HiPEAC info 30

COMPILATION ARCHITECTURE

APPEARS QUARTERLY
APRIL 2012

NETWORK OF EXCELLENCE ON HIGH PERFORMANCE AND EMBEDDED ARCHITECTURE AND COMPILATION

WELCOME TO THE HIPEAC SPRING COMPUTING SYSTEMS WEEK IN GÖTEBORG, SWEDEN, 24-25 APRIL 2012

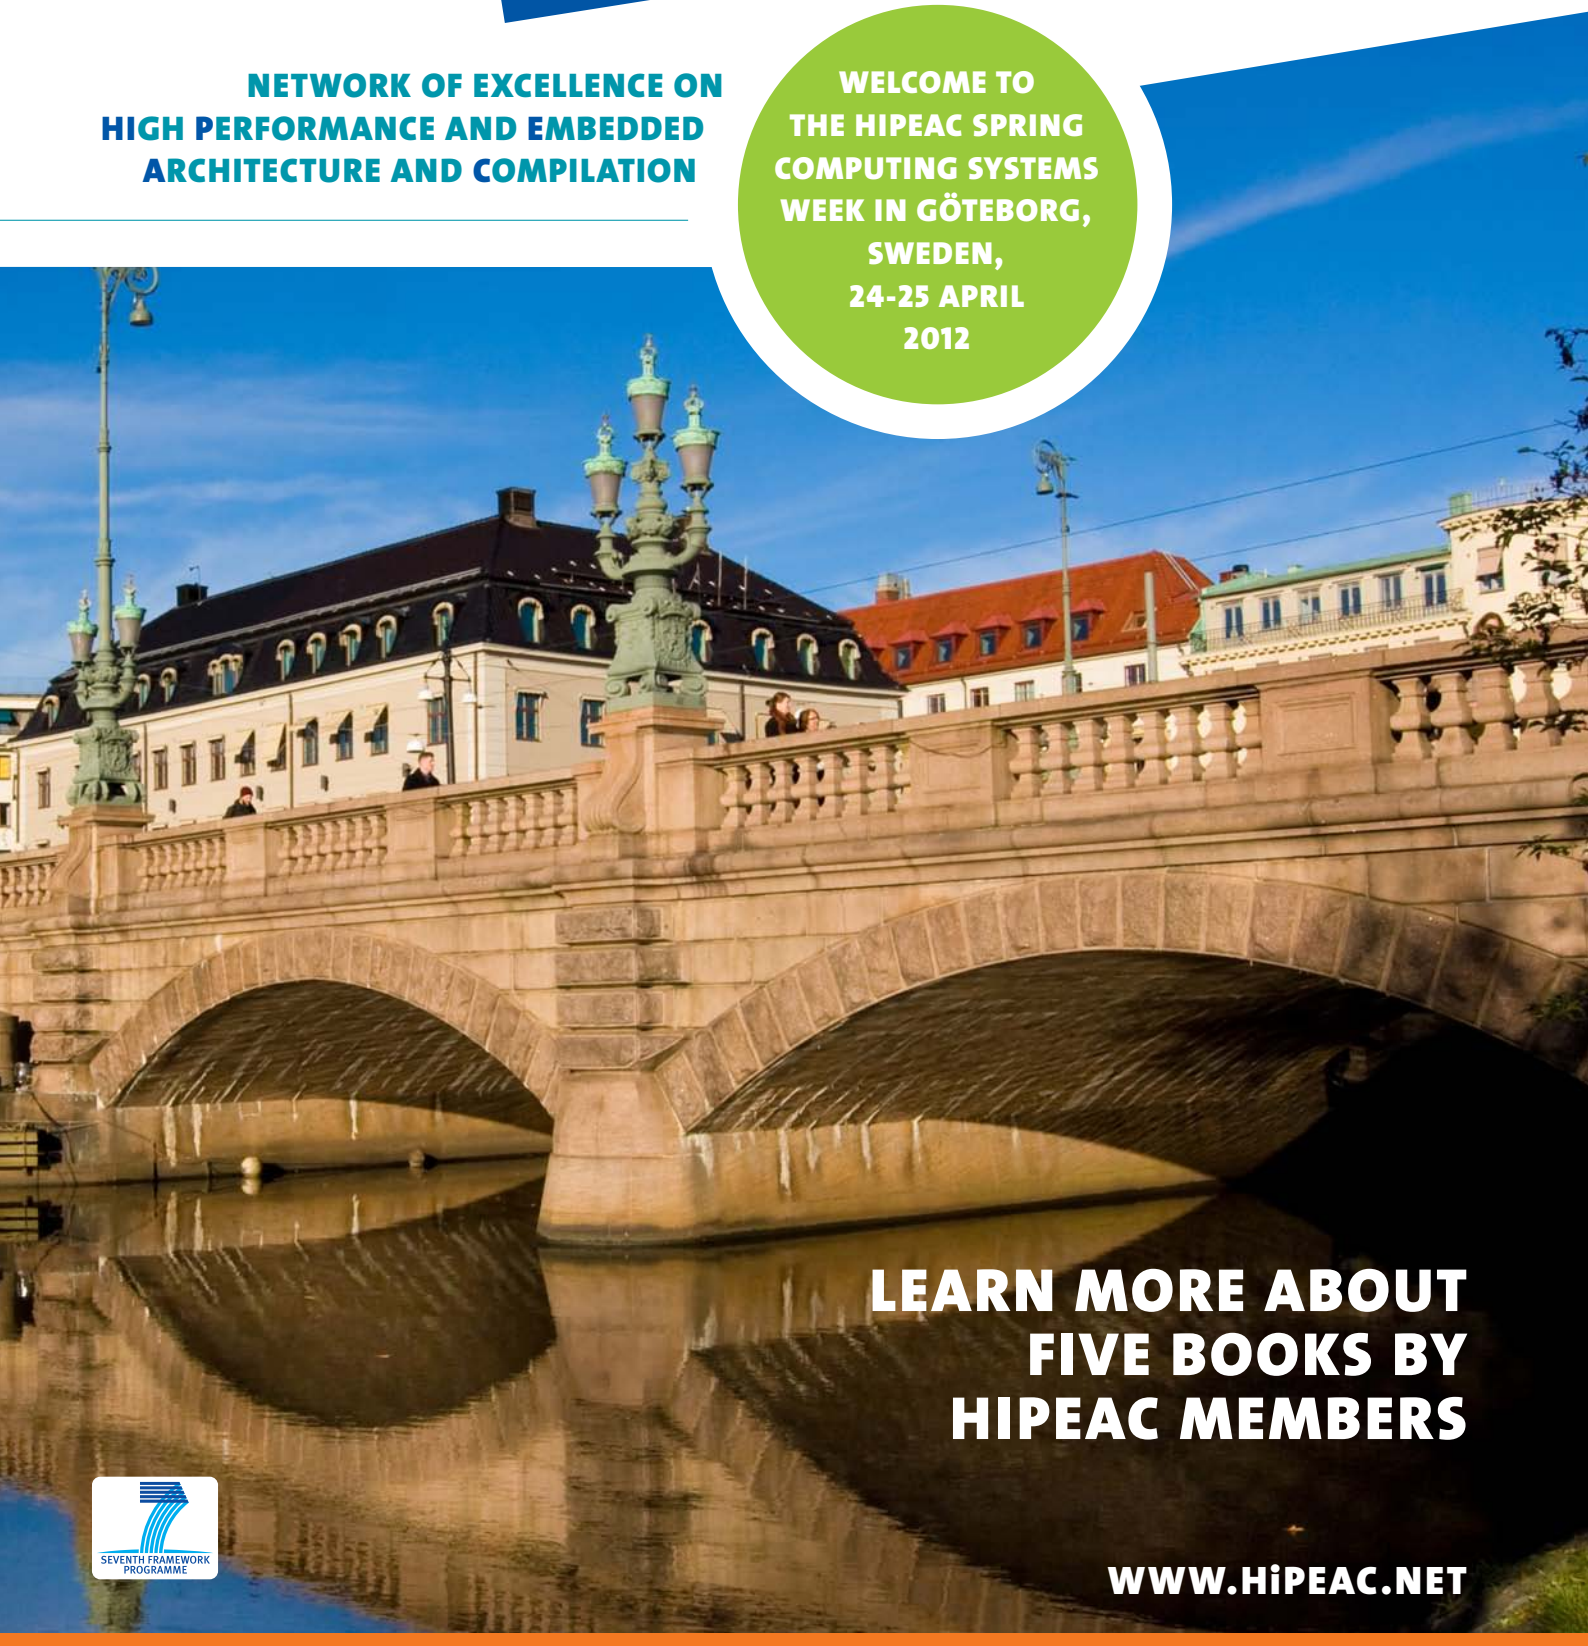

LEARN MORE ABOUT FIVE BOOKS BY HIPEAC MEMBERS

SEVENTH FRAMEWORK PROGRAMME

WWW.HiPEAC.NET

ACACES 2012, FIUGGI, ITALY, 8-14 JULY 2012

MESSAGE FROM THE HIPEAC COORDINATOR

According to a recent study by Cisco, by the end of 2012 there will be more mobile connected devices than there are humans, and 10 billion devices by 2016. Earlier this year, Apple's App Store celebrated its 25 billionth download. Not surprisingly, the malware for mobile devices increases too. According to Juniper Networks, the known malwares for Android went up from 400 in June 2011 to 13302 in December 2011. It is clear that the average user is not yet particularly concerned about the dangers of using mobile devices, and that urgent action is required to improve their security. If we fail to do so, this might seriously hamper the further deployment of mobile technology, and hence weaken one of the drivers for innovation in our economy.

