HiPEAC info 34

COMPILATION ARCHITECTURE

APPEARS QUARTERLY
MAY 2013

NETWORK OF EXCELLENCE ON HIGH PERFORMANCE AND EMBEDDED ARCHITECTURE AND COMPILATION

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

SEVENTH FRAMEWORK PROGRAMME

WWW.HIPEAC.NET

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

MESSAGE FROM THE HIPEAC COORDINATOR

Just as a year has its seasons, HiPEAC has its own schedule of activities. The HiPEAC winter season is dominated by the HiPEAC conference, the summer season by the summer school, and spring and autumn by the computing systems weeks. However, from time to time we have an event which is non-recurring, or which has a different frequency. The HiPEAC roadmap is such a task. We produced roadmaps in 2008, 2010 and 2011 and this newsletter announces the 2013 edition.

I have been involved in the redaction of all four HiPEAC roadmaps so far, and looking back over the last five years of road-mapping, I believe we have learned a lot. We now better understand how to organize the roadmapping process, how to distil a common vision from the input from the HiPEAC community with its more than 1200 members, how to get feedback from external stakeholders, how write it down in a concise and coherent document, and how to communicate back to the community and to the policy makers.
Through being involved in the road-mapping process, I realize how fast our domain is evolving. Every time we start working on the next edition of the roadmap, the plan is to update the previous roadmap, and every time we discover that this is not possible, that insights have progressed, that new challenges have popped up, and that former challenges have become less urgent. Every time, we end up with new emphasis and new recommendations.

I believe that the HiPEAC roadmap is one of the most important deliverables of the HiPEAC network because it is a document that tries to sketch the future of our research area, making abstraction from the daily technical problems we have to deal with. It is a kind of bird's eye view of the big trends that influence our domain, and it makes suggestions where we should invest. I am thankful that the policy makers in the European Commission study our roadmaps in detail, and that they allocate substantial research funding to the major challenges put forward in the roadmap. Since 2004 we have seen a steady increase of the research funding from 25 M€ in the first call, to 72.5 M€ in the 2013 call.
I want to thank everybody who has been involved in the latest roadmapping process, and especially Marc Duranton who has been driving this effort for the HiPEAC community. I hope the roadmap will be a useful resource for the HiPEAC community and beyond. The roadmap is co-owned by all of us, so feel free to share it, to use it in presentations and lectures, and to promote it widely.

Koen De Bosschere

CONTENT

MESSAGE FROM THE PROJECT OFFICER

The European Commission has announced recently the winners of a multi-billion euro competition for Flagship projects in Future and Emerging Technologies (FET). FET Flagships are ambitious large-scale, science-driven, research initiatives that aim to achieve a visionary goal. Following a preparatory phase, Graphene and the Human Brain Project were selected among six candidate proposals by a panel of 25 world-renowned experts. Each project will receive up to €1 billion funding over the next 10 years and will involve hundreds of researchers, companies, research institutes and universities across the EU in a 10-year mission to tackle grand scientific and technological challenges, across research disciplines and national borders.

The mission of the Graphene flagship project is to investigate and exploit the unique properties of graphene, the revolutionary carbon-based material, which has an extraordinary combination of physical and chemical properties. The research effort will cover the entire value chain from materials production to components and system integration, and targets a number of specific goals that exploit the unique properties of graphene.

The Human Brain Project will work in one of the biggest unresolved challenges: understanding the human brain. The project will create the world's largest experimental facility for developing the most detailed model of the brain, for studying how the human brain works and ultimately to develop personalised treatment of neurological and related diseases.

Both flagship projects are of direct interest to computing researchers. Graphene could provide the base for beyond-CMOS nanoelectronics. The Human Brain Project will use advanced High-Performance Computing for data-intensive modelling and simulation; it also has neuromorphic computing as one of its areas of research.

For more information:
www.graphene-flagship.eu
www.humanbrainproject.eu

Panos Tsarchopoulos

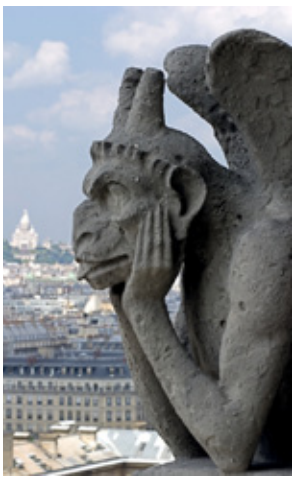

THE HIPEAC 2013 CONFERENCE, BERLIN

With 510 participants, this conference was one of the biggest events ever organized by the HiPEAC network

After the start of the new publication model in Paris last year, the success story has moved to Berlin this year. More than 500 delegates from 38 countries met in the conference centre of the Radisson Hotel in the centre of Berlin from January 21 to 23. Three days of high-quality presentations in the main conference session together with the opportunity to attend a wide variety of satellite events have been offered to all attendees. Poster and industrial presentations completed the event with emphasis on networking opportunities and establishing new collaborations.

HIGHLIGHTS

- 4 keynote talks from industry as well as academia as daily openings and as a wrap-up of the first conference day
- 43 focused paper presentations demonstrated valuable work from researchers from all over the world
- An industrial session with invited presentations from SAP, NVidia and Xilinx
- 24 workshops and 10 tutorials with high-quality presentations, keynote talks, poster sessions and panel discussions
- A student poster session with 45 posters
- 33 projects presented during the European project poster session.
- An amazing traditional German dinner served at the social event in Hofbräuhaus Berlin

Even though the main HiPEAC track forms the focus of the conference, the satellite workshops, tutorials, and poster sessions provided a unique opportunity for networking in Europe. Nearly the complete European communities of computer and embedded system architectures as well as

Over 500 participants attended the HiPEAC conference in Berlin.

Industrial session.

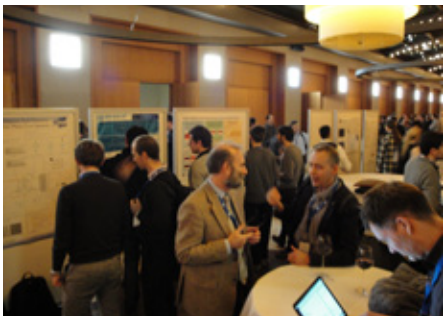

Student poster session.

The conference dinner at Hofbräuhaus Berlin.

compiler development met at the HiPEAC conference for an intensive exchange.

In addition to the poster session, presentations of 26 EC funded projects in workshops and tutorials provided an overview of the European computing systems research landscape. Detailed insights into several projects have been given at dedicated project workshops, together with intensive technical discussions.

A subsequent survey shows that the conference's format with many parallel events, poster session, and exhibitions is highly appreciated. Moreover, several comments expressed the uniqueness of the HiPEAC conference bringing together the research community and industry.

This success is only possible with the generous support from 21 sponsors. These sponsorships are essential to be able to offer attractive registration fees and to support student participation. More information about this year's conference can be found at the conference web site: http://www.hipeac.net/hipeac2013

Next year, the HiPEAC conference will take place in Vienna, Austria. Again, the conference will follow the new publication model where authors of papers accepted for publication in ACM TACO will be invited to present their work at the conference. In parallel to the main session, a number of satellite events including poster sessions will take place. We encourage researchers from the computer architecture, embedded systems, and compiler technology domains to submit their high-quality papers to TACO as well as to propose a workshop or tutorial for the next conference.
See you next year in Vienna!

Sascha Uhrig
TU Dortmund, Germany
Ben Juurlink
TU Berlin, Germany

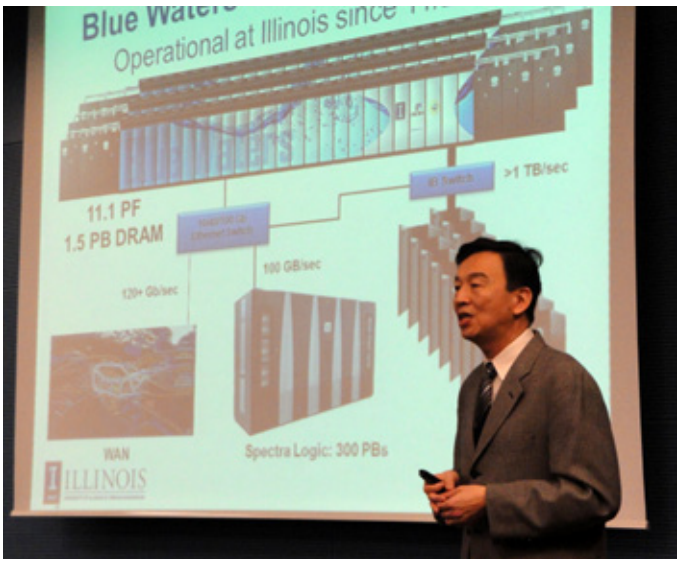

Wen-mei Hwu during the HiPEAC 2013 opening keynote.

Doug Burger gave a very inspirational speech.

