

# HiPEACinfo<sup>14</sup>

## COMPILATION ARCHITECTURE

appears quarterly | April 2008

Network of Excellence on High Performance and Embedded Architecture and Compilation

- 2 Message from the HiPEAC coordinators
- 2 In the Spotlight  
Second Edition "Antiguos Alumnos UPV"
- 3 Message from the project officer
- 4 Hipec Activity
  - HiPEAC Computing Systems Week
  - 5th HiPEAC Industrial Workshop
- 5 Guest Column
  - Grant Martin, Chief Scientist, Tensilica, Inc.
  - ALaRI – University of Lugano
- 7 HiPEAC Start-ups  
Promotion of HiPEAC start-ups
- 9 PHD News
- 12 Upcoming Events



**Welcome to the HiPEAC  
Computing Systems Week  
Barcelona, Spain  
June 2-6, 2008**

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

Paper submission deadline HiPEAC 2009 Conference

July 11, 2008

## Message from the HiPEAC coordinators

Dear colleagues,

The beginning of 2008 was quite exciting for our community. In January 2008, we all witnessed some of the most successful networking events ever organized by the HiPEAC network. The workshops and HiPEAC 2008 conference in Goteborg attracted well over 250 delegates, far more than we ever could have dreamt of.



Mateo Valero



Koen De Bosschere

After the conference, we had the HiPEAC2 kick-off meeting. The meeting facilities for the kick-off meeting were clearly too small for the number of people that showed up. After the successful start of the nine research clusters, we were happy to welcome 26 excellent new members to our network. They enrich the HiPEAC network by providing new expertise and new opportunities for discussion and collaboration.

One of the next opportunities to meet again will be at the Spring Computing Systems Week, which will take place in beautiful Barcelona from June 2 until June 6, hosted by the Barcelona Supercomputing Center, UPC and HP Research Labs. We are also very happy that the Computing Systems Week will also include a multi-core workshop with several international speakers, organized by BSC/Microsoft/HiPEAC. We definitely hope you will be able to attend

this whole week of inspiring events, on which you will find more information in this newsletter. At the Spring Computing Systems Week we will open the call for collaboration grants, which enable students to carry out research at another institution for three months. The cluster meetings provide the ideal opportunity for finding a collaborative partner. For shorter collaborations, we have recently opened the "reimbursement service", where you can enter a request for a short collaboration. More details are on the website under the "Members" section. We hope that the HiPEAC members will take advantage of this low-overhead instrument to start new collaborations in the coming months.

In the next few months, we will also be working hard to further increase the visibility of HiPEAC research by creating a database of recent HiPEAC

research publications, which will be easily accessible by both members and non-members. No single European research idea should remain undiscovered for our community. The database will also help researchers from academia and industry to quickly find a suitable partner for collaboration. We are convinced that over time this database will become one of the more valuable instruments of our network.

It is clear that our community is quickly gaining momentum and has made progress well beyond our expectations when HiPEAC was first started in 2004. As coordinators, we can only be proud of such a community, and we are more than ever committed to working hard to strengthen our community and to create more resources from which your research can benefit.

Koen De Bosschere and Mateo Valero

## In the Spotlight

### Second Edition "Antiguos Alumnos UPV"



José Duato was awarded the Honor Prize by the Association of alumni and the Social Council of the Technical University of Valencia in the second year of this award. Out of 70,000 alumni, this prize recognizes him as one of the most outstanding professionals who graduated from this university.

The very first Honor Prize was awarded to Jaime Mayor Oreja, former minister of the Spanish Government. The prize-giving ceremony was held on January 31st, 2008. Pictures for the event can be found at <http://www.antiguosupv.org/gala15Aniversario/fotos.asp>

# Message from the project officer

Panos Tsarchopoulos  
Panagiotis.Tsarchopoulos@ec.europa.eu

I would like to give a short overview of the conclusions of the three consultation workshops that we have organized to discuss the next work program of the Computing Systems objective. Since the space here is limited, I invite you to take a look at the presentations and reports available on the Cordis web pages for Computing Systems. The next Call for Proposals will be published in November 2008.

## Workshop on Computing Systems

This consultation workshop discussed the major paradigm shift, the massive adoption of multi-core systems by the whole computing systems industry. The workshop identified three major challenges, which impact all market segments of computing systems. These challenges are: how to increase **performance**, how to increase **power efficiency** and how to improve **reliability** of computing systems. The consultation workshop has translated these three challenges into seven research areas:

1. **Parallelization** – automatic parallelization, new high-level parallel programming languages and/or extensions to existing languages taking into consideration that user uptake is a crucial issue.
2. **Continuous Adaptation** – multi-core and/or reconfigurable systems that continuously adapt to a constantly changing environment
3. **Virtualization** – technologies that ensure portability, flexibility and overcome legacy issues for multi-core systems
4. **Customization** – rapid extension and/or configuration of existing systems, architectural templates and tool-chains to optimally address specific application needs and performance/Watt envelopes.
5. **System Simulation & Analysis** – advanced simulation and analysis of complex multi-core systems.
6. **Design Space Exploration** – how to efficiently find the optimal processor configuration for a set of applications.
7. **Technology Impact** – new opportunities and challenges for architectures, tools and compilers created by

advances in semiconductor fabrication technology (for example, 3D stacking).

