



HiPEACinfo¹⁹

COMPILATION ARCHITECTURE

appears quarterly | July 2009

Network of Excellence on High Performance and Embedded Architecture and Compilation

- 2 Message from the HiPEAC coordinator
- 3 Message from the project officer
- 2 HiPEAC Activity:
 - HiPEAC sets flag at DATE 2009
 - HiPEAC 2010 Conference
 - Static Single-Assignment Form Seminar in France
 - HiPEAC Computing Systems Week in Munich
- 6 Community News:
 - Best Paper Award: Design Automation and Test in Europe (DATE)
 - HiPEACinfo Staff Changes
 - Announcing the Open Source Release of TTA-Based Codesign Environment 1.0
 - HiPEAC Mini-sabbatical Program
- 8 Member Profile:
 - ACE Associated Compiler Experts by
 - Professor Thomas Fahringer, University of Innsbruck, Austria
- 10 HiPEAC Start-ups:
 - Sigasi Enters EDA Market with Public Beta Program
- 10 HiPEAC Students
- 13 PhD News
- 16 Upcoming Events



Welcome to
ACACES 2009
La Mola, Barcelona

www.HiPEAC.net

HiPEAC 2010 Conference
Pisa, Italy, January 25–27

Message from the HiPEAC coordinator

Dear friends,

I am writing this column as the first results for the European elections are being published. As a convinced European, I find it a bit strange to see how some parties are questioning the usefulness of Europe, are arguing against immigration and mobility, and are acting against more European integration, while, at the same time, they are blaming Europe for not supporting the troubled European industry, and for the lack of a common foreign policy. I have hardly heard any opinion about the impact of Europe on research, innovation and competitiveness. Perhaps we are to blame for this, and perhaps we should be making more of an effort to ensure success stories of European research programs are made more visible to the general public and are appraised by them. It is definitely an area where integration actually works and which brings benefits to all participants.



Koen De Bosschere

Last month, we received the final review of HiPEAC1. I am proud to quote the conclusions: "In their report, the reviewers conclude that the project was completely successful. The commission is in agreement with the findings of the review report. The commission considers that the HiPEAC NoE has delivered high-quality work and that it has played a crucial role in building a strong European community in the area of architectures and compilers. The success of the HiPEAC conference, the summer school, and the industrial workshops

are some examples of the important achievements."

This conclusion is a very strong motivator for further increasing our efforts to make HiPEAC2 even more successful than its predecessor.

This issue of the newsletter is known as the summer school issue. After the terrible earthquake in L'Aquila, we had to look for an alternative location. Thanks to the efforts of Mateo Valero, we were able to find an excellent place in the vicinity of

Barcelona. This year we will have a record number of participants at the summer school, which is still one of the very successful and visible events organized by the HiPEAC community. It also marks the start of the summer holiday for the HiPEAC community.

Let me therefore wish all participants of the summer school an exciting summer school experience, and all of you a pleasant and relaxing summer holiday.

Koen De Bosschere

HiPEAC Activity

HiPEAC sets flag at DATE 2009



DATE (Design, Automation & Test in Europe) is the premium event in Europe for electronic system-level design and test in spring every year. This year's event was held in Nice in the week of 20–24 April, which, as usual, attracted thousands of attendees from academia and industry. DATE has two integrated parts: the scientific conference, where technical papers and workshops/tutorials are presented and the exhibition, where companies, universities and European projects showcase

the latest results and products. The venue, Acropolis Centre, was well-equipped to accommodate the two parts close to each other, thereby providing the opportunity for the attendees to quickly "switch contexts" between both and thus make the most out of the large program.

DATE '09 offered a special package for dissemination purposes to



Weihua Sheng (RWTH Aachen University) presenting HiPEAC activities to Alexandre Romana (TI France)

EU-funded projects – HiPEAC used this great opportunity by setting up

Message from the project officer

The deadline for the last Call for Proposals in Computing Systems was in April 2009. For this Call we received 30% more proposals than in the previous one, proving that there is a growing European research community in the area of Computing Systems. The quality of the proposals received was again extremely high: three out of every four proposals received were marked above thresholds. The total grant requested by all the proposals received was over €100m. As a reminder, the available funding is €25m.