The newsletter you just picked up is the first genuine HiPEAC3 newsletter. The HiPEAC newsletter is an important communication instrument for our network – the HiPEAC2 newsletter editor Rainer Leupers turned it into a real community magazine. I would like to thank him and his team for this excellent service to our community in HiPEAC2. I also welcome Per Stenström who will produce the HiPEAC3 quarterly newsletter. Together with his team, he will further develop it into an attractive magazine worth reading. We started 2012 with the HiPEAC 2012 conference in Paris. This was by far the biggest event ever organized by the HiPEAC network since its creation in 2004. We registered more than 500 delegates for a three-day event. This is almost three times more than for an average HiPEAC conference. From the survey we learned that the delegates appreciated the new conference model with many parallel activities going on at any time. The three-day event maximized the opportunities to discuss research, to network and to enjoy the city of Paris. Many thanks to Albert Cohen and the other local organizers for organizing this event for our community. In January 2013, the conference will take place in another great European city, Berlin.

HiPEAC is all about promoting research and networking. To further optimize our efforts, we have reorganized the bi-annual computing systems week to make it more community driven. Any HiPEAC member can now propose to organize a so-called thematic session. A thematic session is a small workshop focusing on a particular theme. HiPEAC will support the organizers by taking care of the logistics, and also by giving some financial support if needed. An inspiring thematic session can be the beginning of a future collaboration for the attendees. We hope that this new instrument will help to further structure our community. Eduard Ayguadé is responsible for coordinating the thematic sessions.

In order to better promote the European research in computing systems we have created a press room at the HiPEAC website. At the same time, we are actively building a network of journalists interested in our topics. We will regularly post press releases and stories in the press room. If you have an important HiPEAC-related message to share with the world, feel free to contact us to see how we can help you spreading it. Bilha Mendelson is responsible for the HiPEAC press room.

The HiPEAC network is working hard to strengthen the computing systems community in Europe. If you want to help, don't hesitate to contact us.

Koen De Bosschere

CONTENT

MESSAGE FROM THE PROJECT OFFICER
HORIZON 2020: COMMISSION PROPOSES €80 BILLION INVESTMENT IN RESEARCH AND INNOVATION

The European Commission has recently presented its proposal for Horizon 2020, a package of measures to boost research, innovation and competitiveness in Europe. For the first time, Horizon 2020 brings together all EU research and innovation funding under a single programme. The funding programmes are scheduled to run from 2014 to 2020. The Commission proposal will now be discussed by the Council and the European Parliament.
The European Commission has proposed to focus Horizon 2020 funds on three key objectives: *excellent science, competitive industries and better society*.

Excellent Science: support the EU's position as a world leader in science with a dedicated budget of €24.6 billion.

- support the most talented and creative individuals and their teams to carry out frontier research of the highest quality by building on the success of the European Research Council (ERC);
- fund collaborative research to open up new and promising fields of research and innovation through support for Future and Emerging Technologies (FET);
- provide researchers with excellent training and career development opportunities through the Marie Curie Actions;
- ensure Europe has world-class research infrastructures (including e-infrastructures) accessible to all researchers in Europe and beyond.

Competitive Industries: help secure industrial leadership in innovation with a budget of €17.9 billion.

- build leadership in enabling and industrial technologies, with dedicated support for ICT, nanotechnologies, advanced materials, biotechnology, advanced manufacturing and processing, and space, while also providing support for cross-cutting actions to capture the accumulated benefits from combining several Key Enabling Technologies;
- facilitate access to risk finance;
- provide Union wide support for innovation in SMEs.

Better Society: €31.7 billion will go towards addressing major concerns shared by all Europeans, across six key themes:

- Health, demographic change and well-being;
- Food security, sustainable agriculture, marine and maritime research and the bio-economy;
- Secure, clean and efficient energy;
- Smart, green and integrated transport;
- Climate action, resource efficiency and raw materials; and
- Inclusive, innovative and secure societies.

More detailed information about Horizon 2020 is available at:
http://ec.europa.eu/research/horizon2020/index_en.cfm?pg=home

Panagiotis Tsarchopoulos

MESSAGE FROM THE NEWSLETTER EDITOR

Dear colleagues,

I'd first like to follow Koen De Bosschere and extend my thanks to Rainer Leupers and his team for having taken the HiPEAC Newsletter to the next level. The Newsletter is indeed a valuable source for the community to keep track of all the exciting developments happening inside the HiPEAC network. It is exciting as well as challenging for me to take it from here. Fortunately, I am not doing this on my own. It is my pleasure to introduce Dr. Rubén Titos Gil who will work with me to deliver the Newsletter four times a year. Let me also take this opportunity to encourage the community to continue to contribute to the Newsletter. Without contributions, no Newsletter! We are seeking contributions including, but not limited to events, technical developments, announcements of new books. The deadline for contributions is in the first week of March, June, September, and December and the next deadline is on June 6. I hope that you will appreciate this issue of the Newsletter.

Per Stenström

THE HIPEAC 2012 CONFERENCE

This edition of the conference was the first to offer a much more attractive networking event for the delegates and a much more rigorous paper selection process for contributors.

A BRAND NEW FORMAT FOR A SUCCESSFUL NETWORKING EVENT

The HiPEAC conference is the flagship event of the HiPEAC network. This year's conference took place in Paris on January 23 to 25. Thanks to its new format, it attracted 518 delegates from all over the world, almost tripling its attendance and exposure in comparison with previous editions.

This is thanks to a completely new and innovative conference model. The three-day program was packed with parallel, high-quality technical meetings, and also offered numerous opportunities for networking, establishing collaborations, discussing employment and innovation.

CONFERENCE HIGHLIGHTS

- 3 keynote talks communicating challenges and perspectives from academia, industry, and the European Commission;
- 37 presentations of TACO journal papers selected for the HiPEAC main track;
- 17 workshops and 9 tutorials with a wide coverage of high-performance and embedded system research and innovation, and a wide variety of activities (panels, posters, keynote speakers);
- AN INNOVATIVE, networking-oriented poster session sponsored by IMC, with more than 60 posters, and student presenters competing for 8 Google travel grants and NVIDIA graphics cards;
- A FIRST-OF-A-KIND industry exhibition floor with 19 stands and 50 companies represented, live demonstrations, and a chance for companies to invite students met during the earlier poster session;
- THE NEXT DAY, the exhibition floor opened for stands and posters representing 46 European projects, associated with a special exchange session for project coordinators led by the HiPEAC project officer, Panagiotis Tsarchopoulos;
- THE TRADITIONAL HiPEAC social event, a not-to-be-missed dinner cruise on the Seine river.

NETWORK, EXCHANGE, DISCOVER

The conference created a unique opportunity for the complete portfolio of EU-funded computing systems projects to interact with computing systems researchers, and for more than 50 companies to recruit among the 125 attending graduate students.

The anonymous survey conducted at the end of the conference showed an overwhelming majority of participants planning to attend again next year, many of them planning to submit papers and to organize associated workshops and tutorials.

This success would not have been possible without the generous support from 20 sponsors, most of them also participating in the industry exhibit. Such a mutual-benefit partnership is essential to the sustainability of the HiPEAC network's activities, to offer attractive registration fees for the conference, and to support student participation.

More information and online material about each and every event can be found on the conference web site:
http://www.hipeac.net/conference

We expect HiPEAC 2013 in Berlin to be an even more successful event. The conference model will retain its main highlights, with a continued emphasis on the networking opportunities, and calling for the organization of a wide variety of high quality associated activities. HiPEAC 2013 will strengthen the conference's series as the premier forum for research and innovation in computing systems.

Albert Cohen

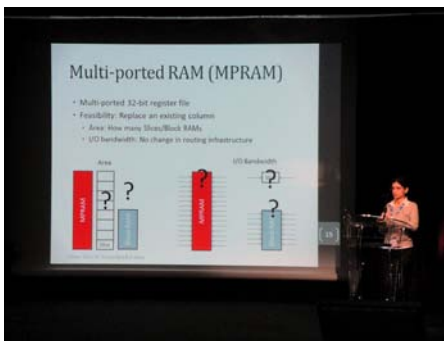

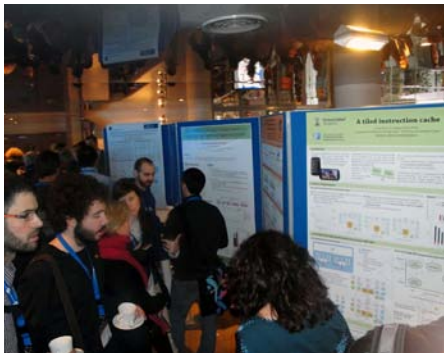

A SUCCESSFUL EXPERIMENT WITH A NEW PUBLICATION MODEL

Papers in computing systems are typically published first in conferences and then extended versions of them are occasionally published in archival journals. While the merit of using conferences as the primary publication vehicle is rapid dissemination, there are several quality issues. First, the review process associated with conferences is not as solid as that of journals. Papers not included in the conferences will have to be revised and submitted again and reviewed by a new set of referees because the review history of papers is not maintained. Second, papers published in conferences are not indexed by citation indexes generally accepted across disciplines and this hurts individual contributors as well as computing systems as a discipline. In contrast, journals have a more rigorous review process in which authors of papers on a clear route to being accepted for publication will be engaged in a constructive dialogue with the reviewers. The downside of journal reviewing, however, has been long turn-around times. But it doesn't have to be that way.

HiPEAC is among the first to have conducted an experiment with a new publication model in which original contributions are submitted to a journal to take advantage of its rigorous review process. The new publication model steers *original* research papers to a journal first and uses the conference to showcase the "catch of the year." This combines the more rigorous and fair review process of journals with the visibility and interaction opportunities of a conference presentation. Here, the journal serves as a publication venue for original research results, while the conference serves as a meeting place for the whole community; a model that works very well across many scientific disciplines.