## Workshop on Reconfigurable Computing

The workshop identified as the highest research priority **the need to enable commercially viable programmability of reconfigurable computing technology**. This requires coherent, integrated (or “integrable”) suites of processes, methods and tools spanning:

- application-level support for reconfigurability that supplements existing design methodologies, including support for verification and validation of reconfigurable behavior and reconfigurability properties of the system so as to satisfy qualification requirements;
- mapping from the output of application design to reconfigurable hardware via intermediate layer(s) of abstraction with standard libraries of functions based on open and widely accepted standards;
- provide a repository of hardware blocks as well as a toolset to generate the hardware, using a HDL, in an automated way; and
- run-time support for reconfiguration, typically through OS extensions for resource allocation, scheduling, and discovery; debugging and monitoring; and fast re-layout of reconfigurable units.

The workshop underlined that research in these topics must recognize the need for compatibility with development paradigms and processes, methods and tools in the applications sectors. In this sense, research in reconfigurable computing should be application-driven. Application sectors, where Europe could gain particular advantage, include embedded health-care, (multi)physical system modeling, biomedical, cognitive radio, portable consumer devices, automotive/avionics, entertainment, and user-driven reconfigurable products

## Workshop on High-Performance Computing

The fact that in the global economy Europe can only maintain its advantage by taking a lead in “knowledge based”

businesses, e.g. design(automotive, aerospace, pharmaceuticals), financial services, etc., the workshop highlighted



**that high-performance computing is a major tool that enables this knowledge-based lead.** The European effort should be focused on the areas that will lead to “high productivity computing”, such as algorithms, robust APIs, standards, and methods of making code more portable.

Focusing on this as a core activity will ensure that as the high-performance computing market develops, this work will naturally “trickle down” to computing systems in general. This will provide not only leverage in that market, but some of those developments will be able to be used as “commodity” building blocks for next generations of high-performance computing, thus creating a “virtuous circle”. Since power consumption and power density will be the limiting factors of future generations of supercomputers, and Europe leads in low-power computing, it is possible to see more European technology at the heart of supercomputers in several computing generations time.

There are challenges from the user viewpoint in scaling, ease of use, portability, reusability and efficiency. Translating these into R&D needs, the workshop identified challenges in **architectural developments, better analysis tools, handling of large data sets, improved and stable APIs, better algorithms and improved levels of abstraction**. To achieve this, more work is needed in multi-disciplinary teams; to be effective, these are usually application driven. The focus should be on improving the performance/cost ratio of existing applications, and extending them into new areas, but also on completely new applications that become possible.

Panos Tsarchopoulos  
Project Officer

## HiPEAC Computing Systems Week June 2-6 2008, Barcelona, Spain

We are glad to announce that the "HiPEAC Computing Systems Week" will be held in Barcelona, Spain, from June 2nd to June 6th, 2008.

This week will include the following three co-located events:

- The HiPEAC Cluster Meetings, hosted by BSC and UPC. These activities will take place on June 2nd and 3rd at BSC-UPC premises in Barcelona.
- The 5th HiPEAC Industrial Workshop, hosted by HP Research Labs. This

event will take place on June 4th at the HP Labs facilities in Sant Cugat del Vallès (near Barcelona). The goal of the workshop is to bring together researchers from academia and industry to discuss tools and methodologies for analysis, compilation, debugging, verification and simulation of parallel programs. More information below.

- The Barcelona Multi-core Workshop, organized and sponsored by Microsoft Research, the Barcelona

Supercomputing Center and the HiPEAC "Programming Models and OS" and "Multi-core Architecture" research clusters. The event will take place on June 5th and 6th at BSC-UPC premises in Barcelona. The workshop will include 2 keynotes, around 10 invited talks and 2 panel sessions on the challenges raised by the multi/many-core architectures of the future.

It will be a whole week of very exciting events, with plenty of time to discuss, meet and do lots of networking. As usual, the program also offers opportunities for other HiPEAC-related meetings, such as research project meetings.

We are looking forward to having you and your cluster collaborators attend the HiPEAC Computing Systems Week. We hope you will have a great experience and much fun in Barcelona!

**Please take note of the following schedule:**

Date/Time	Event	Location
Mon. 2nd June, PM	HiPEAC Cluster Meetings	BSC-UPC, Barcelona
Tues. 3rd June, all day	HiPEAC Cluster Meetings	BSC-UPC, Barcelona
Wed. 4th June, all day	5th HiPEAC Industrial Workshop	HP Labs, Sant Cugat
Thurs. 5th June, all day	Barcelona Multi-core Workshop	BSC-UPC, Barcelona
Fri. 6th June, AM	Barcelona Multi-core Workshop	BSC-UPC, Barcelona

## 5th HiPEAC Industrial Workshop Tools and Methodology for Parallel Programming June 4th, 2008

**Organized by:** HP Labs, Exascale Computing Lab, Barcelona Research Office

**Hewlett-Packard, Sant Cugat del Vallès, Barcelona, Spain**

**<http://www.hipeac.net/industryworkshop5>**

We live in exciting times in the field of parallel architectures. The diminishing IPC returns caused by power and frequency walls have forced a radical change in CPU design. Multi-core is here to stay and has become pervasive throughout the computer industry. In high-performance computing, three out of four "TOP 500" systems are clusters of industry-standard multi-core nodes. In embedded computing, heterogeneous multi-core System-on-Chips have become the dominant standard for high-volume parts.

