HiPEAC info 29

COMPILATION ARCHITECTURE

APPEARS QUARTERLY
JANUARY 2012

NETWORK OF EXCELLENCE ON HIGH PERFORMANCE AND EMBEDDED ARCHITECTURE AND COMPILATION

WELCOME TO THE HIPEAC 2012 CONFERENCE, PARIS, FRANCE, JANUARY 23-25, 2012

HIPEAC3 KICKS OFF IN PARIS

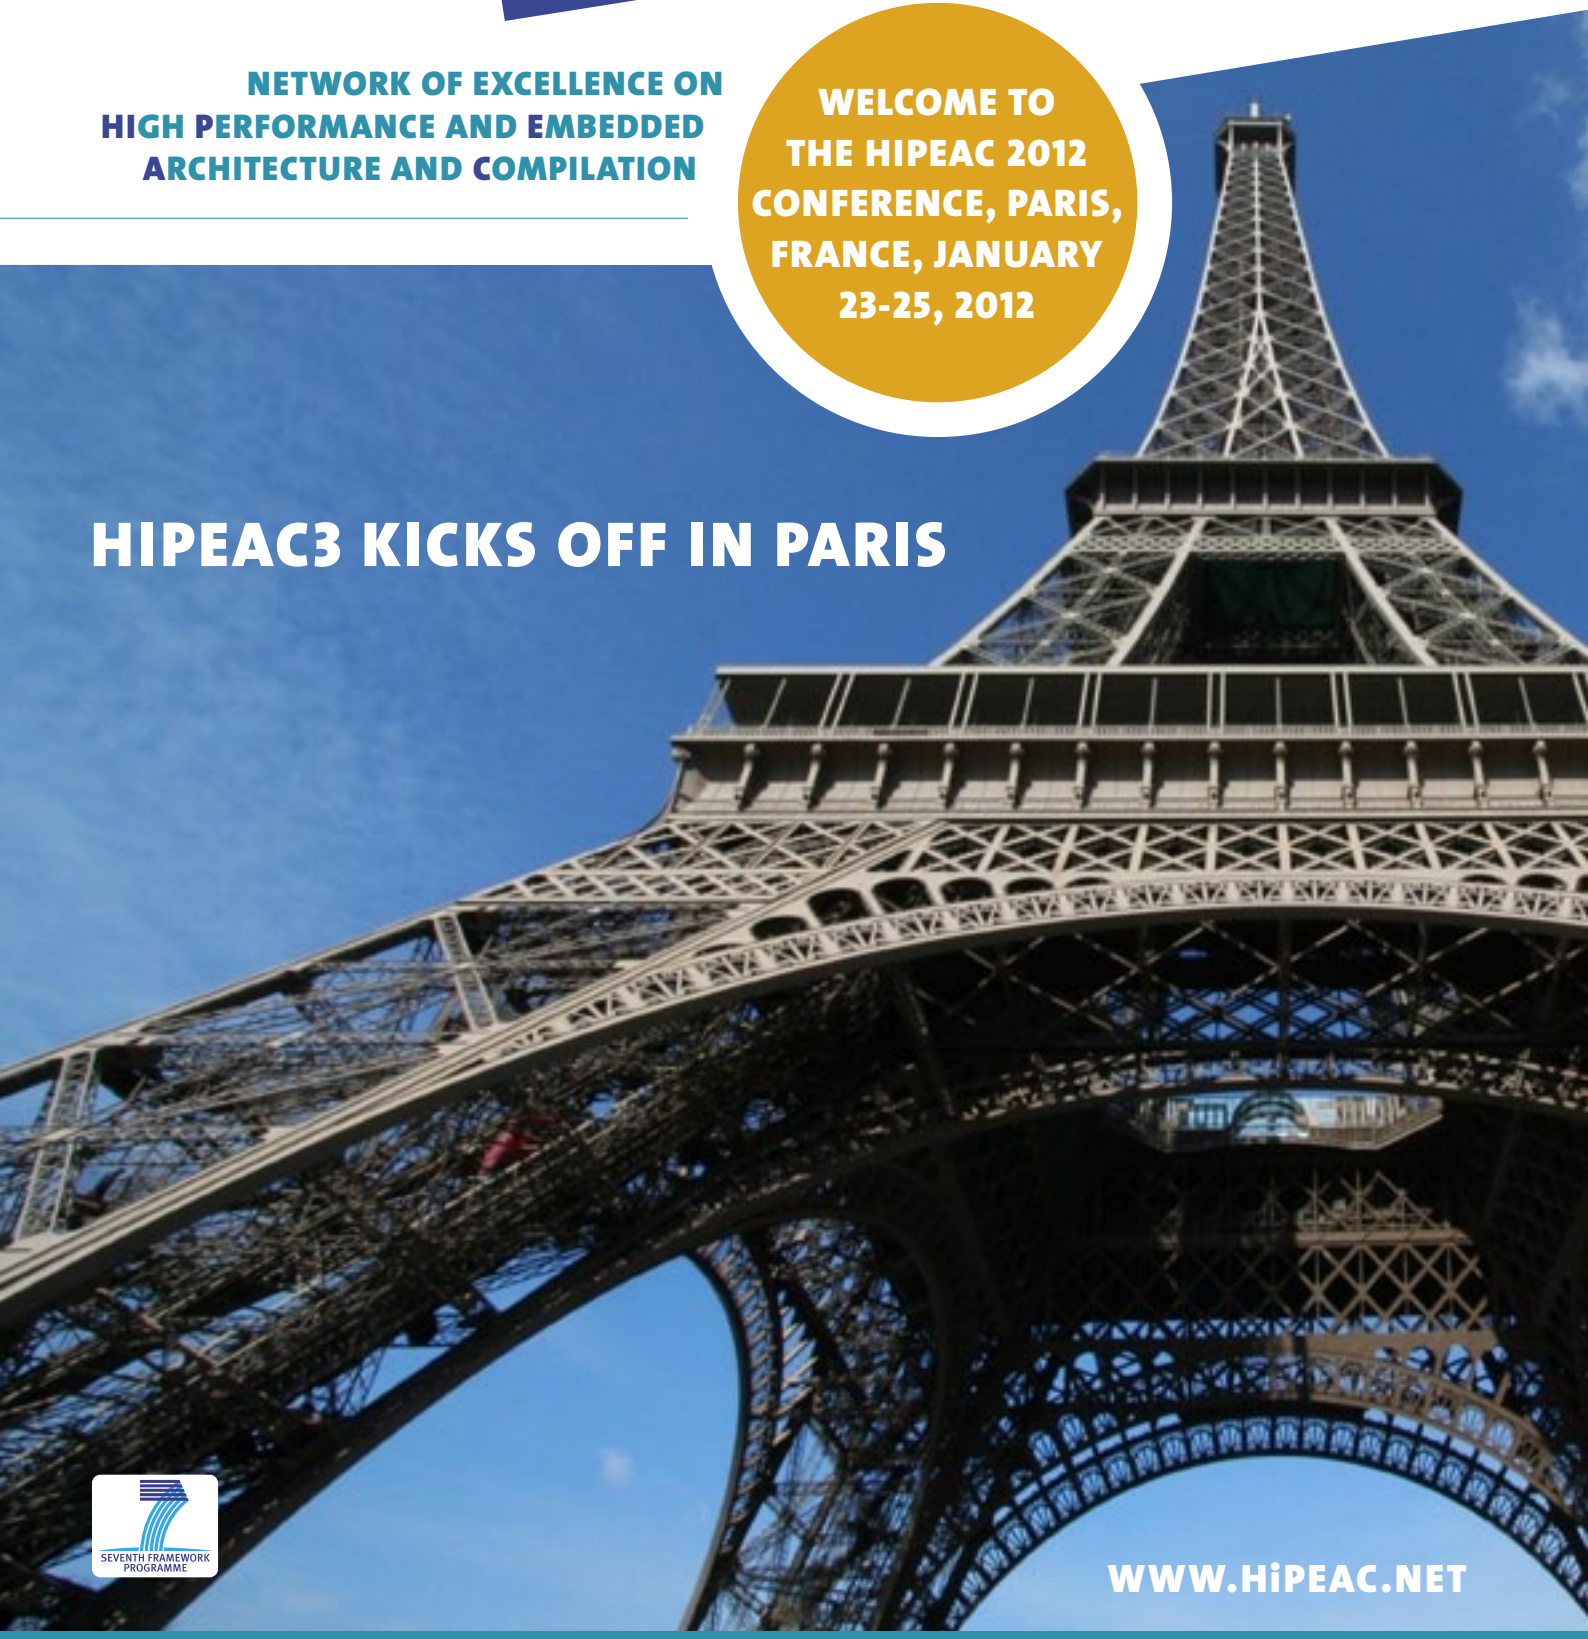

SEVENTH FRAMEWORK PROGRAMME

WWW.HiPEAC.NET

SPRING COMPUTING SYSTEMS WEEK, GOTHEBORG, SWEDEN, 24-26 APRIL, 2012

MESSAGE FROM THE HIPEAC COORDINATOR

First of all, I would like to wish you a happy 2012, both personally and professionally. I hope that 2012 will bring us better news than 2011 – quite a turbulent year for the global economy and for many European countries. I hope that you were not too much affected by it, and that the situation will rapidly improve in 2012, despite the predictions that seem to point in the other direction.

This newsletter marks the transition between HiPEAC2 and HiPEAC3. We decided to redesign its lay-out. We hope that you like it and that it will make our magazine more attractive to read.

One of the good things of 2011 was that the EU trusted us to run the HiPEAC network for another four years, 2012-2015, several years into the next framework program called HORIZON 2020. The transition between HiPEAC2 and HiPEAC3 will be smooth, as most activities will carry on, but sometimes with a slightly different accent. I would like to grasp this opportunity to shed some light on two changes that might impact our community in the short term.

We have seen that the HiPEAC networking events have often helped members to find a better job. To accelerate this effect, we have created a HiPEAC job portal at http://jobs.hipeac.net. At this job portal European employers can post their job openings for computing systems experts. Candidates from all over the world can post their CVs. The job portal will create more visibility for the many European research related jobs in computing, and will attract talent from all over the world, needed for the more than 150 research positions that will be needed to implement the recently approved 45 M€ computing systems projects all over Europe.

A second change I would like to highlight is the new HiPEAC conference. After running the HiPEAC conference in the traditional way for 6 years in a row, we decided to change its concept and to change it from a publication venue into a networking venue. Starting with the 2012 edition, we implemented our 'journal first' publication model. The number of submissions immediately increased by 50%. After two solid review rounds, we were able to accept 37 papers which will all be published as genuine journal papers in ACM TACO, included into the ACM Digital Library, and indexed in all major bibliographical databases. On top of that all these authors got an invitation to present their paper in the main track of the conference. In addition we have over 20 workshops and tutorials running in parallel and there will be daily poster sessions during the coffee breaks. The idea is to have three efficiently packed conference days from which the delegates can compose their personal conference program. The HiPEAC network is committed to further develop this conference model into the yearly premier computing systems event in Europe.

There is no better way to start 2012 than to join us at HiPEAC 2012!

Take care!

Koen De Bosschere

CONTENT

MESSAGE FROM THE PROJECT OFFICER

Panos Tsarchopoulos

Following the evaluation of the last Computing Systems Call for Proposals, the European Commission has launched one Network of Excellence, 14 Research Projects (STREPs) and 3 Support Actions representing 45 million euros of funding. The new projects will provide additional momentum for the European Computing research efforts.

NETWORK OF EXCELLENCE

The **HiPEAC3** Network of Excellence builds on the previous achievements of HiPEAC2 and will work towards steering and increasing the European research in computing systems; improving the quality of the European computing systems research, and creating a visible and integrated pan-European community in computing systems.

