USA, September 8-12, 2008, <http://www.cse.ohio-state.edu/~icpp2008/>

### **IEEE International Symposium on Workload Characterization (IISWC 2008)**

Seattle, USA, September 14-16, 2008, <http://www.iiswc.org/iiswc2008/index.htm>



### **The 14th Annual International Conference on Mobile Computing and Networking (MobiCom 2008)**

San Francisco, USA, September 14-19, 2008, <http://www.sigmobile.org/mobicom/2008/>

### **The 13th ACM SIGPLAN International Conference on Functional Programming (ICFP 2008)**

Victoria, British Columbia, Canada, September 22-24, 2008, <http://www.icfpconference.org/icfp2008/>



### **International Symposium on System-on-Chip 2008**

Tampere, Finland, November 5-6, 2008, <http://soc.cs.tut.fi/>

### **MICRO-41: The 41st Annual IEEE/ACM International Symposium on Microarchitecture**

Lake Como, Italy, November 8-12, 2009, <http://www.microarch.org/micro41/>

### **PDCAT'08: The International Conference on Parallel and Distributed Computing, Applications and Technologies**

Dunedin, New Zealand, December 1-4, 2008, <http://www.cs.otago.ac.nz/pdcat08/>



### **HiPEAC 2009 Conference**

Paphos, Cyprus, January 25-28, 2009, <http://www.hipeac.net/conference>



### **PLDI 2009: Programming Language Design and Implementation**

Dublin, Ireland, June 15-20, 2009, <http://www-plan.cs.colorado.edu/~pldi09/>

## 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)