There is no silver bullet in multi-core programming, and formidable computer science problems underlie the programming of this new breed of parallel architectures. Traditional parallel programming paradigms and tools need to be put to the test and questioned. The heterogeneity of cores and accelerators is creating new challenges for the programming models. Locks and multi-threading memory semantics are being redefined and transactional programming is on the horizon. With new programming paradigms and a shifting architecture target, new tools and analysis methodologies are also granted.

The goal of this workshop is to bring together researchers from academia and industry to discuss tools and methodologies for analysis, compilation, debugging, verification and simulation of parallel programs.

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

- Compiler support for explicit parallelism (threads, transactions)
- Tools to aid programmers to identify, exploit, verify, debug, and tune parallel applications
- Simulation methodology and tools for multi-core/many-core/clustered parallelism
- Characterization of parallel applications and their scalability
- Reliability and fault tolerance for applications running on many-core and distributed architectures
- I/O issues in parallel computing and parallel applications
- Parallel programming languages, algorithms and applications
- Middleware and run time support for parallelism





## Grant Martin, Chief Scientist, Tensilica, Inc.

I am honoured to be asked to contribute a guest column to the HiPEAC newsletter, and to be able to participate in HiPEAC-related activities. I have been working in the electronics industry for almost 30 years (the photo is a few years old; the grey hair in it is not completely up-to-date!) - working at Burroughs, BNR/Nortel, Cadence and now Tensilica. For those of you who do not know Tensilica, we are a configurable, extensible processor IP company that has recently celebrated its 10th anniversary. Our processors can be customised by users to create Application-Specific Instruction set Processors (ASIPs), to be used singly or in a huge variety of combinations in products ranging from toys and games to some of the largest network routers available in the industry, and other complex computational domains. In the middle range of applications using our ASIPs are multimedia-related devices, including audio and video processors in mobile handsets and other devices, and image processing in printers, as well as countless other applications. To complement our configurable processor technology, we have used it ourselves to create a variety of fixed processors, including audio engines (together with

many audio codecs) and video subsystems for decoding and encoding using a variety of standards.

I have a broad background in the design and design automation area, with a particular interest in system-level design, or what is often called these days "ESL". In particular, processor and multi-processor centric design both requires a system-level design approach in configuring processor(s) to applications, and in designing and verifying a complete system that may involve multiple subsystems and a heterogeneous mixture of ASIPs and fixed ISA processors carrying out relevant roles. Here our interests and those of HiPEAC overlap strongly. "High Performance and Embedded Architecture and Compilation" covers many of the areas that we think are important to the evolution of processor-centric design, and since ASIPs can be applied in products ranging from the smallest portable embedded device to the largest routing and super-computing engines targeted to specific applications, much of the HiPEAC work is of interest to us. Participating in the Design Methodology and Tools cluster allows us to bring our experience in ASIPs, ESL modeling and MPSoC appli-



cations into possible collaborations with others, in academia and industry, to share ideas and learn from the work of others. I am particularly keen on establishing new links with people in Europe on common interests in these areas, to better cross-fertilize our work and that of others in these areas. To the extent that being a relatively small company several thousands of kilometers from Europe allows, I hope to use all the communications media available including direct participation to strengthen these common interests.

Grant Martin  
Chief Scientist  
Tensilica, Inc.

## ALaRI – University of Lugano

ALaRI is the Advanced Learning and Research Institute established in 1999 at the University of Lugano (Università della Svizzera italiana - USI), Switzerland, with the mission of promoting research and education in Embedded Systems Design. Aware of the real need for a cross-disciplinary approach to education, ALaRI-USI equips the participants with a unique body of knowledge ranging from electronic engineering to computer science, including interpersonal skills, indispensable in today's industry, such as team work, complex-project manage-

ment, and market sensitivity.

ALaRI-USI research activities focus on topics of scientific interest and industrial applicability, based on real-life design issues. The following list provides the main research lines, with some details and highlights on specific activities:

- **Security & Cryptography:** research in this field has been focused on dedicated hardware implementations and analysis of state-of-the-art cryptographic algorithms, such as the AES, IDEA-NXT or the Secure Hash Standard. Moreover, secure network-



ing protocols such as IPSec have been analyzed and novel architectures of hardware accelerators have been devised.

- **Pervasive computing:** a modeling methodology has been devised to estimate the power consumption of wireless-connected devices at a high (protocol-level) abstraction level, so as to support fast and easy simulation and design space exploration

# ALaRI – University of Lugano

both at node and at network level. Power-related optimizing solutions for protocol management and for security solutions are also a subject of research.

