HiPEAC info 42

COMPILATION ARCHITECTURE

APPEARS QUARTERLY
MAY 2015

NETWORK OF EXCELLENCE ON HIGH PERFORMANCE AND EMBEDDED ARCHITECTURE AND COMPILATION

WELCOME TO THE SPRING COMPUTING SYSTEMS WEEK, MAY 5-7, 2015, OSLO, NORWAY

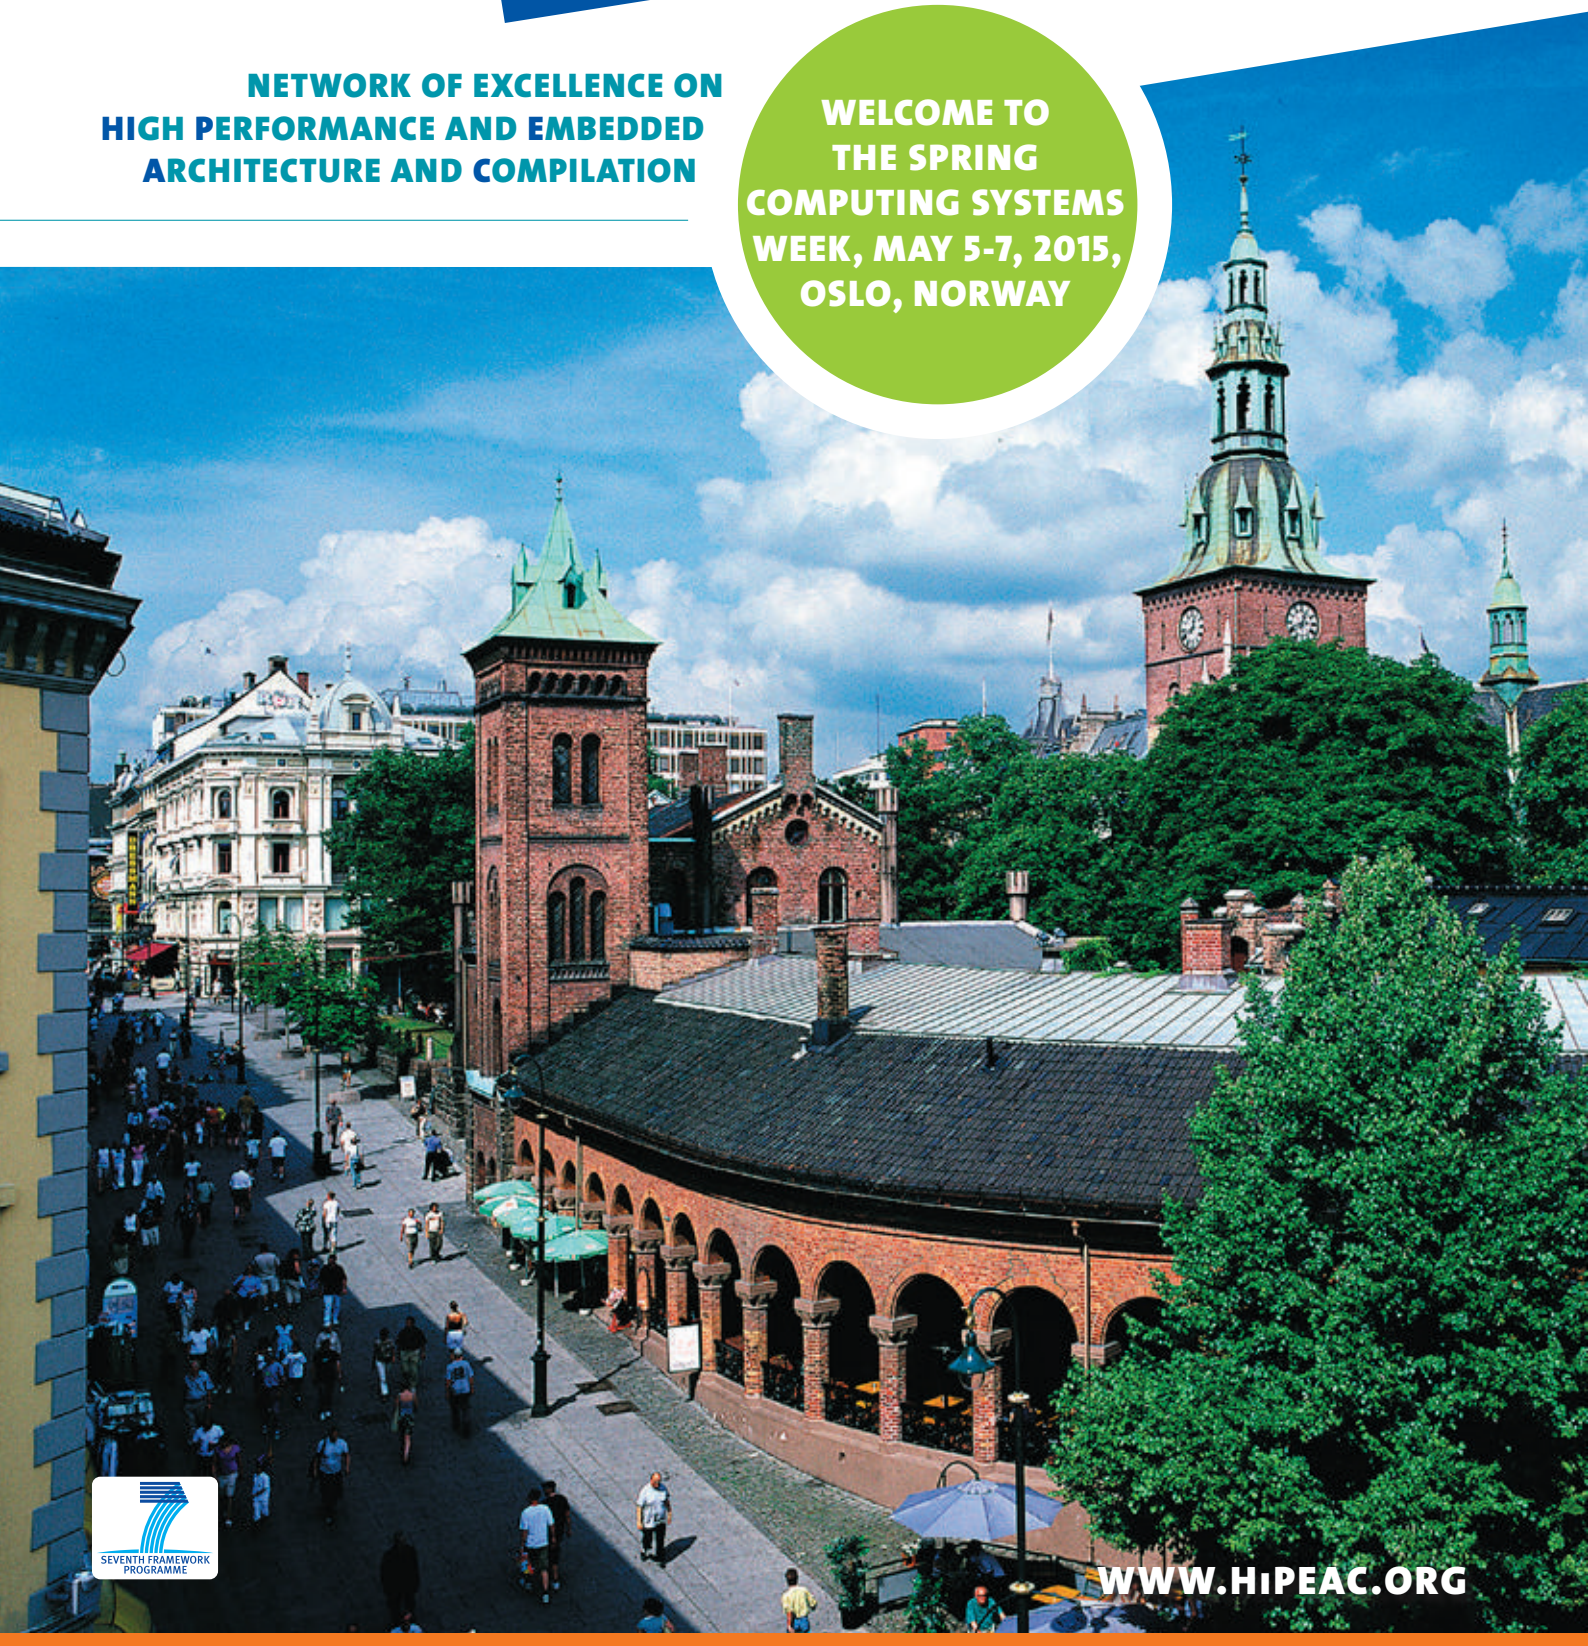

SEVENTH FRAMEWORK PROGRAMME

WWW.HIPEAC.ORG

11TH INTERNATIONAL SUMMER SCHOOL ON ADVANCED COMPUTER ARCHITECTURE AND COMPILATION FOR HIGH-PERFORMANCE AND EMBEDDED SYSTEMS (ACACES 2015), JULY 12-18, 2015, FIUGGI, ITALY