WORKSHOP ON TRANSFER TO INDUSTRY AND START-UPS AT HIPEAC'2013

Improving technology transfer from EU projects to industry

The goal of the first Transfer to Industry and Start-Ups (TISU) workshop was to better understand how we can improve the transfer of scientific output from EU funded projects to industry. To this purpose, three speakers were invited to share their experience and understanding of those issues. Around 35 people participated in the workshop. The speakers were:

- Gerd Ascheid, RTWH Aachen University, "Technology Transfer from Academia to Industry: Is There a Golden Path?"
- Per Stenström, Chalmers University of Technology, "Lessons Learned from Bringing a Research Idea to a VC-backed Start-up"
- Guillem Sague, High-Tech Gründerfonds Management, "Start your Start up: Incubation, Funding and Success (or not?)

Gerd Ascheids' talk pointed out that there is not just a single "olden" path of technology transfer and discussed three different ways with different pro's and con's: (i) by starting a company, which may be a high risk and needs profound management experience; (ii) by close collaboration with SMEs who often need such research projects in order to have the critical mass for investigations; and (iii) by collaborating with multinational corporations who are often interested in acquiring IP rights on research results.

The second part of the workshop consisted of a panel discussion in which a number of propositions were formulated to stimulate discussion. The panel discussed the perception that European researchers are too risk averse, at least as far as entrepreneurial initiatives are concerned, concluding that it is very important to provide the right environment to facilitate technology transfer. Speakers pointed out that while the presence of incubators is pretty common, universities are reluctant to stimulate entrepreneurship and definitely do not see it as an normal part of an academic career path.

The panel also talked about industrial participation, generally finding that it is only motivated by cost reasons and not by the need to innovate. While the TRL (technology readiness level) is the eternal excuse, this is definitely not true for SMEs as they often depend on research projects to be able to do any kind of research. Furthermore, larger corporations are indeed perceived as quite conservative as far as research results are concerned.

Discussions also brought up opinions about how European researchers are only interested in the funding for their research and absolutely not in any kind of economic impact. The panel found that there are important secondary and often unwanted economic effects of research done, without any concrete immediate economic result (such as a patent or a start-up), which should not be neglected.

This workshop will hopefully be the start of a more structured approach to stimulating the HiPEAC members to intensify their efforts in technology transfer. The organisers will try to propose some kind of structural presence of these topics at the HiPEAC gatherings and HiPEAC conference.

K. Bertels, S. Hamdioui,
Delft University of Technology,
the Netherlands

Gerd Ascheid's talk sparked some fruitful discussions.

Prof. Per Stenström shared his start-up experience.

WORKSHOP ON MIXED-CRITICALITY INTEGRATION ORGANIZED BY ACROSS, MULTIPARTES AND ARAMIS PROJECTS

The workshop was co-located with the HiPEAC 2013 conference in Berlin

Entitled "Integration of mixed-criticality subsystems on multi-core processors", the main motivation of this workshop was to introduce the audience to the research carried out in computer architectures that support mixed-criticality integration and the specific challenges from multi-core processors. More than 60 people from 14 countries – representing Europe and Asia – attended the workshop.

Related European research projects under both FP7 and Artemis programs were represented, namely, ACROSS, MultiPARTES, ARAMIS, CERTAINTY and RECOMP. They featured talks focusing on both the practical and scientific perspectives, on topics such as reference architecture, methodology, tools, certification issues and industrial application, among others. The agenda and the material are available in the workshop website (http://www.across-project.eu/workshop2013_program.htm).

A panel discussion was moderated by Salva Trujillo (IK4-IKERLAN) with the following panelists: Alan Burns (University of York), Francisco Cazorla (Barcelona Supercomputing Center), Christian El Salloum (TU Wien) and Bernd Koppenhöffe (EADS Cassidian). The key discussion point was how to assure "sufficient" independence among involved subsystems. In this regard, the definition of "sufficient" was key. Several additional trends were covered, including the appearance of multicore, pressure from consumer electronics, and other hot topics.

Following this success in 2013, a continuation is planned in 2014. A seminar at Dagstuhl is being organized on this topic.

Workshop participants and MultiPARTES poster at HiPEAC 2013.

MULTIPARTES PROJECT MEETING CO-LOCATED WITH HIPEAC 2013 CONFERENCE

The MultiPARTES FP7 STREP project organized a project plenary meeting, collocated with the HiPEAC 2013 conference in Berlin. The main goals of the meeting were to discuss the role of the MultiPARTES project in the context of the mixed-criticality systems scenario and to define the boundaries and synergies among projects, seeking for collaborations.

The conclusion was that MultiPARTES is exploring the boundaries of solutions based on existing COTS solutions via the use of multicore partitioning based on hypervisor technology, whilst other projects are focused on building specific hardware, and others are stricter on the methodological side. It became clear that in every approach there are some limitations. MultiPARTES is exploring such boundaries when existing hardware is being used and it is focuses on industrial demonstration.

Apart from these topics, the current status of the project was reviewed, with particular focus on discussion of the RTD deliverables.

Further pictures and material at:
FP7MultiPARTES@FP7MultiPARTES
https://twitter.com/FP7MultiPARTES

MultiPARTES project meeting in Berlin

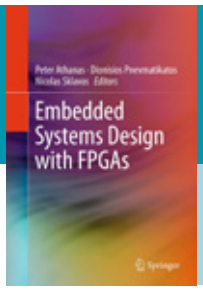

NEW BOOK: EMBEDDED SYSTEMS DESIGN WITH FPGAS

Peter Athanas, Dionisios Pnevmatikatos, Nicolas Sklavos

Peter Athanas (Virginia Tech, USA), Dionysios Pnevmatikatos (Technical University of Crete, Greece) and Nicolas Sklavos (TEI of Patras, Greece), are the co-editors of the newly published book, "Embedded Systems Design with FPGAs", Springer, 2012.

This book describes a variety of methodologies for modern embedded systems design, implements methodologies presented on FPGAs and finally covers a wide variety of applications for reconfigurable embedded systems, including bioinformatics, communications and networking, application acceleration, medical solutions, experiments for high energy physics, astronomy, aerospace, biologically inspired systems and computational fluid dynamics.

This book presents methodologies for modern applications of embedded systems design, using field programmable gate array (FPGA) devices. Coverage includes state-of-the-art research from academia and industry on a wide range of topics, including advanced electronic design automation (EDA), novel system architectures, embedded processors, arithmetic, dynamic reconfiguration and applications.

For more information please refer to: http://www.springer.com/engineering/circuits+&+systems/book/978-1-4614-1361-5

NEW BOOK: COMPILATION AND SYNTHESIS FOR EMBEDDED RECONFIGURABLE SYSTEMS: AN ASPECT-ORIENTED APPROACH

João M. P. Cardoso, Pedro C. Diniz, José Gabriel Coutinho and Zlatko Petrov

The book describes an Aspect-based approach for the compilation and synthesis of applications targeting heterogeneous embedded computing architectures. The approach relies on a domain-specific aspect-oriented language – LARA – that allows designers to control sophisticated design tool chains composed of various compilation and hardware synthesis tools. Using aspect-oriented programming techniques, developers can reuse extensive sections of their designs, exploring a wide range of possible design solution parameters and thus leading to faster and more reliable design solutions. The book includes experimental results of the application of the REFLECT project's aspect-controlled design flow, for various real-life applications from unmanned autonomous vehicle navigation to digital multimedia processing. The results focus on the use of aspects for the development of reusable design strategies for optimization of hardware/software partitioning, hardware synthesis and specialization as well as for the resource -efficiency and safety requirements. The results presented in this book clearly indicate that the Aspect-based approach pursued in the REFLECT project with LARA enhances design reuse, while preserving the original application source code, thus promoting developer productivity as well as architecture and performance portability.

URL at Amazon.com:
http://www.amazon.com/Compilation-Synthesis-Embedded-Reconfigurable-Systems/dp/146144893X/ref=sr_1_1?s=books&ie=UTF8&qid=1362348302&sr=1-1

Project URL: www.reflect-project.eu

João M. P. Cardoso, Universidade do Porto, Portugal
Pedro C. Diniz, INESC-ID, Lisboa, Portugal
José Gabriel Coutinho, Imperial College, London, United Kingdom
Zlatko Petrov, Honeywell Intl. s.r.o., Czech Republic.

KOEN DE BOSSCHERE RECEIVES HERMES AWARD

The award acknowledges De Bosschere's work to foster student-entrepreneurship

The Hermes award is an annual Ghent University award for excellence in service. The 2012 Hermes award was given to the HiPEAC coordinator Koen De Bosschere, for his efforts to stimulate student entrepreneurship at Ghent University. The idea to start a student entrepreneurship project at the university came after a presentation of the Prospekt initiative at the University of Edinburgh, given during the May 2010 HiPEAC innovation event. Today, 60 students from all departments of Ghent University are being coached as beginning entrepreneurs, and their number is increasing every year. The university and the city of Ghent have high expectations for student entrepreneurship in the future.

Koen De Bosschere
Ghent University, Belgium

Prof. De Bosschere poses with the Hermes award.
(© Photo: Hilde Christiaens)

PROFESSOR MATEO VALERO AWARDED AN ERC ADVANCED GRANT

The ERC-funded project proposes a radically new conception of parallel architectures, built using a higher level of abstraction

The most common interpretation of Moore's Law is that the number of components on a chip, and accordingly computer performance, doubles every two years. This experimental law has held since its first statement in 1965 until today. At the end of the 20th century, when clock frequencies began to stagnate at around 3GHz, and instruction level parallelism reached the phase of diminishing returns, industry turned towards multiprocessors and thread level parallelism. Currently, too much of the technological complexity of multicore architectures is exposed to programmers, leading to a software development nightmare that horrifies the entire computing industry.