- **System-level design:** the usage of the Universal Modeling Language to specify the functionality and timing characteristics of electronic devices and software modules has become more and more popular. UML can also be used to identify architectural bottlenecks and possible optimizations early in the design process; efforts have been placed on defining possible methods to enhance the hardware/software co-design of embedded systems using high level specification languages.
- **System-on-chip communication architectures:** on-chip communication in future multimedia and mobile devices will be based on network-like interconnections (Network-on-Chip). ALaRI research in this field focuses on modeling environment to evaluate power consumption, performance and security of NoC-based systems.

ALaRI has been and is active in several European and National programs such as

- **MULTICUB** (FP7) 2008-2010  
MULTI-objective design space exploration of MULTI-processor SoC architecture for embedded MULTI-media applications
- **AETHER** (FP6, IST, FET) 2006-2009  
Self-Adaptive Embedded Technologies for Pervasive Computing Architectures
- **LoMoSA+** (MEDEA+) 2005-2008  
Low-power Expertise for Mobile and Multimedia Applications
- **COOPER** (FP6, IST) 2006-2008  
Collaborative Open Environment for Project-centered Learning
- **Security Design Methodologies for Energy-Efficient Secure Cryptography Coprocessors** (Swiss National Science Fund) 2003-2006
- **ANTITESYS** (FP5, IST) 2002-2005  
Scientific Coordinator  
A Networked Training Initiative for Embedded Systems Design



- **CALIPSO** (Microsoft Embedded Systems RFP) 2003-2004  
IPSec in a Mobile IPv6 Environment
- **Mobile Security** (Gebert Ruff Stiftung) 2002-2003 Technical Coordinator  
Innovative Approaches to the Solution of Security Problems for Mobile Systems

ALaRI teaching staff (31 professors and 10 assistant professors) rely on international experts from renowned EU and US universities, research centres and industries such as Dortmund Universität, EPFL, ETHZ, Humboldt Univ-Berlin, KU-Leuven, PoliMi, RWTH, UPC, and also CSEM, HP, IMEC, NEC Princeton, and others. The teaching program and its final degrees have seen from the beginning the official collaboration of ETH Zürich and Politecnico di Milano.

**At present two Master's degree programs are offered:**

- **Master of Science in Embedded Systems Design, two-year graduate program**

An institutional Master of Science coordinated with the Faculty of Informatics of USI. It welcomes Bachelor school graduates or students with at least 180 credits (ECTS) or three years of study in computer science, telecommunications, electronics, physics and mathematics. The course provides a multi-interdisciplinary approach, including scientific grounding, engineering and management skills. After the first year, two tracks of studies are offered:

- **Design and Research** oriented to academic or industry sector;
- **Business Projects** oriented to management, economics and marketing. The whole program provides 120 credits (ECTS).

- **Master of Advanced Studies (MAS) in Embedded Systems Design, one-year postgraduate program**

This executive program is designed for graduates and participants already in employment. It offers a comprehensive grounding in embedded systems design field, aiming to train and re-skill current and future team leaders. Two solutions of study program are proposed:

- **Full-time program** in 10 months;
- **Part-time program** in 20 months.

The Master of Advanced Studies program provides 70 credits (ECTS).

The master courses start in October and finish in July. The language of tuition is English. Scholarships are available.

During the Master's courses, students are requested to develop and complete a project in collaboration with industrial and academic partners in advanced research topics. The list of industrial partners includes multinational companies such as STMicroelectronics, Intel, Microsoft, Infineon, Hewlett Packard, NEC, CoWare, Synopsys, and others.

For more information, please visit [www.alari.ch](http://www.alari.ch) - or write to [master@alari.ch](mailto:master@alari.ch)

# HiPEAC is promoting start-ups

The creation and consolidation of new companies in the computing domain is among the goals of the HiPEAC NoE. In order to focus our efforts in this direction we have decided to organize specific start-up promotional activities, whose intentions are to exploit the HiPEAC NoE as a facilitating organization for existing and new start-ups.

To begin with, we will promote existing start-ups in the HiPEAC info, starting with this first high-level introduction of the current start-ups, and then we will continue in the next editions with more detailed descriptions on each of them. Other promotional activities include: reserved tracks at our industrial workshops - an excellent opportunity to reach an industrial audience; dedicated initiatives during HiPEAC events, such as panels at the next HiPEAC conferences; invited presentations and possibly specific courses at the ACACES summer school.

In addition to these promotional activities, we will also invite start-up representatives to share their experiences in order to stimulate and facilitate the creation of new ones, focusing on the toughest issues, such as funding. Ideally we would like to establish a business/financial network, parallel and complementary to the main technical one, with reciprocal benefits.

Start-ups will also benefit from the HiPEAC network for training, through active participation to the technical initiatives and the ACACES courses, and from accessing highly qualified and specialized candidates among our HiPEAC students and researchers.

Let us then start this action with the following brief introduction of the current HiPEAC start-ups. Readers are kindly invited to highlight other start-ups operating in the HiPEAC domains that we might have missed.