MESSAGE FROM THE HIPEAC COORDINATOR

Uber was founded as UberCab in San Francisco in 2009. It wants to change the business of commercial transportation of persons by introducing a new business model based on a combination of cloud-based technology and crowd sourcing. Several cities have already banned Uber in order to protect the jobs of the licensed taxi drivers and the tax income they generate for the cities. I am not sure though that cities and taxi drivers are fighting the right enemy. Uber still needs drivers, and taxi drivers could become Uber drivers. The real disruption will come when Uber no longer needs human drivers. Uber recently started a collaboration with Carnegie Mellon University on self-driving cars. According to Carlos Ghosn, CEO of Nissan Co, the commercialization of driver-less cars across Europe is planned for 2020. According to Elon Musk, people – including taxi drivers – could be forbidden from driving cars by 2035. By then, having a transportation contract with Uber might be cheaper than owning a car, and Uber may become, for human transportation, what Google is today for search – meaning that 88% of all (driverless) cars in a city could be (electric) Uber cars continuously driving from destination to destination, without a need for parking lots. Of course, it may also be the case that by then only old people like me will remember the name Uber.

Many of you will read this newsletter when you arrive at the HiPEAC Spring Computing Systems Week in Oslo. This is the second time that we organize a Computing Systems Week in an associated country, i.e. a country that is not part of the EU-28, but that participates in the framework program. I would like to thank our Norwegian colleagues for hosting us this week.

This will be a Computing Systems Week that runs for a whole week, because it is the co-location of two events in one week. The first event is the block review, during which a dozen EU-funded research projects will undergo their review. It will also include other meetings and events where the projects will meet. The goal is to stimulate interaction between projects, and to learn from each other on how to maximize the impact of the projects.

The second event consists of the traditional thematic sessions of the HiPEAC network. These will run on Tuesday, Wednesday and Thursday. I am happy to see that several members took the opportunity to organize a thematic session focused on important HiPEAC themes.

Take care,
Koen De Bosschere

CONTENT

Cover photography: VisitOSLO/Normanns Kunstforlag/Terje Bakke Pettersen

REPORT ON THE HIPEAC 2015 CONFERENCE

With 631 participants, this was the biggest event ever organized by the HiPEAC Network

Ready for HiPEAC2015!

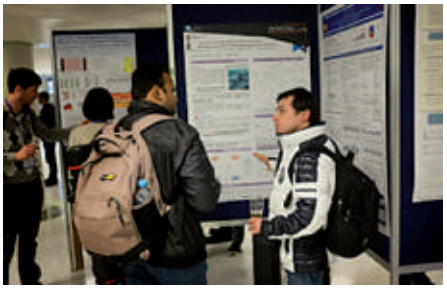

Poster Session on Monday January 19

Keynote of James Larus

Poster Session on Monday January 19

Social Event at the Maritime Museum

More than 600 delegates from 59 countries met in the RAI Conference Center in Amsterdam, the Netherlands, from January 19 to 21, 2015 for the 10th HiPEAC Conference. All attendees joined us for three days of presentations, discussions and networking in a wide variety of events. Posters and industrial presentations completed the event, which focused on creating networking opportunities and establishing new collaborations.

Highlights

- Keynote talks from industry as well as academia
- 38 presentations in the main paper track
- An industrial session with invited presentations from Google, Nvidia, ARM, Ericsson, ASML and many others
- 30 workshops and 4 tutorials with research presentations, inspiring keynote talks, and panel discussions
- Daily poster sessions with research and project results
- An industrial exhibit with 42 participants
- An amazing dinner served at the Maritime Museum of Amsterdam

The survey shows that the conference's format with many parallel events, poster sessions and exhibitions is highly appreciated by the community. Moreover, several comments mentioned the uniqueness of the HiPEAC Conference in bringing together the research community and industry. A large part of the European community of computer architecture and compilation met at the HiPEAC conference for an intensive exchange. This success is only possible with the generous support from the 21 sponsors. The sponsorship they offer is essential to be able to offer attractive registration fees and to support student participation. More information about this year's conference can be found on the conference website: https://www.hipeac.org/2015/amsterdam/.

Next year, the HiPEAC Conference will take place in Prague, Czech Republic from January 18 to 20. Again, the conference will follow the journal-first 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 paper track, a number of satellite events including poster sessions, workshops and tutorials 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 workshops and tutorials for the next conference.

See you next year in Prague!

Thank you to our sponsors for their generous support!

A STUDENT'S IMPRESSION FROM HIPEAC 2015

My name is Tsvetan Shoshkov and I am a PhD student at the Technical University of Sofia, Bulgaria, currently doing research at Leiden University, the Netherlands. I am very glad that I had the opportunity to attend the HiPEAC 2015 conference in Amsterdam, and I would like to thank the organization for the grant I received.
The conference included many interesting workshops, some of which are in the field of my current work. I enjoyed the keynotes and the opportunity to connect with big companies in this area.

I had the opportunity to speak with fellow colleagues during coffee and lunch breaks, and was happy to see people that I had already met at a previous HiPEAC event. My colleagues and I really appreciated the enjoyable and ample food provided during the conference at the RAI Congress Center. The peak point in the social events was the boat trip to the Scheepvaartmuseum (National Maritime Museum) for an official dinner. I really enjoyed the atmosphere in the recently renovated museum, and getting to know the local culture is always useful. I am looking forward to taking part in the conference next year.

Tsvetan Shoshkov, Technical University of Sofia, Bulgaria

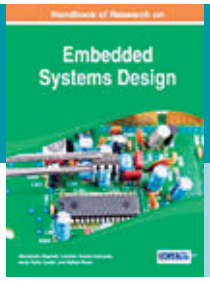

BOOK: HANDBOOK OF RESEARCH ON EMBEDDED SYSTEMS DESIGN

Authors: Alessandra Bagnato (Softeam R&D, France), Leandro Soares Indrusiak (University of York, UK), Imran Rafiq Quadri (Softeam R&D, France) and Matteo Rossi (Politecnico di Milano, Italy)

As real-time and integrated systems become increasingly sophisticated, issues related to development life cycles, non-recurring engineering costs, and poor synergy between development teams will arise.

The Handbook of Research on Embedded Systems Design provides insights from the computer science community on integrated systems research projects taking place in the European region. This premier reference work takes a look at the diverse range of design principles covered by these projects, from specification at high abstraction levels, using standards such as UML and related profiles, to intermediate design phases. This work will be invaluable to designers of embedded software, academics, students, practitioners, professionals, and researchers working in the computer science industry.

URL: http://www.igi-global.com/book/handbook-research-embedded-systems-design/102212

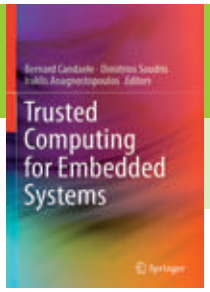

BOOK: TRUSTED COMPUTING FOR EMBEDDED SYSTEMS

Co-Editors: Bernard Candaele, Dimitrios Soudris and Iraklis Anagnostopoulos

To make both hardware and software solutions smarter and more secure, it is necessary to develop and put in place the management of trusted components and other secure technologies.