To get experience with this new publication model ACM Transactions on Architecture and Code Optimization (TACO) and the HiPEAC conference experimented with the model for HiPEAC 2012. To secure as short a review turn-around time as for conferences, we recruited around 80 reviewers that were committed to review a handful of papers each. Submissions for the conference were automatically forwarded to the TACO website. By July 15, 2011, 117 original papers were submitted to ACM TACO as compared to around 80 in previous editions of the HiPEAC conference. This clearly demonstrated a keen interest in this new model.

After the first round of reviews, after only 6 weeks, 7 papers were conditionally accepted with minor revisions, whereas 46 papers had to undergo major revision and the rest were rejected. In a conference model, the acceptance bar would have to be put among the papers in the major-revision category. All authors of the papers in the major-revision category got 5 weeks to improve the papers based on the recommendations by the reviewers. Two papers were not resubmitted. The reviewers then got two weeks to review the revised versions. After this second round 17 papers were finally accepted and another 21 papers were conditionally accepted. In the end, 37 papers were accepted, and 7 papers are still in the TACO queue and might eventually make it into a regular TACO paper. In total, 30 papers benefitted greatly from the additional review rounds that were offered. The 37 accepted papers have all been thoroughly copy-edited, and were published in the ACM digital library in January 2012 – exactly 6 months and 5 days after the submission deadline. The experiment shows that one can afford two review rounds in the same amount of time normally allocated for conference paper reviewing.

Based on the positive experience with this new publication model, it will be used again for the next edition of the HiPEAC conference.

Per Stenström

HIPEAC MINI-SABBATICAL - YIANNAKIS SAZEIDES, UNIVERSITY OF CYPRUS

Report for HiPEAC sponsored Mini-Sabbatical at ARM, Cambridge

During the summer of 2011 I had the great opportunity to spend a three-week HiPEAC sponsored mini-sabbatical at ARM, Cambridge. I am very grateful to HiPEAC for supporting this activity and for ARM for hosting me. I am in particular in debt to Emre Ozer for making this happen. Finally, I like to deeply thank my wife for her support and courage to take care our family while I was away in Cambridge.

My acquaintance and collaboration with Emre started in the context of HiPEAC activities. The mutual interest and technical discussions about reliability lead us into an interesting proposition: the convergence of reliability and security features in future processors. We investigated the basic premises of this hypothesis during the mini-sabbatical by doing an extensive literature review and discussing with a number of ARM engineers about the intertwining of reliability and security. One of the directions that appeared promising is an exploration of methods that can help differentiate between coordinated and stochastic faults. During brainstorming sessions we set the guidelines for what these methods should aimed for and subsequently proposed a number of techniques that satisfied these aims. Due to limited time we decided to focus and develop further only few of these techniques. This effort has lead to the design of three different mechanisms: one of them has been filed for a patent, another is under preparation to be filed for patent in the coming weeks, and the third lead to a paper that has been accepted for publication.

In retrospect the mini-sabbatical at ARM has been a very rich and productive experience. The opportunity for direct exchange with ARM experts on reliability and security was a catalyst towards the developments of our various ideas and techniques. The line of work that started during the mini-sabbatical is further pursued and it is also central to a research proposal under preparation related to reliability and security.

Yiannakis Sazeides

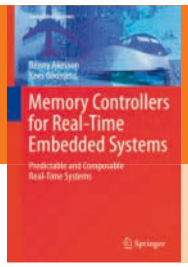

BOOK ON MEMORY CONTROLLERS FOR REAL-TIME EMBEDDED SYSTEMS

Benny Akesson and Kees Goossens, Eindhoven University of Technology

Verification of real-time requirements in systems-on-chip becomes more complex as more applications are integrated. Predictable and composable systems can manage the increasing complexity using formal verification and simulation. This book explains the concepts of predictability and composability and shows how to apply them to the design and analysis of a memory controller, a key component in any real-time system. This book is generally intended for readers interested in Systems-on-Chips for real-time applications. It is especially well suited for readers looking to use SDRAM memories in systems with hard or firm real-time requirements. There is a strong focus on real-time concepts, as well as a brief discussion about memory controller architectures for high-performance computing. Readers will learn step-by-step how to go from an unpredictable SDRAM memory, offering highly variable bandwidth and latency, to a predictable and composable shared memory, providing guaranteed bandwidth and latency to isolated applications. This journey covers concepts for making memories and arbiters behave in a predictable and composable manner, as well as architecture descriptions of hardware blocks that implement the concepts.

BOOK ON SCALABLE MULTI-CORE ARCHITECTURES: DESIGN METHODOLOGIES AND TOOLS

Dimitrios Soudris, National Technical University of Athens, Greece and Axel Jantsch, KTH Royal Institute of Technology, Sweden

Designing an application specific multi-core system-on-chip is a tricky endeavor and requires making dozens of system level decisions with profound impact on performance and cost at an early time in the design phase, when many details are still unknown and performance estimates are notoriously inaccurate. Technology development is progressing and is continuously increasing our raw computing capacity. At the same time, application requirements and standards become more demanding. Together this leads to complex designs that require sophisticated and continuously developed architectures, tools, and methodologies. In particular, since the field during the last decade has embarked to massively increase the number of heterogeneous cores, the picture has become very complicated. Moreover, it is quickly changing with respect to demands on architectures and tools. Hence, the design technology seems to fall more and more behind manufacturing technology, which is also referred to as the design productivity gap. In order to get to grips with this disturbing situation, researchers have tried to develop scalable architectures and tools that continue to be efficient and usable even as the number of cores dramatically grows. The promise of scalable architectures and tools is that we can use them in 2011 with 20 or 50 cores, and we can still use them in 2017 with 200 or 500 cores. Once a scalable architecture for an application domain or a scalable tool for a design problem has been developed, the design research stops running behind the semiconductor manufacturing advances and can truly focus on exploiting the given capabilities of technology. The European FP7 MOSART project (http://www.mosart-project.org/) has set as its goal to develop scalable solutions to architectures, tools and methodologies. Even though not all the work of the project has been fully documented, several important results of our work on architecture and hardware (part I), system level design and exploration (part II), and applications (part III) are reported. After a forward by Dr. Panagiotis Tsarchopoulos, EC, the first part includes techniques and methodologies regarding the memory architecture and management in an NoC platform, the Dynamic Data Management Services in MPSoCs, the ASIP exploration and design and the power management and clocking of NoC-based platforms. The second part includes the tool description of system-level exploration and parallelization of an application onto a Multicore Platform. Lastly, the third part describes two industrial case studies where the developed tool-supported design flow was applied to, namely the Software Defined Radio and the PHY layer of WiMAX.

BOOK ON HARDWARE/SOFTWARE CO-DESIGN FOR HETEROGENEOUS MULTI-CORE PLATFORMS: THE HARTES TOOLCHAIN

Koen Bertels, Delft University of Technology

This book describes the results and outcome of the FP6 project, known as hArtes, which focuses on the development of an integrated tool chain targeting a heterogeneous multi core platform comprising of a general purpose processor (ARM or PowerPC), a DSP (the diopsis) and an FPGA. The tool chain takes existing source code and proposes transformations and mappings such that legacy code can easily be ported to a modern, multi-core platform. Benefits of the hArtes approach, described in this book, include:

- The use of a familiar programming paradigm: hArtes proposes a familiar programming paradigm which is compatible with the widely used programming practice, irrespective of the target platform.
- Enables users to view multiple cores as a single processor: the hArtes approach abstracts away the heterogeneity as well as the multi-core aspect of the underlying hardware so the developer can view the platform as consisting of a single, general purpose processor.
- Facilitates easy porting of existing applications: hArtes provides a migration path where either through manual annotation or the use of the tool chain to apply the necessary modifications, one can test on the real platform how the application behaves and, when necessary, repeat the process if the design objective has not been met.
- Enables development of new applications using powerful toolboxes: the hArtes tool chain provides both high level algorithm exploration tools with subsequent, automatic code generation which can then be fed to other toolboxes in the chain.
- Employs an open tool chain architecture: any development tool can be integrated in the tool chain, so users are not locked into a single vendor technology.
- Allows users easily to retarget to new hardware platforms: the same development tools and environments can be used, no matter what hardware platform one targets.
- Maps from fully automatic to fully manual: developers can choose to opt for a full automatic mapping, semi-automatic or even fully manual. At each step, decisions can be evaluated and overruled if considered inadequate

BOOK ON SCALABLE PARALLEL PROGRAMMING APPLIED TO H.264/AVC DECODING

Ben Juurlink, Mauricio Alvarez-Mesa, Chi Ching Chi, Arnaldo Azevedo, Cor Meenderinck and Alex Ramirez

The book describes the authors' experiences in developing highly efficient and scalable implementations of H.264 decoding. The work they present has been supported by several past and present European projects, in particular SARC and ENCORE, as well as by HiPEAC grants. They were invited by Springer to write a book, because one of their articles was among the top-downloaded articles of the Journal of Signal Processing Systems. The book will appear in the SpringerBriefs series. The main purpose of the series it to distribute emerging, cutting-edge research quickly, primarily through an e-publication, though print will be available. The book has not been released yet, but can already be pre-ordered at Amazon, amongst others.

