Hi PEAC info

COMPILATION ARCHITECTURE

info 36

NETWORK OF EXCELLENCE ON HIGH PERFORMANCE AND EMBEDDED ARCHITECTURE AND COMPILATION

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

SEVENTH FRAMEWORK PROGRAMME

WWW.HiPEAC.NET

THE 9TH INTERNATIONAL CONFERENCE ON HIGH PERFORMANCE AND EMBEDDED ARCHITECTURES AND COMPILERS (HIPEAC 2014), 20-22 JANUARY 2014, VIENNA, AUSTRIA

MESSAGE FROM THE HIPEAC COORDINATOR

I hope you have enjoyed a relaxing summer break with your friends and family. I was pleased to read this summer that the European economy is showing its first signs of recovery after several years of little or no growth. Hopefully, this recovery will drive the market for innovative computing products. Unfortunately, a growing market will bring its own challenges too. In August, a report by Mark P. Mills from Digital Power Group made me aware that the digital economy is consuming an increasing amount of energy, already up to 10% of global electricity generation, or 50% more than the energy used by global aviation. The yearly growth of mobile devices and mobile applications turns out to dramatically outpace the power savings by technology. Since most electricity is still produced from coal, the environmental impact of the digital economy will probably become a more serious issue in the near future.

In July, 151 participants enjoyed the annual ACACES summer school in Fiuggi, Italy. This year's technical program received the highest average appreciation score ever. I would like to thank all the instructors for their excellent contributions. Next year's summer school will be the tenth ACACES summer school, and we want to make it the best we ever got! It will be announced in January 2014.

On October 7-9, we organize the fall Computing Systems Week in Tallinn. Tallinn is the capital of Estonia, and its old town belongs to the list of UNESCO World Heritage Sites. As part of this event, we also organize the second HiPEAC Industry Partner Program, in which we bring industry and academia together to think about common challenges. During the event, we will also update the participants on Horizon 2020, and on the calls that are relevant for the HiPEAC community.

In January 2014, there is the HiPEAC Conference, which will take place in Vienna. This is the third year that we outsource the reviewing process for the conference to ACM TACO, and again we have received many high quality submissions for the conference. We already have 12 papers accepted, and 61 papers got an invitation to submit a revised version. The evaluation of the revised versions will be ready by mid-November. All authors of accepted papers will get an invitation to present their work at the conference. Their paper will be carefully copyedited, published in ACM TACO, and indexed in all major publication databases. The conference itself will be a major networking and recruiting event for the computing systems community in Europe. There will be 30 workshops and tutorials taking place during the conference, and all HiPEAC companies and FP7 computing systems projects will be invited to promote their research and activities.

Koen De Bosschere

CONTENT

BARCELONA COMPUTING WEEK: PUMPS 2013

The fourth edition of PUMPS summer school took place in Barcelona during the week of July 8-12, 2013

PUMPS (Programming and tUning Massively Parallel Systems) is a one week course offered to a very competitive selection of one hundred students from all over the world, in order to learn advanced GPU programming techniques from world-renowned experts in the field, Wen-mei Hwu (U. of Illinois), David Kirk (NVIDIA), Nacho Navarro (BSC/UPC) and Isaac Gelado (BSC).

Organized by the Barcelona Supercomputing Center and the Universitat Politècnica de Catalunya under the CUDA Center of Excellence teaching program, and with the sponsorship of NVIDIA and the HiPEAC Network of Excellence, the PUMPS summer school has become a reference in Europe on GPU and HPC programming.

Summer school co-directors Wen-mei Hwu (U. of Illinois) and Mateo Valero (BSC-UPC) decided that this edition should focus on new advanced GPU programming optimization techniques, so before being accepted to PUMPS, all students had to pass an online course with quizzes and laboratory assignments (supervised by Professor Hwu's team) in order to reinforce their knowledge of CUDA.
During the summer school, the Teaching Assistants from University of Illinois and BSC did a great job by challenging the attendees. This one-to-one mentoring during the afternoon labs made a huge difference and provided the students with a high quality learning experience.
The summer school combined theoretical lectures and hands-on labs, covering scatter-to-gather transformations, non-uniform and sparse data, input binning, interfacing with the host, privatization, tridiagonal solvers and numerical considerations, extending computation to multiple GPUs, FORTRAN interoperability and CUDA libraries, as well as including a one-day training by Rosa Badia (BSC/CSIC) on BSC's OmpSs programming model and Paraver analysis tool. The course›s attendees also enjoyed two interesting sessions by David Kirk (NVIDIA) on computational thinking and the Kepler architecture under CUDA 5.0.


PUMPS 2013 family picture

Continuing with the success of previous editions, PUMPS attendees were encouraged to present their work at the traditional poster session and social event that took place on Wednesday afternoon. A jury composed of the summer school instructors and BSC›s senior researchers evaluated their work and agreed to award to three students the Best Poster Award: Thomas Auzinger (Vienna U.T.) for his work "Analytic Rasterization on GPGPUs", Weifeng Liu (U. of Copenhagen) for his work "Load Balanced SpMM on Heterogeneous Systems", and Jose Manuel Molero (U. of Almeria) for his work "A Batched Cholesky Solver for Local RX Anomaly Detection on GPUs".

More information: http://bcw.ac.upc.edu/

Víctor García, BSC, Spain

ACACES 2013 – TRIP REPORT

Attending ACACES summer school for the second time wasn't exactly what I planned for this summer, but it turned out to be the best plan and it was as fun as the first time!
I'm doing my PhD studies in the field of fault tolerance systems at Chalmers University. Even though the main focus of the school is on computer architecture and FPGAs, which are not much related to the research in my Ph.D. studies, I enjoyed knowing more about them and I believe that these courses gave me good insights for my work. Indeed, there was one course related to reliability and fault tolerance systems in both years.