The ERC-funded project 'Riding on Moore's Law (RoMoL)' proposes a radically new conception of parallel architectures, built using a higher level of abstraction. Instead of expressing algorithms as a sequence of instructions, we will group instructions into higher-level tasks that will be automatically managed by the architecture, much in the same way superscalar processors managed instruction level parallelism. RoMoL envisions a holistic approach where the parallel architecture is partially implemented as a software runtime management layer, with the reminder in hardware. The hardware gains the freedom to deliver performance at the expense of additional complexity, as long as it provides the required support primitives for the runtime software to hide complexity from the programmer. The OmpSs programming model together with the Mercurium compiler and the Nanos++ runtime, which are being intensively developed at the Barcelona Supercomputing Center (BSC-CNS), will provide the perfect environment for this approach (http://pm.bsc.es/).

This regained freedom in hardware design will enable us to revisit a number of previously proposed architectural concepts, in light of the new technology context, for unforeseen impact in performance and energy efficiency. This now becomes possible thanks to the joint development

of the runtime layer with the parallel architecture. This holistic approach towards parallel architectures offers a single solution that could solve most of the problems we encounter in current approaches: handling parallelism, the memory wall, the power wall, and the upcoming reliability wall in a wide range of application domains from mobile up to supercomputers. RoMoL will tackle these crucial challenges for the next five years and will involve around 15 researchers, including PhD students, engineers, and post-doctoral researchers. RoMoL will facilitate collaborations with

high-level research visitors and maximize positive synergies with current research activities in the Computer Sciences Department at the Barcelona Supercomputing Center.

Altogether, this novel approach toward future parallel architectures is the way to ensure continued performance improvements, getting us out of the technological mess that computers have turned into, once more riding on Moore's Law.

Prof. Mateo Valero is a world-recognized computer architect and has been awarded prizes such as the 2007 Eckert-Mauchly Award, the 2009 'Harry H. Goode' Award and two national awards in research. He is also a Fellow of the IEEE, a distinguished fellow at Intel, Fellow of the ACM, and member of some of the most prestigious Academies. Valero is Engineer and Doctor in Telecommunications and a full professor at the Universitat Politècnica de Catalunya – BarcelonaTech since 1983. He has published over 600 articles on HPC architecture, and has more than 500 former PhD students and academic descendants. Valero is the founding director of the Barcelona Supercomputing Centre since 2004. Mateo Valero has also been one of the main actors responsible for the return of computer architecture research as a priority in the EU research agenda, and the founding coordinator of HiPEAC.

For more information on the ERC Advanced Grants 2012:
http://erc.europa.eu/sites/default/files/press_release/files/press_release_adg2012_results.pdf

RWTH AACHEN UNIVERSITY AT THE 2013 MOBILE WORLD CONGRESS

The UMIC Research Cluster was present in Barcelona

North Rhine-Westphalia Stand.

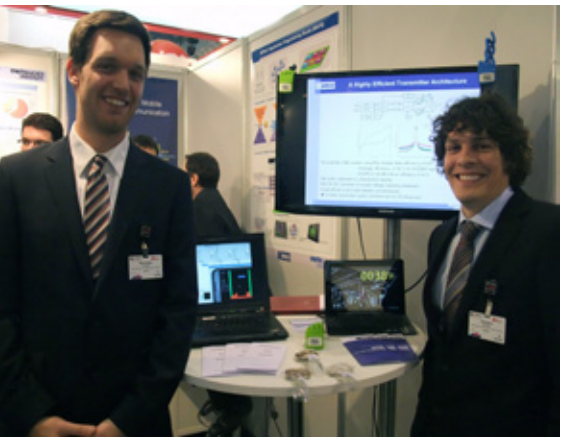

Researchers at the UMIC Stand.

During the Mobile World Congress (MWC) week in Barcelona, the RWTH University was present in two different stands, displaying some of the latest research achievements. From the 25th to the 28th of February, the attendees of the MWC could get an overview of the UMIC Research Cluster (www.umic.rwth-aachen.de) at the official stand of the state of North Rhine-Westphalia, and could delve into the details of several research projects at the dedicated UMIC stand.

From the multiple research initiatives developed within UMIC, four projects were exhibited at the MWC. The first two projects originated from the Chair for Communication and Distributed Systems of the Informatics Faculty, headed by Prof. Klaus Wehrle. They presented an approach for On-demand Content-centric Wireless Networking and a technique for High-coverage Bug Detection with Symbolic Execution. From the chair of Mixed Signal CMOS Circuits of the Electrical Engineering Faculty, Prof. Renato Negra and his team presented a Highly Efficient Transmitter Architecture for 4G base stations. The core of the architecture is a power amplifier with a peak efficiency of 58%. Finally, the chair for Software for Systems on Silicon from the same faculty showcased a framework for Multi-core programming called MAPS. With MAPS, Prof. Leupers and his team demonstrated how applications can be automatically deployed onto parallel, heterogeneous computing platforms.

This is the third time that UMIC has been represented at the MWC, the largest exhibition on mobile communication and mobile Internet technology. As in previous years, the MWC continues to be an exciting fair, providing high visibility to the RWTH Aachen University and new industry contacts for UMIC researchers.

Jeronimo Castrillon
RWTH Aachen University, Germany

MYRMICS: THE BARE-METAL RUNTIME FOR THE 520-CORE HARDWARE PROTOTYPE

Developing a scalable runtime system for dependency-aware task scheduling on heterogeneous many-core architectures

You may remember the 520-core Heterogeneous Hardware Prototype from Crete. In FORTH-ICS, we built a 4x4x4 cube of 64 of our own FPGA-based "Formic" boards and we connected them in a 3D-mesh fashion. We programmed the FPGAs with a non-coherent hardware architecture, which fits eight Xilinx MicroBlaze cores in each board, with private L1 and L2 caches, mailboxes, per-core DMA engines and DMA completion and synchronization primitives. 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, which we use to connect to the Formic cube. We designed this 520-core heterogeneous, non-coherent hardware architecture to be a realistic model of a future single-chip processor.

We recently developed Myrmics, a runtime system, which specifically addresses this class of emerging single-chip many-core architectures. As partners in the European Union ENCORE project (http://www.encore-project.eu), we use Myrmics to execute applications in the OmpSs task-parallel programming model. The programmer writes a serial application, and splits it in small functions (tasks). Compiler pragmas are used to annotate each task with information about which of its arguments are to be read or written. Myrmics uses this information to understand how these tasks depend on one another and which ones can run in parallel. It then schedules them accordingly to available cores and performs appropriate DMA transfers to keep the producers' and consumers' caches coherent. We also enhance the OmpSs model with regions, a programming model concept that allows the user to group arbitrary collections of objects together. Regions enable task-level parallelization of irregular codes that traverse linked data structures using pointers. Myrmics supports hierarchies of regions and task nesting, which we believe to be promising solutions to the scalability problems encountered in current runtime libraries.

In Myrmics, not all cores are created equal. We use the eight powerful ARM cores as "schedulers", which run the main parts of the runtime system. Schedulers handle memory allocation, analyze the complex task dependencies, perform task scheduling and cooperate to balance the load across the system. We use the 512 weaker MicroBlaze cores as "workers", which run only a very thin portion of the runtime system and are reserved to obediently execute the tasks' code. Myrmics is a bare-metal runtime, which runs directly on the hardware prototype without any intermediate operating system.

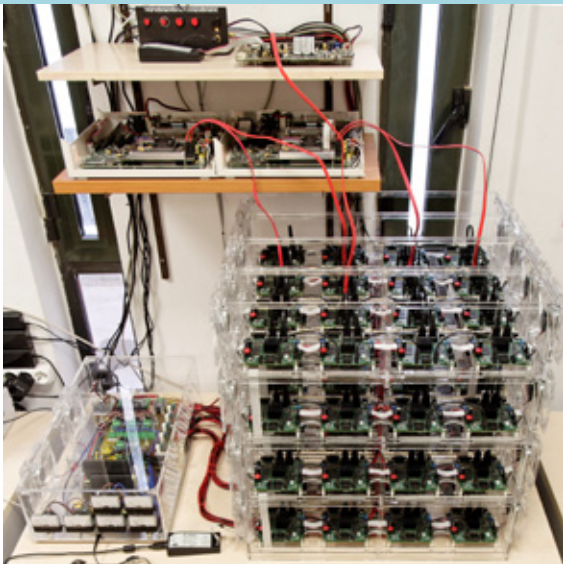

The 4x4x4 cube of 64 Formic boards (bottom right) connected to the two ARM Versatile Express platforms (upper middle).

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 at FORTH-ICS in Scalable Multicore Systems, you can find out more in http://www.ics.forth.gr/carv.

Spyros Lyberis
FORTH-ICS, Greece

HIPEAC MEMBER OLIVIER TEMAM APPOINTED SC CHAIR OF CGO

Olivier Temam, INRIA, has become the new Steering Committee chair of the ACM International Symposium on Code Generation and Optimization (CGO). CGO provides a premier venue to bring together researchers and practitioners working at the interface of hardware and software on a wide range of optimization and code generation techniques and related issues. The conference spans the spectrum from purely static to fully dynamic approaches, including techniques ranging from pure software-based methods to architectural features and support.