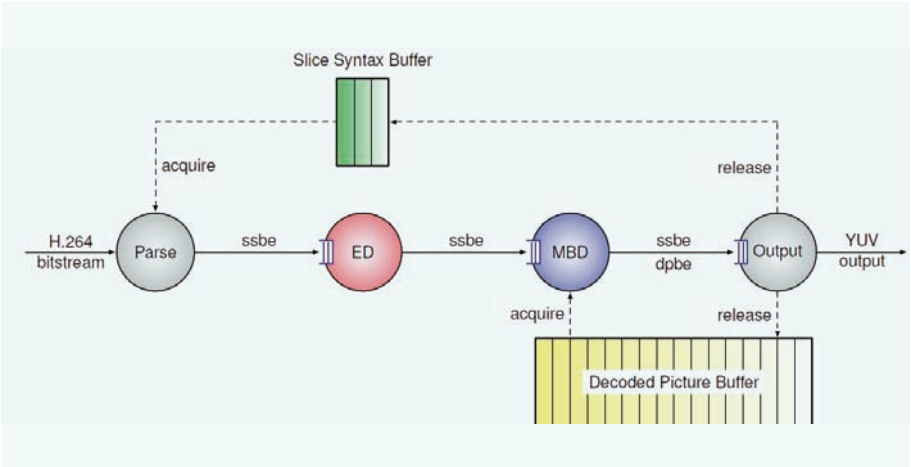

H.264 decoding pipeline

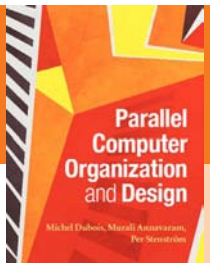

BOOK ON PARALLEL COMPUTER ORGANIZATION AND DESIGN

Michel Dubois, Murali Annavaram and Per Stenström

Two emerging forces in computer architecture today, the advent of parallelism at all levels of the hardware stack in modern computer systems and the move away from performance as the quintessential design concern, motivated this new computer architecture textbook for senior undergraduate and graduate courses.
Complexity, power, and reliability have emerged as first-rate design factors. In fact, the shift from performance to power has driven the pervasive adoption of parallelism in architecture design today. These two trends are interrelated and are covered in this light throughout Parallel Computer Architecture and Design, by Michel Dubois, Murali Annavaram and Per Stenström.

The book focuses exclusively on the functional design aspects of modern computer architectures and the reasons behind design choices, from a functional point of view. This allows the authors to explain and analyse, in great depth, complex concepts such as speculative execution or coherence and memory consistency models. By understanding these basic concepts, students can build intuition, which they can then use later in more advanced courses, in their research, or in their engineering practice.

The text also gives reasoning frameworks to clarify and understand more complex systems, without going into complex theoretical formalism. This means the book is compact and that it is possible to teach this material in one academic semester or two academic quarters at the graduate level. It includes numerous homework exercises, the solutions to which will be provided on the book's website.

For further information please go to www.cambridge.org/comp

LUK VAN ERTVELDE OF GHENT UNIVERSITY FIRST TO WIN NEW SPEC DISTINGUISHED DISSERTATION AWARD

The Research Group of the Standard Performance Evaluation Corp. (SPEC) has awarded one of its first two Distinguished Dissertation Awards to Luk Van Ertvelde of Ghent University in Belgium. The awards will be presented at the International Conference on Performance Engineering (ICPE), which will take place on April 22-25, 2012 in Boston, Mass., USA.

The SPEC Distinguished Dissertation Award was established in 2011 to recognize outstanding dissertations within the scope of the SPEC Research Group in terms of scientific originality, scientific significance, practical relevance, impact, and quality of the presentation.
Luk Van Ertvelde was nominated by Professor Lieven Eeckhout from Ghent University. He receives the award for his dissertation titled "Workload Generation for Microprocessor Performance Evaluation." The dissertation proposes and studies three workload generation and reduction techniques for microprocessor performance evaluation. (1) It proposes code mutation, a novel methodology for hiding proprietary information from computer programs while maintaining representative behavior; code mutation enables dissemination of proprietary applications as benchmarks to third parties in both academia and industry. (2) It contributes to sampled simulation by proposing NSL-BLRL, a novel warm-up technique that reduces simulation time by an order of magnitude over state-of-the-art. (3) It presents a benchmark synthesis framework for generating synthetic benchmarks from a set of desired program statistics. The benchmarks are generated in a high-level programming language, which enables both compiler and hardware exploration.
The SPEC Research Group (RG) serves as a platform for collaborative research efforts in the areas of computer benchmarking, performance evaluation, and experimental system analysis in general, fostering the interaction between industry and academia in the field. The scope of the conducted research efforts includes methodologies, techniques and tools for measurement, load testing, profiling, workload characterization, dependability and efficiency evaluation of computing systems.

NEW HIPEAC MEMBER: DIMITRIS GIZOPOULOS, UNIVERSITY OF ATHENS

Dimitris Gizopoulos is an Associate Professor at the Department of Informatics & Telecommunications, University of Athens where he leads the Computer Architecture Laboratory. Before joining the University of Athens, he led for ten years the Computer Systems Laboratory at the University of Piraeus. He holds a Computer Engineering & Informatics diploma from the University of Patras (1992) and a PhD from the University of Athens (1997).

His research focuses primarily on dependable and fault tolerant computer architecture, and in particular error detection, on-line testing, fault tolerance, and design validation for general purpose microprocessors, embedded processors, as well as multicore/multithreaded processors. Gizopoulos has published 110 papers in peer reviewed transactions/journals (40 papers) and conference proceedings (70 papers). He is also inventor of a US patent on arithmetic units self-test, author of a book and editor of a second (both published by Springer). He is Associate Editor for IEEE Transactions on Computers, IEEE Transactions on VLSI Systems, IEEE Design & Test of Computers Magazine, and Springer's Journal of Electronic Testing: Theory and Applications, as well as guest editor for several special issues in IEEE Transactions on Magazines. He served several times as member of the Steering, Organizing and Program Committees of major IEEE and ACM events (DATE, ITC, MICRO, VTS, ETS, IOLTS, DFTS, ATS, etc) and he has been to General Chair and Program Chair of IEEE IOLTS and DFTS symposia several times. He is member of the executive committee of IEEE Computer Society's TTTC. He is a senior member of IEEE, a Golden Core Member of IEEE Computer Society, and a member of ACM and SIGARCH. He participated in and led several research and development projects funded by the European Union, the Greek Government and private funds and companies and has received donations by major companies.

Gizopoulos is currently coordinator of the HOLISTIC project funded by the Greek government and the European Union that focuses on reliability at different levels of abstraction for multicore architectures. He is also Vice Chair of the ESF-funded COST Action MEDIAN on manufacturable and dependable multicore architectures at nanoscale.

He plans on participating in the HiPEAC community in areas such as dependability/reliability of general purpose and embedded computing systems as well as other architectural aspects (performance, power) for massively parallel and heterogeneous architectures.

Contact information:
Prof. Dimitris Gizopoulos
dgizop@di.uoa.gr
http://www.di.uoa.gr/~dgizop

LOW POWER HIGH PERFORMANCE GPU RESEARCH CONSORTIUM AWARDED €2.8M R&D GRANT

LPGPU
Advanced Lighting and AI for Mobile GPUs

The EU Commission has awarded a European games consortium a grant of €2.8 million to research and develop the next generation of low power high performance Graphics Processor Units (GPUs) to be used in future games consoles and mobile devices.

The consortium includes four European technology companies: Codeplay, an Edinburgh based GPU technology company, Cambridge-based Geomerics, a graphics technology company, AiGameDev.com, an Austria-based game AI company and Think Silicon, a Greek low gate-count graphics semiconductor IP core company. Two European universities, TU-Berlin (Germany) and Uppsala University (Sweden) complete the group.

The research project will run over three years and the key objectives are:

- To enhance Europe's lead in mobile graphics and video games technologies
- To enable the next-generation of advanced graphics technologies for videogames and mobile devices. In particular to achieve power and bandwidth reductions of 2x or more on real-world software;

- To produce applications for power-efficient GPUs including advanced lighting and game AI techniques;
- To provide answers to critical questions about the development of future parallel processors;
- To produce research that will help target the fast-growing smart-phone and graphics processor technology markets;
- To provide an analysis framework to help software developers produce graphically innovative software in the future.

Andrew Richards, CEO at Codeplay said:
"We are delighted the EU has recognised the need to invest in the future of consumer technology. This project will have real tangible, practical benefits and our ambition is that it will allow Europe to take the lead in the development of the next generation of low power high performance mobile and videogame graphics technology."

"European companies lead in the design of mobile phone CPUs and GPUs, and are world leaders in videogame technology. These companies need to make large investments in R&D for graphics, for which it is vital that they have reliable information. To date there has been little academic research in this area. R&D has previously been dominated by high performance programming in sectors such as aerospace, medicine and energy. This project will help to redress that balance by delivering advances in graphics applications, tools and hardware for consumer devices."

Chris Doran, Founder of Geomerics, commented:
"At Geomerics we have tracked the rise in performance of mobile platforms very closely. It is clear that they are fast approaching console performance, and at the high end will overtake this in 2012. We are very excited about the prospect of taking our console technology to mobile platforms, and to contributing to joint R&D on the future direction of graphics technology for mobile devices."

Professor Ben Juurlink at TU Berlin adds:
"This research project will allow the universities in the consortium to extend our leadership position in the field of low-power computer architecture, parallel applications and multi-core architectures. The main market areas for increased processor performance over the next few years are graphics and videogames where the companies in the consortium are world leaders. The project is expected to strengthen their market position and a manifold return on investment is envisaged."

For more information about the project visit http://lpgpu.org/wp/.

IOLANES: SCALING STORAGE I/O ON MULTICORE PLATFORMS

Are you already using SSDs to speed up your workloads? Now you can get ready for the real problems.