However, I think the main goal of the summer school is to give a chance to meet many different students from different countries and universities, and also some of the best teachers in different fields. The friendly environment and quite long meal times (!) made it easy to talk about different topics, including your work and studies, and it is really joyful to get to know different ideas and get feedback on what you do while you get served delicious Italian cuisine! And this was not only about students; I can say the teachers were really open and friendly to talk to during coffee breaks and meal times. I found these chats really encouraging, and I got many useful comments.
Fiuggi is a cozy town famous for its spa. Walking to the old town in the evening and biking to the lake and just wandering in town feel so pleasant that you feel you want to stay there for more days! At

the end, I would like to thank the organizers of the school, who were really helpful and responsible. More than a hundred students attend each year, from different countries with different preferences and needs, and I could tell that the organizers really care that all students have good conditions. I already miss the good gelato and I really wish to get the chance to attend ACACES again!

Fatemeh Ayatolahi, Chalmers University of Technology

ACACES 2013 family photo

HIGH PERFORMANCE COMMUNICATIONS BY TORUS

TORUS develops and commercializes software solutions for reducing the time-to-results of HPC and Big Data applications, focused on, but not limited to, Java codes.

TORUS, promoted and participated by the HiPEAC members Ramón Doallo, Juan Touriño and Guillermo López (associate), commercializes high-speed communication technologies for accelerating HPC and Big Data applications. TORUS targets financial services and the Internet industry

TORUS is a spin-off of the University of A Coruña (Spain) dedicated to commercializing a set of technologies developed over the last ten years by the Computer Architecture Group (http://gac.udc.es/english/), namely two products, Java Fast Sockets (JFS) and FastMPJ. JFS is a high-performance socket implementation that transparently bypasses TCP/IP, the only communication protocol supported by the Java Virtual Machine. JFS provides ultra low latency communications, between and 10 times faster than standard sockets. More specifically, TCP/IP is replaced by a custom high-speed protocol for localhost communications and high-speed networks (InfiniBand and 10/40 Gigabit Ethernet).

FastMPJ, a Java binding of MPI, is fully portable (100% Java) and able to take advantage of high-speed networks (e.g., InfiniBand). FastMPJ outperforms MPI on NUMA servers while providing quite competitive results on InfiniBand clusters. FastMPJ has been able to scale up to 8000 cores in HPC codes and has boosted the performance of Big Data projects.

TORUS products significantly reduce the response time in Big Data and HPC applications, allowing them to take advantage of high-speed Java communications. TORUS technology has been the subject of investment by the Science Fund of Barrie Foundation, allowing the company to verify their competitive advantage and achieve their first commercial contacts in the financial sector.

In 2013, UK Trade & Investment (part of the UK government) awarded TORUS with the prize of the first edition of the UKTI Spain Technology Competition. The Panel of Judges highlighted: "Torus' innovative technology, commitment to R+D+I and clear strategy for the UK market and internationalisation."

Further documentation and software downloads are available at http://www.torusware.com
Academic licenses are free.

Guillermo López Taboada
University of A Coruña and TORUS, Spain

TORUS receiving UKTI Award (from UK Ambassador in Spain).

TETRACOM: TECHNOLOGY TRANSFER MADE EASY

A new FP7 Coordination Action specifically focused on academia-industry technology transfer

TETRACOM *(Technology Transfer in Computing Systems)* is a new FP7 Coordination Action specifically focused on academia-industry technology transfer (TT). Starting in September 2013, it will stimulate European TT over the duration of 36 months. At the same time, it serves as a pilot project for possible new TT instruments in Horizon 2020. The total project volume is two million Euros. TETRACOM is open to the entire HiPEAC community, and the eight founding partners cordially invite you to participate!

Generally, European academia-to-industry TT in Computing Systems is considered insufficient today. While great scientific results are generated, in particular within EU funded collaborative projects, and impressive TT activities do happen, there is no systematic way of stimulating TT so far at a small to medium scale. Occasionally, outstanding R&D results are commercialized via start-up companies. However, the effort in setting up a company, including all legal and financial matters, is often experienced as a major hurdle. In addition, the majority of promising individual solutions for highly specific problems, which do not justify a start-up on their own, rarely make it into industrial practice. Reasons for this include the following:

- Academic researchers are seldom well connected to relevant industries or they are inexperienced in academia-industry TT scenarios, thereby missing opportunities.
- There is too much complexity and bureaucracy involved in legal negotiations between universities and companies.

- There is insufficient time or budget to bridge the wide gap between a research demonstrator and a real-life industrial product or design flow.

TETRACOM provides a new key instrument to improve this situation, called Technology Transfer Projects (TTPs). TTPs provide a novel and systematic incentive for small to medium scale TT at the European level. As an important support measure, the TTPs will be backed by Technology Transfer Infrastructures (TTIs), such as regular workshops, trainings, and TT consultation services by experts, which will be widely announced. For instance, one of the first events will happen during the HiPEAC CSW on October 8 in Tallinn.

All HiPEAC community members can apply for TTP funding in TETRACOM using an efficient proposal scheme. In short, the "TTP algorithm" works as follows:

1. All TTPs are based on bilateral academia-industry TT partnerships. One academic partner A teams up with one industry partner B, who is interested in taking up a specific technology or IP developed by A for internal use, evaluation, or productization.

2. The total volume of the intended TTP is between 10k-200k Euros, and the total TTP duration is between 3 and 12 months.
3. Partner A, assisted by B, submits a lightweight three-page TTP proposal to TETRACOM, which will be efficiently evaluated by experts according to several well-defined and public criteria.
4. Following a positive evaluation, TETRACOM can provide funding of up to 50% of the total TTP volume. This funding will be received only by partner A, but it will indirectly also benefit partner B, of course.