BSC VIDEO WINS INTERNATIONAL SCIENCE VISUALIZATION CHALLENGE PROMOTED BY NSF AND SCIENCE MAGAZINE

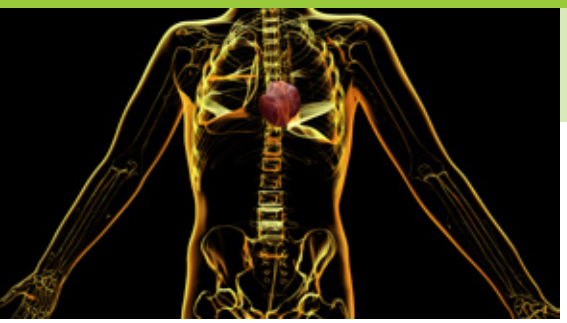

The National Science Foundation (NSF) announced the winning proposals from among 200 projects submitted from 18 countries

The Science magazine and the National Science Foundation (NSF) today announced the winners of the 10th annual international science and technology visualization challenge. The visualization team from the Barcelona Supercomputing Center – Centro Nacional de Supercomputación (BSC-CNS) was the winner in the category on visualization videos, with the video 'Alya Red: a computational heart', which explains the centre's project to simulate a human heart. The prize is the publishing of a report on this project and the other winning projects in the magazine's February issue, which will see the light tomorrow.

The jury of the event received over 200 proposals on science visualization, submitted by researchers, illustrators, photographers, programmers and graphic arts specialists from 18 countries; among these proposals was Alya Red.

At a later stage, a committee made up of members from the magazine and NSF made a selection of the most prominent works and published them on a website for the audience to cast their votes. According to the NSF, a total of over 3,150 individual votes were cast. The Alya Red video received both the jury's award and the audience award.

The BSC-CNS has a science visualization team that, in close collaboration with researchers, prepares videos to help interpret and visualize numerical data. This team is a combination of scientists, engiwneers and graphic designers who aim to produce representations of data that present the scientific results in an effective and visually attractive way.

Alya Red: a computational heart explains the research carried out by scientists from the centre and their collaborators, paying special attention to the complexity of the issue and describing the modern techniques used to tackle it. The video explains the motivation, means and goals of the project. The underlying point is the long standing issue of the heart and the use of supercomputing to try understanding it, which is reflected through a combination of original medical images and 3D representations of computer simulations.

ALYA RED, ONE OF THE STAR PROJECTS OF THE BSC-CNS

Alya Red, a project on Biomechanics, takes its name from the Alya System, the simulation tool developed exclusively at BSC-CNS. Its objective is to develop a computer model to simulate how a human heart works, and it is being developed by a multidisciplinary team of doctors, bioengineers and researchers in supercomputing and medical imaging. Currently, thanks to Alya Red, the scientists working on the project can simulate ventricular models coming from real geometries.

Another objective is to create a new tool to help provide a better understanding of how the cardiovascular system works to doctors carrying out clinical and pharmaceutical research. This tool will be a technological simulation infrastructure linked to high performance computing (HPC).

HIPEAC STUDENT GRANTED IBM PHD FELLOWSHIP AWARD

Pablo Llopis is one of the recipients of this award, which aims to support research, curriculum innovation and educational assistance

Pablo Llopis Sanmillán, a Ph.D. student of the Computer Architecture and Technology Area at the Computer Science and Engineering Department of Universidad Carlos III of Madrid, has been awarded the 2013 IBM Ph.D. fellowship. The IBM Ph.D. Fellowship Awards Program is an intensely competitive worldwide program, which honors exceptional Ph.D. students who have an interest in solving problems that are important to IBM and fundamental to innovation in many academic disciplines and areas of study. Awardees are granted a stipend for one academic year, and are encouraged to participate in an internship at an IBM Research facility during their studies. Motivated by the increasing importance of energy efficiency and growing data sets, Pablo's research is centered on efficient virtualized data sharing, which will enable scientific applications to achieve higher performance and efficiency when run on cloud-based HPC environments.

Pablo Llopis Sanmillán
Universidad Carlos III of Madrid, Spain

INTERNSHIP REPORT: KIRILL KONONENKO

Towards Dynamic Neutralization of Data Leakage at IBM Research Lab in Haifa

As a PhD candidate at Darmstadt Technical University in Germany, near Frankfurt-am-Main, I am working in the areas of cryptography, cyber-security, mobile, and cloud computing. The focus areas of my research include the detection and mitigation of data leakage in timing channels in statically compiled languages, such as C/C++.

Thanks to the HiPEAC internship grant I was able to work from September to December 2012 on the topic of dynamic code optimization at the IBM Research Lab in Haifa, Israel. At the IBM lab in Israel, I was hosted by Dorit Nuzman, Michael Vinov, Revital Eres, and Sergei Dyshel from the Post-link Analysis and Optimization Technologies group. The experience I gained during the internship allowed me to design a hybrid static/dynamic system for time-channel mitigation, and also to extend the Libjit-Linear-Scan library for the compilation of program code in CIL language into the IR used by the IBM JIT compiler.

Today, we use optimizing compilers to compile C/C++ cryptographic applications into binary code. Such applications are highly vulnerable to attacks by hackers, since compiler optimizations can result in time channels, which are variations in the execution time of cryptographic programs that may reveal full or partial information about the analyzed data, such as cryptographic keys or user passwords. This raises the challenge of automatically transforming critical code into a more secure state.

As a foundation for a unified approach to the secure execution of programs, I use the .NET virtual execution environment and the CIL bytecode. Usually, execution of the native CIL bytecode that is not within the CLI managed runtime is undesirable. But, in certain cases it allows non-standard optimizations and non-standard implementation of the functionalities of the CIL virtual machine. In my research, CIL bytecode is transformed into an Intermediate Representation (IR) of a specialized library for dynamic optimization, Libjit-Linear-Scan. Libjit-Linear-Scan supports methods of fast compilation and static analysis, based on linear scan. It generates rules in DKAL language that are saved near the binary code, for later use during runtime recompilation.

The Libjit-Linear-Scan IR is then compiled into the IR of the IBM JIT compiler, to allow subsequent dynamic recompilation based on the information computed by Libjit-Linear-Scan. Based on a heuristic, this framework will automatically select the optimum way to remove the vulnerability and the most robust technique for neutralizing the timing channel at runtime. The framework uses a combination of mitigation techniques, for example, insertion of timing delays in the code, code flattening, and code parallelization. Dynamic instrumentation and monitoring is used to limit the application of these transformations to vulnerabilities that actually materialize at runtime, rather than to all potential vulnerabilities (as would be done by static mitigation).

I would like to express my gratitude to HiPEAC for providing this excellent opportunity. This allowed me to see how research is conducted in industry, meet interesting people, and establish long-term contacts and collaborations. I think all students would greatly benefit from the internships and collaboration sponsored by HiPEAC.

Kirill Kononenko
Darmstadt Technical University, Germany

INTERNSHIP REPORT: LOIS OROSA NOGUEIRA

Building a flexible, configurable and expandable system-level functional simulator at Recore Systems

I am a PhD student at the University of Santiago de Compostela, working on hardware support for transactional memory systems and parallel debugging, under the supervision of Elisardo Antelo and Javier D. Bruguera.

Thanks to the HiPEAC internship grant, I had the opportunity to spend four months in Enschede (Netherlands), working on System-level simulation of multicore SoCs at Recore Systems. Recore is a semiconductor company that develops advanced digital signal processing platform chips and licenses reconfigurable semiconductor IP.

Simulation in computer architecture is indispensable for researching, development and testing of new architectures, and software development and debugging before the actual hardware is available. Furthermore, from the point of view of a hardware company, this results in a big reduction in cost.

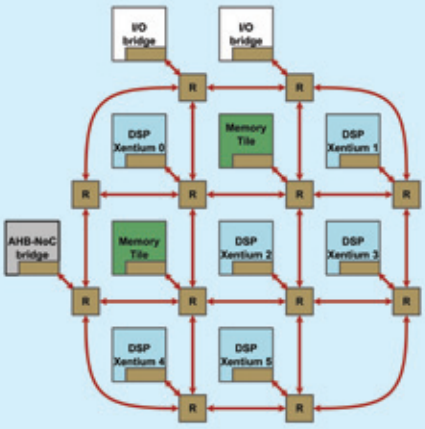

Simulated architecture

The main target of this project was to develop a flexible, configurable and expandable functional simulator to build new configurations and architectures that matches the needs of Recore and that could run unmodified code built for the company's FPGA prototype.
The shared memory simulator should be composed of several Digital Signal Processors (DSPs), one General Purpose Processor (GPP), and several peripherals (Memory Tiles, DMA controller, UART, etc.).
My starting point was a single -ore DSP Instruction Set Simulator (ISS) developed by the company, and from there I built a complete functional system. The first steps were to add some functionality to the ISS to match the characteristics of the prototype, and to develop the peripherals.