The interest is so big that industrial companies have started to define, develop and validate trusted hardware and firmware. This book describes, but it is not limited to, the results of the TOISE project funded by ENIAC framework. The purpose of this book is threefold. Firstly, to be used as an undergraduate or graduate level textbook as an introduction to trusted embedded systems, providing students and practicing engineers with the fundamentals as well as the details in the many facets of security problems. Secondly, to be used as reference for researchers in the field. Thirdly, to be used as a guide for professionals in analyzing and designing state-of-the-art trusted embedded systems. The book consists of three parts. The first part is an introductory chapter presenting information about existing solutions and programming interfaces in the trusted domain. The second part presents four real application use-cases, which depict the need to design trusted embedded systems. Finally, based on the second part, the third part presents the building blocks on which the presented application use-cases were built.

URL: http://www.springer.com/gp/book/9783319094199

ADRENALINE: A TOOL FOR BOOSTING OPENVX APPLICATIONS ON EMBEDDED MANY-CORE ACCELERATORS

As industry pushes towards advanced embedded vision devices, acceleration of computer vision kernels becomes a must. Adrenaline is a novel tool for hardware/software design space exploration of many-core accelerators, with a strong focus on embedded vision applications. Adrenaline consists of two main components, a heterogeneous SoC simulator and an OpenVX programming model, a cross-platform standard for imaging and vision. It models a SoC composed of a host processor and a many-core accelerator, providing several knobs to tune the accelerator architecture (hardware FPU, L1/L2 memory size, number of cores, DMA bandwidth/latency). On the software side we provide an OpenVX run-time with specific optimizations for the target many-core. An OpenVX graph can run partly on the host and partly on the accelerator, where an OpenVX graph node can be further parallelized among available cores. Our tool is intended to provide design exploration capabilities to a wide range of end users: platform architects can explore different hardware configurations for a target application; application developers can explore different partitioning schemes for applications; SDK vendors and researchers can explore various optimizations, scheduling policies and algorithms for the implementation of the OpenVX support layer. An FPGA-based version of the tool will be released soon, so stay tuned!

URL: http://www-micrel.deis.unibo.it/adrenaline/

Giuseppe Tagliavini, University of Bologna

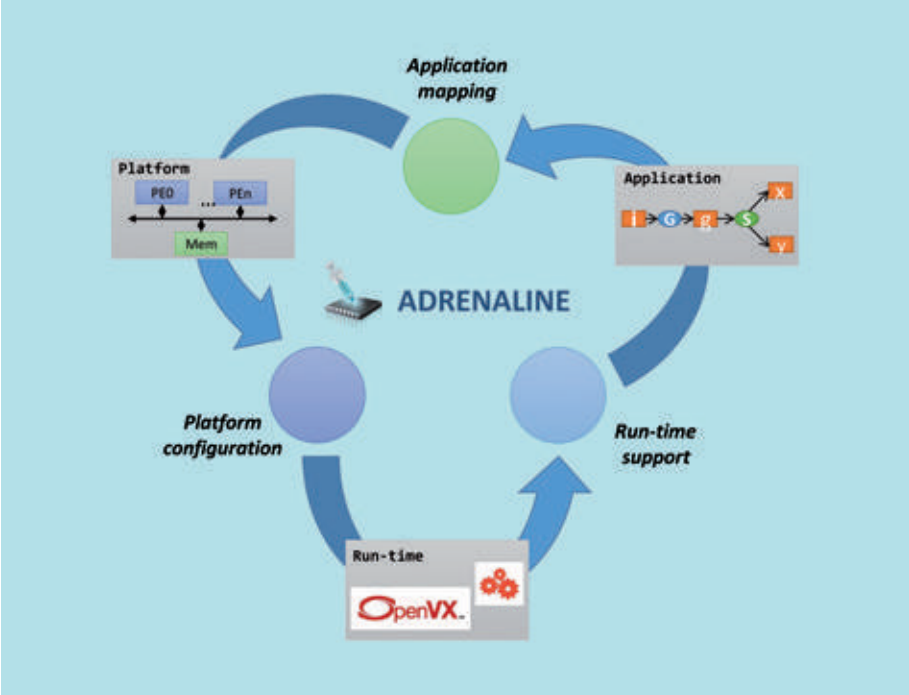

SILEXICA RECEIVES EMBEDDED INDUSTRY AWARD

SLX multicore programming solution honored by international expert jury at Embedded World 2015 as the most innovative tools product of the year.

The prestigious embedded AWARD 2015 was handed over on Tuesday, February 24, to Silexica´s managers by Prof. Dr.-Ing. Matthias Sturm, chairman of the selection committee, during a morning press tour through Embedded World, the world´s largest trade show in embedded system HW/SW technologies. The award acknowledges the high potential of the groundbreaking SLX multicore toolsuite for solving many of today´s and tomorrow´s multicore software development issues. “We are especially delighted to receive this important recognition among more than 500 international tools exhibitors here in Nuremberg”, said Maximilian Odendahl, CEO at Silexica. “Disruptive innovation is for sure a strength of technology start-ups, and our customers particularly appreciate the high software productivity gains delivered by the SLX tools. We are excited about fast growing market opportunities, given the strong trend towards multicore as an enabling technology, as is visible in so many products at Embedded World.” (Image by Uwe Niklas)

URL: http://www.silexica.com/

Rainer Leupers, Chief Scientist, Silexica, Germany

PARALLWARE: THE OPENMP-ENABLING SOURCE-TO-SOURCE COMPILER

As a new member of the ETP4HPC, Appentra Solutions is working to tackle the programmability issues of modern heterogeneous hardware platforms through the innovative & disruptive Parallware technology.

PRODUCT

Appentra solutions is pleased to announce that since march 2015 it is a member of the European Technology Platform For High Performance Computing (www.etp4hpc.eu/members/appentra). The ETP4HPC provides a framework for stakeholders, led by industry, to define research priorities and action plans on the HPC strategy of the EC.

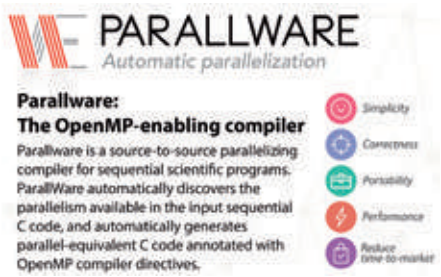

Parallware is the first commercial OpenMP-enabling source-to-source compiler to succeed in automatically adding OpenMP capabilities to numerical simulations programs. Boosted by the worldwide exclusive Parallare technology, the compiler discovers the parallelism available in sequential programs written in the C programming language, and automatically adds OpenMP 3 pragmas targeting multicore systems. For more details on the technology and the compiler consult www.appentra.com.

During 2015, Parallware will be presented in the following prestigious international events in the high-performance technical computing (HPTC) market:

1. 2015 Rice Oil & Gas HPC Workshop (http://www.appentra.com/democratization-hpc-oilgas-industry-automatic-parallelization-parallware/)
2. Finalist in NVIDIA's Early Stage Challenge at GTC'2015 (http://www.appentra.com/appentra-finalist-nvidias-early-stage-challenge/)
3. PRACE Scientific and Industrial Conference 2015 (PraceDays15), 26-28 May in Dublin.
4. OpenMPCon 2015 conference, 28-30 September in Aachen.

Manuel Arenaz, CEO, Appentra, Spain

THE UNIVERSITY OF CANTABRIA GRANTS DOCTOR HONORIS CAUSA TO MATEO VALERO

On March 19th the University of Cantabria granted an Honorary Doctorate to the director of BSC-CNS, Mateo Valero, who was nominated by the Department of Electronic and Computer Engineering and the Faculty of Science for his "extraordinary contributions to the field of engineering and computer architecture, in which he has played a pioneering role at both national and international levels. It will also recognise his creation and fostering of supercomputer research centres, being director of the Barcelona Supercomputing Centre since its inception."