RESEARCH PROJECTS

ALMA will provide a full design path and associated tools from Scilab-based specifications to C code running on state-of-the-art multi-core architectures. ALMA will extend both the underlying architectures and Scilab to allow user driven parallelization of Scilab high-level system models.

AutoTune will develop automatic online tuning plug-ins and integrate them with existing on-line performance analysis tools. The resulting framework for performance and energy efficiency tuning will be able to tune serial and parallel codes with or without GPU kernels and will return tuning recommendations that can be integrated into the production version of the code. The whole tuning process, consisting of automatic performance analysis and automatic tuning, will be executed online, i.e., during a single run of the application.

CARP will work on the design of a novel language, PIL (Portable Intermediate Language), for productive accelerator programming. Advanced compilation tools will generate low-level code from PIL programs targeting OpenCL. Portable performance will be achieved through profile-based auto-tuning, enabling compiled code to run efficiently across a range of accelerator platforms.

DeSyRe will develop new design techniques for future reliable Systems-on-Chips. The primary project objectives are to reduce the power and performance overheads of fault-tolerance as well as to improve yield decreasing the number of defective chips. The developed techniques will be used to design new advanced embedded systems targeting future high-tech medical devices for new treatments.

FASTER will facilitate the use of reconfigurable technology by developing novel techniques for optimizing and verifying static and dynamic aspects of a reconfigurable design, while minimizing run-time overheads on speed, area and power consumption. The project will also provide a powerful run-time system that will be able to run on multiple reconfigurable platforms and manage the various aspects of parallelism and adaptivity with reduced overhead.

FlexTiles will define and develop an energy-efficient programmable heterogeneous manycore platform with self-adaptive capabilities. The manycore platform will be associated with an innovative virtualisation layer and a dedicated tool-flow to improve programming efficiency. The project will deliver SystemC simulators and FPGA demonstrators to demonstrate the capabilities of the manycore.

LPGPU seeks to achieve power and bandwidth reductions for real-world software running on next-generation GPUs, as well as GPU architecture designs that are capable of advanced real-time graphics techniques at power levels suitable for battery-powered devices like portable games consoles and mobile phones.

MultiPARTES aims to support mixed criticality for trusted embedded systems based on multi-core open source virtualization. This will be done through the definition, demonstration and validation of a complete methodology and tools. The project results will be demonstrated in applications from aerospace, offshore wind turbines, video surveillance, railway and automotive.

ParaPhrase will develop a new pattern-based approach to programming parallel applications, developing high-level design and implementation patterns as well as software virtualisation abstractions, and applying refactoring to achieve parallelism dynamically in heterogeneous systems. The project will demonstrate applicability of its results in a collection of representative high-performance industrial applications like large-scale databases, video streaming, 3D modelling, advance data mining, machine learning, and renewable energy production.

parMERASA targets a timing-analysable system of parallel hard real-time applications running on a scalable multi-core processor. The parMERASA multi-core processor and system software will provide temporal and spatial isolation between tasks and scale up to 64 cores. Parallelisation of hard real-time applications will be demonstrated in the domains of avionics, automotive and construction machinery.

RELEASE will develop a scalable concurrency-oriented programming infrastructure and its associated tool set to build reliable general-purpose software on massively parallel machines and datacenters. The project builds on the Erlang language and Open Telecom Platform (OTP) libraries that have concurrency and robustness as their main design principles.

T-CREST will design time-predictable resources (processor, interconnect, memories) and the corresponding compiler

infrastructure. Worst case execution time aware optimization methods will be developed along with detailed timing models allowing the compiler to benefit from the known behavior of the hardware. Two industrial use cases (avionics and railway) will be used for validation.

TOUCHMORE targets the development of automatically customisable software tool chains for heterogeneous multi-cores, starting from a high level modelling language (UML/SysML) and taking into consideration energy and system variability constraints. Project results will be demonstrated in applications from automotive infotainment and intelligent driving.

vIrtical targets the full vertical development of effective embedded virtualization with a specific focus on heterogeneous accelerator-based systems. A virtualization-ready System-on-Chip platform and the associated programming models will be developed, tackling all system layers: applications, programming model, hypervisor and hardware.

SUPPORT ACTIONS

The Support Actions **EU-INCOOP**, **RISC** and **SCC-Computing** will work on analysing international research agendas, identifying common research challenges and preparing concrete initiatives for research collaboration in Computing Systems with India, Latin America and China respectively.

Panos Tsarchopoulos

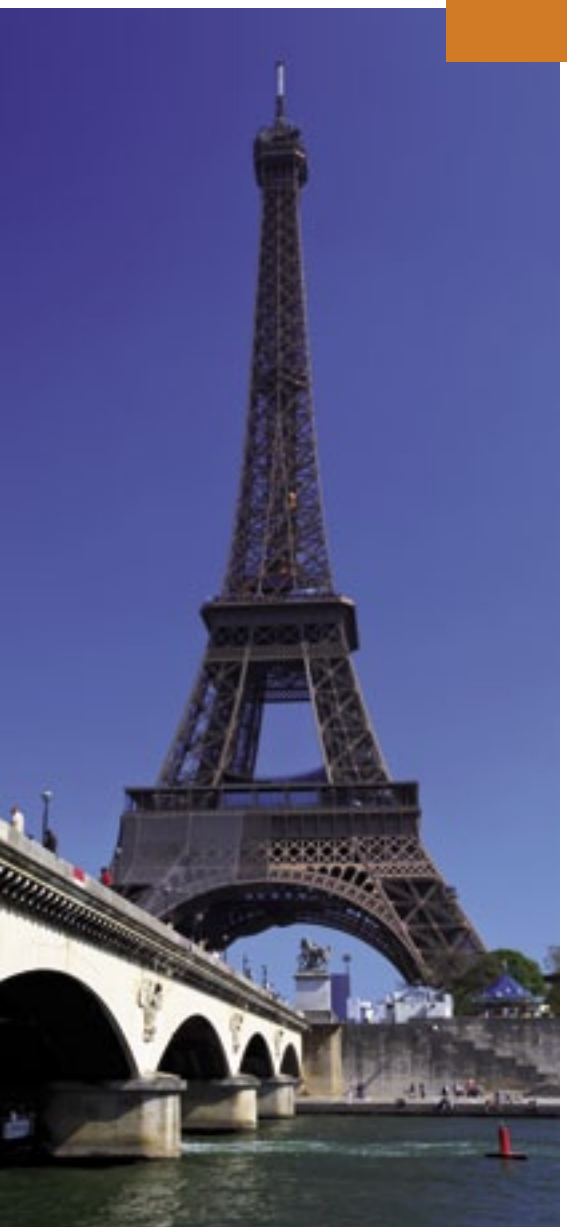

MESSAGE FROM THE NEWSLETTER EDITOR

Dear colleagues,

HiPEAC Info 29 is the last newsletter issue produced by RWTH Aachen University. With the transition to HiPEAC3 some partners´ responsibilities will shift. It is my pleasure to announce that Per Stenström from Chalmers University will take over as newsletter editor from the next issue on.
If my counter is correct, the team at RWTH has produced 16 newsletter issues in the past years. I´d like to acknowledge the help by Jeronimo Castrillon and Anastasia Stulova from RWTH, who enthusiastically managed the day-to-day newsletter business and worked very hard in peak times to keep the publishing deadlines. Both invested enormous effort in quarterly generating a nice-to-read document for you out of the vast raw material received. Further thanks go to Klaas Millet from Ghent for providing quick assistance whenever needed and to the layout team of Wallace & Sanders for their reliable services and short response times. Moreover, various proofreaders helped in quality assurance, including Marisa Gil, Tom Crick, Leigh Murray, and Igor Boehm. Last but not least, I am grateful for all the active HiPEAC members and students that – often on a regular basis – supplied very interesting contributions. This helped a lot to generate a true HiPEAC "community spirit" via the newsletter next to the various face-to-face project meetings.

It was a great pleasure to work with many of you during the HiPEAC2 period, and I am absolutely sure that Per Stenström will continue to provide you with enjoyable "HiPEAC Infos" in the future. In HiPEAC3 I will continue to help on the Steering Committee and to coordinate the "Design and Simulation" research activities. Besides, my major new responsibility is the network´s membership program. I am especially looking forward to support HiPEAC´s mission to expand into the new EU member states. If you have contacts there that might lead to potential new members, please write me at leupers@ice.rwth-aachen.de.

Rainer Leupers,
RWTH Aachen University

HIPEAC'S COMPUTING SYSTEM WEEK BARCELONA MULTICORE WORKSHOP

BSC, HiPEAC and Microsoft Research organized the third Barcelona Multicore Workshop (BMW) which took place in Barcelona on 2-3 November 2011 as part of the HiPEAC Computing Systems Week.
Building on the success of the highly successful 2008-2010 Barcelona Multicore Workshop series, BMW2011 was attended by more than 220 researchers and company decision makers from Europe as well as USA and Asia.

Traditionally BMW workshops examine the most important issue faced by the software community: how to program multi-core processors in the most productive way. Additionally, the overriding important issue faced by the hardware community is how to design the multi-cores so as to maximize the potential performance within a limited power envelope. Those issues describe the same challenge facing both communities: how to maximize the effectiveness of many-core processors.

The proposed solutions to the above problem require a multidisciplinary HW-SW participation. BMW2011 brought together prominent researchers active across the field – computer architecture, programming languages and formal foundations in the course of a two day workshop. The leitmotif this year was on High-Performance Computing. The workshop consisted of a combination of invited talks, and two panels: one on the Fifth Anniversary of the BSC - Microsoft Research Centre (www.bscmsrc.eu) and another on European Research Roadmaps.

The HiPEAC Computing Systems Week continued with a general assembly on topics such as the upcoming HiPEAC conference in Paris, HiPEAC's roadmap, and information about the achievements of HiPEAC2 and the soon to commence HiPEAC3, followed on the next day and a half by the traditional cluster meeting sessions. On Thursday evening the attendees were able to socialize and exchange ideas at the social dinner, where they were able to try some traditional Spanish dishes and wine.

The sessions included a inter-cluster workshop on Languages and Tools for Heterogeneous and GPU-based Multicores organized by the programming models and operating systems, compilation and binary translation, virtualization clusters, and the task force on applications. This inter-cluster session opened with an interesting keynote by Anton Lonkhmotov (ARM) where he presented the current take of ARM on accelerator programming.
Following the keynote, Christoph Dubach (U. of Edinburgh) talked about the Liquid Metal project and their unified language Lime that can be used for programming hybrid computers comprised of multi-core processors, GPUs and FPGAs. Rosa M. Badia (BSC) presented the work done at the BSC on dataflow programming models, specially the CMP-oriented OmpSs. To conclude, Carlos Alvarez (UPC) gave an overview of BSC's research on reconfigurable architectures.

The Design and Simulation cluster meeting included a set of very interesting talks from HiPEAC members regarding system design and integration, modeling and system simulation with emphasis on FPGAs, application-specific computing and hardware acceleration. The first speaker, João Cardoso (U. do Porto), presented an overview of the LARA domain-specific language in the context of the FP7 REFLECT project, and showed its flexibility and effectiveness for exploring the design of flexible hardware cores on FPGA-based platforms. Phillip Coussy (U. de Bretagne-Sud) introduced GAUT, a high-level synthesis tool able to extract the parallelism existing in functions written in pure C code, and generate RTL code to be synthesized on an FPGA. Phillip also showed the GAUT graphical user interface, that helps the user to guide and visualize the GAUT operation. The third talk, given by Christian de Schryver (TU Kaiserslautern), motivated the need for hardware acceleration on financial market simulations, and introduced their approach to create application-specific hardware accelerators for that purpose. Ola Dahl (Linköping U.), who recently joined the HiPEAC community, gave an overview of his previous work at ST-Ericsson on the field of system design and high-level simulation, and shared with us his research interests to start collaborations within HiPEAC. Finally, Juan Eusse (RWTH Aachen U.), showed a hybrid simulation methodology for fast forwarding some parts of the target application, and estimate their execution time based on the time taken on the host on previous executions.