Marco Cornero  
Director Compiler,  
Operating Systems  
and Applications,  
STMicroelectronics  
marco.cornero@st.com  
HiPEAC partner

### Acumem

<http://www.acumem.com/>

Contact: Erik Hagersten



Acumem was founded during the spring of 2006 by Prof. Erik Hagersten of Uppsala University and currently employs 9 people. Acumem explores some of the performance-modeling techniques developed by Hagersten and his PhD students. It also has long-running collaboration with Prof. Kaxiras' research group at the University of Patras.

The new patented technology developed by Acumem is targeted at performance problems related to the move of software to multi-core architectures. Their first product, the Acumem Virtual Performance Expert (VPE), was launched during the Supercomputing Conference 2007. VPE captures a small "application fingerprint" from a running application, performs detailed performance analysis on the fingerprint, detects Slowspots in the studied application and suggests modifications to the source code needed to avoid the Slowspots. This will turn a normal programmer into a performance expert, creating the virtual performance expert.

Several of the leading computer companies, such as Sun Microsystems, HP and AMD, have already adopted Acumem's technology and have established partnership agreements. For example, VPE is included in the multi-core tools offered by HP ([www.hp.com/go/multi-coretoolkit](http://www.hp.com/go/multi-coretoolkit)). Future developments at Acumem will package additional performance-analysis capabilities and insights acquired by the fingerprint technology into new tools.

### CAPS

<http://www.caps-entreprise.com/>

Contact: François Bodin



Founded in 2002 by members of an INRIA research team, CAPS develops and commercializes innovative software for high-performance applica-

tion tuning in the domains of HPC and embedded systems. CAPS offers a whole range of development tools and services enabling its customers' applications to optimize the performance of multi-core processors used by the last generation hardware.

CAPS mission is in keeping with the innovative and fast moving multi-core market and helps industries with high-level HPC issues, such as oil and gas, defense, finance and life sciences to enable their software developers to make the most out of multi-core processors while preserving their legacy source.

Built over five years of advanced research and development, CAPS provides high quality and cost-effective programming tools that leverage the computing power of evolutionary, many-core hybrid platforms.

Among these environments: HMPPTM, a toolkit based on a set of compiler directives that comes with development tools and Dynamic Services to simplify the use of hardware accelerators in conventional, general-purpose applications.

A leading innovator in parallel-programming tools, CAPS is also actively involved in many French and European Research and Development projects concerned with the development of multi-core compiling technologies and optimization methods.

### Nema Labs – enabling a smooth transition to multi-cores

<http://www.nemalabs.com/>

Contact: Per Stenström

Nema Labs is an exciting venture-capital-backed start-up, addressing the industry-wide software crisis created by the rapid adoption of multi-core

microprocessors. By providing unique software development tools built on world-class proprietary technology (patents pending), Nema Labs will enable multi-core performance and rapid software innovation for software vendors. The technology combines advanced





compile-time and run-time analysis to identify code segments that can run in parallel. Our tools require no multi-core programming skills and can be used by any programmer to create efficient and reliable code for multi-cores. Nema Labs tools are compatible with current software-development tool chains.

Nema Labs was founded by a research team from Chalmers University of Technology lead by Professor Per Stenström. Our offices are located in Gothenburg, Sweden and we have operations in Silicon Valley. For more information and contact information, please consult [www.nemalabs.com](http://www.nemalabs.com).

### Nanochronous Logic

<http://www.nanochronous.com/>

Contact: Christos P. Sotiriou



Nanochronous Logic, Inc. is a start-up company with headquarters in San Jose, California and an R&D centre in Crete, Greece. It develops Design-for-Variability, Design-for-Manufacturability EDA tools for ASIC/SoC circuits implemented in nanoscale standard-cell libraries.

The company was founded in November 2006 and is currently working with selected strategic customers to evaluate its first two industrial tools, NanoSync and NanoVerify.

Nanochronous Logic's vision is to provide advanced algorithms, tools and methodologies for tackling the problems of the nanometer era, namely (i) high dynamic power consumption and high leakage power, (ii) multiple frequency/voltage operating points and multiple voltage domains on the same IC, (iii) process (static) and voltage, temperature (dynamic) variations, (iv) timing yield uncertainty and (iv) high electromagnetic emissions.

The core technology focuses on a novel timing methodology to create designs that are variation-aware. Utilizing the novel concept of creating timing partitions for a design, an IC's global timing can be optimally tuned to the presence of post-fabrication variations. Timing partitions essentially localize timing constraints and can borrow cycle time from other timing partitions.

With this mode of operation, local variations do not directly impact the IC's global frequency of operation. Actual conditions (in terms of P, V, T), which vary from static SSTA model predictions and dynamic variations (very difficult to model), can be diagnosed and take into account post-fabrication using this timing methodology. Each die is rendered capable, through a fast functional test, to self-diagnose its attainable clock speeds. Continuous DFVS capability is another significant benefit, whereby through merely controlling a die's voltage, its internal timing circuits can self-tune the attainable frequency as a function of the voltage, thus enabling timing-power calibration for optimal power savings. Advanced power-saving techniques, such as MSMV and PSO, fit directly into this technology framework, yielding a complete timing-power solution.