Three calls for TTP proposals will be issued in TETRACOM over the project duration. On top of this, a lot of exciting TTI events will be organized. Please make sure to add www.tetracom.eu to your bookmarks and stay tuned for the first call for TTP proposals to appear in early 2014. In case of questions, please do not hesitate to contact the project coordinator at leupers@ice.rwth-aachen.de.

Rainer Leupers,
RWTH Aachen University, Germany

Tetracom founding partners

MANY-CORE RESEARCH THRIVES IN SCOTLAND

Glasgow University awarded funding for three new research projects

In a recent funding round, the UK Engineering and Physical Sciences Research Council (EPSRC) committed £5 million to a "Many-core Architectures and Concurrency in Distributed and Embedded Systems" (MACDES) research priority. After a competitive process, EPSRC elected to fund five out of sixteen shortlisted proposals. Of these five successful proposals, three are led by researchers at the University of Glasgow's School of Computing Science. HiPEAC members Phil Trinder and Jeremy Singer lead two of the projects. Phil Trinder said, "It's great to see the concentration of Glaswegian and Scottish Parallelism research paying dividends in this way. Effective working relationships will enable us to capitalise on synergies between these projects."

Overall, this funding outcome is a triumph for Scottish research pooling, since four of the five successful projects involve collaboration among Scottish universities (Edinburgh, Glasgow and Heriot-Watt). The projects range from energy-efficiency improvements to adaptive parallel execution models, all targeting future many-core computing systems.

Jeremy Singer
University of Glasgow, UK

ACM DISTINGUISHES MATEO VALERO AND FABRIZIO GAGLIARDI FOR THEIR ROLE IN THE ADVANCEMENT OF COMPUTING

The award recognises Valero's leadership of initiatives in high-performance computing research and education

Mateo Valero, director of the Barcelona Supercomputing Center, has been recognised with the Distinguished Service Award given by the Association of Computing Machinery (ACM). The award recognises Valero's leadership of initiatives in high-performance computing research and education, including the HiPEAC European Network of Excellence and the Barcelona Supercomputing Center (BSC-CNS). This year's edition has also recognised Fabrizio Gagliardi, previously in Microsoft Research and now at BSC, with a *Presidential Award*.

The award is on the basis of "the value and degree of services to the computing community," and is one of the most internationally prestigious in the computing field. The ACM has distinguished Valero for being a firm advocate for the advancement of computer architecture and high-performance computing research, and in nurturing others to pursue these research fields. Valero has worked on many fronts: making a strategic case for major funding initiatives at the regional, national and EU levels; serving as a founding coordinator of HiPEAC; ensuring that such initiatives benefit

scientific research more broadly; fostering academia-industry cooperation; and hosting numerous visitors and conferences in the region of his institution, Universitat Politècnica de Catalunya – Barcelona Tech. Initiatives in which Valero has played a leading role include: the European Center for Parallelism of Barcelona (CEPBA, 1990), the Catalan Center for Computation and Communications (1995-2000), the CEPBA-IBM Research Institute (CIRI, 2000-2004) and the BCN-CNS. The last of these hosts the MareNostrum supercomputer, which is a national treasure that serves scientific and industrial users and has a worldwide reputation.

With the awards, the ACM wants to recognise the achievements of scientists devoted to computing who have contributed to elevating the role of this discipline as an essential tool to push technological innovation forward.

Gagliardi, as the chair of ACM Europe Council, has sharpened ACM's visibility and its professional and educational activities throughout Europe. He also played a leading role in creating the first Heidelberg Laureate Forum, a collaborative initiative between the

Mateo Valero, director of BSC-CNS (in the center), receiving the Distinguished Service Award. ACM President Vint Cerf to the right and John White, ACM Executive Director, to the left.

computing and math disciplines to bring young researchers together with winners of the highest scientific awards, including the ACM Turing Award. The awards ceremony took place in San Francisco on Saturday 15 June.

BSC'S PROPOSAL FOR HANDLING DATA DEPENDENCES IN TASK BASED PROGRAMMING MODELS INTEGRATED INTO OPENMP 4.0

The integration recognizes BSC's efforts and tradition of innovation in programming models for parallel architectures

Barcelona Supercomputing Center's (BSC) proposal for handling data dependences in task-based programming models has definitively influenced the latest version of OpenMP released in July. Task-to-task synchronization in OpenMP is now supported through the specification of task dependency.

This feature is already available in the open-source OmpSs programming model, developed and distributed by the Spanish supercomputing center (http://pm.bsc.es/ompss). OmpSs has been demonstrated in real applications, including those used in European projects such as Mont-Blanc, DEEP, TEXT, and in the Intel-BSC Exascale lab collaboration.

The integration in the OpenMP standard, the most used programming model for shared memory in High Performance Computing, is recognition of BSC's innovative work in programming models for parallel architectures. BSC has 15 years' experience in developing support for parallel programming models. Its researchers have been involved in OpenMP since the beginning, through cOMPunity, participating in the definition of the tasking model, and lately with the inclusion of task dependences.

Jesus Labarta, Director of BSC Computer Science's department says "BSC proposals for expressing data flow and dependences between tasks allow unprecedented amounts of parallelism to be exposed and exploited under very asynchronous execution models. We are committed to programming models and intelligent runtimes research that is a key factor for future Exascale systems".

BSC's research and development in programming models is performed by a team of more than 20 researchers and students. The topics covered by the team include compiler support, support for heterogeneous architectures, including GPUs and MIC, support for distributed architectures and energy efficiency in Exascale platforms.