Lois Orosa Nogueira
University of Santiago de Compostela, Spain

INTERNSHIP REPORT: ALEJANDRO RICO

Power and performance evaluation of the ARM Cortex-A15 and comparison to ARM Cortex-A9, at ARM Cambridge

I am a PhD student at Universitat Politècnica de Catalunya - BarcelonaTech (UPC) and a junior researcher at Barcelona Supercomputing Center (BSC). The research for my PhD focuses on simulation methodologies for large multi-core systems under the supervision of Alex Ramirez and Mateo Valero. Apart from the work for my PhD, I also collaborate in a couple of EU-funded projects at BSC. One of these projects is the Mont-Blanc project, which aims to develop supercomputing technology out of low-power embedded parts.

The motivation for Mont-Blanc is that the increasing size of supercomputing systems is leading to power consumption figures reaching several MW (10^6 W). The performance of the current top supercomputers is under 20 PetaFLOPS (10^15 floating-point operations per second) with a power consumption of around 10 MW. The next target figure for supercomputer performance is one ExaFLOPS (10^18 FLOPS), but supercomputing experts claim that the power consumption of such a system should not exceed 20 MW. That requires an improvement in energy efficiency, beyond that of current supercomputers, which are generally built around general-purpose processing parts, of two orders of magnitude. Nowadays, the most energy-efficient processors are found in mobile phones and tablet computers.

At the beginning of the internship, our group at BSC had already set up a cluster of 256 ARM Cortex-A9 processors, a processor widely used in mobile phones and tablets. The results of this experience were encouraging, but they showed some room for improvement. At the same time, the new ARM Cortex-A15 processors showed promise for better energy efficiency while delivering higher performance. Thus, we wanted to evaluate how much better Cortex-A15 is, compared to Cortex-A9, for high-performance computing (HPC) applications. At that time, however, there was no product available that featured Cortex-A15 processors.

In order to carry out this comparison, I spent three months in the R&D division of ARM in Cambridge under the supervision of Chris Adeniyi-Jones. We employed a test chip featuring Cortex-A15 processors and two platforms using Cortex-A9, including our cluster at BSC. To add another energy-efficiency point for our evaluation, we also evaluated Cortex-A7, a lower-power lower-performance processor. The evaluation was performed in two phases. The first consisted of an architecture-level evaluation using small code kernels to stress certain parts of the micro-architecture. The objective was to get a fundamental understanding of the capabilities of each platform in order to get an upper bound in performance and see how these compared among the multiple platforms. In this first phase, we specifically evaluated the peak floating-point performance and memory bandwidth using code kernels in assembly and the popular benchmarks STREAM and LMBench. For the second phase of the evaluation, we used a set of benchmarks that make a broader use of the micro-processor but still characterise specific processing capabilities. Examples are Dhrystone (for integer performance), the ATLAS scientific library (for floating-point performance), HPLinpack (as a reference for supercomputing rankings) and the Mont-Blanc micro-benchmarks, which cover a set of algorithms employed in many HPC applications. The results showed a large improvement of Cortex-A15 over Cortex-A9 and Cortex-A7 for double-precision floating-point codes, due to its better floating-point unit, and they also showed better performance on control-intensive codes due to its higher issue width and better support for speculative execution. We also identified that work is needed to further improve their floating-point capabilities, and the software stack, including the operating system, drivers and scientific libraries, needs to be optimized for HPC rather than for mobile computing.

I want to thank HiPEAC for providing the frame and means for this collaboration. Also, I really recommend ARM R&D at Cambridge for an internship. I enjoyed my time there and I learned a lot during those three months.

INTERNSHIP REPORT: FOIVOS S. ZAKKAK

Power modelling of multi-core processors at Vector Fabrics

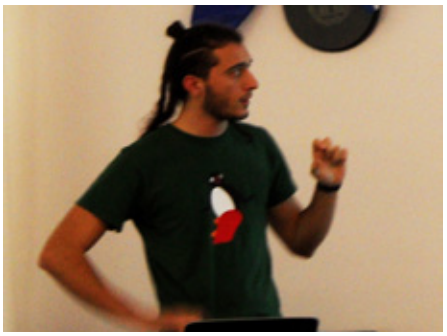

I am currently in my first year of the PhD in the Computer Science Department of the University of Crete. I am also working as a graduate research assistant at the Institute of Computer Science (ICS) of the Foundation for Research and Technology - Hellas (FORTH). Thanks to the HiPEAC funding for industrial internships I was able to spend foumor months working on power modelling of multi-core processors, an interesting and hot topic. During my internship at Vector Fabrics B.V., I had the chance to meet and work with some highly skilled developers in a great environment. This was the first time for me to work for a commercial product and the experience was good. Having a strong background on task parallelism and techniques to statically resolve dependencies between tasks, I found their work really interesting and relative to my research interests.

My work at the host's site focused on developing a new energy consumption model that I would later integrate to the company's product Pareon. The first week was spent on getting familiar with the environment and the company's software. In the second week, being more familiar with the software and its capabilities, I started creating the energy consumption model and studied some related work. After discussions with my mentor and other staff, I came up with a simple cycle-accurate energy consumption model working for both parallel and sequential code segments. The next step was calibration. I calibrated the model for both Intel Core i7-3770 and OMAP-4430 (ARM Cortex-A9). In the first case I used measurements from the performance counters and the Intel RAPL driver. For the OMAP-4430 calibration I used a modified Pandaboard, wired to a multimeter of 0.1 mV precision. The verification was encouraging with low error ranges so I proceeded to the model's integration in the Pareon tool. At the end, Pareon was able to print a breakdown of the energy consumption by computation, by data fetching and by the rest of the system (idling, leakage etc.).

This was an overall great experience and I would love to have more similar experiences before the end of my PhD.

Foivos S. Zakkak
FORTH-ICS, Greece

INTERNSHIP REPORT: ZEUS GÓMEZ

Porting the Barrelfish operating system to the Gem5 simulator at ARM Cambridge

My internship in the R&D department of ARM Ltd in Cambridge turned out to be very interesting. My main focus was to port the Barrelfish operating system, developed by the University of Zürich and Microsoft Research, to the Gem5 simulation model. This port had been partially done in the past, but using an old version of Barrelfish and an old version of the simulator. It was also using an old ARM development board that is no longer supported. So mainly the work was to update the operating system to use a new board for the simulator, new simulator version and new compiler version, merging it to the currently supported ARM boards in the operating system.

It didn't take long to complete this part, so taking into account the heterogeneous design of Barrelfish, we wanted to measure of performance of thread migration in the operating system, using traces and simulator counters. This would provide a comparison with commercial operating systems, such as Linux. In this respect, Linux is able to migrate threads from core to core in a quick way, although two kernel traps need to be performed; i.e. on both the source and destination cores. In Barrelfish, as each core is "isolated" from the other, a thread migration has to be performed in two stages. The first is to extend the virtual address space and process dispatcher to the new core and the second stage to do the actual migration. The first stage is not necessary in a shared-memory operating system like Linux as all processes share the memory space across all cores. But in Barrelfish this first step is makes it a slow process. Once the first stage is done, both Barrelfish and Linux perform equally well.

Staying in Cambridge for some months was wonderful as I could enjoy a very nice atmosphere at ARM. Those are a great group of researchers. I learnt how research is carried out in a company and how to work with a state-of-the-art simulator in an industrial environment. I really recommend staying there for an internship.

HIPEAC TECHNOLOGY TRANSFER AWARDS 2012

The HiPEAC technology transfer awards were granted last year for the first time

HiPEAC member	Affiliation	Partner	Title
Marco Santambrogio	Politecnico di Milano, Italy	Maxeler	PDR-X: Design Flow to Support Dynamic Partial Reconfiguration on Maxeler Architectures
Muhammad Shami	KTH, Sweden	Huawei	Multi-Mode Hardware Accelerator using Dynamically Reconfigurable Resource Array
Umberto Bondi	Alari/ Università di Lugano, Switzerland	Dolphin Engineering	Smart Vineyard –wireless Sensor network-based adaptive MAnagement system to foRecasT VINEYARD pests
Lieven Eeckhout	UGent, Belgium	Intel	Sniper: Fast and Accurate Manycore Simulation
Alex Ramirez	BSC/UPC, Spain	Seco	CARMA (CUDA on ARM) development kit
Piero Foglia	Università di Pisa, Italy	Isac	Small-DCS

PDR-X: DESIGN FLOW TO SUPPORT DYNAMIC PARTIAL RECONFIGURATION ON MAXELER ARCHITECTURES

Maxeler systems combine standard Intel CPUs with high performance Dataflow Engines (DFEs). Maxeler DFEs are composed of one or more FPGAs, memory, interconnect and other specialized resources and are optimized for performing high throughput computation on large datasets. For this reason, the Maxeler system can be viewed, at a lower level, as the result of the interaction between a GPCPU and one or more FPGAs. This work addresses the Maxeler platforms with the objective of designing a new methodology to support Partial Reconfiguration, due to the work done by the group at Politecnico di Milano, in the Maxeler design flow. The concept of Partial Reconfiguration (PR) enhances the flexibility and the possibilities offered by FPGAs: with PR it is possible to modify the behavior of a determined portion of the device while the rest of the board keeps on working (e.g., to spatially-share the usage of the FPGA among multiple applications, interleave different computational kernels, or to tolerate faults at runtime).