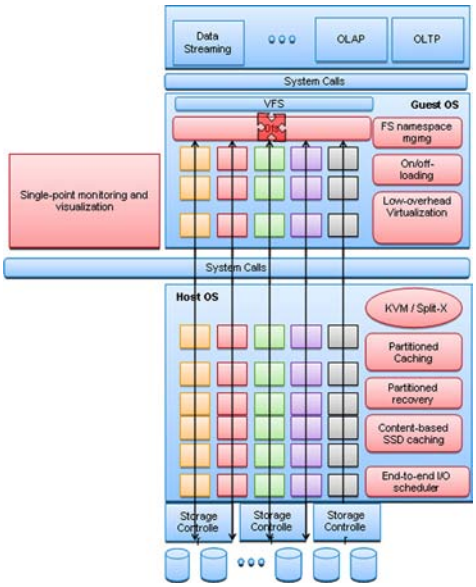

IOLanes partitioned I/O architecture

Storage I/O in modern and future servers is becoming a main bottleneck and design constraint. As modern infrastructures strive to consolidate data and content, improve dependability and integrity without compromising application performance, a number of questions emerge: How much I/O can modern servers support? How much will they need to support in 10-years? What happens to the I/O stack as the number of cores increases? What is the cost of virtualization? How can we handle mixed storage devices, namely SSDs and disks? How can we deal with workload interference in multi-tenant workloads and virtualization?

With emerging storage device technologies, such as solid state disks (SSDs), systems that are capable of millions of I/O operations are expected to become commonplace. This trend shifts the bottleneck from the I/O devices to the I/O path. Recently, industry has started to also shift focus from merely using faster I/O devices, such as SSDs instead of disks, to innovations in the I/O path for achieving scale-out for big data applications.

Today's system software in the I/O path exhibits high overheads and poor scaling when increasing the number of cores and storage devices; shared structures, replication of functionality, synchronization requirements, and workload interference are on the way of supporting current and future I/O intensive applications that end-up consuming many times more cycles to perform each I/O operation when the number of cores increases even in low ranges, as for instance, from one to eight. By looking at the cycles used per I/O operation IOLANES is able to characterize application requirements, evaluate improvements, and project future needs. Early results show that today's I/O stack does

not scale beyond 4-6 cores and that it exhibits high overheads, especially in virtualized environments.

IOLANES (www.iolanes.eu, 2010-2012) is a EU-funded research project targeted at understanding and improving the I/O performance in modern hardware that employs multicore architectures by adapting or redesigning the I/O stack and by providing system-level support that will allow future storage systems to take advantage of multicore CPUs in new ways.

IOLANES re-architects the systems software in the I/O path to reduce overheads from application to devices, allow I/O to scale as physical resources increase, reduce interference with multi-tenant workloads, and reduce effort for leveraging multicores when performing storage I/O. IOLANES designs I/O subsystems that will be capable of sustaining one order of magnitude more load than todays' systems, reducing the required infrastructure cost and energy to support large workloads and data sets. IOLANES evaluates these mechanisms using OLTP, OLAP, and data streaming applications and examining their impact on I/O efficiency. To better understand complex interactions IOLANES provides a single-point for performance monitoring and visualization of the I/O path from applications to devices.

IOLANES is supported by partners Foundation for Research and Technology – Hellas (FORTH) - Coordinator, Barcelona Supercomputing Centre (BSC), University of Madrid (UPM), INTEL Performance Labs Ireland, IBM Research Labs in Haifa, and Neurocom S.A. IOLANES is funded by the E.C. under the 7th Framework program and is part of the portfolio of the Embedded Systems Unit - G3 Directorate General Information Society.

Angelos Bilas,
FORTH-ICS, Greece

THE 520-CORE HETEROGENEOUS HARDWARE PROTOTYPE

An FPGA-based board that targets a combination of SRAM, DRAM and GTP high-speed serial connectors at a low cost

Formic is designed to be a building block for large systems and that's exactly what we did recently in FORTH-ICS. We built a 4x4x4 cube of 64 Formic boards, interconnected in a 3D-mesh fashion. Each board uses six GTP links (the small black cables in the picture, two per dimension) to connect to its X, Y and Z neighbors. The FPGAs contain our prototype non-coherent hardware architecture, which fits eight Xilinx MicroBlaze cores in each board. Every core has private L1 and L2 caches. The cores communicate among them, as well as with the other boards, using a crossbar-based network-on-chip. The cube is physically made of plexiglas and uses large, silent fans to cool the system efficiently. Long JTAG and I^2C chains are present for debugging and monitoring. A modified PC power supply unit measures the cube consumption, which averages 460 Watts.

The 512-core cube prototype is connected to two ARM Versatile Express platforms. Each such platform features a quad-core ARM Cortex A9 processor coupled with an FPGA daughterboard. We have developed logic for this FPGA which connects the four ARM cores to the cube network-on-chip through its compatible GTP links (the red longer cables in the picture). Our logic also enhances the ARM cores with the special features of our prototype architecture, such as Mailboxes, per-core DMA engines and DMA completion-monitoring counters. The total system is a heterogeneous, non-coherent 520-core architecture of 8 "large" and 512 "small" cores. We see this system as a realistic model of a future single-chip processor, where a small number of strong cores will handle the complex parts of the programs, such as scheduling and control, and a large number of weaker cores will perform the bulk of the work, such as the data processing. As partners in the European Union ENCORE project (http://www.encore-project.eu) we are developing a runtime system for this architecture which executes applications in the OMP-SS task-parallel programming model in this fashion. The programmer splits a sequential program in small tasks with annotated I/O arguments and the runtime system automatically parallelizes these tasks. The control-intensive part of the runtime runs in the fast ARM cores and deals with preserving the data dependencies, with scheduling the tasks onto the worker cores and with keeping the system memory coherent using DMA transfers. The tasks themselves run onto the slow, but many, MicroBlaze worker cores in the cube.

For more information about the Formic board and our FPGA hardware prototype, please visit http://formic-board.com. If you are interested in the research we do in FORTH-ICS in Scalable Multicore Systems, you can find out more in http://www.ics.forth.gr/carv

Spyros Lyberis,
FORTH-ICS, Greece

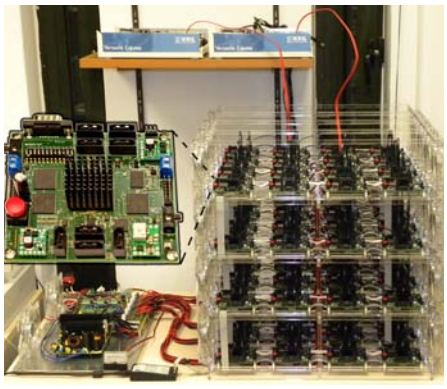

The 4x4x4 cube of 64 Formic boards (right), the two ARM Versatile Express platforms (top) and a closer look at the Formic board (left, inset)

DRAMPOWER: OPEN-SOURCE DRAM POWER AND ENERGY ESTIMATION TOOL

Karthik Chandrasekar, Benny Akesson, Kees Goossens, TU Delft and TU Eindhoven

DRAMPower is an open-source tool for accurate DRAM power and energy estimation.

The DRAMPower open-source tool for fast and accurate DRAM power and energy estimation was released in February 2012. This tool is based on the DRAM power model developed jointly by TU Delft and TU Eindhoven and is targeted at users and designers of DRAM memory controllers. The tool conforms to the JEDEC specifications for DDR2 and DDR3 memories and employs datasheets provided by DRAM memory vendors for architecture, timing, current and voltage specifications.

This tool enables users and designers of DRAM memory controllers to evaluate different memory controller policies on the basis of their energy consumption. This tool implements the basic DRAM power model proposed in our paper titled, "Improved Power Modeling of DDR SDRAMs", which was presented at the 14th Euromicro Conference on Digital System Design (DSD) 2011. Currently, the tool supports common memory operations like read, write, refresh, activate and precharge/auto-precharge. In the near future, we plan to extend this tool to support the Power-down and Self-refresh modes as well. Support for LPDDR and LPDDR2 memories will also be added to the tool soon.

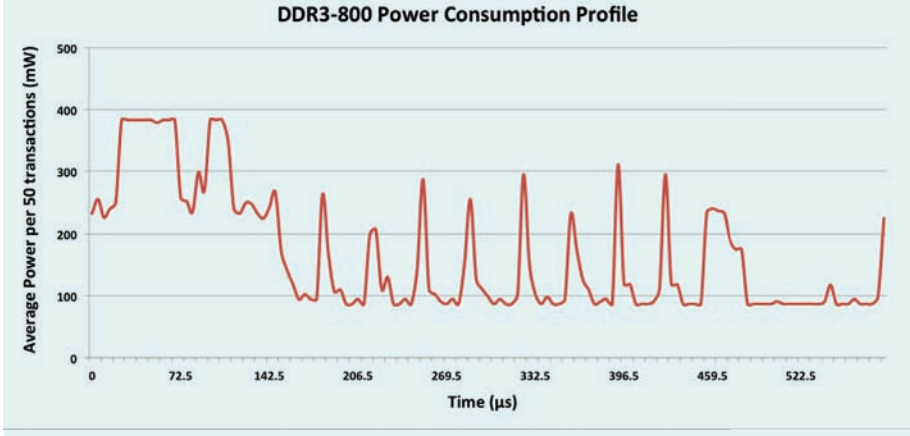

Memory power profile of a sample trace using Micron's 1Gb DDR3-800 device"

The tool accepts a DRAM command trace (sent from the memory controller to the memory) as input and performs a cycle-by-cycle examination of the same, thus considering all the command scheduling decisions of the memory controller. This results in an accurate study of the impact of the memory controller policies and helps evaluate them from the power perspective. This detailed analysis gives a precise estimate of power and energy consumption of the different memory operations as opposed to methods/tools using only the minimum timings of the memory device to derive power and energy estimates.

We have released this tool under the DRAMPower Public License (DPPL 1.0), which is based on the Simple Public License (SimPL-2.0). This license gives the end-users and developers the flexibility to employ, develop and re-distribute the source code with minimal obligations.

For more information, or to download the tool, please refer to the official DRAMPower website: http://www.es.ele.tue.nl/drampower

Further questions about the tool and the software package can be directed to Karthik Chandrasekar (k.chandrasekar@tudelft.nl).