Ramon Beivide (Professor of Architecture and Computer Technology) had the honour of delivering the laudation of Mateo Valero and he extolled the "energy and time" Valero invested to achieve his technical objectives, which has led to his personal projection and that of the group he leads in Barcelona. But he believes that his success is due not only to his "dedication" but also his personal traits such as his "bonhomie and good nature, which he treasures".

Mateo Valero used his speech to assert that "if an advanced country wants to generate good ideas, it needs a proper ecosystem". He also noted that "if an environment is to grow rich in opportunities, then governments, companies and researchers need to collaborate, moving in the same direction, so that the sum is greater than the parts. If just one of these actors fails, then the end result will be unsatisfactory."

Barcelona Supercomputing Center

BSC RESEARCHER DAVID CARRERA SELECTED FOR ERC STARTING GRANT

Barcelona Supercomputing Center (BSC) researcher David Carrera has been awarded one of the most prestigious grants offered by the European Commission to study how data centres can learn how to become more energy efficient independently. It has been estimated that servers around the world currently consume as much energy as the whole of Italy; as such, their efficiency is one of the main problems facing computer science today.

David Carrera, who holds a PhD in Computer Science, has been awarded the grant for his project Hi-EST: Holistic Integration of Emerging Supercomputing Technologies. In this project, the BSC researcher and Universitat Politècnica de Catalunya-BarcelonaTech associate professor will set out to ensure that services use the information which they themselves generate about their activity to learn how to autoconfigure their settings and thereby achieve maximum efficiency while consuming minimal energy. This process is to be carried out using artificial-intelligence techniques which help the servers decide which types of hardware should be used for each task and where data should be saved, among other parameters which improve their efficiency. Carrera's research interests are focused on the performance management of data-centre workloads, the topic of his doctoral thesis. He has carried out work for IBM and Microsoft and leads various projects relating to server configuration and big-data processing at BSC.

Barcelona Supercomputing Center

BEST PAPER PRESENTATION AWARD HIPEAC 2015

Nvidia, gold sponsor of the HiPEAC Conference, offered development boards to the HiPEAC2015 Best Paper Presentation Award winners. The following papers were selected:

- Efficient Data Mapping and Buffering Techniques for Multilevel Cell Phase-Change Memories, presented by Justin Meza
- Topology-Aware and Dependence-Aware Scheduling and Memory Allocation for Task-Parallel Languages, presented by Andi Drebes

Congratulations to the presenters and the authors. They will receive a Jetson TK1 from Nvidia!

Jetson TK1 is Nvidia's embedded Linux development platform featuring a Tegra K1 SoC (CPU+GPU+ISP in a single chip). Jetson TK1 comes pre-installed with Linux4Tegra OS (basically Ubuntu 14.04 with pre-configured drivers). There is also some official support for running other distributions using the mainline kernel.

HIPEAC TECHNOLOGY TRANSFER AWARDS 2014

HiPEAC not only wants to encourage scientific excellence through publications in the HiPEAC award conferences, but it also wants to stimulate uptake of research results by industry. To reward and celebrate the transfer of research results into industry (be it through technology licensing, by providing dedicated services to an existing company, or through the creation of a new company), HiPEAC has created the so-called technology transfer awards.

The following members received a HiPEAC Technology Transfer Award in December 2014.

Grigori Fursin	Transferring cTuning technology to ARM to systematize benchmarking and combine it with predictive analytics
Christian Weis and Norbert Wehn	LTE-IP --- A Turbo Decoder IP for the next generation of mobile devices (5G)
Olav Lysne and Sven-Arne Reinemo	Fabriscale - Scalable and reliable fabric management
Ben Juurlink	H.265/HEVC on Embedded GPUs
Andrea Bartolini, Christian Pinto and Luca Benini	Energy-Optimal workload allocation for microcontroller platforms based on heterogeneous multicores

Congratulations!

BSC CELEBRATES ITS 10TH ANNIVERSARY

BSC was created on April 1 2005, catapulting Spain into the elite of international supercomputing, the result of an already long existing collaboration between the UPC, the Spanish Ministry and the Generalitat. This had begun with the creation of CEPBA, the European Center of Parallelism of Barcelona, in 1991, and continued with the CEPBA-IBM Research Institute founded in 2000.

In our ten years of existence BSC has achieved much more than could reasonably been expected back then. We have enabled more than 3,000 scientific projects to be run on the MareNostrum; we have driven the design of cutting-edge hardware and software technologies; we have helped to create the Spanish Supercomputing Network (RES) and the pan-European distributed supercomputing infrastructure, PRACE; we have won funding for more than 90 EU projects; we have created long-term joint research centres with some of the best IT companies in the world (IBM, Microsoft, Intel and NVIDIA); we have helped Repsol to become a global leader in its industry; we have worked with Iberdrola to create greener energy; we have achieved the Severo Ochoa recognition; we have helped to create the ETP and the Public Private Partnership for HPC; we have published hundreds of papers; we have supervised 140 PhD thesis; we have closely collaborated internationally especially with Latin America and particularly through the creation of RISC, the Ibero-American Supercomputing Network; and we have grown to become an internationally respected centre with more than 420 people.

During the past ten years almost 300 people have passed through the BSC and have used the centre as a springboard from which to launch themselves into other excellent organisations. With the launch of the BSC Folks platform, we want to make sure that we keep in touch with these people and continue the mutually beneficial relationship with them. In the past five years, during a time of grave economic crisis, BSC has created 150 jobs and our ambition is to further strengthen our role as a generator of wealth and wellbeing.

We realise that our achievements have only been possible though the bright ideas, rigour, hard work and dedication of a huge number of people. We also realise that not everything is perfect, and that as directors, we need to continuously improve. We can do that using input from different sources including the "comité de empresa", the scientific advisory board, and through contact and conversation with all of you, and we believe that the implementation of the HRS4R action plan will be an important step in this direction.

The 10th anniversary event on April 1 was an opportunity to celebrate our achievements, and to look forward to the challenge of continuing to position BSC as an essential facilitator and practitioner of excellent data intensive science and innovation. Above all, it was a chance to get together and to enjoy ourselves.

The quote "the future is not what it used to be" now seems truer than ever, and the pace of technological and social change is often so fast that we feel overwhelmed. BSC's task over the next ten years will be to provide science and society with the necessary tools to cope with this uncertain future. Judging from the results of the past ten years, we have every confidence that we will continue to be an example of how collaboration between scientists, public institutions and companies can produce both excellent science and a positive economic impact.

Mateo Valero, BSC Director and Francesc Subirada, BSC Associate Director

FIRST PRIZE IN THE ICFPT 2014 DESIGN CONTEST

The team from Universidad de Zaragoza, Spain, composed of Javier Olivito, Alberto Delmas and Javier Resano, has won the ICFPT 2014 Design Contest held in Shanghai, China

The International Conference on Field-Programmable Technology (ICFPT, www.icfpt.org) each year organizes a design competition for programmable logic (typically FPGAs). The goal of this event is to demonstrate that programmable logic can be used to develop custom processors that efficiently deal with complex computational problems.

In recent years the organizers have selected different board games as targets for this competition. Board games are interesting problems since they involve complex tasks such as data processing, where a strong player must process thousands or even millions of boards per second; pattern recognition to identify important game situations; design space exploration to evaluate the consequences of each possible move in the future and select the best one.

The target of this year was a board game called Blokus Duo, which is a popular game that has received several international awards. Basically Blokus Duo is a two-player game played on a square 14x14 grid board. Each player has 21 different-shaped blocks, and can place them with eight different rotations. The objective is to occupy as many squares with your blocks as possible while trying to reduce the number of squares of your opponent. To this end players must follow some simple rules. The Figure presents a game won by the blue player, who has placed 68 squares while green has only 66. Although the rules are simple, this is in fact a complex game for two main reasons. On the one hand, it is difficult to know whether a player is winning or not since the score of a game can change a lot at the end. On the other hand, the game tree to explore is huge since in a given board there may be thousands of possible moves to analyze.