The Interleaving of Power and Reliability cluster meeting saw an extended session, the first part with talks by Avi Mendelson (Microsoft) on the Teraflux approach to reliability and fault-tolerance in future large-scale CMP architectures, Ioannis Sourdis (Chalmers) on the FP7 project Desyre, and Grigorios Magklis (Intel) on co-designed virtual machines; the second part with presentations by Jaume Abella (BSC) on reliability and power in safety-critical systems, Yanos Sazeides (U. Cyprus) on power and reliability trade-offs, and Stefanos Kaxiras (Upsala).

Other highlights of the CSW were: the I/O, Network, and Storage Virtualization cluster meeting, where Muli Ben-Yehuda (Technion and IBM Research) talked about bare-metal peformance for x86 I/O virtualization, and András Vajda (Ericsson) discussed how to scale the data-center network; the Reconfigurable cluster session with talks by Pedro Diniz (INESC-ID) on the REFLECT project: using aspects and strategies for design-space exploration, Dionisios Pnevmatikatos (FORTH) on the FP7 FASTER project: tools for efficient reconfiguration, and Carlos Alvarez (UPC) with an overview of BSC research on reconfigurable architectures; the Education Task Force with a presentation of Ghent University's approach to encouraging students to venture and create new business enterprises; and the MinIR Compilation meeting, where INRIA, Kalray and STMicroelectronics presented their MinIR compiler exchange format.

Alex Rico & Victor Garcia, UPC

HIPEAC MINI-SABBATICAL - AYAL ZAKS

I first met Prof. David August during the 1st HiPEAC ACACES Summer School in 2005, where we both taught courses. Over the years my Compiler Technologies group at IBM tracked and corresponded with David's Liberty group at Princeton University (http://liberty.princeton.edu/) on advanced compilation techniques, including the automatic extraction of pipeline parallelism (Decoupled Software Pipelining). I am again very grateful for HiPEAC's support in bringing us together again, this time for a three months mini-sabbatical at Princeton during the summer of 2011.

The dozen or so Ph.D. students under David's supervision, collectively called Liberators, were absolutely fantastic to work with. Not only was each enthusiastic and open for collaboration about his or her own research topics and work, but what struck me right away was their mutual assistance to one another in providing feedback, advice and support when needed. I immediately plunged into intensive collaboration on a joint paper on dynamic speculative auto-parallelization whose submission deadline was merely two weeks away. Despite the time shortage, I managed to work out some of the math and contribute to the writing. This work continues, with several new directions identified for potential follow-up research and collaboration.
Continuing to the next submission deadline, I joined another line of research, lead researcher and paper preparation, on adapting the execution of automatically parallelized programs, this time with a longer horizon (enough for me to also contribute figures). As we prepared the paper and approached the original deadline, we decided to postpone its submission to another, more ambitious venue. Collaborative work continued after my return from the mini-sabbatical, and a joint paper was finally submitted recently.

Towards the end of my stay, I corresponded closely with another team member working in the area of speculative privatization for automatic parallelization. Several high-level thought-provoking aspects were raised and discussed. Here too, collaborative work continued after my return from the mini-sabbatical, and a joint paper was recently submitted.

Throughout my stay, I enjoyed discussions with each of the group members, learning new lines of work, some which are closer to the programming languages and semantic aspects, others which are closer to runtime and hardware aspects, and caught up on plenty of prior-art reading. Ongoing discussions are being conducted to continue the collaboration after my return, towards additional joint work and publications.

Ayal Zaks, IBM

Day trip to the Catskills with Liberty team members

FOURTH SWEDISH WORKSHOP ON MULTICORE COMPUTING (MCC-2011)

23-25 November 2011, Linköping, Sweden

Christoph Kessler opening MCC-2011 at Linköping University

MCC is an annual workshop series by the network Swedish Multicore Initiative (SMI), started in 2008, as an event to bring together researchers and practitioners from academia and industry to discuss recent work in the area of multi-core computing. Previous MCC workshops have been held at Blekinge Institute of Technology in Ronneby (2008), Uppsala University (2009) and Chalmers University of Technology (2010). MCC-2011 was hosted by Linköping University, organized by Prof. Christoph Kessler at the Department of Computer and Information Science.

As can be seen from the submissions, the workshop has gathered great interest not only in Sweden but also in other countries, especially from Denmark, Finland and Norway. In total, 31 papers were submitted to MCC-2011. Each paper was peer-reviewed by three to four program committee members or external reviewers. Based on the reviews and the PC discussion, we finally decided to accept 16 papers for regular presentation and 15 papers for poster presentation.

In this year, MCC featured for the first time an extra half-day with two tutorials, centered around the topic of portable programming of heterogeneous multicore systems: Dr. Samuel Thibault from INRIA / University of Bordeaux, France, gave a tutorial about the StarPU runtime system, and Dr. George Russell from Codeplay Inc., Edinburgh, UK, presented the Offload-C++ language and compiler. The tutorial program was made possible by the EU FP7 project PEPPHER, which co-located its end-of-year meeting with MCC at Linköping University.

Also, MCC-2011 offered no less than four keynote presentations: Dr. David Moloney from Movidius Ltd., Dublin, Ireland, talked about "Green Multicore" for designing power-efficient multicore architecture for mobile platforms, Dr. Victor Pankratius from KIT Karlsruhe, Germany, about "Ubiquitous Auto-tuning", Prof. Philippas Tsigas from Chalmers University of Technology about "Design challenges for scalable concurrent data structures", and Dr. David Black-Schaffer from Uppsala University about "GPUs: The Hype, The Reality, and The Future".

The technical program was complemented by a panel discussion about "Teaching Multicore Programming – When, What, How?", contrasting perspectives from industry and academia.

Program details can be found at http://www.ida.liu.se/conferences/mcc2011

BOOK ON MULTIPROCESSOR SYSTEMS ON CHIP: DESIGN SPACE EXPLORATION

By Torsten Kempf, Gerd Ascheid, Rainer Leupers

Latest generation smartphones support a wide range of applications from the domain of multimedia, entertainment and infotainment. The various requirements of these devices have created one of the most challenging assignments in engineering today. Supported applications have high computational requirements and high-energy efficiency is requested due to the limited capacity of batteries. Furthermore, short time-to-market, extremely short life-cycles and low cost put a particular pressure on system architects during the design process. As an optimal choice for implementation heterogeneous Multi-Processor System-on-Chip (MPSoC) platforms are considered. However, system architects have to apply new technologies in order to take the right design decisions and to select the best trade-off between the contradicting requirements.

For an effective design process, it is of vital importance to quickly identify suitable implementation candidates and to have a simple and quick analysis of the platform characteristics subject to the application requirements of a single design point.

The book "Multiprocessor Systems on Chip: Design Space Exploration" answers these challenges by combining analytical- and simulation-based models. The methodology and tools described within the book allow evaluation of a single design before any time- and cost-intensive development has to be carried out. Furthermore, it enhances the exploration by enabling a smooth transition from high-level models down to a detailed system implementation. This comprises an analytical implementation model and an abstract simulation-based model, including a virtual processing unit and advanced task modeling. In addition, to the unique exploration methodology the book provides a comprehensive overview of design space exploration techniques and their use cases.

BOOK ON MULTI-OBJECTIVE DESIGN SPACE EXPLORATION OF MULTIPROCESSOR SOC ARCHITECTURES

By Cristina Silvano, William Fornaciari, Politecnico di Milano, Italy; Eugenio Villar, University of Cantabria, Spain

This book serves as a reference for researchers and designers in Embedded Systems who need to explore design alternatives. The MULTICUBE project (an EU Seventh Framework Programme project) has focused on this problem for the past three years and is the basis for this book. It provides a design space exploration methodology for the analysis of system characteristics and the selection of the most appropriate architectural solution to satisfy requirements in terms of performance, power consumption, number of required resources, etc.

This book

- Focuses on the design of complex applications, where the choice of the optimal design alternative in terms of application/architecture pair is too complex to be pursued through a full search comparison, especially because of the multi-objective nature of the designer's goal, the simulation time required and the number of parameters of the multi-core architecture to be optimized concurrently.
- Describes the MULTICUBE Design Space Exploration methodology, which provides a multi-level system specification and modeling framework to provide static and dynamic evaluation of the system-level metrics
- Provides a common tool interface composed of several layers that are connected through standardized interfaces
- Offers a short path to real design space exploration, through use of industrial design flows for examples and tools;
- Includes optimizations in areas such as multi-processor architectures, multi-media, low-power architectures, system-level simulation and profiling and run-time management of resources.

ACACES 2012: 8TH – 14TH JULY, 2012, FIUGGI, ITALY

8TH INTERNATIONAL SUMMER SCHOOL ON ADVANCED COMPUTER ARCHITECTURE AND COMPILATION FOR HIGH-PERFORMANCE AND EMBEDDED SYSTEMS

We are proud to announce the eighth HiPEAC Summer School, which will take place in downtown Fiuggi, a historical hill town near Rome, during the second week of July. We start on Sunday evening with an opening keynote. The 12 courses start on Monday, spread over two morning and two afternoon slots. There are three parallel courses per slot, from which the participants can take one course. The courses have been allocated to slots in such a way that it will be possible to create a summer school program that matches your research interests.
The following world-class experts will present the topics of this year's Summer School.

Instructor	Title
Murali Annavaram - University of Southern California, USA	Reliability: the next frontier for systematic benchmarking and monitoring tools
David August - Princeton University, USA	Next generation automatic parallelization
Katherine Compton - University of Wisconsin-Madison, USA	Resource management in reconfigurable computing systems
Natalie Enright Jerger - University of Toronto, Canada	Networks-on-Chip: communication challenges for many-core architectures
Antonio Gonzalez - Intel & UPC, Barcelona, Spain	Processor microarchitecture
Ted Huffmire - Naval Postgraduate School, USA	Hardware support for trustworthy systems
Krishna Palem - NTU, Singapore & Rice University, USA	What to do about the end of Moore's law (probably)?
Keshav Pingali - University of Texas, Austin, USA	Multi-core parallel programming using the Galois system
Ravi Rajwar - Intel, USA	Transactional memory and synchronization
Josh Simons - VMware, USA	Virtualization and High Performance Computing
Jürgen Teich - University of Erlangen-Nuremberg, Germany	Domain-specific and resource-aware computing on multi-core architectures
Thomas Wenisch - University of Michigan, USA	Designing efficient data centers

On Wednesday afternoon, participants are given the opportunity to present their own work to other participants during a huge poster session; and finally, on Friday evening there will be a farewell dinner and party. The accommodation will be provided by a consortium of hotels in Fiuggi, all located closely together. There will be abundant Italian food, and the town of Fiuggi will provide plenty of opportunities to socialize in the evenings. At the end of the event, all participants will receive a certificate of attendance detailing the courses they took. If you are a student member of HiPEAC, you can apply for a grant that covers the registration fee. In this newsletter, you will find a summer school poster. Please post it at some visible place in your department. You can find more information about the summer school at http://www.hipeac.net/summerschool. We look forward to seeing you there!