SNIPER MULTI-CORE SIMULATOR V2.0 RELEASED

Lieven Eeckhout, Ghent University

Version 2.0 of the Sniper simulator was recently released with a notable new feature. It is now possible to simulate multi-programmed workloads together in Sniper using our new trace format. Along with this new feature, we have also released a number of bug-fixes and improvements. It is available for download at http://snipersim.org and can be used freely for academic research.

Sniper is a next generation parallel, high-speed and accurate x86-64 simulator. This multi-core simulator is based on the interval core model and the Graphite simulation infrastructure, allowing for fast and accurate simulation and for trading off simulation speed for accuracy to allow a range of flexible simulation options when exploring different multi-core architectures. Using this methodology, we are able to achieve good accuracy against hardware for up to 16-threaded applications.

The Sniper simulator allows one to perform timing simulations for multi-threaded, shared-memory applications and multi-programmed workloads with 10s to 100+ cores, at a high speed when compared to existing simulators. The main feature of the simulator is its core model which is based on the interval core model, a fast mechanistic core model. The interval model raises the level of abstraction in architectural simulation which allows for faster simulator development and evaluation times; it does so by 'jumping' between miss events, called intervals. On recent multi-core hardware, we see simulation speeds up to 3 MIPS.

This simulator, and the interval core model, is useful for uncore and system-level studies that require more detail than the typical one-IPC models. As an added benefit, the interval core model allows the generation of CPI stacks, which show the number of cycles lost due to different characteristics of the system, like the cache hierarchy or branch predictor, and lead to a better understanding of each component's effect on total system performance.

Visit http://snipersim.org for more information.

MADES PROJECT: SYSML/MARTE HIGH LEVEL METHODOLOGY FOR REAL-TIME AND EMBEDDED AVIONICS SYSTEMS

A holistic model-driven approach for the design, validation, simulation, and code generation of complex embedded systems

The MADES FP7 project aims to develop novel model-driven techniques for the design, validation, simulation, and code generation of complex real-time and embedded systems for avionics and surveillance embedded systems industries, in order to vastly improve the current practice in the field.

The project proposes an effective subset of existing standardized UML profiles for embedded systems modeling: SysML and MARTE, while avoiding incompatibilities resulting from simultaneous usage of both profiles. As both of these profiles provide numerous concepts and supporting tools, they are in turn difficult to be mastered by system designer. For this purpose, the MADES language subset is proposed along with a specific set of unique diagrams for expressing different aspects related to a system.

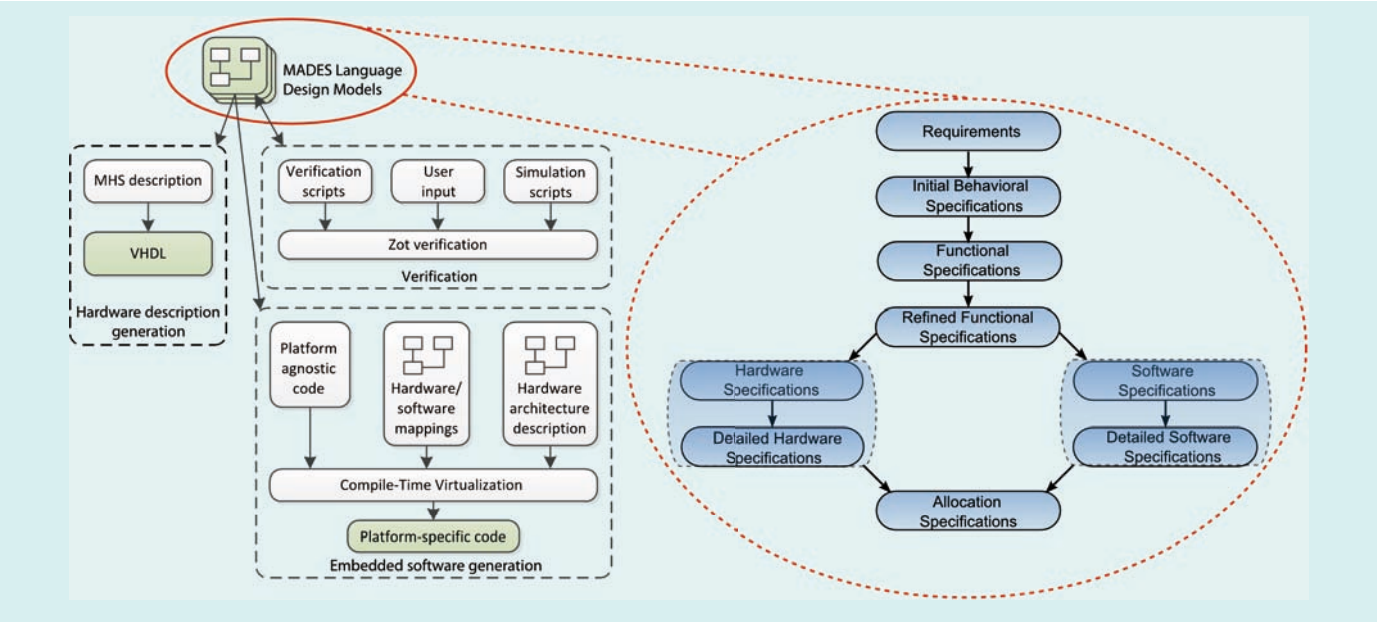

A Global overview of the MADES methodology and the underlying SysML/MARTE language subset

Initially, the high level models are created in Softeam's open source Modelio environment, that is capable of fully supporting SysML and MARTE profiles along with the dedicated MADES diagrams introduced in the course of the project, for example, dedicated diagrams for expressing the hardware or software aspects of a functional system in question. After specification of the design models that include user requirements, related hardware/software aspects and their eventual allocation along with schedulability analysis; underlying model transformations developed by University of York are used to bridge the gap between these abstract design models and subsequent design phases, such as verification, hardware descriptions of

modeled targeted architecture and generation of platform specific embedded software from architecturally neutral software specifications. For implementing model transformations, MADES uses the Epsilon platform, that enables model transformations, code generation, model comparison, merging, refactoring and validation; while Politecnico Di Milano's Zot tool is used for verification of key properties of designed concepts (such as meeting deadlines, etc.) and that of model transformations integrated in the design flow.

The aim of the project is twofold: to specify and implement a ground based radar processing unit and an on-plane radar control unit case studies by two major industrial organizations in the Avionics and Defense domains (TXT e-solutions and EADS)following the high abstraction level design MADES methodology, and to also provide effective guidelines to help the general embedded systems community in the usage of a combined SysML/MARTE methodology, and to improve upon the current OMG standards. For this purpose, the MADES project is also being presented at the OMG Workshop on Real-time, Embedded and Enterprise-Scale Time-Critical Systems in April 2012.

MADES project members at first successful project evaluation in Brussels, 2011

Alessandra Bagnato,
Project Coordinator,
TXT E-Solutions,
alessandra.bagnato@txtgroup.com

DESYRE FOR EXTREMELY RELIABLE ULTRA LOW-POWER SYSTEMS

For pacemakers and other implantable medical devices three are the key factors: extreme reliability, small size, and long longevity. In the EU project DeSyRe, researchers tackle these issues with a new approach: building a reliable system on unreliable components.

Project title:
DeSyRe (on-Demand System Reliability)
Project coordinator:
Ioannis Sourdis
Chalmers University of Technology
sourdis@chalmers.se
Partners:
Chalmers
Bristol
EPFL
FORTH
Imperial College London
Neurasmus B.V.
Recore Systems B.V.
Yogitech SpA
Project website:
www.desyre.eu
Project start date:
1st October 2011
Duration:
36 months

To counter the increasing fault-rates expected in the next technology generations, DeSyRe develops new design techniques for future Systems-on-Chips to improve reliability while at the same time reducing power and performance overheads associated with fault-tolerance. Ioannis Sourdis, Assistant Professor in Computer Engineering at Chalmers University of Technology, is the project leader of DeSyRe (on-Demand System Reliability).

"We focus on the design of future high reliable Systems-on-Chips that consume far less power than other designs for high reliability systems," he says. "This approach allows by design devices that combine high reliability with small batteries and state-of-the-art longevity. It is perfect for safety-critical applications such as in implantable medical devices, for example pacemakers or deep brain stimulators that treat Parkinson's disease".

Research in reliable systems typically focuses on fail-safe mechanisms that use various redundancy schemes, in which sensitive subsystems are entirely doubled as a fail-safe. Checking for faults in the subsystem increases the energy consumption and decreases the performance of chips, as testing all subsystems cost time and energy. The DeSyRe consortium takes a different approach, and separates the System-on-Chip (SoC) into two different areas: one which is extremely resistant to faults, and one area with fault-prone processing cores. The cores on the fault-prone area are interchangeable and the task of one core can easily be transferred to any of the other cores in case of a diagnosed malfunction. The fault-free part of the chip is responsible for monitoring the operation of the fault-prone part by performing sanity-checks of the processing cores, and for assuring that each core correctly handles an assigned sub-task.

"It sounds perhaps counterintuitive to design a highly reliable System-on-Chip on the basis of components that may fail, and yet this is exactly what we propose to do. Since our subsystems consist of small, interchangeable processing cores, we can test and exclude individual cores while the function of the whole systems stays intact", says Gerard Rauwerda, CTO of Recore Systems, one of the industry partners of DeSyRe.

The researchers expect this type of fault-

tolerance to reduce energy consumption by at least ten to twenty percent compared to other redundancy schemes, while at the same time minimizing penalty on performance.

To counter the increasing fault-rates expected for the next technology generations, DeSyRe develops new design techniques for future SoCs to improve reliability while at the same time reducing power and performance overheads associated with fault-tolerance. The DeSyRe consortium brings together leading European experts in the field of fault-tolerant and self- repairing designs, both from academia and industry.