ABOUT OPENMP

The OpenMP Application Program Interface (API) is a multi-platform shared-memory parallel programming model for the C, C++ and Fortran programming languages. It is a portable, scalable model that gives shared-memory parallel programmers a simple and flexible interface for developing parallel applications for platforms ranging from multicore systems and SMPs to embedded systems.

Incorporated in 1997, the OpenMP ARB is the non-profit corporation that oversees the OpenMP specification and produces and approves new versions of the specification.

*Further information can be found at www.openmp.org.*

Rosa M. Badia, BSC Team leader of the Computer Sciences Dept., giving a programming models demo at SC12, (USA).

Example of task dependency graph

UMIC BOOTH AT 50TH DESIGN AUTOMATION CONFERENCE

High-level power estimation showcased by RWTH Aachen University

UMIC booth at 50th DAC with NoC power demo; Stefan Schürmans, Prof. Anupam Chattopadhyay, Prof. Rainer Leupers (left to right)

The UMIC research cluster (Ultra High Speed Mobile Information and Communication) of RWTH Aachen University was represented at the 50th Design Automation Conference (DAC), which took place in Austin, Texas during June 2 to 6.

UMIC presented its work in the area of future mobile communication technologies, including further increasing of data rates over wireless links and possible future mobile applications enabled by those higher data rates. Additionally, two demonstrations were shown, in order to give insight into ongoing research at UMIC. Both demos focused on power estimation at a high abstraction level. This topic is important, since low power consumption is a critical concern for new multiprocessor systems on chip (MPSoCs), which are used in mobile phones. High-level design decisions taken early in the design process have a much bigger impact on power consumption than later decisions at lower levels. Thus, it is highly desired to be able to obtain approximate power numbers already at a high level.

The first demonstration showcased the power extension of simulators for application specific instruction set processors (ASIPs). Simulators created from a description in the instruction set description language LISA using Synopsys Processor Designer usually model only functionality and timing. The UMIC researchers made it possible to extend those simulators to also output a power estimate for the processor.

SystemC models of communication architectures found in MPSoCs are typically also limited to functionality and timing, and do not model power consumption. Therefore, the second demonstration was about a methodology to back-annotate available power information to existing SystemC models (see DAC program, session 19.6, Creation of ESL Power Models for Communication Architectures using Automatic Calibration). At the booth, a SystemC simulation of a network on chip (NoC) was running and estimating the NoC power consumption.

Stefan Schürmans
RWTH Aachen University, Germany

BEST PAPER AWARD FOR INVASIC TEAM AT ASAP'13

Symbolic Parallelization of Loop Programs

Professor Teich is honored by receiving the best paper award at the 24th IEEE International Conference on Application-specific Systems, Architectures and Processors (ASAP) in Washington DC, USA, for the contribution on "Symbolic Parallelization of Loop Programs for Massively Parallel Processor Arrays" by Jürgen Teich, Alexandru Tanase, and Frank Hannig.

Jürgen Teich,
Friedrich-Alexander-Universität Erlangen-Nürnberg, Germany

Professor Teich (second from the right) receiving the best paper award.

HIPEAC MANAGEMENT STAFF CHANGE

Vicky Wandels

As of 10 June, Jeroen Borghs has left the HiPEAC Network as Project Management Assistant. Jeroen managed all financial and administrative aspects of HiPEAC. He also helped with the logistics of several activities including the annual summer school, the HiPEAC conference and the computing systems week. The HiPEAC Network thanks Jeroen for his excellent work during the past years and wishes him success in his further career.

Jeroen has been replaced by Vicky Wandels. Vicky holds a Master degree in languages from the University of Ghent and has worked at KULeuven for several years. If you have any questions, do not hesitate to contact her at vicky@hipeac.net.

NEW BOOK: "ENERGY-EFFICIENT COMMUNICATION PROCESSORS: DESIGN AND IMPLEMENTATION FOR EMERGING WIRELESS SYSTEMS"

Robert Fasthuber, Francky Catthoor, Praveen Raghavan, Frederik Naessens

The book presents an integrated template-based design approach for the implementation of efficient communication processors. These processors are needed to enable emerging high-performance mobile wireless communication devices, such as smartphones with significantly improved user experience.

The first part of the book reviews current trends and motivates why future designs need to be jointly optimized for high performance, high energy efficiency, high programmability, high reusability and high technology scaling-friendliness. A state-of-the-art review shows that existing design approaches, e.g. based on the ASIC or ASIP architecture style, optimize only for one or several of these design criteria, but not for all together.

For this reason, the main part of this book proposes a new design approach that can cope with the joint requirements. The core of this design approach is the proposed Domain-Specific Instruction set Processor (DSIP) architecture template. To achieve high efficiency, innovative architectural concepts, such as Software SIMD and distributed instruction buffers, have been included and the template has been designed with future process technologies and algorithm characteristics in mind.

To demonstrate the feasibility and high efficiency of the proposal, the book includes three relevant case studies, i.e. on the design of a MIMO detector, a FIR filter and a FFT, such as are needed for emerging 4G and 60 GHz communication systems. All three implementations are a factor of two to three times more energy efficient than state-of-the-art programmable solutions.

The main text is supplemented with an appendix on semi-custom back-end design, which is an important part of the overall proposed design approach.

Robert Fasthuber, Francky Catthoor, Praveen Raghavan, Frederik Naessens, EPFL, Switzerland / KUL, Belgium / IMEC, Belgium

EU-US COLLABORATION UNDER THE TERAFLUX PROJECT

Roberto Giorgi visited new Teraflux partner University of Delaware in mid August

TERA^FLUX

Dr Roberto Giorgi, Principal Investigator of the EU TERAFLUX project and Professor at the University of Siena, came to visit Dr Guang Gao, Endowed Distinguished Professor and his research group at the University of Delaware, to have a dedicated working week towards their common goal of research for extreme-scale high-performance computing – in the context of the TERAFLUX consortium.