Koen De Bosschere
Summer school organizer

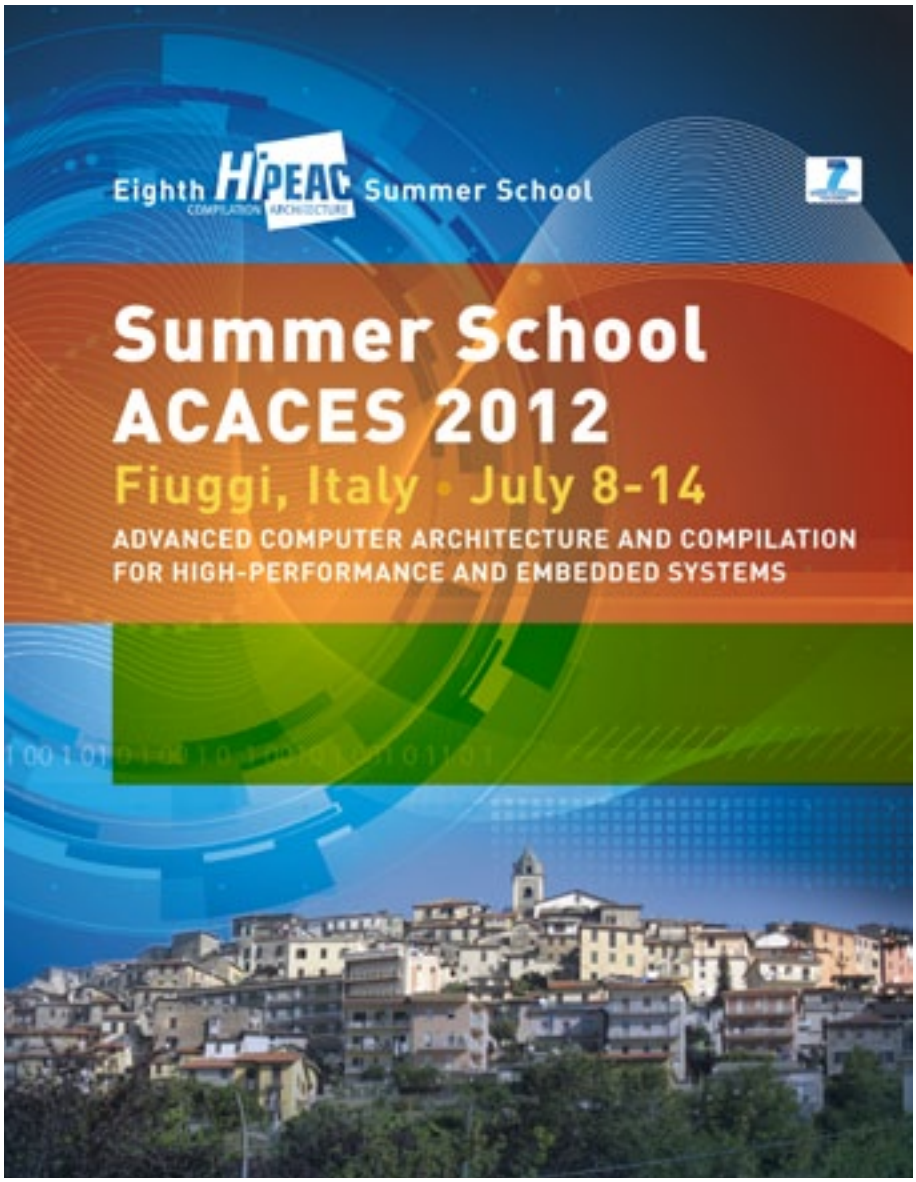

ICE OBTAINS THE BEST PAPER AWARD AT SOC 2011

Weihua Sheng presenting the paper at SoC

The authors from the ICE Institute of RWTH Aachen University (Weihua Sheng, Stefan Schürmans, Maximilian Odendahl, Rainer Leupers and Gerd Ascheid) received the best paper award of the International Symposium on System-on-Chip (SoC), held in Tampere, Finland, October 31st to November 2nd 2011. The paper entitled "Automatic Calibration of Streaming Applications for Software Mapping Exploration" was recognized for its contributions in the area of fast and accurate virtual prototyping for multi-core software development. Nowadays, software engineers face a huge challenge to map the parallelized applications to multi-core platforms efficiently. The paper introduces a tool-flow to create fast multi-core virtual prototypes supporting fully functional execution of software with good timing accuracy. Supported by the UMIC Research Centre of RWTH Aachen University, the work presented has been developed in the context of MAPS (MPSoC Application Programming Studio), which is a long-term R&D investment by ICE to enable programming of real-life complex MPSoC platforms.

FLEXSOC AN EXPOSED DATAPATH ARCHITECTURE TOOLCHAIN

The FlexSoC toolchain has been developed to enable research on exposed datapath architectures, in which the compiler has complete control of the datapath. This opens up a wide range of new optimizations, while putting more demands on the compiler. In an exposed datapath, the instructions, or rather micro operations, are statically scheduled by the compiler. It is therefore up to the compiler to handle matters like data forwarding between units and to decide what values are stored in the register file.

As it is the compiler's task to create the final static schedule of micro operations, the compiler has the possibility of optimizing the interactions between these. For example, if a produced value is only temporarily needed and can be directly forwarded to another unit, this value does not need to be stored in the register file. Alternatively, if a load/store instruction has an offset of zero, the address does not need to be computed, allowing the ALU to be used for other useful operations. This gives opportunities to both save power (for example, no unnecessary writes to the register file) and improve performance (for example, perform an ALU operation in parallel with a load/store operation) without requiring more hardware resources than an in-order datapath.

A wide control word might increase the code size and the required instruction bandwidth. To address this problem, the FlexSoC toolchain can be used to generate an application domain specific instruction decompressor that is reconfigurable during runtime. This enables the creation of application specific or even application phase specific compressed instructions.
The FlexSoC toolchain consists of a number of tools developed at Chalmers University of Technology. The toolchain consists of compilers (FlexComp), based on GCC or LLVM, with code generating backends based on SAT solvers, functional simulators (FlexSim), VHDL code generator (FlexGen) for generating exposed datapath cores (FlexCores), instruction decoder generator (FlexSize), a reference datapath implementation (MIPS), and an extensive set of scripts and makefiles for automatic compilation, simulation, and synthesis of benchmarks and hardware implementations.

For more information and to download the toolchain please visit http://www.flexsoc.org.

Magnus Själander, Chalmers

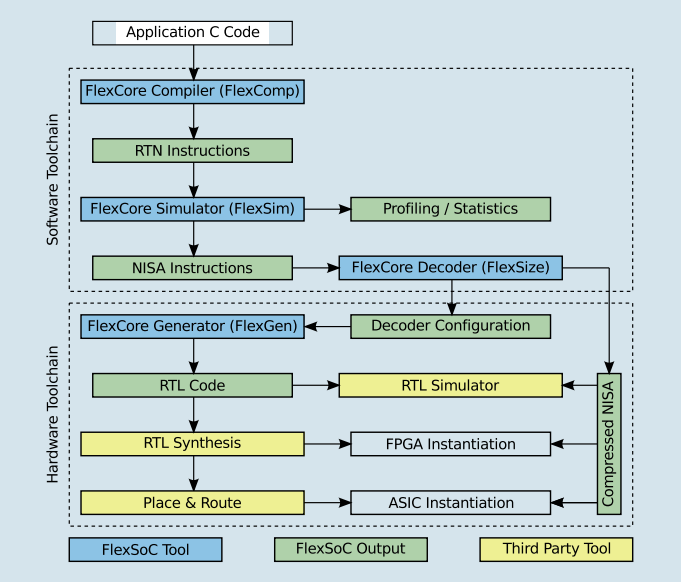

FlexSoC toolflow

ENERGY EFFICIENT ACCELERATION OF ASSET SIMULATIONS USING FPGAS

The recent happenings on the financial markets all around the world have sensitized politicians, institutes and all of us to the inherent risks of our current economic system. As a consequence and to obviate such incidents in the future, the G-10 and the EU commission have introduced the Basel III and Solvency II regulations for finance and insurance institutes. These regulations impose strict conditions on the institutes and force them to asses their financial risks much more frequently than in the past.

However, risk assessments are highly computational tasks. They are mainly based on simulation the future behavior of the underlying asset prices. However, accurate simulations require sophisticated models such as the industry standard Heston model that in many cases lack closed-form solutions and require numerical solving methods, for example finite difference or Monte Carlo methods.

Nowadays, those number crunching simulations are performed on huge CPU- or GPU clusters with up to several thousands of nodes. For instance, an associate project partner from the insurance domain has to compute 2,000 exotic option prices every five minutes. They are currently using an IBM CPU cluster with 2,000 nodes that consumes approximately 2.9 GJ of energy every hour. This is an equivalent of nearly 11 t of CO_2 per day.
Not only due to increasing energy costs, financial and insurance institutes are forced towards more energy efficient

simulation platforms. Recent GPUs can provide an immense speedup of number crunching applications: investigations in a joint cooperation of the Microelectronic Systems Design Research Group and the Stochastic Control and Financial Mathematics Group within the Center for Mathematical and Computational Modelling (CM)2 at the University of Kaiserslautern in Germany have shown that an accelerated system with an Nvidia Tesla C2050 GPU can speed up a Monte Carlo asset simulation by a factor of around 5.5, reducing the energy per task to 44% at the same time, compared to a fully loaded 8-core Intel Xeon server running at 3.07 GHz. Although GPUs require a much higher programming effort, they are currently state-of-the-art in industry for those applications.

However, 44% remaining energy still mean 4.8 t of CO_2 per day. Former investigations from the reconfigurable hardware community have shown devices like field programmable gate arrays (FPGAs) have an enormous potential for energy saving. The first FPGA accelerators for Heston model simulations have been presented very recently at the SC11 supercomputing conference and the ReConFig hardware conference in November 2011. We have shown by real measurements that a floating point Monte Carlo Heston solver on a state-of-the-art FPGA can save up to 98% of the energy per task, compared to GPU clusters. First commercial FPGA accelerators for financial simulations are now available, for example by Maxeler Technologies.

With the increasing variety of heterogeneous platforms, different methods and algorithms as well as numerical characteristics, comparing and fairly evaluating accelerators become a challenge. For that reason, we have developed a platform independent and comprehensive benchmark set for asset price simulations based on the Heston model. It is freely available for download as an executable version on http://www.uni-kl.de/benchmarking. The benchmark set consists of a bunch of Heston parameter sets from real market scenarios, and provides high-precision reference prices. By providing the metrics priced options per seconds and energy per priced option for a given accuracy, accelerators become transparently comparable over architectural and algorithmic borders.

Altogether, FPGAs have a huge potential for bringing down the energy needed for number crunching simulations, particularly in the finance and insurance business. However, deploying application to this devices is still very challenging and end-user oriented tools are mandatory to push FPGAs forward. A lot of ongoing research projects are currently trying to bridge this gap, so we can dare to look forward to seamlessly integrating FPGAs in future simulation clusters easily, making number crunching somewhat greener.

Christian de Schryver & Norbert When, University of Kaiserslautern

AN OPEN-SOURCE LTE UPLINK BENCHMARK

The LTE Uplink Receiver PHY benchmark is an open-source LTE benchmark developed at Chalmers University of Technology during a research project in collaboration with Ericsson AB. The benchmark is a realistic implementation of the baseband processing in an uplink for an LTE mobile base station.

The research community has up until now lacked a shared, freely available benchmark model of baseband processing. This has prevented comparison of system designs and strategies. Given the potential market share that can be controlled by the baseband providers with the best price/performance trade-offs, companies have been strictly guarding their intellectual property. Nonetheless, in light of the rapid expansion of the number, needs, and distribution of consumers of mobile connectivity, an argument can be made for working together as a community to meet future baseband processing requirements.

The LTE Uplink Receiver PHY benchmark is organized as a software pipeline in which modules can easily be replaced to model different algorithms. Although the benchmark's performance cannot rival state-of-the-art proprietary software and platforms, it provides a reasonable baseline for comparing hardware designs or resource management strategies. In short, it realistically captures the dynamic behavior of an LTE baseband uplink as viewed by the base station.