### QuviQ

<http://www.quviq.com/>

Contact: John Hughes



QuviQ AB develops and markets

QuickCheck, a tool for automated software testing which generates test cases randomly from an executable formal specification of the system-under-test. When a failing case is found, QuickCheck begins to simplify it systematically, separating the "signal" causing the test failure from the "noise" that is always present in random data, thereby producing a failing-case report to the user that often makes the cause of the problem self-evident. Thus QuickCheck both reveals errors that would often be missed by conventional testing, and simplifies their diagnosis.

QuickCheck has been applied extensively in the telecoms industry, to products such as media gateways and radio base stations, where the demands on quality and robustness are particularly high. QuickCheck builds on functional programming technology, in particular Ericsson's Erlang/OTP programming language and system.

QuviQ was founded in May 2006 by John Hughes and Thomas Arts, of

Chalmers University and the University of Gothenburg in Sweden. Building on successful initial projects with Ericsson, the company now serves customers in five countries and on two continents. Further information can be found at [www.quviq.com](http://www.quviq.com).

### SDS

<http://www.splitted-desktop.com/fr/>

Contact: Olivier Temam



SDS is a French start-up company in the domain of computing systems. The rationale for SDS is that current computing systems are too complex to be used by a large share of the population. The root cause of this situation is that operating systems were initially developed with a bottom-up approach, starting from the hardware and progressively abstracting it to the user over time. Even though operating systems have made drastic improvements in terms of usability, they are still a far cry from most users' view and needs.

Many users just want to see the computer as a simple consumer electronics device offering a set of rich services, and a user may want to see several such devices spread within the home. Therefore, the computing system must be easy to use, be cheap, be powerful, be connected, and be ubiquitous. Therefore SDS is working on (1) redesigning the operating system from a user perspective from the top down, and (2) designing a hardware platform that fulfills the aforementioned constraints. One of the first innovations of SDS is the design of a full hardware system that can offer PC-class performance (using recent AMD x86 processors) without a fan, thanks to patented power dissipations mechanisms, special motherboards and encasing. As a result, the system consumes 20W on average, is totally silent, and provides instant-on access. The first prospective customers of SDS are ISPs (Internet Service Providers) which are increasingly looking to extend the array of services offered in their Internet boxes. More information at <http://www.splitted-desktop.com/en/>



## Algorithmic and Scheduling Techniques for Heterogeneous and Distributed Computing

**By Cyril Banino-Rokkones**  
([banino@yahoo-inc.com](mailto:banino@yahoo-inc.com))  
**Prof. Lasse Natvig**  
**Norwegian University of Science and Technology (NTNU),**  
**Trondheim, Norway**  
**March 2007**

The computing and communication resources of high-performance computing systems are becoming heterogeneous, are exhibiting performance fluctuations and are failing in

an unforeseeable manner. The Master-Slave (MS) paradigm, that decomposes the computational load into independent tasks, is well-suited for operating in these environments due to its loose synchronization requirements. The application tasks can be computed in any order, by any slave, and can be resubmitted in case of slave failures. This thesis provides models, techniques and scheduling strategies that improve the scalability and performance of MS applications. In particular, we claim

that deploying multiple masters may be necessary to achieve scalable performance. We address the problem of finding the most profitable locations on a heterogeneous grid for hosting a given number of master processes, such that the total task throughput of the system is maximized. Furthermore, we provide distributed scheduling strategies that better adapt to system-load fluctuations than traditional MS techniques.

## Architectural Techniques to Improve Cache Utilization

**By Haakon Dybdahl**  
([dybdahl@google.com](mailto:dybdahl@google.com))  
**Prof. Lasse Natvig**  
**Norwegian University of Science and Technology (NTNU),**  
**Trondheim, Norway**  
**May 2007**

To compensate for slow main memory, the area of a chip dedicated to the last-level cache is substantial, but its performance can be rather low. In this PhD dissertation, we present three approaches to improving the utilization of the cache. The first extends cache

functionality to include write-backs for destructive-read DRAM. In the second approach, several different schemes are proposed that improve the use of the last-level cache. The first scheme uses a run-time heuristic to bypass memory accesses that are transient. The second scheme uses a LRU-based replacement policy augmented with frequency of access to protected cache blocks from eviction. This is efficient in a chip multiprocessor (CMP) with shared last-level cache, since it can stop a single processor from filling the cache with unnecessary blocks. In the

third approach the last-level cache in a CMP is dynamically partitioned. Novel mechanisms for determining the best partition sizes for each processor are described. A performance monitoring mechanism is proposed that can stabilize cache schemes that are not robust. The performance gains of the new schemes are shown relative to conventional architectures, state-of-the-art techniques and compared to upper-bound oracle schemes.

## Counting Integer Points in Polyhedra and Applications to Program Optimization

**By Rachid Seghir**  
([rachid.seghir@gmail.com](mailto:rachid.seghir@gmail.com))  
**Prof. Catherine Mongenet**  
**and Dr. Vincent Loechner**  
**Université Louis Pasteur de**  
**Strasbourg, France**  
**December 2007**

The polyhedral model is a well-known framework in the field of automatic program optimization. Iterations and array references in affine loop nests

are represented by integer points in bounded polyhedra, or (parametric) Z-polytopes. In this thesis, three new counting algorithms have been developed: counting integer points in a parametric Z-polytope, in a union of parametric Z-polytopes and in their images by affine functions. The result of such a counting is given by one or many multivariate polynomials in which the coefficients may be periodic numbers. These polynomials, known as Ehrhart

quasi-polynomials, are defined as subsets of the parameter values called validity domains or chambers. Many affine loop-nest analysis and optimization methods require such counting algorithms. We applied them in array linearization, which achieves memory compression and improves spatial locality of accessed data. Besides program optimization, the proposed algorithms have many other applications, such as in mathematics and economics.