TERAFLUX AND A NEW US PARTNER ACROSS THE ATLANTIC

The Computer Architecture and Parallel Systems Laboratory (CAPSL) led by professor Guang Gao has become a new partner of the European Union-funded TERAFLUX project (http://teraflux.eu). TERAFLUX aims to provide a solid solution from the hardware to the software for tomorrow's large-scale computations. Future large-scale systems require higher energy efficiency – a power reduction of a factor of at least 100 compared with current technology. In addition, exploiting massive parallelism and increasing perfor-mance are both necessary to include new predictive and analytical capabilities that will ideally underpin all mathematical models going forward and make effective use of unprecedented levels of concurrency. Examples range from faster and more complex weather models in order to anticipate catastrophes and take informed decisions under critical conditions, such as tornados and hurricanes, to applications that analyse complex molecular chains that can result in new treatments for cancer or Parkinson, among others.
TERAFLUX is originally composed of ten partners coming from the EU: both from academia (U. of Augsburg, U. of Manchester, U. of Sienna, U. of Cyprus, INRIA, Barcelona Supercomputing Center) and industry (CAPS Entreprise, Hewlett Packard, Microsoft R&D, THALES SA). The TERAFLUX consortium and its members are coordinated by Dr Roberto Giorgi, Associate Professor at the University of Siena, Italy. Since mid-2012, the consortium has expanded to include its first US partner, the University of Delaware.
By joining the TERAFLUX consortium, the University of Delaware helps in building a bridge over the Atlantic between Europe and the United States of America, in order to research new ways of exploiting extreme-scale architectures for the next twenty years.

INITIAL ACHIEVEMENT THROUGH EU/USA COLLABORATION

From the angle of the University of Delaware, several contributions have been made in the area of energy efficiency in modern parallel architectures. The EU/USA collaboration has provided new scalable energy modeling methodologies for modern many-cores. The novel modeling techniques have been the foundation for a methodology that is able to partition an application into several pieces for parallel computation and minimize the energy consumption of the whole amount of work, targeting the operations that are energy hungry. This collaboration has already resulted in several research papers accepted for publication in top conferences and journals (HiPC, LCPC, PACT, Computing Frontiers, WHIST , NPC, etc.).

The collaboration of the University of Delaware and their counterparts in EU through TERAFLUX in the field of data flow is particularly oriented to fine-grained data flow inspired execution models extended to address the increasing amount of resources available in new architectures. Furthermore, new challenges in terms of energy, performance and resiliency have led to proposed self-aware systems and exploring of the major hardware requirements and types of interfaces and information that will be needed for exascale runtime systems. Under this umbrella, several research papers have been published in top conferences and journals (PACT, Euro-Par, COOL chips, Computing Frontiers, ISPA, DFM, etc.).

EDUCATION AND EXCHANGES

This cross-Atlantic collaboration, between the European Union and United States, has benefited several graduate students, post-doctoral researchers and faculty from both sides, allowing interchange of experiences and technology. Professor Gao has visited several institutions across Europe (e.g. Italy, France, England) while Professor Giorgi visited the University of Delaware. Students and Researchers have benefited through paper presentations and interchanges at different levels in Europe and US. The impact is going beyond, through paper presentations in Asia and Oceania (e.g. China, Korea, Japan, India and Australia)

SUMMARY AND OUTLOOK

By joining the TERAFLUX consortium, the University of Delaware is helping to forge lasting ties between the United States and European Union's top research institutions in the field of future supercomputing systems. Such collaboration between Europe and the United States is paramount to the dissemination of ideas and helping to improve the world of supercomputing in general, providing a broader impact to the world in the long run.

Prof. Roberto Giorgi (bottom, left) during his visit at University of Delaware)

FP7 EUROSERVER PROJECT

Project name: EUROSERVER
Coordinator: Yves Durand, CEA-Leti, France
Partners:
CEA-Leti, France
STMicroelectronics, France
ARM Ltd, UK
Eurotech S.p.A., Italy
TU Dresden, Germany
BSC, Spain
FORTH, Greece
Chalmers, Sweden
OnApp Ltd, UK
Start date: September 2013
Duration: 36 months

Data Centres (DCs) are a key resource for innovation and leadership of industry in Europe. They drive the Information Society through hosted cloud applications. To sustain the ever-increasing demand of storing and processing data, DCs need to improve their capabilities and scale in size. With current server technology, however, DC scaling is limited by the IT equipment's density and energy consumption. To keep up with data growth, in the face of power distribution constraints and steadily increasing energy costs, the IT equipment must become dramatically smaller and more efficient in power and energy. This moves the focus of server design and the interests of industry from performance to power/energy efficiency and total cost of ownership (TCO).
The basic components of future servers and their integration into a full system must be reconsidered from the ground up. The processor, the memory hierarchy, I/O, the system interconnects and the systems' software all require fundamental changes to match the application's performance in less space and with drastically lower energy costs.

EUROSERVER is addressing these challenges in a holistic manner; we advocate the use of state-of-the-art low-power ARM processors in a new server system-architecture that uses 3D integration to scale with both the numbers of cores, and the memory and I/O, all managed by new systems software providing transparent system-wide virtualization and efficient resource use by cloud applications. The EUROSERVER prototype will demonstrate how the proposed approach can lead to **10x DC Energy Efficiency by 2020.**

ENERGY-EFFICIENT ARCHITECTURE

EUROSERVER will use 3D stacking to increase the density of computing cores. Placing cores on a silicon board – interposer –, together with I/O and memory will significantly reduce the number of chips and improve fabrication yield. Through effective data communication and processing, it will also improve the inter-chip efficiency providing cross-allocation of memory and I/O resources across all the system, while containing load and traffic within smaller "islands".