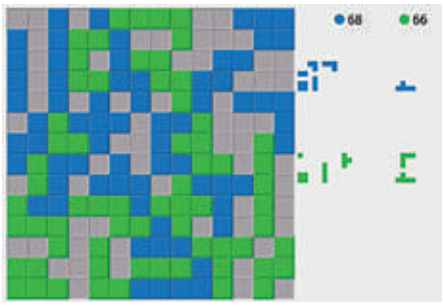

This year, fourteen teams from different universities and research centers participated in the competition in Shanghai, China (Dec. 10-12, 2014). Each competitor presented a custom design implemented in an FPGA. The best designs were selected and the finalists played among them in one of the sessions of the ICFPT 2014 conference.

In the finals the winner was the design from the Computer Architecture Group (GaZ) from Universidad de Zaragoza. The team was composed of Alberto Delmás (last-year student on Computer Science engineering), Javier Olivito (PhD student) and Javier Resano (Associate Professor of the Computer Science Department and member of the I3A research institute).

Javier Resano, Javier Olivito and Alberto Delmás, Universidad de Zaragoza, Spain

SCALABLE RESOURCE-EFFICIENT SYSTEMS FOR BIG DATA ANALYTICS

Data is generated at an ever-increasing rate, and added value and cost savings can be obtained by analyzing big data streams. The analysis of large data sets requires scalable, high-performance, and resource-efficient computer systems. The research project "Scalable resource-efficient systems for big data analytics" at Blekinge Institute of Technology combines existing expertise in machine learning, data mining, and computer engineering to create new knowledge in the area. The value will be demonstrated and evaluated in two application areas, i.e., decision support systems and image processing.

The needs and interests of our eleven industrial partners are grouped into six industrial challenges. Based on these challenges, and in cooperation with our partners, we have defined seven initial sub-projects grouped into four research themes:

- Theme A: Big data analytics for decision support
- Theme B: Big data analytics for image processing
- Theme C: Core technologies (data science)
- Theme D: Foundations and enabling technologies

The project is financed by the Knowledge Foundation in Sweden with approx. EUR 4 million during 2014-2020. 11 companies participate: Arkiv Digital AD, Compuverde, Contribe, Ericsson, Indigo IPEX, Medical Management Innovation, Noda Intelligent Systems, Scorett Footware, Telenor Sweden, Sony Mobile Communications, and Wireless Maingate Nordic.

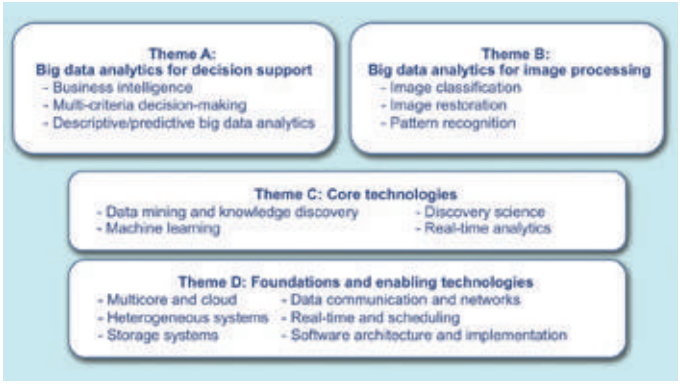

URL: http://www.bth.se/bigdata/

Håkan Grahn, Blekinge Institute of Technology

THE INTO-CPS H2020 PROJECT

Integrated Tool Chain for Model-based Design of Cyber-Physical Systems

Consortium Members:
Aarhus University (DE)
Newcastle University (UK)
University of York (UK)
Linköping University (SE)
Verified Systems International (DE)
Controllab Products (NL)
ClearSy (FR)
TWT GmbH - Science & Innovation (DE)
Kongskilde Industries (DK)
United Technologies (IR)
Softeam (FR)

Project coordinator:
Aarhus University (DK)

URL: http://into-cps.au.dk/

The aim of the INTO-CPS project is to create an integrated "tool chain" for comprehensive Model-Based Design (MBD) of Cyber-Physical Systems (CPSs). The tool chain will support the multidisciplinary, collaborative modelling of CPSs from requirements, through design, down to realisation in hardware and software. This will enable traceability at all stages of development. INTO-CPS will support the holistic modelling of CPSs, allowing system models to be built and analysed that would not have been possible using standalone tools. We will integrate existing industrial-strength tools with high Technology Readiness Levels (TRL 6–9) in their application domains, based around Functional Mockup Interface (FMI)-compatible co-simulation. The project focuses on the pragmatic integration of these tools, making extensions in areas where a need has been recognised. The tool chain will be underpinned by well-founded semantic foundations that ensure the results of analysis can be trusted. The tool chain will provide powerful analysis techniques for CPSs, including connection to SysML; generation and static checking of FMI interfaces; model checking; Hardware-in-the-Loop (HiL) and Software-in-the-Loop (SiL) simulation, supported by code generation. The tool chain will allow for both Test Automation (TA) and Design Space Exploration (DSE) of CPSs.

Remarks: The INTO-CPS technologies will be accompanied by a comprehensive set of method guidelines that describe how to adopt the INTO-CPS approach, lowering entry barriers for CPS development.

Alessandra Bagnato, SOFTEAM - Research and Development

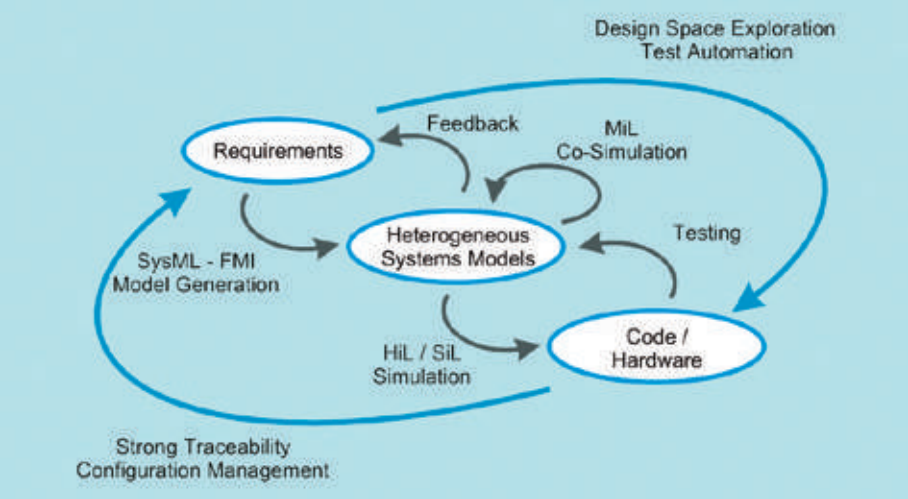

EU PROJECT AUTOTUNE STARTS DEMONSTRATION CENTRE

Consortium Partners:
Technische Universität München (DE)
Leibniz Supercomputing Centre (DE)
Universitat Autònoma de Barcelona (ES)
University of Galway (IE)
University of Vienna (AU)
IBM (DE)

Project coordinator:
Technische Universität München (DE)

The European Commission (EC) funded project in the field of high-performance computing (HPC) prolongs its optimization services after the official project end in April 2015, with a dedicated centre hosted at Leibniz Supercomputing Centre in Garching (Germany).

The European consortium forming the "Automatic Online Tuning" (AutoTune) project (FP7 project no. 288038) announces the opening of the AutoTune Demonstration Centre (ADC) after the end of the project in April 2015 (www.autotune-project.eu). The purpose of the centre is to bring together users, developers, vendors, and supercomputing experts and educate them on the benefits of the automatic online optimisation of scientific codes as well as plugins and methods developed within the AutoTune project. Thereby, the consortium members want to exploit the project results and make them available to a larger interested HPC audience. This shall enable scientists and commercial HPC users alike to make better usage of supercomputing resources, regarding not only the compute performance but also energy efficiency. The ADC will be hosted at Leibniz Supercomputing Centre in Garching, near Munich (Germany) and ICHEC (Ireland), and will be open to users working with the consortium. Users will also have access to online documentation, best practice guides, discussion forums and individual support, as well as training sessions.