Maxeler's MaxCompiler design flow allows the programmer to write a Java program describing the DFE configuration, integrate with CPU code in C and automatically configure and manage the FPGAs contained in the DFE at runtime. The Maxeler flow does not allow specification of a partially reconfigurable design - every DFE program represents a complete chip configuration. Supporting PR in the Maxeler flow would allow the design of a more flexible application, where parts of the DFE behavior can change at run time without affecting the remainder of the DFE. But flexibility is not the only improvement brought by PR: its employment would allow the system to overcome limitations due to the limited FPGA capacity, by time-multiplexing the device. An initial implementation has already been completed, and work is ongoing to integrate a first public version of PR functionality into MaxCompiler in early 2013. The technology transfer between PoliMi and Maxeler to support PR design in Maxeler architecture will:

- Enhance the Maxeler design tool chain to open future design possibilities for the system that can be realized with such a technology. Maxeler recognize the benefits of this technology transfer and this is demonstrated by the integration in their official tool chain of the work done in collaboration with PoliMi to support PR.
- Open future possibilities for EU projects that can use Maxeler solutions as runtime reconfigurable systems for their designs. This point is supported, as an example, by the work done in the context of the FASTER EU project that will use this technology to prototype their solutions.

More information can be found at the following links:
FASTER webpage: http://www.fp7-faster.eu/
FASTER YouTube Channel: https://www.youtube.com/user/FASTERprj

MULTI-MODE HARDWARE ACCELERATOR USING DYNAMICALLY RECONFIGURABLE RESOURCE ARRAY

KTH collaborated with Huawei to design a runtime reconfigurable multi mode accelerator to implement 256/512/1024/2048-point FFT for 4G LTE and a multi-finger CP for 3G UMTS. The requirement for total gate count and active mode power consumption, at a clock speed of 450 MHz, was set to be less than 500K gates and 150mW respectively. The 2048-point FFT had to complete in 5000 cycles, and the multi-finger CP algorithm had to compute for a maximum of 1024 fingers in 256 Chips (time), when the chip rate is 1*3.84 Mbps.

Dynamically Reconfigurable Resource Array (DRRA), shown in Figure 1, is a CGRA fabric developed in KTH, and is used to implement both the FFT and the CP algorithm in mutually exclusive way. Only the configware for the respective algorithms are loaded to run the algorithms. All the design and timing constraints were met. The FFT algorithm has been verified against a 64-bit floating point Matlab model and CP against a C# model built in parts in consultation with the Huawei team in Kista, Sweden. A list of publications relevant to the topic can be found on http://web.it.kth.se/~shami/.

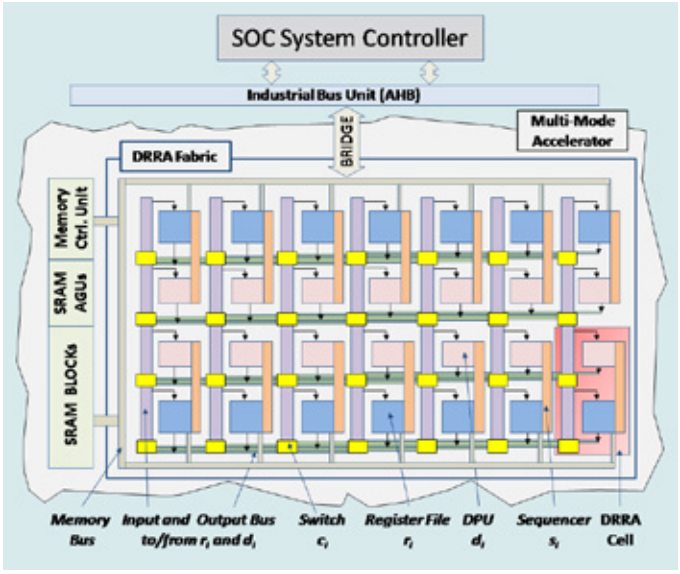

12 DRRA cells are used to implement the two kernels: FFT and Correlation Pool. One Kernel is executed at a time.

SNIPER: FAST AND ACCURATE MANYCORE SIMULATION

Sniper is a fast yet accurate x86 multi/many-core simulator. Its key features include analytical core models to raise the level of abstraction in architectural simulation, improving both development and simulation time, and a parallel simulation engine to scale simulation performance with simulation host core count. Sniper allows for simulating both multi-threaded and multi-program workloads on multi/many-core architectures with up to hundreds of cores at a simulation speed up to several MIPS. Sniper has been validated against Core2, Nehalem and Xeon Phi hardware, and is especially valuable for uncore and system-level design studies involving long-running workloads.

Sniper is a result of the collaboration between Ghent University and the ExaScience Lab, one of the four Intel European Exascale Labs, and is currently being used as simulation platform for Intel's pathfinding activities.

More information about Sniper can be found at http://snipersim.org/.

Lieven Eeckhout, Wim Heirman, Trevor E. Carlson, Stijn Eyerman, Kenzo Van Craeynest
Ghent University, Belgium
Ibrahim Hur, Alexander Isaev
Intel, Belgium

CARMA (CUDA ON ARM) DEVELOPMENT KIT

The Heterogeneous Architectures research team of the Barcelona Supercomputing Center (BSC), headed by Dr Alex Ramirez, has been awarded the 2012 HiPEAC Technology Transfer Award for its development of an innovative platform combining, for the first time in the industry, an ARM platform with a programmable GPU. BSC defined and coordinated the design of the CARMA platform, while the company SECO provided the board design and product manufacturing and validation. NVIDIA contributed to this platform with the silicon technology and software stack support for the ARM platform. The research project has also been part of the PRACE initiative to explore a set of prototypes to test and evaluate promising new technologies. The result of this project was a commercial product called CARMA Development Kit. The product was released in summer 2012 and was a huge sales success. "This award recognizes the fact that Europe produces technologies that will allow computers to achieve higher performance with lower power consumption, and at a lower cost," says Mateo Valero, BSC Director.

Caption: The CARMA Development kit.

SMALL-DCS

The new demands for flexibility and reconfigurability in manufacturing and process industries has determined the transition from centralized automation systems to systems with distributed intelligence, an important class of which is represented by Distributed Control Systems (DCS). The new ISO-IEC 61499 standard for DCS introduces advanced software engineering concepts in the design of such systems. Its diffusion is currently limited, however, since its adoption requires the designer to abandon the widely used IEC 61131 standard. To overcome this limitation, we are designing an architecture that allows the integration of systems based on IEC 61131-3 into an IEC 61499 DCS.

ISAC s.r.l., http://www.isacsrl.eu, an SME operating in Tuscany (Cascina, Italy), was interested in enhancing their PLCs lines with DCS capabilities, but it was mandatory for ISAC to continue to conform to the IEC 61131 standard for the specification of the control law. An implementation of our architecture can satisfy ISAC requirements: IEC 61499 code for the distribution can be generated transparently to the user, which can program control law in the IEC 61131 language. As a consequence, ISAC supported the further development of our research, and the results are going to be transferred to ISAC products. Together with ISAC, we are planning to integrate our solution in IEC 61131 development environments.

IEC 61499 and IEC 61131-3 control logic coexist in the same control environment. Interactions between IEC 61131-3 programs and IEC 61499 applications occur through PLC Interfaces (PIs), which are realized using interface function blocks performing PLC data exchanges (PDEs). Execution semantics of each standard are not modified, so IEC 61131-3 programs are mapped to tasks while IEC 61499 applications have event-driven execution.

Pierfrancesco Foglia,
Università di Pisa, Italy

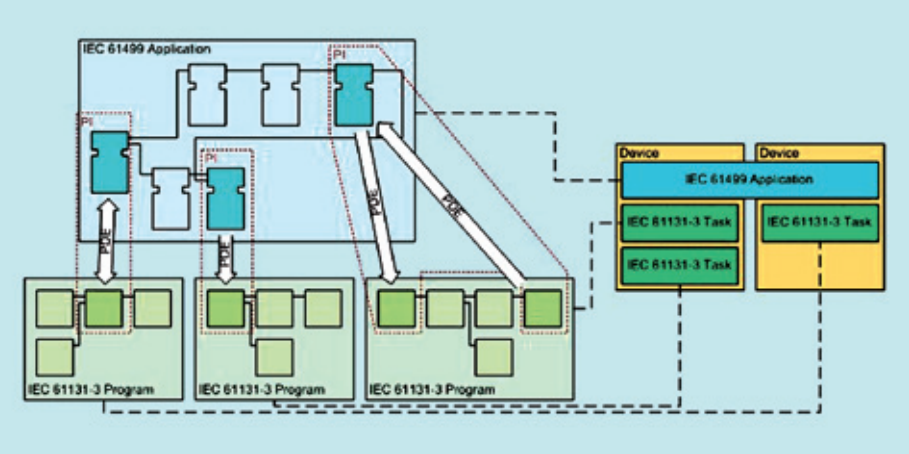

The proposed architecture.

VECTORIZATION MADE EASY: THE SARC APPROACH

Computational science and engineering has only increased in importance over the last decades, driven by the advances in high-performance computing. Many of the computational challenges are related to efficient manipulation of large data structures. At the same time, single-processor performance growth recently stagnated due to power and thermal dissipation constraints, and the additional resources provided by the technology are turned into more (often specialized) on-chip processing elements. As a result, innovative approaches are emerging to translate technology improvements into performance advances.