For more info: www.desyre.eu

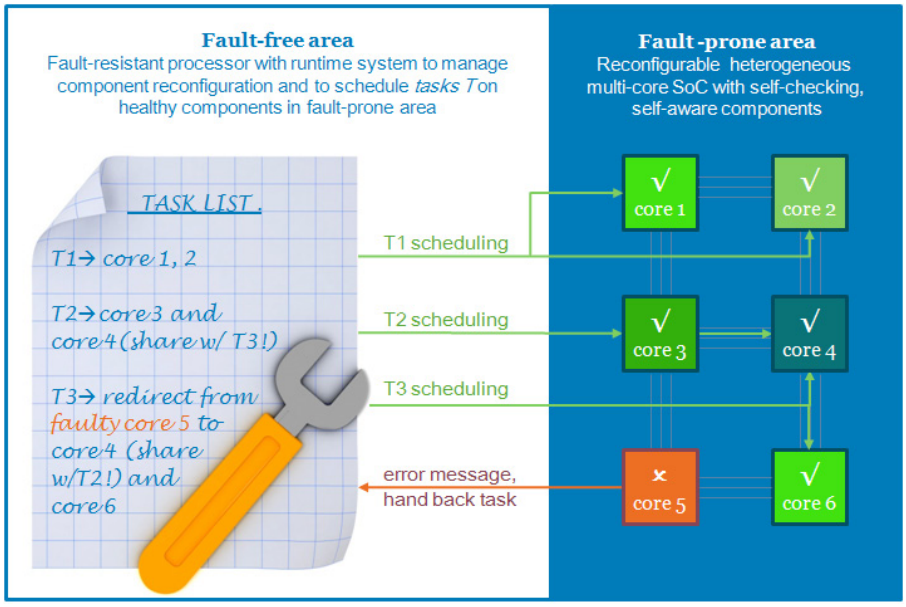

DeSyRe design for fault-tolerant Systems on Chip

EUROPE INVESTS IN DESIGNING A NEW ENERGY-EFFICIENT EXASCALE MACHINE

The Mont-Blanc project brings together a purely European consortium gathering industrial technology providers and supercomputing research centres

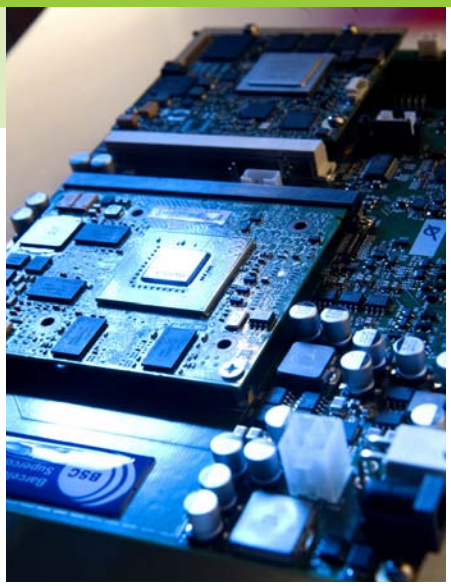

MONT-BLANC

Energy efficiency is already a primary concern for the design of any computer system and it is unanimously recognized that future Exascale systems will be strongly constrained by their power consumption. This is why the Mont-Blanc project, which was launched in October 2011 in Barcelona, has set the following objective: to design a new type of computer architecture capable of setting future global High Performance Computing (HPC) standards that will deliver Exascale performance while using 15 to 30 times less energy.

This new project is coordinated by the Barcelona Supercomputing Center (BSC) and has a budget of over 14 million Euros, including over 8 million Euros funded by the European Commission, has 3 objectives: to develop a fully functional energy-efficient HPC prototype using low-power commercially available embedded technology; to design a next-generation HPC system together with a range of embedded technologies in order to overcome the limitations identified in the prototype system, and to develop a portfolio of Exascale applications to be run on this new generation of HPC systems.

With energy efficiency being a key issue, supercomputers are expected to achieve 200 Petaflop/s (PF) in 2017 with a power budget of 10 MW, and 1000 PF (1 Exaflops/s) in 2020 with a power budget of 20MW. That means an increase in energy-efficiency of more than 20 times compared with the most efficient supercomputers today. "First, we must take into account that not all energy is used for computing within the cores. In current systems the processors consume the lion's share of the energy, often 40% or more. The remaining energy is used to power up the memory, interconnection network, and storage system. Furthermore, a significant fraction is wasted in power supply overheads, and in thermal dissipation (cooling), which do not contribute to performance at all", says Alex Ramirez, leader of the Mont-Blanc project. The system architecture in Mont-Blanc will rely on energy-efficient ARM processors, also used in embedded and mobile devices. It is expected to achieve from 4 to 10 times increase in energy-efficiency compared with current technologies.

The project brings together leading European technology companies such as Bull, ARM and Gnodal as well as some of the most important supercomputing centers in Europe: JSC and LRZ (Germany), CNRS and GENCI (France), CINECA (Italy) and BSC-CNS (Spain).

INTERNSHIP REPORT - DIEGO CABALLERO

Optimizing OpenCL kernels in a LLVM compilation and runtime environment

I am a Ph.D. student at Universitat Politècnica de Catalunya - Barcelona Tech (UPC), and I am also associated with the Barcelona Supercomputing Center (BSC) as a research student in the Parallel Programming Models team. My advisors are Xavier Martorell and Alejandro Duran. My research interests include computer architecture and programming models, ranging from new multi- and many-core architectures, SIMD and short-vector units, vectorization techniques and new programming models focused on high productivity, performance and portability.

The HiPEAC internship allowed me to collaborate with ARM, one of the worldwide leading companies in the industry of computer architecture for embedded systems and mobile devices. This fact, independently from the internship activities, gave me a privileged and valuable view of how a private enterprise of this magnitude works.

During my stay at ARM Cambridge I did research on applying automatic transformations to OpenCL kernels, both at compile- and run-time, such as vectorization, kernels fusion and data layout transformations. I carried out these activities on an ARM production compiler based on the LLVM compiler infrastructure. This was my first contact with a compiler infrastructure based on LLVM. As soon as I got used to the LLVM environment, I started implementing a new vectorization pass initially able to automatically vectorize simple OpenCL kernels. The powerful LLVM capabilities in compiler analysis and code transformations allowed me to easily improve the vectorizer with more complex features to vectorize a wider number of kernels. LLVM was a new discovery for me, providing new possibilities to my Ph.D. and even to some projects at BSC.

I also conducted benchmarking activities on the new Mali architecture with the intention of getting information about performance and the memory hierarchy behavior. For this task, I had the privilege of having early access to the architecture which informed about new tendencies in this kind of devices.

This internship was especially valuable in the professional relationships. I had the opportunity of meeting outstanding professionals and Ph. D. students. In addition, this experience helped to establish the communication between the ARM Media Processing Division and the Parallel Programming Models Group at BSC, which will possibly result in future collaborations.

The evaluation of the whole internship was excellent. I want to thank to HiPEAC for offering me this opportunity and to ARM for choosing me. In my opinion, ARM has all the features to lead an internship to success. I encourage HiPEAC to continue collaborating with them and Ph.D. students to join them, as one of the most appropriate places to conduct a research internship.

Diego Caballero, UPC/BSC, Spain

EVENT TRIP REPORT - JOSÉ L. ABELLÁN

The poster session of the HiPEAC Conference in Paris gave students a great opportunity for networking

My name is José L. Abellán. I am a fourth-year PhD student in the Computer Engineering Department at the University of Murcia (Spain). My advisors are Juan Fernández (now at Intel Barcelona Research Lab) and Manuel E. Acacio. My research interests include multicore architectures, parallel programming models, efficient synchronization operations and efficient cache-coherence protocol designs.

Since I expect to finish my PhD before summer, I am fully committed to a lot of tasks including not only those of writing the thesis, performing final experiments or attending conferences, but also it is also time to look for a qualified job. Conscious about this important concern among people from both academia and industry, this year the organizers of the HiPEAC Conference had the great idea to incorporate a special poster session as a recruitment event for the companies and projects that attended the conference. Here is when I decided to do a poster with the work done along my PhD, go to Paris and present my poster in this special session. Perhaps I could find a job that would meet my expectations.

The HiPEAC organizers usually offer a number of grants to attend the conference primarily aimed at students to pay registration fees. This year it was not an exception and I applied for a student grant during my registration process. Fortunately, my request was accepted and I had not to pay anything to present my poster, to attend the main conference presentations or other events such as associated workshops, tutorials or special sessions.

In the end, despite I didn't find a suitable job for me, my experience in the poster session was very positive because on the one hand, people were very interested in

my poster and I enjoyed a lot discussing all their questions. On the other hand, it provided me with very useful feedback that is helping me to finish my PhD.
Finally, I would like to thank the HiPEAC organizers for many reasons. Firstly, because they gave me a student grant to attend the conference. Secondly, for the huge effort shown to help people like me to find a job by improving the traditional HiPEAC conference model. Finally, because they have also enabled the HiPEAC jobs website (http://jobs.hipeac.net/) where, almost every day, you can find new job openings, or even you can also submit your CV, as I did.

Jose L. Abellán, University of Murcia, Spain
jl.abellan@ditec.um.es

INTERNSHIP REPORT - PETAR RADOJKOVIC

On the Evaluation of the Impact of Shared Resources in Multithreaded COTS Processors in Time-Critical Environments

I am a PhD student at computer architecture department of Universitat Politècnica de Catalunya, Barcelona, Spain. For the last five years, I have also worked as a researcher at Computer Architecture and Operating Systems Interface group at Barcelona Supercomputing Center. The main area of my research is the interference between tasks that execute concurrently on multicore/multithreaded processors.