Additionally, an in-depth documentation of the AutoTune project is provided in the book "Automatic Tuning of HPC Applications - The Periscope Tuning Framework

(PTF)" edited and published by Michael Gerndt (TU Munich), Eduardo César (Universitat Autònoma de Barcelona) and Siegfried Benkner (University of Vienna). PTF (periscope.in.tum.de) was developed within the AutoTune project. The goal is to close the gaps in the application tuning process and thus to simplify development of efficient parallel programs on a wide range of architectures. The framework is unique, since it combines analysis and tuning of multiple aspects into an online automatic tuning framework. The book will be published in April and is available via Shaker Verlag. It will go into detail on different tuning concepts as well as on various tuning plugins like for instance MPI Parameters Plugin or a Parallelism Capping Plugin.

AutoTune is a European Commission (EC) funded research project made up of an international consortium of scientific institutions and industrial companies coordinated by Technische Universität München (TUM). The project started mid October 2011 and ends in April 2015. It disposed of an overall funding of 2.3 MEuro and pursued the target of automatically optimizing applications in the area of high-performance computing. Next to TUM, the Leibniz Supercomputing Centre (LRZ) of the Bavarian Academy of Sciences, the Universitat Autònoma de Barcelona (UAB), the Centre for High-End Computing (ICHEC) at the University of Galway as well as the University of Vienna are partners in the project, with IBM as associated partner.

Michael Gerndt, Technische Universität München

EU-FUNDED RESEARCH TO CREATE SECURE-BY-DESIGN ARCHITECTURES

Consortium Members:
Foundation for Research and Technology - Hellas (GR)
Vrije Universiteit Amsterdam (NL)
Chalmers University of Technology (SE)
Technische Universität Braunschweig (DE) Neurasmus BV (NL)
OnApp Limited (UK)
IBM Haifa Research Lab (IL)
Elektrobit Automotive GmbH (DE)

Project coordinator:
Foundation for Research and Technology - Hellas (GR)

Project URL:
http://www.sharcs-project.eu/

The Horizon 2020 SHARCS project aims at designing, building and demonstrating secure-by-design system architectures that deliver end-to-end security for their users. The new technologies developed will be directly usable in applications and services that require end-to-end security.

Tremendous technological achievements such as Medical IoT, Smart Cars, Smarter Cities, Smarter Grids etc. have led society as a whole and individual citizens to rely ever more on critical systems which sense and control the physical environment. Such "Cyber Physical Systems" (CPS) use a blend of embedded devices and traditional computing systems, and a variety of communication channels. The adoption of CPS necessitates revisiting of the security stack to ensure that the new generation of devices and services encompass the lessons learned as part of the ICT cybersecurity battles of the last decades.

To address the pervasive security problems, the SHARCS project will push security mechanisms down the system stack, from software to hardware. Hardware security mechanisms have the advantage of being hard, if not impossible, to be bypassed by attackers. The primary reason for this is that hardware is not typically modifiable as is the case with software. Also, hardware security mechanisms, in most cases, have the advantage of being more efficient in terms of performance, simplicity, and power usage. SHARCS will also implement a series of bottom-up design and implementation concepts to explore the entire system stack, and the interactions between the various components. Some of the challenges that will be addressed by SHARCS include: how to utilize the new functionality, how maintain backwards compatibility with legacy applications, and how to handle the interactions between different administrative domains (e.g., one that has high security requirements and one with low security expectations). SHARCS also investigates how to dynamically recover from errors and how to minimize the burden on software developers and users.

It received EU funding of € 3.1 million over three years (1/1/2015 to 31/12/2018).

The main objectives of the SHARCS project are the following:

1. Extend existing hardware and software platforms towards developing secure-by-design enabling technologies.
2. Leverage hardware technology features present in today's pro cessors and embedded devices to facilitate software-layer security.
3. Build methods and tools for providing maximum possible security-by-design guarantees for legacy systems.
4. Evaluate acceptance, effectiveness and platform independence of SHARCS technologies and processes.
5. Create high impact in the security and trustworthiness of ICT systems by:
   - Producing new technologies and tools for security-by-design;
   - Bringing the SHARCS outcomes to market via partnerships with industry;
   - Developing recommendations for EU regulators for a minimal set of requirements (hardware and software) for devices used in critical applications;
   - Forming a standardization proposal for security-by-design technologies;
   - Communicating and disseminating SHARCS results broadly and effectively to industry, academia and the general public.

Vassilis Prevelakis, Technische Universität Braunschweig

EC FP7 STREP HARPA PROJECT

Harnessing Performance Variability

Project Coordinator:
Prof. William Fornaciari, Politecnico di Milano
william.fornaciari@polimi.it
Partners:
Politecnico di Milano (IT), IMEC (BE), ICCS/NTUA (GR), UCY (CY), IT4I (CZ), THALES (FR), HENESIS (IT)
Project website:
http://www.harpa-project.eu

The goal of HARPA is to enable next-generation embedded and high-performance heterogeneous many-core processors to cost-effectively confront variations and yet provide Dependable-Performance: correct functionality and timing guarantees throughout the expected lifetime of a platform under thermal, power, and energy constraints. This will be done by utilizing both proactive techniques (in the absence of hard failures) and reactive techniques (in the presence of hard failures).

The HARPA project will deliver a comprehensive framework, along with its associated methodology and supporting tools, that ensures performance dependability in the many-core systems of the near future. Specifically, the HARPA Engine – the core deliverable of this project – will have direct applicability in existing systems, and it will pave the way for more extensive architectural support in future SoC platforms. The components to be developed in this project will be validated using various important use-cases, which have been selected due to their significance in emerging application trends.

The HARPA project aims to address several scientific challenges in this direction: (i) develop performance dependable heterogeneous multi-core architectures applicable to both embedded system environments and HPC setups; (ii) develop monitors and knobs in hardware and software designs; (iii) develop system software designs that support high-performance dependability requirements; (iv) develop run-time designs that support high-performance dependability requirements; (v) develop and analyze low level methodologies to tradeoff conflicting metrics, i.e. power, performance, reliability and thermal constraints; (vi) develop software/hardware interfaces to provide a fluent communication flow; and (vii) design new application design guidelines to improve performance dependability.

Prof. William Fornaciari
william.fornaciari@polimi.it

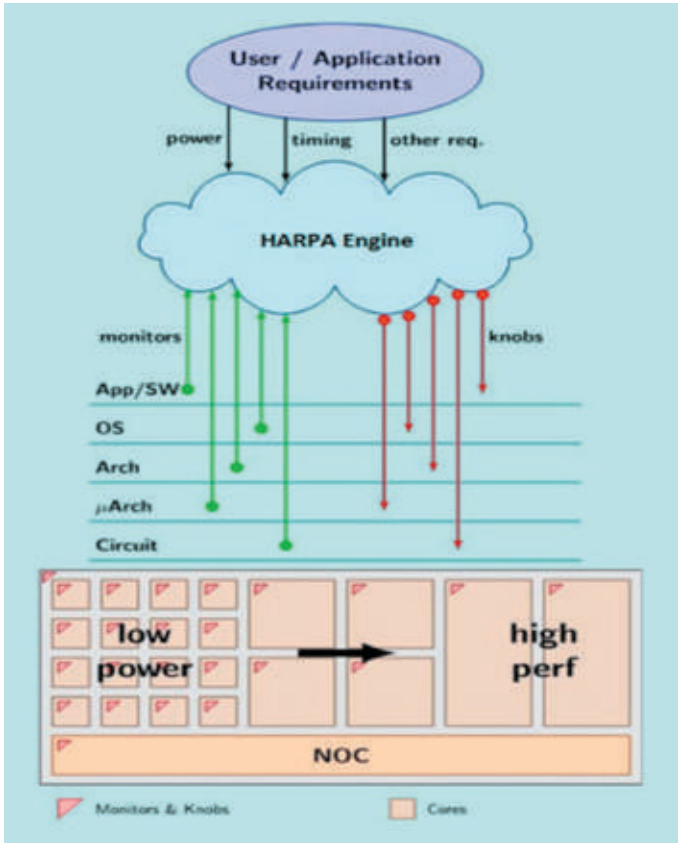

Technical Approach: HARPA Engine Overview

Figure 1 shows the main concepts of the HARPA architecture and the key modules of an architecture that can provide performance dependability guarantees. Note that this generic framework applies to both embedded systems and high-performance genewral-purpose systems. The main elements that distinguish a HARPA-enabled system are: (i) monitors and knobs, (ii) user requirements, and (iii) the HARPA Engine. Conceptually, the HARPA Engine mainly consists of a feedback loop, where the different metrics (performance, timing, power, temperature, errors and manifestations of time-dependent variations etc.) of the system are continuously monitored. The HARPA engine actuates the knobs to bias the execution flow as desired, based on the state of the system and the performance (timing/throughput) requirements of the application.