The Scalable ARChitecture (SARC) project (FP6 – FET, contract #27648) worked out a scalable architecture applicable to a wide range of application domains. SARC proposed a heterogeneous multicore architectural template based on several high-performance master processors and multiple clusters with application-specific

From left to right, Prof. Yale Patt, Prof. Georgi Gaydadjiev, Dr. Catalin Ciobanu, Prof. Per Stenstrom and Dr. Peter Hofstee.

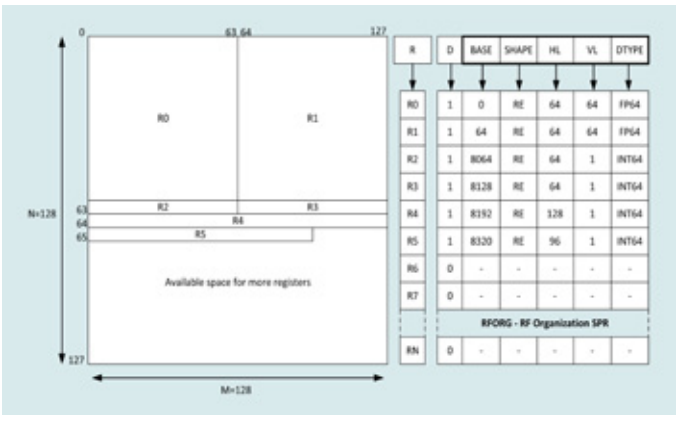

The Polymorphic Register File

accelerators for, especially, multimedia, bioinformatics and scientific computing.
Dr. Cătălin Ciobanu's work, successfully defended on March 8, 2013 (see photo), has focused on the SARC Scientific Vector Accelerator targeting data parallel workloads typically computed using the Single Instruction Multiple Data (SIMD) approach. SIMD processors, however, are notorious for turning performance programmers into low-level hardware experts which hurts productivity. Moreover, legacy programs often require rework to follow (micro)architectural evolutions. In his dissertation, Dr. Ciobanu addresses the problems of SIMD-accelerator programmability, code portability and performance-efficient data management. The proposed Polymorphic Register File (PRF) provides a simple programming interface, allowing programmers to focus on algorithm optimizations rather than complex data transformations or low-level details. The overall PRF size is fixed, while the actual number, dimensions and sizes of its individual registers can be readjusted at runtime. The figure above depicts an example of a 128KB PRF.
The HDL implementation achieves clock frequencies of 100 to 300MHz for FPGA and 500 to 900MHz for ASIC, demonstrating the PRF practical usability. In specific cases, PRF systems can execute by up to three orders of magnitude fewer instructions than the Cell scalar core, and can outperform the Cell SIMD engine by up to three times. For separable 2D convolution, PRFs outperform the NVIDIA C2050 GPU for masks of 9 x 9 elements and bigger.
The research will continue at Chalmers, and the results will be used in the scope of the Facilitating Analysis and Synthesis Technologies for Effective Reconfiguration (FASTER) (FP7 – contact # 287804) project. Dr. Ciobanu's PRF FPGA implementation provides FASTER designers with powerful tools to handle 2D scientific data, allowing low-overhead runtime adjustment of additional PRF parameters, e.g., the overall storage size and aspect ratio, the number of read/write ports and vector lanes.

For more information, contact Dr. Ciobanu (email:catalin@chalmers.se)

RESEARCH PROJECT FOR ENERGY-EFFICIENT RISK MANAGEMENT ACCELERATION STARTED

FPGAs can save up to 90% of energy in finance and insurance

The precise and significant simulation of risk scenarios is essential for rating current and future portfolios and investments in the financial and insurance sectors. Even large CPU or GPU clusters can have a workload of many hours for such simulations, resulting in huge energy consumption.
The ESR project is an interdisciplinary cooperation between partners from the sectors of asset management, financial mathematics, and hardware design. The main objective is to make current reconfigurable hardware (field programmable gate arrays, FPGAs) easily applicable for users in the finance and insurance markets. It has already been shown that this approach can save up to 90% of energy for simulations, compared to state-of-the-art CPU or GPU implementations. In order to achieve that goal, a holistic algorithm / hardware / software co-design flow is mandatory to design the optimal implementation over all design levels.
The technical challenge is to keep the high flexibility that is always required because of fast changing product descriptions, models and algorithms while using efficient and therefore dedicated accelerators. During the project the consortium develops a demonstrator for accelerating selected algorithms that are relevant for industry. This demonstrator will form the foundation for a future commercial platform system.
The project consortium consists of:

- University of Kaiserslautern: Microelectronic Systems Design Research Group
- Karlsruhe Institute of Technology (KIT): Information processing technologies
- Assenagon GmbH, Munich: Asset Management
- Creonic GmbH, Kaiserslautern: Hardware design for FPGAs
- cronologic GmbH & Co KG, Frankfurt: FPGA boards and integration
- Fraunhofer Institute for Industrial Mathematics (ITWM), Kaiserslautern

The ESR project is funded by the Federal Ministry of Education and Research and is accompanied by the German Aerospace Center (DLR). The duration of the project is three years from August 2012 until July 2015.

Project website: www.esr-projekt.de

Christian de Schryver
University of Kaiserslautern, Germany

POWER CONSUMPTION: THE MOST IMPORTANT ISSUE IN COMPUTING

Energy-efficient ICT infrastructures are a survival issue of our local and global economy. Until the year 2030, if current trends continue, the electricity consumption caused by the internet is predicted to grow by a factor of 30, reaching much more than that of the total world today. Energy prices will grow substantially. The next generation of oil and gas seismic simulations will require two or more orders of magnitude more computations. Already during the past 11 years the price of crude oil increased by a factor of 9. What's next until the year 2030? Assuming only a factor of 10 would lead here to an overall unaffordable electricity cost increase by a factor of 300.

The strong tendency of growing wireless access and the growth of the cloud computing business also massively increases the electricity cost for communication. This altogether would lead to a collapse of our ICT infrastructures and the end of the global economy. A future peta- or exa-flop supercomputer will need its own power plant. At the Top 500 list from 2008 thru 2012 the energy efficiency has improved from 228 MFlops/Watt to 280 MFlops/Watt, yielding only an average factor of 1.05 per year. To avoid collapse, we must reinvent not only data centers and supercomputing, but also computer engineering and computer science totally. Compared to the current status of research on power-efficient high performance ICT, we are urgently forced to achieve results being better by several orders of magnitude. This is possible only by a radically new structure of the entire R&D landscape of both industry and academia and only by drastically reformed ICT-related education, since many fundamental backgrounds have dramatically changed.

We have to detect much more efficiency-promising subareas. To obtain many synergy effects for more parallelism we must introduce connected thinking everywhere by removing the barriers between abstraction levels and between paradigm domains: communication, instruction streams, data streams, configware, hardware and heterogeneous systems. We must introduce thinking into new patterns of cross-cutting design trade-offs, involving power, performance, reliability and quality of service, allowing power to play more active role, being not only an optimization criterion but also a constraint. To obtain a lot of synergy effects we must radically reconsider the interactions between design flow, dynamic reconfiguration, algorithms, programming, languages, innovative compilers and runtime systems and much more. All this is the scope of the PATMOS conference at Karlsruhe, Germany, September 9 - 11, 2013. See http://www.itiv.kit.edu/patmos-vari2013/

Reiner Hartenstein
TU Kaiserslautern, Germany

CUSTOM FLOATING-POINT ARITHMETIC FOR INTEGER PROCESSORS: ALGORITHMS, IMPLEMENTATION, AND SELECTION

Jingyan Jourdan-Lu
Ecole Normale Supérieure de Lyon/STMicroelectronics
Grenoble, France
Advisors: Claude-Pierre Jeannerod, Christophe Monat, Jean-Michel Muller
November 2012

Media processing applications typically involve numerical blocks that exhibit regular floating-point computation patterns, which, on processors without hardware floating-point units, can be turned profitably into custom operators. This thesis addresses the design of such custom operators as well as the compilation techniques to select them in application codes. Specific algorithms targeting high instruction-level parallelism are proposed and implemented for squaring, scaling, two-dimensional dot products, simultaneous sine/cosine, etc. On DSP benchmarks, supporting the selection of these custom operators allows speed ups of up to 1.59x, compared to the sole usage of basic ones (+, -, X, / and √).

PROFILING METHODS FOR MEMORY CENTRIC SOFTWARE PERFORMANCE ANALYSIS

David Eklöv
Uppsala University, Sweden
Advisors: Prof. Erik Hagersten and Prof. David Black-Schaffer
December 2012

This thesis presents several novel profiling methods for memory-centric software performance analysis of cache and bandwidth usage. The goal is to provide general, high-level, quantitative information describing how the profiled applications utilize the resources in the memory hierarchy, and thereby help software and hardware developers identify opportunities for optimization. Our techniques have minimal data collection overhead, are not dependent on custom hardware and/or operating system extensions, and provide accurate and easy to interpret information. These methods provide high-level performance metrics (CPI, bandwidth, cache) as a function of memory resource allocation for multicores..