## Dynamic and Software Optimization of Data Accesses

**By Jean Christophe Beyler**  
(beyler@icps.u-strasbg.fr)  
**Prof. Philippe Clauss**  
**Université Louis Pasteur, Strasbourg,**  
**France**  
**December 2007**

This thesis concerns the development of dynamic approaches for the control of the hardware/software couple. More precisely, this work has the main goal of minimizing program execution times on uni- or multi-processor architectures, by anticipating memory

accesses through dynamic pre-fetch of useful data in the cache memory and in a way that is entirely transparent to the user. The developed systems are entirely software and do not use any dedicated hardware. They consist of a dynamic analysis phase, where memory access latencies are measured, a phase of binary optimizing transformations when they have been evaluated as efficient, and where data pre-fetching instructions are inserted into the binary code, a dynamic analysis phase of the optimizations efficiency, and finally a

canceling phase for transformations that have been evaluated as inefficient. Every phase applies individually to every memory access, and eventually applies several times if memory accesses have behaviors that are varying during the execution time of the target software. On several benchmark programs compiled either with gcc or icc, speed-ups ranging from 2% to 163% are obtained, with an average overhead of only 3% of the execution time. The manuscript is in English and can be obtained from the author's website.

## Algorithms and Architectures for Motion Estimation Stage in Hybrid Video Encoders Based on H.263 and H.264/AVC

**By Sebastian Lopez**  
(seblopez@iuma.ulpgc.es)  
**R. Sarmiento and J.F. Lopez**  
**University of Las Palmas GC, Spain**  
**January 2008**

A new Variable Block Size - Adaptive Cost Block Matching (VBS-ACBM) algorithm is proposed. This algorithm is able to compute the motion vectors for each of the block sizes and levels of sub-

pixel precision demanded by H.263 and H.264/AVC. Thanks to a novel adaptive strategy, VBS-ACBM is able to obtain optimum compression rates with low computational cost, independent of the characteristics of the input video sequence and the compression requisites established by the user. This means degradation in the performance of the video encoder under changing characteristics, which are typically

present in other algorithms, can be avoided. A set of architectural solutions are introduced when implementing VBS-ACBM in video encoders under real-time constraints, covering integer pixel motion estimation and sub-pixel motion vector refinement processes. Significant improvements for both cases are obtained when compared with state-of-the-art designs.

## Simulation & Modeling of MPSoC for an Early Performance and Energy Estimation

**By Rabie Ben Atitallah**  
(rabie.ben-atitallah@inria.fr)  
**Prof. Jean-Luc Dekeyser,**  
**Prof Smail Niar,**  
**INRIA Lille, France**  
**March 2008**

Multiprocessor system on chip (MPSoC) simulation in the first design steps has an important impact on reducing the time-to-market and power consumption of the final product. However, MPSoC are becoming increasingly complex and heterogeneous, thus making

early performance and power estimation hard to obtain. In this thesis, we propose a framework composed of several simulation levels. This enables early performance evaluation in the design flow. The proposed framework is useful for design space exploration and enables rapidly adequate Architecture/Application configurations to be found. In the first part of this thesis, we present an efficient simulation tool composed of three levels that offer several performance/energy trade-offs. The three levels are differentiated by the accu-

racy of architectural descriptions based on the SystemC-TLM standard. In the second part, we focus on the MPSoC energy consumption. For this purpose, we enhanced our simulation framework with flexible and accurate energy consumption models. Finally in the third part, a compilation chain based on a Model Driven Engineering (MDE) approach is developed and integrated in the Gaspard environment. This chain allows automatic SystemC code generation from high level MPSoC modeling.

## Three Pitfalls in Java Performance Evaluation

**By Andy Georges**  
(Andy.Georges@elis.UGent.be)  
**Prof. Koen De Bosschere,**  
**Prof. Lieven Eeckhout**  
**Ghent University, Belgium**  
**April 2008**

Evaluating the performance of a virtualized language environment, such as the Java platform, is not easy due to the complex behavior of the Java application and the Java virtual machine and

the interactions they have while executing. In this work, we address three pitfalls that have not been taken into account when researchers conduct performance analysis of Java applications. First, performance results from one VM are not representative for other VMs. The same holds for the input to the Java application. In general, extrapolating performance leads to mistakes. Second, research has not dealt adequately with the non-determinism that is present

when executing Java applications. Furthermore, techniques proposed to (partially) remove said non-determinism are not used in a statistically rigorous manner, resulting in misleading or even incorrect conclusions when comparing performance. Finally, average performance numbers provide little information. In this dissertation, we uncover these pitfalls and formulate solutions on how to deal with them.