The benchmark can be downloaded from http://sourceforge.net/projects/lte-benchmark/

Magnus Själander, Chalmers

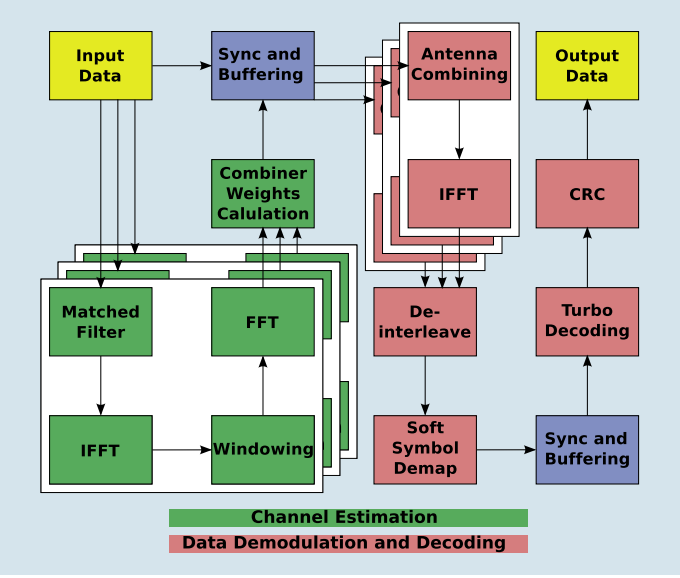

LTE benchmark block diagram

DOMINIK GREWE RECEIVES NVIDIA GRADUATE FELLOWSHIP AWARD

Dominik Grewe, a PhD student from the Compiler and Architecture Design group at the University of Edinburgh, has been awarded an NVIDIA Graduate Fellowship. This year, 11 students worldwide have been selected from over 200 applicants from 25 countries with Dominik being the only recipient from a European university. The NVIDIA Graduate Fellowship Program supports research on GPU-based heterogeneous computer systems. The award includes financial as well as technical support for the students.

Dominik's research interests lie in the area of compiler optimisations for GPU-based heterogeneous systems. His focus is on mapping and optimising parallel tasks on such systems using machine learning. The support from NVIDIA will allow Dominik to expand his research to a broader range of hardware, as well as foster future collaboration with academic and industrial institutions. In fact, he will join NVIDIA's compiler team for an internship in Seattle next year.

For more information about the NVIDIA Graduate Fellowship Program and how to apply visit: http://research.nvidia.com/relevant/graduate-fellowship-program

1st MAPS USER GROUP WORKSHOP ORGANIZED BY ICE, RWTH AACHEN UNIVERSITY

MUG 2011 workshop attendees and tutors

On September 28-29, the 1st MAPS User Group Workshop (MUG 2011) has been organized by ICE in Aachen. MAPS (MPSoC Application Programming Studio) is a programming tool suite for heterogeneous multicore architectures that has been developed in the context of RWTH´s UMIC research cluster (www.umic.rwth-aachen.de). It uses both sequential C and a C language extension (CPN) for describing applications in the form of process networks, and it performs optimized temporal and spatial task-to-processor mapping for embedded MPSoC platforms. The workshop was organized as a training event, interleaving short presentations with hands-on sessions, primarily targeted to industrial participants. The workshop consisted of four tutorials covering different aspects of the MAPS tool infrastructure such as its programming model, mapping and scheduling, sequential C code partitioning and integration with the state-of-the-art ESL (Electronic System Level) and silicon vendor tools. During the tutorials, the participants had full access to the MAPS tools which are integrated in an Eclipse-based environment. A team of MAPS tutors was also on-site assisting the participants and having 1:1 discussions with them. Moreover, various demonstrations were shown during the workshop including a case study of applying MAPS tools to the TI OMAP3 System-on-Chip platform.

The main goal of MUG 2011 was to introduce MAPS for the first time to a wider non-academic audience, and to receive feedback for future enhancements. As a result, MUG attracted more than 20 personally invited participants from 15 international companies, including various HiPEAC member companies. Many interesting discussions and comments were received both on-line in the tutorial sessions and off-line during breaks, the social event, and dinner. Overall, MUG 2011 was very well received and provided valuable feedback to the MAPS team for its future roadmap.

Rainer Leupers, RWTH Aachen University

TU DELFT RECEIVES THE BEST PHD STUDENT PAPER AWARD AT ICSTCC

The paper entitled "Leakage-enhanced 3D-Stacked NEMFET-based Power Management Architecture for Autonomous Sensor Systems", authored by Marius Enachescu, George R. Voicu, and Sorin Cotofana, has been granted the Best PhD Student Paper Award by the committee of the 15th International Conference on System Theory, Control, and Computing (ICSTCC) held in Sinaia, Romania from October 14 to October 16, 2011. The authors are with the Computer Engineering Laboratory, Faculty of Electrical Engineering Mathematics and Computer Science, Delft University of Technology, The Netherlands. In this paper, the authors deal with the increase of leakage power consumption, induced by CMOS fabrication technology advance into the deep sub-100 nm region, in autonomous sensor systems. To this end, they make use of emerging nano-devices, i.e., Nano-Electro-Mechanical FETs (NEMFETs), and Through-Silicon Via (TSV) based 3D chip stacking in an attempt to substantially diminish the leakage power. They propose a novel NEMFET-CMOS hybrid 3D-Stacked power management approach to alleviate the leakage overhead associated with the use of CMOS devices as sleep transistors, always-on cells, in isolation (ISO) cells, and in the Power Management (PM) controller. Their proposal relies on: (i) moving the ISO cells and the PM controller on the NEMS die, and (ii) redesigning them in the NEMS technology to take advantage of the NEMFET ultra low leakage power. The paper also explores the practical implications of such an approach and evaluates the performance and energy efficacy of this power management architecture in a real-life scenario, i.e., an openMSP430 processor running a heartbeat rate monitoring application. The results are promising and clearly indicate that 3D stacked hybrid integration can pave the way towards the industrial utilization of NEMS and CMOS technologies in the design and realization of energy effective computation platforms.

BARCELONA SUPERCOMPUTING CENTER NAMED NVIDIA CUDA CENTER OF EXCELLENCE

The Barcelona Supercomputing Center (BSC) in association with Universitat Politecnica de Catalunya (UPC) has been awarded by NVIDIA as a CUDA Center of Excellence (CCOE). The announcement took place during the International Conference for High Performance Computing, Networking, Storage and Analysis (SC11) held in Seattle (USA), November 14th 2011.

BSC and UPC currently offer a number of courses covering CUDA architecture and parallel computing. BSC became the first NVIDIA CUDA Research Center in Spain, in 2010. BSC/UPC Prof. Mateo Valero is Principal Investigator for the Center, and Prof. Nacho Navarro is acting director of the actual BSC CUDA Research Center. During the next three years, BSC plans to build an Education Program on Massively Parallel Programming, as a continuation of the Programming and Tuning Massively Parallel Systems (PUMPS) summer schools held in 2010 and 2011 and co-organized with the HiPEAC NoE. The center will foster multi-GPU and cluster-aware programming environments for GPUs promoting a unified resource management like GMAC, and will enable the forthcoming Exascale supercomputing era with GPU acceleration using the task-based StarSs programming model and its OmpSs implementation. BSC is already starting to build a new GPU-based cluster prototype system to explore the potential of low-power GPU clusters as high-performance platforms.

"We feel honored to become a CUDA Center of Excellence and for sure this will strengthen even more the well settled collaboration with NVIDIA", said Mateo Valero, BSC Director. "Our aim is that Barcelona becomes a hub for training in CUDA for scientists and faculty members from European academic and research institutions."

The CUDA Center of Excellence program recognizes rewards and fosters collaboration with leading institutions at the forefront of parallel computing research. A world leader in Computer, Life and Earth sciences as well as computational applications in Science and Engineering, BSC joins a network of 13 elite institutions worldwide that have demonstrated a unique vision for improving the technology and application of parallel computing, and are empowering academics and scientists to conduct world-changing research.

Mateo Valero (BSC) with Jen-Hsun Huang (NVIDIA)

BSC-MICROSOFT CENTRE RESEARCH GROUP WINS TWO BEST PAPER AWARDS

ICPE 2011

The paper titled: "RMS-TM: A Comprehensive Benchmark Suite for Transactional Memory Systems" won the Best Paper Award in the 2011 International Conference on Performance Engineering (ICPE). Authored by Gokcen Kestor, Vasileios Karakostas, Osman S. Unsal, Adrián Cristal, Ibrahim Hur and Mateo Valero, the paper introduces RMS-TM, a Transactional Memory (TM) benchmark suite composed of seven real-world applications from the Recognition, Mining and Synthesis (RMS) domain. In addition to featuring current TM research issues such as nesting and I/O and system calls inside transactions, the RMS-TM applications also provide a mix of short and long transactions with small/large read and write sets with low/medium/high contention rates. These characteristics, as well as providing lock-based versions of the applications, make RMS-TM a useful TM tool. The evaluation with selected Software TM and Hardware TM systems shows that the RMS-TM benchmark suite is also scalable, which is useful for evaluating TM designs on high core counts.

GLSVLSI 2011

The paper titled: "Circuit Design of a Dual-Versioning L1 Data Cache for Optimistic Concurrency" won the Best Paper Award in the 2011 Great Lakes Symposium on VLSI (GLSVLSI) held in Lausanne, Switzerland in May. Authored by Azam Seyedi, Adria Armejach, Osman S. Unsal, Adrian Cristal, Ibrahim Hur and Mateo Valero, the paper proposes a novel L1 data cache design with dual-versioning SRAM cells for chip multiprocessors that implement optimistic concurrency proposals. The paper proposes a complete circuit design of a 32-KB dual-versioning L1 data cache with 45-nm CMOS technology at 2GHz processor frequency and 1V supply voltage. The paper also introduces three well-known use cases that make use of optimistic concurrency execution and that benefit from the proposed design: Lock Elision, Speculative Multithreading and Transactional Memory. The authors evaluate one of the use cases to show the impact of the dual-versioning cell in both performance and energy consumption.

A NEW WORLD-CLASS OPEN STANDARD FOR DIRECTIVE MANYCORE PROGRAMMING

CAPS and Pathscale announced last June that OpenHMPP – a non-profit association of academic, research, and industry partners which mission is to work with and for the HPC community to make HMPP directives a world-class open standard – was now live!

Since 2010, CAPS and Pathscale have been coordinating their effort to make HMPP directives an Open Standard denoted OpenHMPP. This initiative aims at creating a community effort for proposing a directive manycore programming interface and thus capitalizing on existing experiences with HMPP.

"CAPS has always been committed in delivering to users solutions that will secure their software investment. We think the directive approach is definitely the best way to meet this issue. Now, with the open standard and the software companies adopting HMPP, our customers feel confident in their choice." Declares Laurent Bertaux, CAPS CEO.

The HMPP directives, initially developed by CAPS, have been designed and used for more than 3 years by major HPC leaders. As a high level of abstraction for programming GPUs in scientific applications, HMPP preserves customers' software assets by keeping applications portable. The HMPP directive-based programming model offers a powerful syntax to efficiently offload computations on hardware accelerators and to optimize data movement.

HMPP directives describe remote procedure call (RPC) on an accelerator device (e.g. GPU) or more generally a set of cores. The directives annotate C or Fortran codes to describe two sets of functionalities: the offloading of procedures (denoted codelets) onto a remote device, the optimization of data transfers between the CPU main memory and the accelerator memory.