STREAM-BASED COMPUTING AND FINE-GRAINED PARALLELISM: FROM ALGORITHMS TO RECONFIGURABLE HARDWARE

Frederico Pratas
Universidade Técnica de Lisboa, Portugal
Advisors: Prof. Leonel Sousa, and Associate Prof. Pedro Trancoso
December 2012

This thesis is focused on exploiting fine-grained parallelism and stream-based computing to accelerate a set of applications, mainly from Chemistry and Bioinformatics. The application behavior is analyzed for multicore CPUs, Cell/BE and GPUs. The results show that parallel efficiency and performance are highly application-specific. Based on these results, a new systematic method for developing stream-based applications and accelerators for streaming platforms is also proposed. Moreover, a multi-stream multi-core reconfigurable accelerator architecture is also proposed, which provides extra hardware-to-software flexibility and adaptability by supporting reconfiguration and reprogrammability.

METHODS TO ENHANCE CONTENT DISTRIBUTION FOR VERY LARGE SCALE ONLINE COMMUNITIES

Juan M. Tirado
Universidad Carlos III de Madrid, Spain
Advisors: Prof. Jesus Carretero and Prof. Florin Isaila
January 2013

In this thesis, we propose a novel framework that employs automatically-generated workload forecasting models to design control solutions. Prediction models enable the designer to determine with enough anticipation the workload to be served by the system. These predictions are employed to determine the system size and to design locality-aware data distribution policies that exploit temporal data access patterns. Additionally, we combine the models with adaptive solutions that permit the system to adapt if the predictions become inaccurate. In particular, we propose two novel error-based metrics. Our evaluations show improvements in content locality and resources utilization even under unexpected workload peaks..

FINE-GRAINED CONTROL FLOW ERROR CHECKING FOR HARD REAL-TIME SYSTEMS

Julian Wolf
University of Augsburg, Germany
Advisor: Prof. Dr. Theo Ungerer
January 2013

Safety and reliability are key requirements in many areas of embedded systems. In this thesis, we propose a mechanism to detect logical and temporal control flow errors in hard real-time systems. The fine-grained operation with very low detection latency enables an immediate reaction to any misbehaviour and possibly the execution of a fall-back solution within the required deadlines. In addition to a detailed description of the developed technology, this work presents a reference implementation for a real-time capable processor model. Based on evaluations using fault injection, we show both the effectiveness of the detection mechanism and its low overhead.

UPCBLAS: A NUMERICAL LIBRARY FOR UNIFIED PARALLEL C WITH ARCHITECTURE-AWARE OPTIMIZATIONS

Jorge González Domínguez
University of A Coruña, Spain
Advisors: Prof. María José Martín Santamaría, Prof. Juan Touriño Domínguez
February 2013

This PhD thesis describes UPCBLAS, a parallel library for numerical computation using the PGAS Unified Parallel C (UPC) language. The routines are built on top of sequential BLAS functions, and exploit the particularities of the PGAS paradigm, taking into account data locality in order to achieve good performance. This thesis also presents Servet, a suite of benchmarks focused on detecting a set of parameters with high influence on the overall performance of multicore systems. UPCBLAS routines use the hardware parameters provided by Servet to implement optimization techniques that improve their performance.

SYSTEM-LEVEL POWER ESTIMATION METHODOLOGY FOR MPSOC BASED PLATFORMS

Santhosh Kumar Rethinagiri
University of Valenciennes/ INRIA Lille Nord Europe, France
Advisors: Prof. Jean-Luc Dekeyser, Prof. Smail Niar and Assoc. Prof. Rabie Ben Atitallah
March 2013

My PhD thesis proposes a tool based on efficient hybrid system-level power estimation methodology for MPSoC platforms. Here, a combination of Functional Level Power Analysis (FLPA) and system-level estimation technique is used to compute the power consumption of the whole system. This combination leads to a hybrid power estimation, which gives a better trade-off between accuracy and speed. Based on the proposed novel methodology, a Power Estimation Tool at System-Level (PETS) is developed. My areas of expertise are virtual platform development (SystemC, C++, TLM), Just-In-Time compilation, and multimedia applications. My interests are MPSoC, scheduling, compilation and low power-aware design.

CUSTOMIZABLE REGISTER FILES FOR MULTIDIMENSIONAL SIMD ARCHITECTURES

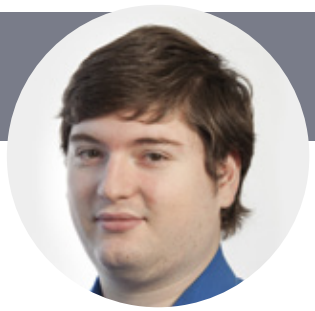

Cătălin Bogdan Ciobanu
Delft University of Technology, The Netherlands
Advisors: Prof. dr. ir. G. Gaydadjiev, Prof. dr. ir. H. J. Sips
March 2013

We proposed a Polymorphic Register File (PRF) allowing SIMD programmers to focus on algorithm improvements instead of low-level data manipulations. The number and dimensions of all PRF registers are adjustable at runtime. With clock frequencies of 100 to 300MHz on FPGA and 500 to 900+MHz on ASIC, our PRF can be used in real systems. We show how a PRF-based vector processor outperforms the Cell SIMD engine by up to 3X for matrix workloads. In addition, depending on the size of the vector registers, PRF reduces the number of committed instructions by up to three orders of magnitude. For separable 2D convolution, PRF systems outperform a state-of-the-art NVIDIA GPU for masks of 9 x 9 elements and bigger.

SPECULATIVE STATE AND CONTENTION MANAGEMENT IN HARDWARE TRANSACTIONAL MEMORY

Anurag Negi
Chalmers University of Technology, Sweden
Advisor: Prof. Per Stenstrom
March 2013

The Thesis presents analyses and designs that further current understanding of the behavior and relative performance of various hardware transactional memory (HTM) designs. HTM is expected to provide significant boost in productivity to programmers writing multithreaded applications in shared memory systems. The thesis proposes and evaluates 1) designs that can adapt their behavior based on workload characteristics; 2) the use and impact of write buffering techniques in HTM designs; 3) a cache designs that exploits cache associativity to support multiple speculative contexts; and 4) a prefetching technique that speeds up transactions and improves high contention performance.

PROGRAMMING HETEROGENEOUS MPSOCS: TOOL FLOWS TO CLOSE THE SOFTWARE PRODUCTIVITY GAP

Jeronimo Castrillon
RWTH Aachen University, Germany
Advisor: Prof. Dr. Rainer Leupers
April 2013

Embedded systems have evolved from single-processor platforms running single applications to complex heterogeneous Multi-Processor Systems-on-Chip (MPSoCs), capable of executing multiple applications. This hardware complexity, together with stringent application requirements, makes programming a daunting task. This thesis presents tool flows for improving the programmers' productivity, including a parallelism extraction flow for C applications, a mapping and scheduling flow for parallel applications and a flow for analyzing multiple applications at design time. An evaluation on heterogeneous platforms revealed a significant productivity increase with a low efficiency loss.

UPCOMING EVENTS

The 8th IEEE Symposium on Industrial Embedded Systems (SIES 2013)
19-21 June 2013, Porto, Portugal http://www.cister.isep.ipp.pt/sies2013/

The 40th International Symposium on Computer Architecture (ISCA 2013)
23-27 June 2013, Tel Aviv, Israel http://isca2013.eew.technion.ac.il/

International Conference on Localization and GNSS (ICL-GNSS 2013)
25-27 June 2013, Torino, Italy http://www.icl-gnss.org/

International Conference on Embedded Computer Systems: Architectures, Modeling, and Simulation (SAMOS XIII)
15-18 July 2013, Samos, Greece http://www.samos-conference.com/

The 9th International Conference on Wireless and Mobile Communications (ICWMC 2013)
21-26 July 2013, Nice, France http://www.iaria.org/conferences2013/ICWMC13.html

The 18th International European Conference on Parallel and Distributed Computing (Euro-Par 2013)
26-30 August 2013, Aachen, Germany http://www.europar2013.org/

The 23rd International Conference on Field Programmable Logic and Applications (FPL 2013),
2-4 September 2013, Porto, Portugal, http://www.fpl2013.org/

The 22nd International Conference on Parallel Architectures and Compilation Techniques (PACT 2013)
7-11 September 2013, Edinburgh, Scotland http://www.pactconf.org/

23th International Workshop on Power and Timing Modeling, Optimization and Simulation (PATMOS 2013)
9-11 September 2013, Karlsruhe, Germany http://www.itiv.kit.edu/patmos-vari2013/

International Conference on Parallel Computing (ParCo 2013)
10-13 September 2013, Munich, Germany http://www.mac.tum.de/parco2013/

The 5th Mini-Symposium on Parallel Computing with FPGAs (ParaFPGA 2013)
10-13 September 2013, Munich, Germany http://parafpga.elis.ugent.be//

International Symposium on System-on-Chip 2013 (SoC 2013)
23-24 October 2013, Tampere, Finland http://soc.cs.tut.fi/

The 9th International Conference on High Performance and Embedded Architectures and Compilers (HiPEAC 2014)
20-22 January 2014, Vienna, Austria 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 34
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. FP7/ICT 287759
website: http://www.hipeac.net.
subscriptions: http://www.hipeac.net/newsletter

design: www.magelaan.be

AUTUMN COMPUTING SYSTEMS WEEK, 7-9 OCTOBER 2013, TALLINN, ESTONIA