## Analytical Performance Analysis and Modeling of Superscalar and Multi-Threaded Processors

**By Stijn Eyerman**  
(Stijn.Eyerman@elis.UGent.be)  
**Prof. Lieven Eeckhout**  
**Ghent University, Belgium**  
**May 2008**

Analyzing performance on superscalar out-of-order and multi-threaded processors is challenging. While simulation produces accurate performance numbers, it is very slow and provides limited

insight into the factors that determine overall performance. In this PhD dissertation, we present a simple, yet accurate, model to analyze processor performance, called interval analysis. Using interval analysis, we develop an analytical model for estimating performance without needing detailed simulations, and we show its use in processor design through a pipeline depth and width study. We also use interval analysis to

develop a novel, hardware-performance counter architecture for constructing accurate cycle component stacks. Extending this counter architecture to SMT processors enables per-thread monitoring of single-threaded progress during multi-threaded execution; this enables better quality of service on SMT processors.

## Embedded Systems: Low-power Techniques and Applications

**Paolo Bennati**  
(bennati@dii.unisi.it)  
**Prof. Roberto Giorgi**  
**University of Siena, Italy**

Embedded systems are at the core of every modern electronic product, ranging from consumer electrical appliances to electronic games to communications equipment (including cell phones). An embedded system is a combination of computer circuitry and software built into a product that performs a narrow range of pre-defined tasks.

Requirements for embedded systems are very stringent and sometimes conflicting. Some of these requirements are cost pressure, long life-cycle, real-time requirements, reliability and low power consumption (especially for portable devices). It is difficult to successfully apply traditional computer design

methodologies and tools to embedded systems domain. The design of embedded systems is getting more and more challenging and it is attracting many people from industry and academia.

The contributions of this thesis are twofold. In the first part, the problem of power consumption, i.e. becoming a main constraint, is addressed and the filtered leakage-saving cache architecture is presented. Due to the rapid growth in magnitude of leakage current, static power reduction has become predominant and techniques for its reduction are required. Chip size is dominated by memory and in particular by caches. Therefore, caches represent a critical aspect that is worth analyzing. The filtered leakage-saving cache architecture aims to achieve low-leakage capabilities without degrading the performance. A tiny filter cache is

placed between CPU and L1 cache with leakage-saving capabilities (i.e., drowsy or decay). This simple solution further reduces the total leakage power and improves the performance making program execution faster.

In the second part, the BlueSign system is presented. It is an example of a complex application for a system with scarce computation capabilities, designed for specific tasks, such as a Personal Digital Assistant (PDA). The BlueSign is a system for the visualization of information in sign language for deaf. It is based on an avatar (a three-dimensional animated model). Signs have been coded in such a way that they are understandable to a computing system. The coding of the signs and the rendering engine has been developed to guarantee an efficient system that can run on a PDA device.



## Upcoming events


**HiPEAC Computing Systems Week,**  
Barcelona, Spain, June 2–5, 2008.



**5th HiPEAC Industrial Workshop,**  
Barcelona, Spain, June 4, 2008.



**ICS 2008, 22nd International Conference on Supercomputing,**  
Island of Kos – Aegean Sea, Greece, June 7–12, 2008



**PLDI 2008, Programming Language Design and Implementation,**  
Tucson, USA, June 7–13, 2008.

**LCTES 2008, Languages, Compilers, and Tools for Embedded Systems,**  
Tucson, USA, June 12–13, 2008. In conjunction with PLDI 2008.

**DAC 2008, 45th Design Automation Conference,**  
Anaheim, USA, June 8–13, 2008.



**SASP 2008, 6th IEEE Symposium on Application Specific Processors,**  
Anaheim, USA, June 8–9, 2008. In conjunction with DAC 2008

**SIES'2008, IEEE Third Symposium on Industrial Embedded Systems,**  
Montpellier - La Grande Motte, France, June 11–13, 2008.

**ICDCS 2008: The 28th International Conference on Distributed Computing Systems,**  
Beijing, China, June 17–20, 2008.

**ISCA 2008, 35th International Symposium on Computer Architecture,**  
Beijing, China, June 21–25, 2008



**MPSoC' 08, 8th International Forum on Application-Specific Multi-Processor SoC,**  
Aachen, Germany, June 23–27, 2008.



**HPDC 2008, High Performance Distributed Computing,**  
Boston, USA, June 23–27, 2008.

**ASAP 2008: 19th International Conference on Application specific, Systems, Architectures and Processors,**  
Leuven, Belgium, July 2–4, 2008.



**ACACES 2008, Fourth HiPEAC Summer School,**  
L'Aquila, Italy, July 13–19, 2008



**SAMOS VII: International Symposium on Systems, Architectures, Modeling and Simulation,**  
Samos, Greece, July 21–24, 2008, [http://sam0s.et.tudelft.nl/samos\\_viii/](http://sam0s.et.tudelft.nl/samos_viii/)



**International Symposium on System-on-Chip 2008,**  
Tampere, Finland, November 5–6, 2008, <http://soc.cs.tut.fi/>

**HiPEAC 2009 Conference,**  
Paphos, Cyprus, January 25–28, 2009



### 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)