The OpenHMPP consortium will oversee the HMPP directives specification, produce and approve new versions of the specification. In accordance with its vision involving participation and knowledge sharing, OpenHMPP works towards providing programmers with a portable, scalable directive-based programming model for developing parallel applications on many-core platforms.

A set of core partners (CAPS, GENCI, ICHEC, INRIA, PATHSCALE, etc.) are currently in discussion to set up the non-profit organization that will own the OpenHMPP brand and run its evolutions. This organization will gather technology providers, application developers as well as ISVs and will work within working groups such as: HMPP directives, Dataflow, Tracing event...

The OpenHMPP consortium is now open to any interested company to join. OpenHMPP first user club was held on November 14th in Seattle, at Pan Pacific Hotel.

CAPS Press Contact:
Estelle Dulsou
estelle.dulsou@caps-entreprise.com
Website: www.caps-entreprise.com

ABOUT CAPS ENTREPRISE

CAPS is a major supplier of solutions dedicated to application migration and deployment on manycore processors.
CAPS global solution for manycore leads the developer to performance by providing top-of-the-range technology (HMPP hybrid compiler and wizard), code porting methodology and ecosystem (third software tools, expertise, training, etc.).
Its directive-based & multi-target HMPP™ compiler enables developers to safely move to hybrid CPU / GPU model and quickly get performance by leveraging the computing power of stream processors without the pain associated to GPU programming. HMPP™ is offered within CAPS DevDeck™ package: an ALL-IN-ONE multi-level suite for manycore application definition, porting and optimization with tools (HMPP compiler, development tools such as HMPP Wizard, debugging & profiling software and scientific libraries), methodology and resources (tutorials, use cases, etc.).

FP7 FASTER PROJECT: FACILITATING ANALYSIS AND SYNTHESIS TECHNOLOGIES FOR EFFECTIVE RECONFIGURATION

Project coordinator:
Dionisios Pnevmatikatos
FORTH-ICS
pnevmati@ics.forth.gr
Partners:
FORTH-ICS
Chalmers
Imperial College London
Politecnico di Milano
Gent University
Maxeler
ST Microelectronics
Synelixis
Project website:
www.fp7-faster.eu
Project start date:
September 1st, 2011
Duration:
36 months

Extending product functionality and lifetime requires constant addition of new features to satisfy the growing customer needs and the evolving market and technology trends. Software component adaptivity is straightforward but not enough: recent products include hardware accelerators – for reasons of performance and power efficiency – that also need to adapt to new requirements. Reconfigurable logic allows the definition of new functions to be implemented in dynamically instantiated hardware units, combining adaptivity with hardware speed and efficiency.

The FP7 FASTER project will facilitate the use of reconfigurable technology by providing a complete methodology that enables designers to easily implement and verify applications on platforms with general-purpose processors and acceleration modules implemented in the latest reconfigurable technology. The FASTER tool-chain input will be based on high-level programming languages with an initial decomposition described using existing formalisms (such as OpenMP). This input will be transformed to the corresponding task graph, which in turn will be partitioned in space and time using new algorithms derived from graph theory. FASTER will support both region- and micro-reconfiguration (a technique that reconfigures very small parts of the device), an ability that opens up a new range of application opportunities for run-time reconfiguration.
FASTER will develop novel techniques for optimizing and verifying static and dynamic aspects of a reconfigurable design, while minimizing run-time overheads on speed, area and power consumption. FASTER will also provide a powerful run-time system that will be able to run on multiple reconfigurable platforms and manage the various aspects of parallelism and adaptivity with reduced overhead.

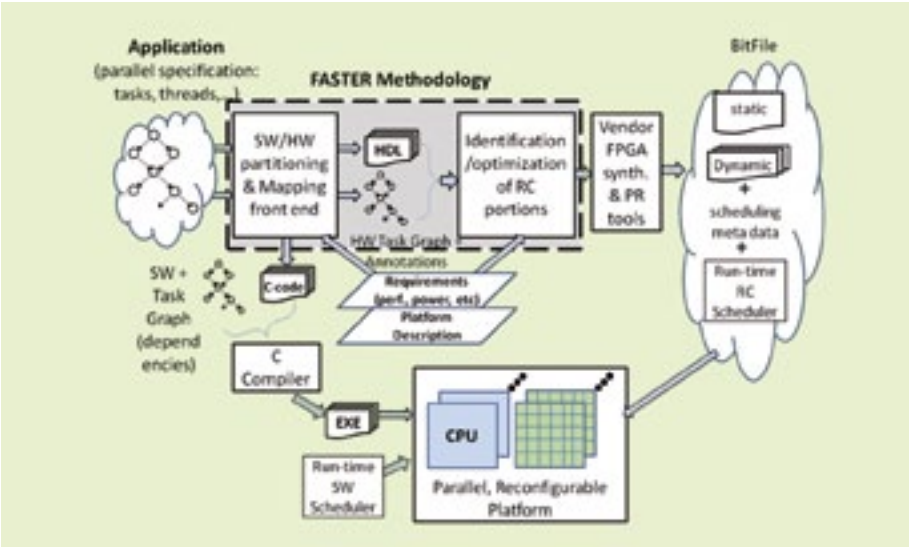

FASTER tool chain

INTERNSHIP REPORT - TASSADAQ HUSSAIN

I am doing Ph.D. degree in Computer Architecture at UPC-Barcelona Tech, under the supervision of Professor Eduard Ayguadé. The focus of my work is on investigating and evolving the high performance computing systems for data and compute bound applications. To improve performance of HPC application and I am working over efficient memory access patterns, which helps the multi processors/accelerators architecture to make best use of the memory.

To provide the understanding for the conceptual framework, I got a HiPEAC Industrial Internship Grant that bestowed me the opportunity to work closely with High Level Synthesis designers at PLDA Italia to develop efficient memory architecture for Reverse Time Migration kernel for FPGA.
During my visit in PLDA Italia, I used a High Level Synthesis (HLS) environment, the Hardware Compiling Environment (HCE), for HPC applications that allows specifying algorithms to be mapped onto FPGAs as standard C programs.
To validate functionalities and features of HCE environment, the computing intensive part of the Reverse Time Migration (RTM) kernel is selected.

In HCE environment, RTM kernel is described in the abstract level (C program) and mapped on FPGAs. The proposed design environment provides ease in designing HPC application, and flexibility to reuse design at the abstract level without going into Hardware/Software details. The design methodology reduces production time to market by addressing full hardware design flow including structural design exploration, RTL design, functional verification, board level simulation, and system integration. The RTM design not only optimizes the "compute intensive" part of the algorithm but also improvise on-chip/off-chip units.
The resulting system has full control and computes units on the same chip which advances the performance of system by removing access and grant time.
The resulting design, when implemented in an Altera Stratix IV EP4SGX230 and EP4SGX530 devices, achieves 11.2 and 22 GFLOPS respectively. Despite low operational frequency of the HCE based design, generated RTM design provides production quality yields up to 7.9x faster than CPU based approach.
I would like to thank HiPEAC for kindly supporting my trip to Italy. It was a great opportunity to learn and investigate design methodology for HPC applications.
This work has been accepted in IEEE International Conference on Field-Programmable Technology FPT IIT New Delhi, India 12-14 December 2011, as titled "Implementation of a Reverse Time Migration Kernel using the HCE High Level Synthesis Tool". The authors of the paper are TASSADAQ Hussain, Miquel Pericas, Nacho Navarro, Eduard Ayguade.

Tassadaq Hussain, UPC

COLLABORATION GRANT REPORT - JOSÉ L. ABELLÁN

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

Thanks to the HiPEAC collaboration grants program, last summer I worked with Prof. Davide Bertozzi and his research group in the Electronic Department at the University of Ferrara (Italy). This research group has a solid background in hardware design and characterization through a complete industrial design flow.

The work performed during our collaboration was focused on improving the performance of barrier synchronizations for many-core CMPs. Traditional software-based barrier implementations tend to produce hot-spots in terms of memory and network contention as the number of processors increases. This is even worse when considering future many-core CMPs which will contain several dozens of cores. Therefore, some sort of hardware acceleration becomes essential in order to meet the desirable performance.

Aimed at developing an efficient hardware-based barrier for future many-core CMPs, prior to our collaboration, my advisor, co-advisor and I had developed a novel mechanism, namely GBarrier, that provides extremely fast, scalable and power-efficient barrier operations. Unlike other hardware proposals, GBarrier relies on a dedicated on-chip interconnection network that decouples completely barrier synchronization from any kind of memory-and-traffic-related activity. Moreover, the very high efficiency of GBarrier is achieved through a simple synchronization protocol

and the use of state-of-the-art technology that enables almost speed-of-light communications. This full-custom technology is not within reach of a standard cell design methodology, hence the goal of the collaboration was to determine whether the level of efficiency achieved by the GBarrier mechanism could be kept up with current cost-effective industrial design flows.

To do so, we started by implementing our GBarrier scheme through the mainstream industrial toolflow. Also, other two schemes were proposed and implemented due to their lightweight and performance-efficient connectivity patterns: CBarrier, a centralized implementation that uses a master-slave paradigm, and TBarrier, that is based on a tree connectivity pattern. In particular, all these schemes were evaluated in terms of the hardware overhead that they entail. For that, we implemented the three designs using Verilog/VHDL, and went through a physical synthesis process mapping them onto a 45nm silicon technology. Such characterization was carried out in terms of critical path degradation and on-chip area overhead. We also performed a comprehensive exploration in terms of layout sizes and varying clustering granularities. Inter-cluster barriers were designed using globally-asynchronous locally-synchronous (GALS) communications. This was an issue that has never been adequately investigated before. The results of the study performed during the collaboration were included in the paper "Design of a Collective Communication Infrastructure for Barrier Synchronization in Cluster-based Nanoscale MPSoCs", which has recently been accepted at DATE'2012 conference.

Personally, I am very grateful to HiPEAC for supporting the collaboration. I would also like to thank Prof. Bertozzi and his group for a very pleasant internship, especially to my new friends: Daniele Ludovici, Hervè Tatenguem and Alessandro Strano.

José L. Abellán, University of Murcia

INTENSHIP REPORT - JORGE ALBERICIO

Quantitative exploration of memory hierarchy in a context of many-core computing system

This is the second time I am doing an internship at ST Microelectronics Grenoble granted by HiPEAC to work with Arthur Stoutchinin. In both internships, we have studied the memory hierarchy of a many-core System-On-Chip designed by ST Microelectronics in collaboration with the CEA, called "Platform P2012". The platform includes multiple processors with a shared memory. Last year I studied the inclusion of a special cache adapted for prefetching low-temporal-locality data in image and video target applications. This year the objective is to quantitatively compare two configurations for the memory hierarchy. The two different memory organizations that we are comparing are: a tightly-coupled and a loosely-coupled shared memory architectures.

In a tightly-coupled shared memory architecture, a number of processors directly share a level-1 memory. This configuration leads to a reduced hardware area but suffers from more complicated design and reduced scalability. In a loosely coupled shared memory architecture, a number of processors share level-2 memory via a level-1 data cache. This configuration has a potentially larger area and requires maintaining cache coherence at level-1 caches. On the other hand, the loosely coupled architecture allows higher operating frequencies and inclusion of Memory Management Unit (MMU) with comparatively large number of cores (>8).

The tightly coupled architecture seems well suited for control intensive parallel applications requiring small number of cores such as controlling video and image hardware pipelines. On the other hand, for compute intensive applications requiring many cores, the overall performance of the loosely coupled architecture may be superior if the cache miss rate and coherence overhead as well as HW area overhead are reasonably low. In this project we have investigated and compared performance-to-area ratio of the two architectures on some compute-intensive target applications. In particular, we have estimated silicon area for both architectures and also the performance of the target applications using simulation. We have also characterized the cache coherence overhead of the loosely coupled architecture in the target applications and we have studied and proposed a low-cost hardware coherence scheme.