The concepts that are to be developed within the HARPA context address equally both the HPC and ES domains. More specifically, from the HPC domain we will use Disaster and Flood Management Simulation, while from the ES domain, we will use Radio frequency spectrum sensing application, Face Detection Application, Object Recognition and Beesper Landslide Multimodal Monitoring. In particular, HARPA use cases will be demonstrated in two HPC platforms: (i) Intel Xeon, and (ii) x86-64 multi-core plus a GPU, as well as two embedded platforms: (i) Freescale i.MX 6Quad, and (ii) ODROID XU-3 (Octa Core Linux Computer Samsung Exynos5422 Cortex-A15 2.0Ghz quad core and Cortex-A7 quad core)

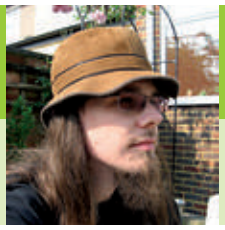

INTERNSHIP REPORT: ERIK TOMUSK

Host: ARM Research and Development (Cambridge, UK)
Title: "Prototype Hardware Accelerators and Demonstration of the Benefits of the ACAI Framework"

Through a HiPEAC internship, I was able to spend four months at ARM Research & Development in Cambridge, UK. The internship revolved around an ongoing hardware accelerator interface research project. While I was at ARM, I was the primary person working on the project, but I worked closely with the project's owner and other engineers involved with the project.

The internship involved working with the entire hardware-software stack. My tasks included making changes to the hardware and synthesizing to FPGA, debugging cache coherency and virtual memory, patching and building the Linux kernel, updating and writing device drivers, updating and writing user-space libraries, and developing software test cases. A primary focus of the work was debugging and improving both functional and performance aspects of the system. I spent a significant amount of time understanding the distributed virtual memory behavior of the system. In some cases, idiosyncrasies of the Linux kernel's memory management would cause tests to fail. In other cases, correctly configured hardware would significantly improve performance. I also developed test automation infrastructure to lower the overhead of testing the system, and I improved software support for test case development. Finally, I worked with several bus protocols and implemented a bus interface in hardware as part of integrating a new accelerator with the system. For ARM, the internship led to a better understanding of performance considerations and usage constraints of the accelerator interface. At the end of the internship, I documented the lessons learned and presented them to the R&D group.

I would like to thank the team at ARM for a positive working environment and a very interesting project, and HiPEAC for making the internship possible.

Erik Tomusk, University of Edinburgh

INTERNSHIP REPORT: CLÁUDIO MAIA

Host Institution: Thales Research and Technology (Paris, France)
Title: Safety Critical Runtime for Multicore Embedded Systems

This HiPEAC Industrial PhD Internship was conducted at Thales Research and Technology (Paris, France), from August 2014 until December 2014, under the supervision of Daniel Gracia Pérez. The internship had as its main objective the research of alternative models to maximize the performance of safety-critical solutions executing on multiprocessor platforms in the avionics domain. The avionics domain is constantly developing state-of-the-art solutions that need to take into account the demanding performance requirements of new applications, or requirements arising from the integration of several applications that need to coexist in the same hardware platform. With this increasing demand for performance, uniprocessor platforms are no longer the best candidates as target platforms for new solutions (due to the physical limitations of processor chips), and have been naturally replaced by more powerful multiprocessor platforms. Nevertheless, this move towards the use of multiprocessor solutions brings along new challenges and problems.

One of the most challenging problems specially arising in safety-critical systems - systems in which a failure of the system under control may lead to catastrophic consequences and therefore determined unacceptable by design - is to find ways of minimizing interference between different applications, possibly those executing at different criticality levels, without affecting their performance. The most commonly adopted solution in the avionics domain is to isolate applications, in terms of either time or space, as a means to obtain zero interference between them, and therefore achieve a higher level of predictability. However, such an approach leads to under-utilized platforms due to the worst-case execution time estimation guarantees that need to be provided. Maximizing parallelism is important for avionic systems, not only because it allows one to execute more complex applications but also because at the same time it increases the utilization of the platform itself. Nevertheless, parallelism comes at the cost of decreasing the system's predictability.

Under the internship two solutions were proposed to increase the utilization when using parallel platforms. First, a model was proposed to take advantage of the idle time left available during runtime, so that offline guarantees can still be met. A simulator was partially built in order to test the feasibility of the proposed approach. The second solution is an offline framework (on which work is still on-going) that allows one to analyse task behaviour with respect to its schedulability.

I would like to thank HiPEAC and Thales Research and Technology for the opportunity of doing this internship, as it was a very valuable experience at all levels. Moreover, I would like to thank all the people that contributed to my integration in the company, as well as those that shared their knowledge.

Cláudio Maia, School of Engineering – Polytechnic of Porto

INTERNSHIP REPORT: MARINA ZAPATER

Host Institution: Performance and Energy-Aware Computing Laboratory (PeacLab) at Boston University
Title: Workload and cooling aware optimization techniques to reduce the energy consumption in Data Centers

Energy efficiency in Data Centers has been and continues to be an important research challenge. Data Centers are easily found in every sector of the worldwide economy and their power consumption has exhibited an unsustainable increase in recent years, mainly due to data globalization and computer ubiquity.

The focus of my research is the development of thermal-aware optimization policies to reduce the energy consumption and environmental impact of data centers, contributing to placing them on a more scalable curve. I am currently a PhD candidate at Universidad Politécnica de Madrid (UPM), and my research is in collaboration with Universidad Complutense de Madrid (UCM), in Spain, where I am also a part-time lecturer. On Summer 2012 I visited Prof. Ayse K. Coskun at the Performance and Energy-Aware Computing Laboratory (PeacLab) at Boston University, where we collaborated on the development of leakage and cooling-aware energy optimization policies at the server level. On Fall 2014, thanks to a HiPEAC collaboration grant I was able to visit the PeacLab group again for a three-month period. During my stay at the PeacLab group we have been able to extend our previous work on the server level to the Data Center scope. In particular, we have examined the trade-offs in terms of power for state-of-the-art highly efficient data center scenarios. To create and calibrate our models, the PeacLab group provided access to data from the "Massachusetts Green High Performance Computing Center" (MGHPCC) a high-density cooling-optimized data center that supplies research computing to Boston University, Harvard, MIT, Northeastern University, and UMass. Our work relies on these models to propose workload allocation and cooling management strategies to further reduce the overall energy costs of these facilities.

This HiPEAC collaboration grant has been a great opportunity to strengthen the bonds between the three institutions involved. As a PhD student, visiting a research group that aims to deliver such high-quality research, and that has so many collaborations, is a very rewarding experience. Research at the PeacLab group is strongly aligned with my work and I believe that this project has brought new perspectives to the development of my PhD thesis. I am very grateful to HiPEAC for granting me this opportunity, and I am certain that the collaboration between UPM, UCM and BU will keep on going in the future. I would like to thank my advisors, Prof. José Luis Ayala, and Prof. José Manuel Moya for making this collaboration possible, my host Prof. Ayse K. Coskun for her time and guidance, as well as all the members of the PeacLab group at Boston University.