LOW-POWER 64-BIT PROCESSOR

ARM processors, which dominate the mobile and embedded markets, have the potential to dramatically change the landscape of DC servers. EUROSERVER will use 64-bit ARM processors ideal for the DC workloads that exhibit a high amount of application concurrency.

EFFICIENT SYSTEM SOFTWARE

EUROSERVER proposes a new architecture for system platform software to allow efficient virtualization and sharing of resources across chiplets. The systems software will efficiently manage shared resources and processors and dynamically assign workloads to the appropriate group of resources, reducing workload interference and achieving high resource utilization with no performance compromises.

REAL APPLICATIONS

EUROSERVER will assess and evaluate the results and potential impact against a set of use-case scenarios including DC, Telecommunications, Transportation and Security that will demonstrate the effectiveness and efficiency of the new design and architecture against typical industry scenarios.
The EUROSERVER consortium brings together world-class industrial and academic leaders in their own fields having broad access to the required technologies and specialized knowledge. EUROSERVER creates the critical mass required to deliver competitive "More than Moore" solutions in Europe and to bring new levels of density and energy-efficient computing worldwide.
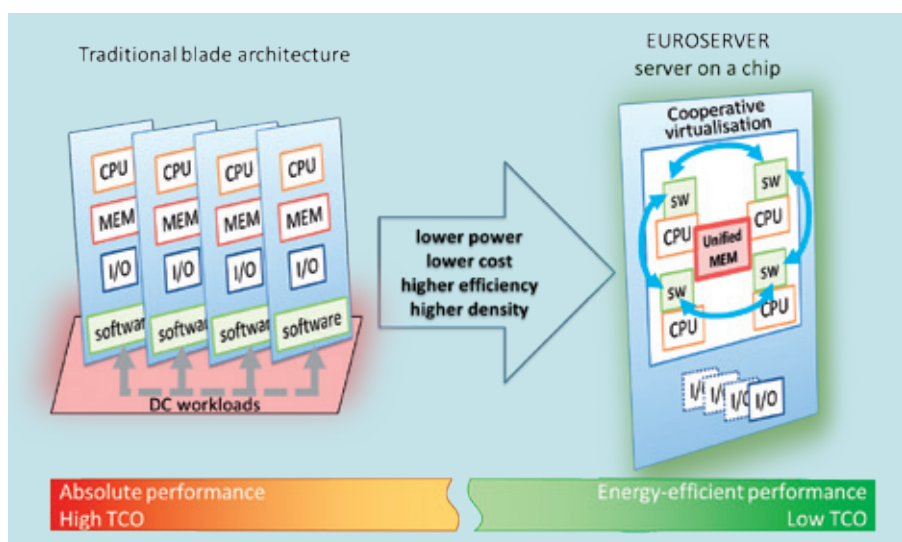
FP7 SCORPIO: SIGNIFICANCE-BASED COMPUTING FOR RELIABILITY AND POWER OPTIMIZATION

A new computing paradigm that exploits uncertainty to design systems that are energy-efficient and scale gracefully under hardware errors

Project name: SCoRPiO (Significance-Based Computing for Reliability and Power Optimization)
Coordinator: CERTH, Greece
Partners:
CERTH, Greece
EPFL, Switzerland
RWTH Aachen University, Germany
The Queen's University of Belfast, UK
IMEC, Belgium
INRIA, France
Start date: June 2013
Duration: 36 months
**Website:** http://www.scorpio-project.eu/

SCoRPiO
Significance-Based Computing for
Reliability and Power Optimization

SCoRPiO is a FET-Open project that seeks to dramatically change the way that next generation integrated circuits are designed and fabricated, by relaxing the requirement that hardware is always correct. It exploits the observation that there are many applications in which errors are *acceptable* by the users. For example, video applications can tolerate small variations in pixel values and audio applications can tolerate slight distortions, provided that they produce an overall acceptable outcome.

SCoRPiO will research methods for characterizing the *significance* of various parts of the program for the quality of the end result, and their tolerance to faults and imprecision. Based on this information, computations and data can then be steered to either low-power and less-reliable, or higher-power and fully-reliable functional

and storage components of the underlying computer platform. In addition, it becomes possible to aggressively reduce the system's power footprint by powering hardware modules even below nominal values.