Jorge Albericio, University of Zaragoza

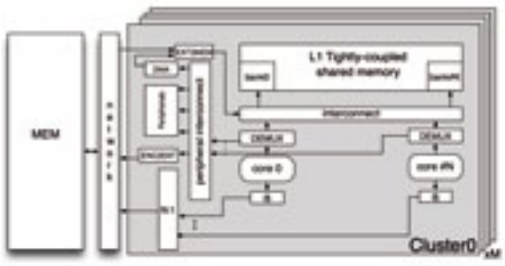

A tightly-coupled shared memory architecture

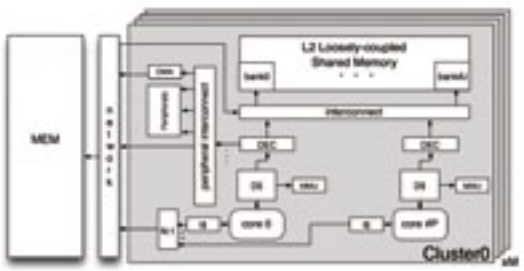

A loosely coupled shared memory architecture

FAULT TOLERANCE CONFIGURATION FOR UNCOORDINATED CHECKPOINTS

By Leonardo Fialho (leonardofialho@gmail.com)
Advisor: Dr. Dolores Rexachs
Universitat Autònoma de Barcelona, Spain
July 2011

Parallel computers are growing in complexity and in number of components. The components miniaturisation and concentration are the major root causes of the failures increasingly seen on these computers. Thus, in order to achieve the execution end, parallel application should use a fault tolerance strategy.

A widely used strategy is the rollback-recovery, which consists of saving the application state periodically. In the event of a fault occurring, the application resumes it execution from the most recent saved state. These fault tolerance protocols include an overhead on the parallel application execution.

Using a coordinated checkpointing protocol it becomes easy to estimate the application execution time, as well as to calculate the frequency in which checkpoints should be taken. In fact, there are very precise models to estimate the application execution time and the checkpoint interval nowadays.

However, the use of the coordinated checkpointing may not be the best solution to provide fault tolerance on the next-generation parallel computers. In other words, the current paradigm of fault tolerance for parallel applications is not suitable for the future parallel computer.

Fault tolerance protocols such as uncoordinated checkpointing permits that each process of the parallel application saves its state independently of other processes. The combination of uncoordinated checkpointing with logging of message-passing events avoids the inconvenience of this sort of protocol, such as the domino effect and orphan messages. This is the emergent paradigm of fault tolerance for scalable parallel applications.

For instance, there is no model suitable to estimate the execution time of a parallel application protected by uncoordinated checkpointing. As well as there is no convenient model to calculate the frequency in which those checkpoints should be taken.

The objective of this thesis is to define suitable models that can be used with each paradigm: the coordinated and the uncoordinated. These models should provide an estimation of the application wall time clock running under each fault tolerance paradigm, as well a methodology to define the value of the variables used to calculate the checkpointing interval.

The main motivation of this work is to provide at the same time the knowledge necessary to face the emergent fault tolerance paradigm and make it suitable to be used by parallel applications users.

CUSTOMIZABLE MEMORY SCHEMES FOR DATA PARALLEL ARCHITECTURES

By Chunyang Gou (c.gou@tudelft.nl)
Advisor: Prof. Georgi Gaydadjiev
Delft University of Technology, the Netherlands
September 2011

Memory system efficiency is crucial for any processor to achieve high performance, especially in the case of data parallel machines. Processing capabilities of parallel lanes will be wasted, when data requests are not accomplished in a sustainable and timely manner. Irregular vector memory accesses can lead to inefficient use of the parallel banks/modules/channels and significantly degrade overall performance even when highly parallel memory systems are employed. This problem is also valid for many regular workloads exhibiting irregular vector accesses at runtime. This dissertation identifies the mismatch between the optimal access patterns required by the workloads and the physical data layout as one of the major factors for memory access inefficiency. We propose customizable memory schemes to address this issue in data parallel accelerators. More specifically, this thesis extends traditional approaches by proposing two new parallel memory schemes that alleviate bank conflicts for commonly used access patterns. We also propose a framework to capture and convey the access pattern information to the proposed parallel memory schemes. Furthermore, we describe techniques that dynamically adjust the instruction sequencer of a multithreaded vector architecture and customize the access patterns to improve on-chip, local memory efficiency. Last, we identify and exploit new locality type to dynamically adjust off-chip memory access granularity of manycore data parallel architectures, in order to improve main memory efficiency. We implemented our proposals as extensions of contemporary data parallel architectures and our evaluation results demonstrate that memory efficiency and overall system performance can be improved at minimal hardware cost, while at the same time programming overhead can be greatly reduced.

OPTIMIZING DYNAMIC DATA STRUCTURES FOR DYNAMIC APPLICATIONS IN EMBEDDED SYSTEMS

By Christos Baloukas (cmpalouk@ee.duth.gr)
Advisor: Prof. Dimitrios Soudris
National Technical University Of Athens, Greece
September 2011

Modern embedded systems and services exhibit intensive data transfer and storage requirements. The range of applications includes network and multimedia along with intensive 3D graphics applications. The developed methodologies and tools target the optimization of the dynamic data structures of the application's source code, in order to minimize the number of data and consequently memory accesses. The proposed methodology optimizes the dynamic data structures based on the particular access pattern that the application uses to access the data in each structure. The access pattern is categorized into several cases like FIFO, sequential or random access etc. Each of these categories is served optimally by a different access pattern.

The first step is to break the access pattern down to a set of characteristics that can uniquely describe each access pattern. That way each pattern can be identified automatically with the use of certain profiling techniques. In this thesis all activity of dynamic structures is recorded and later analyzed to extract the dominant access pattern for each structure. Furthermore, the nature and functionality of a data structure is also broken down to a set of decisions like the number and the kind of headings available, direct or linked organization and single or multiple layers. Each set of choices represents a different data structure implementation. Based on the particular access pattern a certain set of decisions is selected that needs to be further explored to find the optimal implementation. That way the original vast search space is now reduced to a handful of implementations. These implementations can be simulated and the optimal one can then be extracted.

The results presented in this thesis offer solutions to: a) characterizing the dynamic data structures' behavior b) dynamic data management and c) exploration of alternative data structures' implementations.

Results of my PhD work were published in: 2 books, 2 book chapters, 5 journal papers, 8 conference papers and 3 PhD Forums.

DYNAMICALLY AND PARTIALLY RECONFIGURABLE EMBEDDED SYSTEM ARCHITECTURE FOR AUTOMOTIVE AND MULTIMEDIA APPLICATIONS

By Naim Harb (naim.harb@gmail.com)
Advisors: Prof. Smail Niar,
Université de Valenciennes et du Hainaut Cambrésis, Valenciennes, France, Prof. Mazen Saghir, Texas A&M University at Qatar, Doha Qatar
September 2011

The work in this dissertation targets the exploration of the DPR recent FPGAs feature in the advantage of automotive and multimedia systems. We target a Driver Assistant System (DAS) system as our automotive base system. On the other hand, we have selected the H.264 encoder as our multimedia target system.
Starting with the automotive application, in this part of the work, a dynamically reconfigurable filtering hardware block for tracking applications in DAS is presented. Our system shows that there will be no reconfiguration overhead because the system will still be functioning with the original configuration until the new reconfiguration is completed. The free reconfigurable regions can be implemented as improvement blocks for other DAS system functionalities. Two approaches were used to design the filtering block according to driving conditions: one was related to the number of targets while the second was based on targets' proximity and danger level.
Regarding the H.264 multimedia system standard, we propose a dynamically reconfigurable H.264 motion estimation computational unit whose architecture can be modified to meet specific energy and image quality constraints. By implementing 16 reconfigurable regions, we are able to support multiple configurations each with different levels of accuracy and energy consumption. Image accuracy levels were controlled via application demands, user demands or support demands. Using a reconfiguration heuristic, our system can support up to 35 different configurations. With a maximum saving of up to 51% in energy consumption, the system can support all block sizes in an H.264 encoder.

RUNNING STREAM-LIKE PROGRAMS ON HETEROGENEOUS MULTI-CORE SYSTEMS

By Paul Carpenter (paul.carpenter@bsc.es)
Advisors: Eduard Ayguade, Alex Ramirez
UPC, Barcelona, Spain
October 2011

This thesis develops compiler and run-time techniques to map stream programs to multiprocessor machines. Writing high-performance parallel software is difficult, time-consuming, and error prone, leading to high cost and a long time-to-market. We believe that money and time will be saved if parallel software is written using domain-specific high-level languages, or DSLs. These languages improve productivity and they reveal implicit parallelism, enabling the compiler to reorganise the program's structure and data layout. The programmer can then concentrate on whatever the program is supposed to do.

Stream languages are DSLs that represent the program as independent kernels that communicate only through point-to-point channels, known as streams. This representation is natural for many important applications, including those involving audio and video encoding and decoding, 3D graphics, and software radio.

The first part of the thesis is related to the ACOTES compiler (Advanced Compiler Technologies for Embedded Streaming). It introduces a new static partitioning algorithm, which determines which kernels should be fused together, and how they should be mapped to processors. The algorithm uses a convexity constraint to control the length of the software pipeline. It also takes account of constraints imposed by the compiler: the ACOTES compiler, for instance, can only fuse certain combinations of kernels. The thesis also introduces a static queue sizing algorithm, which allocates buffers to balance latencies and hide variations in computation time.

The second part of the thesis investigates dynamic scheduling of stream programs, using the Nanos++ run-time. Previously known online, non-preemptive, non-clairvoyant schedulers either suffer from poor locality or they periodically run out of parallelism. The thesis proposes two low-complexity dynamic schedulers for stream programs.

The thesis also introduces StarssCheck, a tool to find bugs in OmpSs task-level programs. Such tools are important, because a new programming language will only be widely adopted if users can get their programs to work.

MULTICORE ARCHITECTURES FOR BIOINFORMATICS APPLICATIONS

By Sebastian Isaza (sebastianisaza@gmail.com)
Advisors: Georgi Gaydadjiev and Henk Sips
Delft University of Technology, the Netherlands
October 2011

Abstract: In this dissertation, we address the challenges of performance scaling for bioinformatics applications on multicore architectures. In particular, we focus on sequence alignment, one of the fundamental tasks in bioinformatics. Due to the exponential growth of biological databases and the computational complexity of the algorithms used, high performance computing systems are required. Recently, computer architecture has shifted towards the multicore paradigm in an attempt to sustain the performance scalability that single-core processors did provide in the past. Although multicore architectures have the potential of exploiting the task-level and data-level parallelism found in bioinformatics workloads, efficiently harnessing systems with hundreds of cores requires deep understanding of the applications and the architecture specifics. This thesis presents a study of two sequence alignment applications that are modeled, characterized, mapped and optimized targeting two multicore architectures. More precisely, we use the Cell BE and the SARC architectures, the latter developed within the project this thesis was part of. The targeted applications, i.e., HMMER and ClustalW, are used for pairwise alignment and multiple sequence alignment. We first propose an analytical model to predict the performance of applications parallelized under the master-worker scheme. We use HMMER and the Cell BE processor for our experimental case study and for validation of our model. Results show the high accuracy of the model and the scaling behavior of the application phases. Next we investigate the optimal mapping of ClustalW on the Cell BE, identify a number of limitations in the architecture and propose few instruction-set extensions to accelerate the main ClustalW kernel. Last, we study ClustalW and HMMER scalability on the SARC architecture with up to one thousand cores. We also investigate the impact of different input types on ClustalW performance.