As a part of the HiPEAC mobility program, in winter 2010/2011, I spent three months at Embedded System Lab at Thales Research &Technology, Palaiseau, France. During the internship, I was exploring the possible application of multithreaded Commercial-Off-The-Shelf (COTS) processors in time critical environments.
COTSprocessors are now commonly used in real-time embedded systems. The characteristics of these processors fulfill system requirements in terms of time-to-market, low cost, and high performance-per-watt ratio. However, multithreaded processors are still not widely used in real-time systems because their timing analysis is complex. In multithreaded processors, simultaneously running tasks share and compete for processor resources, so the timing analysis has to take into account the possible impact that the inter-task interferences have on the execution time of the applications.

During the internship, we developed a method that quantifies the slowdown that simultaneously running tasks may experience due to collision in shared processor resources. To that end, we designed benchmarks that stress specific processor resources and we used them to: (1) Estimate the upper limit of the slowdown that simultaneously-running tasks may experience because of collision in different shared processor resources, and (2) Quantify the sensitivity of time-critical applications to interaction with other applications in these resources. We used the presented method to determine if given multithreaded processors are good candidates for systems with timing requirements. Finally, we showed that measuring the slowdown that real applications experience when simultaneously running with resource-stressing benchmarks is an important step in measurement-based timing analysis. This information can be used during the incremental verification of multithreaded COTS architectures.

The work done during the internship is summarized in the paper "On the Evaluation of the Impact of Shared Resources in Multithreaded COTS Processors in Time-Critical Environments" that is published at ACM Transactions on Architecture and Code Optimization (TACO) in January 2012.

I am grateful to HiPEAC for the support of my internship at Thales Research & Technology. I would also like to thank Sami Yehia, my supervisor at Thales, and the colleagues from the Embedded Systems Lab group for a pleasant collaboration.

Petar Radojkovic, BSC, Spain

PERFORMANCE ANALYSIS AND TARGET-SPECIFIC CODE OPTIMIZATIONS FOR THE GRID ALU PROCESSOR

By Ralf Jahr
Advisor: Prof. Dr. Theo Ungerer
University of Augsburg, Germany
December 2011

The post-link optimizer GAPtimize is introduced to exploit the features of the Grid Alu Processor (GAP) for quicker execution of sequential instruction streams without recompilation. Each of the four main issues restraining the performance of the GAP is addressed by at least one of the optimizations implemented in GAPtimize. In an automatic design space exploration (ADSE) with heuristic algorithms, the most efficient and effective configurations for hardware parameters of the GAP are identified. The search is extended to GAPtimize proving that it can optimize both performance and area effectiveness of the GAP. The impact of GAPtimize can be raised by selecting parameters adaptively.

REAL-TIME SCHEDULING FOR SIMULTANEOUS MULTITHREADED PROCESSORS

By Jörg Mische
Advisor: Prof. Dr. Theo Ungerer
University of Augsburg, Germany
February 2012

The microarchitecture of a single threaded superscalar in-order processor is enhanced to support simultaneous multithreading with one completely isolated thread. This thread is executed as if it were the only thread in a single threaded system. Hence established methods for Worst Case Execution Time analysis can be applied. A hardware module, directly connected to the issue stage of the processor pipeline, provides sophisticated scheduling algorithms for a nearly arbitrary number of threads (time slicing for hard real-time, control of the instructions per cycle for soft real-time). The SystemC and the sythesizable VHDL model support the TriCore instruction set.

PERFORMANCE COUNTER-BASED STRATEGIES TO IMPROVE DATA LOCALITY ON MULTIPROCESSOR SYSTEMS: REORDERING AND PAGE MIGRARTION TECHNIQUES

By Juan Angel Lorenzo
Advisors: Dr. Juan C. Pichel and Dr. Francisco F. Rivera
University of Santiago de Compostela, Spain
January 2012

We approach the study of Precise Event-Based Sampling (PEBS) techniques to improve the performance of applications on a NUMA system. We demonstrate that a low-cost, PEBS profiling can support strategies to improve the performance of an important group of scientific codes in runtime. Two different contexts have been considered. Firstly, the development of hardware counter-based strategies to assist reordering techniques for irregular codes in order to reduce their cost and improve their behaviour. And secondly, the development of novel hardware counter-guided, dynamic page migration algorithms that take advantage of the new features provided by the PEBS.

VLSI MICRO-ARCHITECTURES FOR HIGH-RADIX CROSSBARS

By Giorgos Passas
Advisors: Prof. Manolis Katevenis and Prof. Dionisios Pnevmatikatos
FORTH-ICS & University of Crete, Greece
March 2012

This thesis shows that a 128x128x64Gb/s crossbar, interconnecting 128 1mm² "user tiles" in a single hop, fits in just 2mm2 in 32nm CMOS. The crossbar is 32bits wide, runs at 2GHz, and consumes 8Watts. This result shows that, contrary to common belief, high-radix crossbars are competitive to popular network-on-chip topologies for chip multiprocessors. Moreover, in the domain of system-area routers, crossbar speedup is inexpensive, and thus combined input-output queuing is competitive to crosspoint queuing. We made high-radix crossbars feasible by developing novel VLSI micro-architectures, shaping the intricate wiring of crossbars into regular layouts, and routing wires --not packets. For more details: http://www.ics.forth.gr/~passas.

COMPILER ASSISTED RUNTIME ADAPTATION

By Vlad-Mihai Sima
Advisors: Prof. Koen L.M. Bertels and Prof. Henk J. Sips
Delft University of Technology, the Netherlands
January 2012

The dissertation addresses the problem of runtime adaptation of the application to its execution environment. The work focuses on heterogeneous multicore architectures. It addresses three aspects of application optimizations: hardware software mapping, memory allocation and parallel execution. The applications used to validate the algorithms are real life applications from the multimedia field. One of the results of the work performed is a toolchain targeting heterogeneous architectures. One of the most important tools developed specifically for this toolchain is the DWARV C-to-VHDL compiler. More information can be found at: http://ce.et.tudelft.nl/~vladms

POWER-EFFICIENT TIGHTLY-COUPLED PROCESSOR ARRAYS FOR DIGITAL SIGNAL PROCESSING

By Dmitrij Kissler,
Advisor: Prof. Dr.-Ing. Jürgen Teich
University of Erlangen-Nuremberg, Germany
December 2011

In this thesis, we consider highly area- and power-efficient, massively parallel, tightly-coupled processor arrays, which can be used as hardware accelerators in mobile embedded systems for sophisticated digital signal and image processing. Our research fully proved the need and the benefits of deploying such efficient accelerators in modern high-performance embedded systems: The prototype implementations of an processor array in 90 nm CMOS ASIC technology with 4 and 24 processing elements revealed power efficiency values ranging from 98 MOPS/mW to 124 MOPS/mW. The corresponding chip area lies between 0.2 mm² (4 PEs) and 2.2 mm² (24 PEs).

PERFORMANCE EVALUATION OF APPLICATIONS FOR HETEROGENEOUS SYSTEMS BY MEANS OF PERFORMANCE PROBES

By Alexandre Strube
Advisor: Prof. Dr. Emilio Luque
University Autónoma of Barcelona, Spain
July 2011

This work describes a novel way to select the best computer node out of a pool of available heterogeneous computing nodes. Today's systems allocate resources based on history of executions, which takes long time to be effective, and benchmarks are hardly representative of any application at all. My solution, called "Probe", performs the important computations in the same way as the application, but without taking the time of a full execution. It predicts application execution time with over 90% accuracy, well under 5% of the original time. It allows one to choose the right execution nodes in grids, clouds, workflows engines and other environments very quickly and precisely.

UPCOMING EVENTS

The 7th IEEE Symposium on Industrial Embedded Systems (SIES 2012)
June 20-22, 2012, Karlsruhe, Germany http://sies2012.ira.uka.de

The 8th International Conference on Wireless and Mobile Communications (ICWMC 2012)
June 24-29, 2012, Venice, Italy http://www.iaria.org/conferences2012/ICWMC12.html

International Conference on Localization and GNSS (ICL-GNSS 2012)
June 25-27, 2012, Munich, Germany, http://www.icl-gnss.org/

The 23rd IEEE International Conference on Application-specific Systems, Architectures and Processors (ASAP 2012)
July 9-11, 2012 Delft, The Netherlands, http://asap-conference.org/

International Conference on Embedded Computer Systems: Architectures, Modeling, and Simulation (SAMOS XII)
July 16-19, 2012, Samos, Greece, http://samos.et.tudelft.nl/

The 18th International European Conference on Parallel and Distributed Computing (Euro-Par 2011)
August 27-31, 2012, Rhodes Island, Greece. http://europar2012.cti.gr/

The 22nd International Conference on Field Programmable Logic and Applications (FPL 2012)
August 29-31, 2012, Oslo, Norway, http://www.fpl2012.org/

International Symposium on System-on-Chip 2012
October 11-12, 2012, Tampere, Finland, http://soc.cs.tut.fi/

Conference on Design and Architectures for Signal and Image Processing (DASIP 2012)
October 23-25, 2012, Karlsruhe, Germany http://www.ecsi.org/dasip/

The 8th International Conference on High Performance and Embedded Architectures and Compilers (HiPEAC 2013)
January 21-23, 2013, Berlin, Germany, http://www.hipeac.net/conference

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

HiPEAC info 30
COMPILATION ARCHITECTURE

HiPEAC info is a quarterly newsletter published by the HiPEAC network of excellence, funded by the 7th European Framework Programme (FP7) under contract no. IST-217068.
website: http://www.hipeac.net.
subscriptions: http://www.hipeac.net/newsletter

design: www.magelaan.be

AUTUMN COMPUTING SYSTEMS WEEK, GHENT, BELGIUM