Just as parallelism became a central aspect of modern computing, having a profound impact on the worlds of software and hardware, in the same vein SCoRPiO has the ambition to elevate *significance* as such a foundational property of next-generation computing systems, and show that it needs to be actively supported by all system layers. The ideas researched in SCoRPiO will allow the semiconductor industry to progress beyond 22nm nodes, but may also turn out to be more generally valid irrespective of the technology and material properties of the underlying hardware.
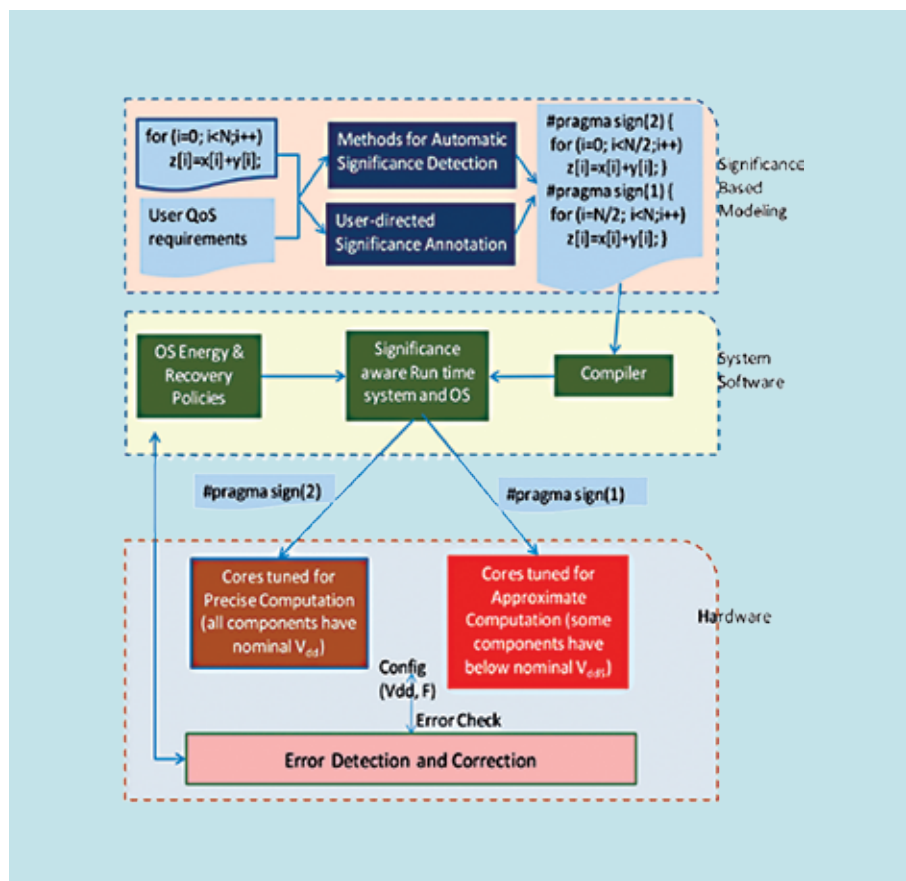
FP7 FET MINECC PROJECT ENTRA

Promoting Energy Transparency to Enable Energy-Efficient Software Development

Project name: ENTRA: Whole-Systems Energy Transparency
Coordinator: Roskilde University, Denmark
Partners:
Roskilde University, Denmark
University of Bristol, UK
XMOS Ltd, UK
IMDEA Software, Spain
Start date: October 2012
Duration: 36 months
Website: http://entraproject.eu/

ENTRA research team

The research project "ENTRA: Whole-Systems Energy Transparency" (grant agreement 318337) promotes energy transparency as the key enabler for the development of energy-efficient software. The ENTRA project, led by Roskilde University, Denmark, is funded by the EU 7th Framework Programme, Future and Emerging Technologies (FET) under the Minimizing Energy Consumption of Computing to the Limit (MINECC) objective. The consortium also includes the University of Bristol (UK), XMOS Ltd (UK) and the IMDEA Software Institute, Madrid (Spain). The ENTRA project runs for three years from 1st October 2012, with a budget of EUR 2.1 million.

REDUCING ENERGY CONSUMPTION OF COMPUTING IS AN INCREASINGLY VITAL DESIGN GOAL.

Energy consumption and the environmental impact of computing technologies have become major global concerns. The growth of energy consumption in cloud computing and Internet traffic is unsustainable at current energy efficiency levels. Furthermore, it is increasingly important to optimize the energy consumption of battery-powered devices, especially in pervasive computing applications.

THERE ARE SIGNIFICANT GAINS FROM OPTIMIZING SOFTWARE FOR ENERGY EFFICIENCY.

While hardware can be designed to save a modest amount of energy, the potential for savings are far greater at the higher levels of abstraction in the system, with the greatest savings expected from energy-efficient software. This is because energy is consumed by the hardware performing computations, but control over the computation ultimately lies in the software and algorithms running on the hardware. Energy-efficient software development requires a deep understanding of how energy is consumed during computation, which is fundamentally at odds with the trend, over the last decades, to de-couple software engineering from the operations of the hardware, in the interests of portability, understandability and software reuse. The ENTRA project aims to close the gap between software and hardware, by making visible how algorithms and their encoding in software impact the energy consumption of a computation.

ENTRA researchers are developing advanced program analysis and energy consumption modelling techniques to obtain predictions of the energy consumption of a computer system early in the software design phase. These predictions help software developers gain a better understanding of how much energy is consumed by the programs they write; they enable the toolchain to optimize code for energy efficiency, and can eventually be used to certify that the energy consumption of the final product is within specified bounds.

Further details about the ENTRA project are available at http://entraproject.eu/

INTERNSHIP REPORT - HARRY WAGSTAFF

System simulator analysis at Freescale Semiconductor

As part of a HiPEAC sponsored industrial internship, I spent three months at Freescale Semiconductor's Network and Multimedia Solutions Group in East Kilbride, Scotland. Freescale's main markets include automotive processors, as well as platforms for network and cellular infrastructure, and this is the second time that this group has hosted an intern from my research group at the University of Edinburgh. My PhD research focuses on high-speed instruction set simulation and tool autogeneration, so my work at Freescale focused on working on and analysing Freescale's internal system simulator.

While much of my own work focuses on core simulation, Freescale's simulator includes a large range of device models, in order to support the complex accelerators and network processing hardware that they provide. The internship began with some engineering work on the simulator as a 'warm-up' task (involving implementing some semihosting-like functionality in the simulator, as well as several bug fixes), followed by an analysis of the simulator, focusing on areas in which the simulator could be improved or extended. This involved drawing on knowledge gained during my research, the abundance of papers published on the subject of system simulation, as well as ideas and requests from potential users of the simulator inside Freescale, in order to extract improvements which make sense in the context of the simulator, and which might make good opportunities for future internships or direct collaboration. I also contributed to a user guide for the simulator from the perspective of an applications engineer, which gave me the opportunity to learn about many of the hardware structures and techniques used in network processing platforms.

This internship was also the first time I have used the Power ISA, as well as several other technologies, which has given me lots of ideas for new areas to explore in my future research. I am grateful to HiPEAC for providing the opportunity to participate in this internship, as well as to Freescale for hosting it.