THE GPU AS A PROCESSOR FOR NOVEL COMPUTATION: ANALYSIS AND CONTRIBUTIONS

By José M. Cecilia (chemacecilia@gmail.com)
Advisors: Prof. José M. García and Dr. Manuel Ujaldón
University of Murcia, Spain
October 2011

Graphics Processing Units (GPUs) have been at the leading edge of increasing chip level parallelism over the last decade, evolving into sophisticated massively parallel coprocessors. The advent of programming models such as CUDA or OpenCL has facilitated the application of GPUs into many computational domains. However, developers have to deal with a new parallel programming paradigm that is quite different from the traditional ones, and therefore, applications need to be redefined to leverage all GPU capabilities.

Novel Computation techniques have been designed to provide novel solutions for the challenges of the new era. Most of them are inspired in the nature to define massively parallel computational devices. They have been traditionally implemented in sequential architectures that compromises their effectiveness.

The overall goal of this dissertation is to use the GPU as a processor for bridging the gaps between the theoretical definition of those novel computational methods and their practical implementation on silicon-based architectures. Concretely, we focus our analysis on two different bioinspired research topics, Membrane Computing (P systems) and Swarm Intelligence (Ant Colony Optimization, ACO).

We discuss the simulation of P systems solving the Satisfiability (SAT) problem on GPUs. For an efficient handling of the exponential workspace created by the P systems computation, we enable different data policies to increase memory bandwidth and exploit data locality through tiling and dynamic queues. We report speed-up factors exceeding four orders of magnitude when running our simulations on the Tesla S2050 server.

The parallelization strategies on GPUs for the different stages of ACO are also presented, contributing with a novel data parallelism scheme tailored to the GPU, novel GPU programming strategies, and a new mechanism called I-Roulette to replicate the classic Roulette Wheel while improving GPU parallelism. Our implementation leads to factor gains exceeding 20x when compared to its sequential counterpart version.

HARDWARE TECHNIQUES FOR HIGH-PERFORMANCE TRANSACTIONAL MEMORY IN MANY-CORE CHIP MULTIPROCESSORS

By Rubén Titos-Gil (ruben.titos@gmail.com)
Advisors: Manuel E. Acacio and Jose M. García
University of Murcia, Spain
November 2011

One of the key challenges for the computer architects of the multicore era is to understand which abstractions can enhance the productivity of parallel software development and then introduce the appropriate hardware support to realize it. Transactions are a good candidate for such an abstraction, and this thesis focuses on the hardware mechanisms that provide optimistic concurrency control with stringent guarantees of atomicity and isolation, with the intent of achieving high-performance across a variety of workloads, at a reasonable cost in terms of design complexity.

This thesis identifies key inefficiencies that impact the performance of several scalable hardware implementations of transactional memory (TM), and proposes techniques to overcome such limitations. In this dissertation we consider both eager and lazy approaches to hardware TM (HTM) system design, and address important sources of overhead that are inherent to each policy. This thesis presents a hybrid-policy, adaptable HTM system that combines the advantages of both eager and lazy approaches in a low complexity design, by selecting the appropriate policy at the granularity of cache lines.
Furthermore, this thesis investigates the overheads of the simpler, fixed-policy HTM designs that leverage a distributed directory-based coherence protocol to detect data races over a scalable interconnect. For eager systems, we propose a mechanism to prevent the directory controller from becoming a bottleneck during situations of high contention. For lazy systems with early conflict detection, we present a solution that unburdens transactional execution from the penalty of accessing the directory in the common case to guarantee correctness, while providing true commit parallelism for non-conflicting transactions. Lastly, the thesis also demonstrates that common structural optimizations such as store buffers play a major role in determining the overall performance of an HTM implementation, bridging the performance gap between eager and lazy designs.

PREDICTABLE MULTI-PROCESSOR SYSTEM ON CHIP DESIGN FOR MULTIMEDIA APPLICATIONS

By Ahsan Shabbir (a.shabbir@tue.nl)
Advisor: Prof. dr. Henk Corporaal
TU Eindhoven, the Netherlands
November 2011

The design of multimedia systems has become increasingly complex due to consumer requirements and large number of applications. The concurrent execution of these applications causes interference and unpredictability in the performance of these systems. Conventional design techniques use simulation but simulation methods do not scale. In this thesis, predictable components and analysis based techniques have been presented which can generate MPSoCs capable of meeting throughput constraints of multiple applications. A communication assist has been presented as a predictable component that provides guarantees on the transfer of data.

Based on the communication assist, a complete design flow has been presented which can synthesize the generated MPSoC platform onto FPGAs. Further in the thesis, run-time mechanisms are developed to handle dynamic situations. These run-time distributed resource managers are more scalable as compared to existing state of the art resource management schemes.

A TILED CACHE ORGANIZATION

By Darío Suárez Gracia (dario@unizar.es)
Advisors: Prof. Teresa Monreal Arnal and Prof. Víctor Viñals Yúfera
Universidad de Zaragoza, Spain
November 2011

The gap between logic and DRAM speed has widen with technology scaling. As a result, current processors include a complex memory hierarchy to minimize the cost of accesses to main memory. The most common organization comprises several on-chip cache levels. The Last Level Cache (LLC) is optimized for density (size) and the first level cache for latency. As larger LLCs are incorporated, we can note a growing latency gap between them, what we call the on-chip inter-level latency gap. This work takes this fact, recognizes it as a potential problem, and proposes a new cache able to deal with it.

This dissertation proposes Light NUCA (Non-Uniform Cache Architecture), a tiled cache organization made of small caches connected with very specialized Networks-in-Cache. L-NUCA improves performance and reduces energy by capturing temporal locality at a finer granularity than other cache organizations. Further, we have proved with a layout in 90 nm that its regular organization has potential for easy verification and a reduced time-to-market.

With an accurate in-house simulation framework and layout-extracted consumptions, we show that the proposed organization provides benefits in several environments, namely, high-performance uniprocessors, high-performance low-power embedded, and simultaneous multithreading. Since the simulation of simultaneous multithreading workloads is very costly in terms of time, we have proposed a statistics-based mechanism to select representative combinations of benchmarks.

For the high-performance low-power embedded domain, we have extended Light NUCA with several techniques to reduce energy consumption without impacting performance. As the rest of the design, these techniques leverage the Networks-in-Cache, and are very easy to implement. This Light Power NUCA adapts the cache latency to variations in the working set of programs. Nevertheless, its power consumption is independent of the cache hit rate. To provide an adaptive behavior, we have proposed a learning based controller to detect when the cache is not providing blocks and drop them. To further reduce energy, the same mechanism changes the cache access from parallel to serial in the tag and data arrays.

FILTERING DIRECTORY LOOKUPS IN CMPS

By Ana Bosque (anica@unizar.es)
Advisors: Prof. José M. Llabería, Prof. Pablo Ibáñez, Prof. Víctor Viñals
University of Zaragoza, Spain
November 2011

Nowadays, most computer manufacturers offer chip multiprocessors (CMPs) due to the always increasing chip density. These CMPs have a broad range of characteristics, but all of them support the shared memory programming model. As a result, every CMP implements a coherence protocol to keep local caches coherent.

Coherence protocols consume an important fraction of power to determine which coherence action to perform. Specifically, on CMPs with write-through local caches, a shared cache and a directory-based coherence protocol implemented as a duplicate of local caches tags, we have observed that energy is wasted in the directory due to two main reasons.

Firstly, an important fraction of directory lookups are useless, because the target block is not located in any local cache. The power consumed by the directory could be reduce by filtering out useless directory lookups.

Secondly, useful directory lookups (there are local copies of the target block) are performed over target blocks that are shared by a small number of processors. The directory power consumption could be reduced by limiting the directory lookups to only the directory entries that have a copy of the block.

Along this thesis we propose two filtering mechanisms. Each of these mechanisms is focused on one of the problems described above: while our first proposal focuses on reducing number of directory lookups performed, our second proposal aims at reducing the associativity of directory lookups. Several implementations of both filtering approaches have been proposed and evaluated, having all of them a very limited hardware complexity. Our results show that the power consumed by the directory can be reduced as much as 30%.

CONTROL BASED DESIGN OF COMPUTING SYSTEMS

By Martina Maggio (maggio.martina@gmail.com)
Advisors: Prof. Alberto Leva
Politecnico di Milano, Italy
December 2011

Two trends nowadays appear evident: one is a steadily increasing complexity of architectures and applications, the other is an attempt to make systems as self-adaptive as possible, in order to have them maintain certain levels of performance and/or reliability also in varying environmental conditions. More powerful hardware leads designers to devise more advanced functionalities, which in turn increases the overall complexity. More complex systems can be in a larger number of conditions, and present to the outside world a more articulated interface, so that making them capable of adapt themselves is more and more difficult.

This work attempts to start synchronising the use of control-related methods to control computing systems, by fostering the use of the systems and control theory also in the design of computing systems. The remark is initially made that the main problem is that, when a feedback loop is closed around a computing system that is already fully designed and functional, the controlled object includes some "core physical phenomenon", that normally is quite simple, plus a number of programming superstructures that were designed and introduced with little if any knowledge of dynamics. As a consequence of that, the suggestion is here made and motivated to use the systems theory as a common language and design framework, aiming specifically at concentrating the modelling effort on the mentioned core phenomenon, so that anything to control it emerge naturally as the implementation of a controller model in the system-theoretical sense.

With this innovative approach, it is shown that not only many relevant computing system functionalities can be viewed as feedback controllers in nature, but also that adopting that point of view allows for often surprisingly simple solutions, that can be implemented into existing systems and also provide innovative guidelines to streamline the design of future ones.

UPCOMING EVENTS

The 18th IEEE International Symposium on High-Performance Computer Architecture (HPCA-18)
25-26 February 2012, New Orleans, LA, http://www.hpcaconf.org/hpca18

The 17th International Conference on Architectural Support for Programming Languages and Operating Systems (ASPLOS 2012)
3-7 March 2012, London, UK, http://research.microsoft.com/asplos_2012

The International Conference on Compiler Construction (CC 2012)
24 March - 1 April 2012, Tallinn, Estonia, http://conferences.inf.ed.ac.uk/cc2012

The International Symposium on Code Generation and Optimization (CGO 2012)
31 March - 4 April 2012 San Jose, California, USA, http://www.cgo.org/cgo2012/

The IEEE International Symposium on Performance Analysis of Systems and Software, (ISPASS 2012)
1-3 April 2012, New Brunswick, NJ, http://ispass.org/ispass2012/

The European Conference on Computer Systems (EuroSys 2012)
10-13 April 2012, Bern, Switzerland, http://eurosys2012.unibe.ch

The 9th European Dependable Computing Conference (EDCC 2012)
8-11 May 2012, Sibiu, Romania, http://edcc.dependability.org/

The 17th International Conference on Reliable Software Technologies (Ada-Europe 2012)
11-15 June 2012, Stockholm, Sweden, http://www.ada-europe.org/conference2012

The 15th Euromicro Conference on Digital System Design (DSD'12): Architectures, Methods & Tools
5-8 September 2012, Izmir, Turkey, www.univ-valenciennes.fr/dsd2012/

The 26th ACM International Conference on Supercomputing (ICS 2012)
25-29 June 2012, Venice, Italy, http://www.ics-conference.org/

contributions If you are a HiPEAC member and would like to contribute to future HiPEAC newsletters, please contact Per Stenström at per.stenstrom@chalmers.se

HiPEAC info 29
COMPILATION ARCHITECTURE

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

design: www.magelaan.be

THE 7TH INTERNATIONAL SUMMER SCHOOL ON ADVANCED COMPUTER ARCHITECTURE AND COMPILATION FOR HIGH-PERFORMANCE AND EMBEDDED SYSTEMS (ACACES), JULY 8-14, 2012, FIUGGI, ITALY