Marina Zapater, Universidad Politécnica de Madrid & Universidad Complutense de Madrid

INTERNSHIP REPORT: PATRICIA ARROBA GARCÍA

Host Institution: The Cloud Computing and Distributed Systems (CLOUDS) Laboratory at The University of Melbourne - Title: Energy-Aware Automatic Optimization of Resource Allocation in CloudSim

Computational demand on data centers is increasing due to growing popularity of Cloud applications. Nowadays, this industry consumes about 2% of the worldwide energy production and the carbon footprint generated by cooling systems is expected to overtake airline industry emissions by 2020.

Furthermore, as Cloud applications expect services to be delivered as per Service Level Agreement (SLA), power consumption in data centers has to be minimized while meeting this requirement whenever it is feasible. Also, Cloud workloads vary significantly over time, making optimal allocation and DVFS configuration not a trivial task. Thus, Cloud providers need to implement an energy-efficient management of physical resources to meet the growing demand of their services while ensuring Quality of Service (QoS). Understanding the relationship between power, cooling, DVFS and workload consolidation is crucial to enable energy-efficient management at the data center level.

My research, at Universidad Politécnica de Madrid (UPM) in tight collaboration with Universidad Complutense de Madrid (UCM), proposes new trade-offs between energy, workload consolidation and performance that help on combining DVFS with power and thermal-aware strategies. In this scope, the key contribution of my internship at CLOUDS Lab focuses on addressing the energy challenge in Cloud data centers from a proactive perspective. We propose a consolidation policy that jointly minimizes consumption while maintaining QoS. We have performed an extensive evaluation on the CloudSim toolkit, a state-of-the-art simulator developed at CLOUDS Lab, using real Cloud traces and an accurate power model based on data gathered from real servers. Our results demonstrate that combining DVFS awareness with workload management provides substantial energy savings around 37.86% for scenarios under dynamic workload conditions.

I would like to thank HiPEAC for giving me the excellent opportunity to collaborate with a group that is clearly an international reference in Cloud computing optimizations. This internship has allowed us to start a relationship between UPM, UCM and CLOUDS Lab to perform future research in collaboration. I also need to thank members of the CLOUDS Lab at The University of Melbourne for their comments and support, especially Prof. Rajkumar Buyya for his supervision and guidance.

Patricia Arroba García,
Universidad Politécnica de Madrid

EVALUATION OF HIGH PERFORMANCE COMPUTING PLATFORMS FOR DRUG DISCOVERY

Ginés Guerrero
Advisors: José M. Cecilia and José M. García
Institution: Universidad de Murcia
Graduation Date: June 2014

In my dissertation we evaluate the current landscape of computation using as a case study a high-impact problem for society, which is virtual screening. The study has covered all processing levels, starting with an extensive analysis of the various available alternatives at the chip level, through their evaluation in a cluster environment, to scale to cloud computing and volunteer computing levels.

This study concludes that GPUs are at the leading edge for the development of scientific applications with massively parallel computing patterns and high computational demands. However, migration to GPUs may require an application redesign and even rethinking, but this is actually part of computational thinking, which is essential to the development of scientific applications.

Alternatives to the use of clusters need to be evaluated, such as the cloud computing and volunteer computing for a larger-scale executions. Cloud computing can be an interesting option if the computation you wish to perform is executed periodically, as non-use of local resources implies that the economic investment is not justified. On the other hand, the option of using a volunteer computing platform looks interesting for developing certain HPC applications, as it offers a huge amount of hardware resource at no cost.

ENERGY CHARACTERIZATION METHODOLOGIES FOR SMT/CMP PROCESSOR SYSTEMS

Ramon Bertran
Advisors: Marc Gonzalez Tallada, Nacho Navarro Mas
Institution: Universitat Politècnica de Catalunya
Graduation Date: December 2014

Computer systems performance and affordability is limited by energy and power consumption. Power density limits the operational frequency of computer systems, affecting their performance. Similarly, the energy consumption of computer systems limits the operational time of battery-powered devices and increases the operational costs of data-centers. These energy-related limitations, known as the power wall, motivate researchers to look for new solutions to overcome them.

We must understand in detail the power consumption of computer systems in order to propose solutions that mitigate these energy-related limitations. Given the complexity of today's systems, systematic energy characterization methodologies are required to understand energy-related issues on today's computer systems.

In this thesis, we present various systematic energy characterization methods as well as the software frameworks required to implement them. This includes a systematic method for producing counter-based power models; an energy accounting method on shared virtualized systems using counter-based power models; a systematic instruction-wise power/performance profile generation method; a systematic maximum power stressmark generation method; and Microprobe, a microbenchmark generation framework.

MULTIPROCESSING SYSTEMS DEVELOPMENT FOR IMPLEMENTATIONS OF APPLICATIONS

Calliope-Louisa Sotiropoulou
Advisor: Assoc. Professor Spiros Nikolaidis
Institution: Aristotle University of Thessaloniki
Graduation Date: March 2015

The main objective of this thesis is to identify the optimum heterogeneous multiprocessing design architecture for a given application. This objective is tackled though two different approaches: a) formulation of a theoretical model that can be used to optimize each system according to the application's specification and the platform's characteristics, and b) direct approach through finding the optimum architecture for two characteristic and high performance applications with hard real-time requirements. The developed theoretical model formulation is based on Integer Linear Programming and it can be used to identify the optimum architecture and task assignment for a given application and a specific platform. The first application used to implement a working example of an optimized MPSoC system is a real-time machine vision system for real-time flow detection on microfluidic Lab-on-Chips. The system is designed to follow a 60 fps camera with 1 Mpixel resolution. The second application used is a high performance 2D pixel clustering implementation for streaming data. The first application used to implement a working example of an optimized MPSoC system is a high performance machine vision system for real-time flow detection on microfluidic Lab-on-Chips.

UPCOMING EVENTS

ACM International Conference on Computing Frontiers 2015 (CF15)
18-21 May 2015, Ischia, Italy
http://www.computingfrontiers.org/2015/

20th International Conference on Reliable Software Technologies (Ada-Europe 2015)
22-26 June 2015, Madrid, Spain
https://ae2015.dit.upm.es

International Conference on Embedded Computer Systems: Architectures, Modeling and Simulation (SAMOS XV)
20-23 July 2015, Samos Island, Greece
http://samos-conference.com/

The 2015 International Conference on High Performance Computing & Simulation (HPCS 2015)
20-24 July 2015, Amsterdam, the Netherlands
http://hpcs2015.cisedu.info/

International Symposium on Low Power Electronics and Design (ISLPED'15)
22-24 July 2015, Rome, Italy
http://www.islped.org/2015/index.html

Euro-Par 2015
24-28 August 2015, Vienna, Austria
http://www.europar2015.org/

22nd European conference on circuit theory and design (ECCTD2015)
24-26 August 2015, Trondheim, Norway
http://www.ntnu.edu/ecctd2015/

ParCo2015
1-4 September 2015, Edinburgh, UK
http://www.parco.org/

International Conference on Field-programmable Logic and Applications (FPL 2015)
2-4 September 2015, London, UK
http://www.fpl2015.org/

IEEE 9th International Symposium on Embedded Multicore/Many-core Systems-on-Chip (MCSoC-15)
23-25 September 2015, Turin, Italy
http://mcsoc-forum.org/2015/

2015 IEEE Nordic Circuits and Systems Conference (NORCAS)
26-28 October 2015, Oslo, Norway
http://www.norcas.org/

contributions If you are a HiPEAC member and would like to contribute to future HiPEAC newsletters, please visit https://www.hipeac.org/publications/newsletter/

HiPEAC info 42
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: https://www.hipeac.org/
subscriptions: https://www.hipeac.org/publications/newsletter/

design: www.magelaan.be