Harry Wagstaff
University of Edinburgh, Scotland

PhD NEWS

PARALLEL PROGRAMMING MODELS FOR ARCHITECTURAL SYNTHESIS

Muhsen Owaida
University of Thessaly, Greece
Advisor: Assoc. Prof. Nikolaos Bellas
October 2012

Recent trends in computing industry favor heterogeneous platforms to exploit parallelism at multiple granularity levels. Programming models such as OpenCL have been designed targeting parallel heterogeneous systems. However, using FPGA platforms as part of a heterogeneous system requires a tremendous effort to translate OpenCL code into RTL descriptions. This dissertation introduces a methodology to automatically synthesize hardware accelerators from OpenCL applications targeting FPGA platforms. We map OpenCL kernels on FPGA accelerators using architectural templates that explicitly decouple computation from memory communication, dramatically reducing development time.

MYRMICS: A SCALABLE RUNTIME SYSTEM FOR GLOBAL ADDRESS SPACES

Spyros Lyberis
FORTH-ICS, Greece
Advisors: Prof. Dimitrios
S. Nikolopoulos, Prof. Angelos Bilas
July 2013

In this thesis we implement Myrmics, a runtime system for future, heterogeneous, non-coherent processors. Myrmics uses novel, distributed algorithms and policies for hierarchical memory management, dependency analysis and task scheduling. We evaluate Myrmics on a heterogeneous 520-core FPGA prototype, modeled faithfully after current predictions for many-core processors. Our prototype runs code 50,000 times faster than software simulators. Our experimental results show that Myrmics scales well compared with reference, hand-tuned MPI baselines, while automatic parallelization overheads remain modestly low (10-30%). We verify that many of our proposed algorithms and policies are promising.

EVALUATION FRAMEWORK FOR TASK SCHEDULING ALGORITHMS IN DISTRIBUTED RECONFIGURABLE SYSTEMS

M. Faisal Nadeem
Delft University of Technology,
The Netherlands
Advisors: Assoc. Prof. Dr. J.S.S.M.
Wong and Prof.dr. K.L.M. Bertels
(Promoter)
August 2013

This dissertation presents the design of a simulation framework to evaluate the performance of reconfigurable nodes in distributed systems. The framework incorporates partial reconfigurable functionality of the nodes, in which they can execute multiple tasks simultaneously. For resource management, we introduce simple but essential data structures for information update and maintenance of nodes. We also present a generic scheduling algorithm for task distribution, utilizing the partial and full reconfigurability of nodes. For a given set of parameters, we conducted several experiments to generate and evaluate various performance metrics in order to verify the functionality of our framework.

RUN-TIME ADAPTABLE VLIW PROCESSORS -- RESOURCES, PERFORMANCE, POWER CONSUMPTION, AND RELIABILITY TRADE-OFFS

Fakhar Anjam
Delft University of Technology,
The Netherlands
Advisors: Prof. Dr. K.L.M. Bertels
and Assoc. Prof. Dr. J.S.S.M. Wong
August 2013

In this thesis, we propose a dynamically adaptable soft core VLIW processor called ρ-VEX, which is based on the VEX ISA. The processor can adapt itself to different and changing characteristics, e.g., parallelism, among and within applications. This is achieved through parameterization of the hardware design to reduce (area) overhead and maximize flexibility. Examples of parameterization are issue-width, execution units, register file size, interrupt/exceptions support, fault tolerance, and custom operations. These options enable users to trade-off between hardware area/resources, performance, power/energy consumption, and reliability. The processor is available under an academic license.

Conference on Design and Architectures for Signal and Image Processing (DASIP 2013)

8-10 October 2013, Cagliari, Italy   http://www.ecsi.org/dasip

International Symposium on System-on-Chip 2013 (SoC 2013)

23-24 October 2013, Tampere, Finland   http://soc.cs.tut.fi/

ICT 2013 - Create, Connect, Grow

6-8 November 2013, Vilnius, Lithuania   https://ec.europa.eu/digital-agenda/en/ict-2013

Compiler, Architecture and Tools Conference

18-19 November 2014, Haifa, Israel   http://software.intel.com/compilerconf2013

The 46th International Symposium on Microarchitecture (MICRO 2013)

7-11 December 2013, Davis, USA   http://www.microarch.org/micro46/

The 9th HiPEAC Conference

20-22 January 2014, Vienna, Austria   http://hipeac.net/HiPEAC_conferences

19th International Conference on Architectural Support for Programming Languages and Operating Systems (ASPLOS 2014)

1-5 March 2014, Salt Lake City, USA   http://www.cs.utah.edu/asplos14/

The 16th Design, Automation and Test in Europe Conference (DATE 2014)

24-28 March 2014, Dresden, Germany   http://www.date-conference.com

The International Conference on Compiler Construction (CC 2014)

5-13 April 2014, Grenoble, France   http://www.etaps.org

The European Conference on Computer Systems (EUROSYS 2014)

13-16 April 2014, Amsterdam, The Netherlands   http://eurosys2014.vu.nl/

The Design Automation Conference (DAC 51)

1-5 June 2014, San Francisco, USA   http://www.dac.com/

35th Programming Language Design and Implementation Conference (PLDI 2014)

9-11 June 2014, Edinburgh, UK   http://conferences.inf.ed.ac.uk/pldi2014/

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

HiPEAC *info* 36
COMPILATION ARCHITECTURE

design: www.magelaan.be

THE 9TH INTERNATIONAL CONFERENCE ON HIGH PERFORMANCE AND EMBEDDED ARCHITECTURES AND COMPILERS (HIPEAC 2014)
20-22 JANUARY 2014, VIENNA, AUSTRIA