In many cases, the proposals addressed more than one of the research areas of the Call. The area that attracted most interest was the area of “*parallelisation and programmability*”. Proposals in this area addressed, among others, issues related to multi-core-oriented software architectures,

automatic identification of parallelism in sequential programs, program-mability breakthroughs for hardware platforms with hundreds of cores, and configurable hardware architectures facilitating parallel programming.

The area of “*methodologies, techniques and tools*” also attracted many proposals, which addressed, among others, static analysis and compiler transformations, virtualisation of computational demanding applications, instruction-set virtualisation, design-space exploration for customisation of parallel systems, and reconfigurable heterogeneous architectures.

Proposals in the area of “*system simulation and analysis*” mainly addressed ways of reducing the simulation time of complex multi-core architectures.

In the area of “*technology implications*”, proposals mainly addressed the challenges of 3-dimensional integration technology for new parallel multi-core platforms.

Over the next couple of months, our Cordis website will provide more details about the selected projects.

Panos Tsarchopoulos
Project Officer



Panos Tsarchopoulos
Panagiotis.Tsarchopoulos@ec.europa.eu

a booth in the exhibition area in order to present its research mission and increase its visibility. The booth preparations were carried out by Ghent University and RWTH Aachen University – the booth was decorated with posters, latest newsletter copies and leaflets. It also used a running video presentation displayed on a large TFT LCD to attract more attention. Fortunately, the HiPEAC booth was located directly across from a main meeting-point, and so even more traffic was generated to the booth during the two coffee breaks. The booth was well-received and well-attended throughout the event. As well as greeting the numerous HiPEAC veteran members, we

also had many new visitors everyday asking for project information, wanting to learn about the research clusters and expressing interest in joining the HiPEAC events.

The HiPEAC booth at DATE gave the project great visibility to a large community. Follow-up communications to spread the word on HiPEAC’s influence further will take place in the near future and hopefully we will get new members onboard. It was a wonderful experience running the booth,



Felix Engel (RWTH Aachen University) at HiPEAC University Booth in DATE

talking to the people, and (after the booth hours) enjoying the breathtaking beauty of the Côte d’Azur!

HiPEAC 2010 Conference

HiPEAC invites its members to the 5th International Conference on High-Performance Embedded Architectures and Compilers which provides a forum for computer and compiler designers in the field of high-performance architecture and compilation for embedded and general-purpose systems, with a special emphasis on cross-cutting research that can be applied to both. The conference aims

to aid the dissemination of advanced scientific knowledge and to promote the forging of international contacts among scientists from academia and industry. HiPEAC 2010 will take place on January 25-27, 2010, in Pisa, Italy.

Submission of papers is currently open. The deadline for abstracts is July 3. Please check <http://www.hipeac.net/conference> for detailed



information about the topics of interest, the rules for paper acceptance, and further important dates. ■

Static Single-Assignment Form Seminar in France

On 27–30 April, 2009, Laboratoire d'Informatique du Parallélisme de Lyon, Saarbrücken University and STMicroelectronics jointly organized a workshop exclusively focussed on Static Single-Assignment form (SSA). This workshop was sponsored by LIP, HiPEAC, STMicroelectronics and INRIA, and took place in Autrans (France) in the beautiful Vercors range of the Alps. This event was a unique opportunity for several worldwide compiler experts to engage in a couple of days of presentations, discussions and active work sessions.

In fact, STMicroelectronics, the INRIA/LIP Comsys laboratory and Saarbrücken University have a long-standing research relationship in the domain of highly optimizing compilers. Several important advances in the field of SSA have been achieved by these groups:

First, with PSI-SSA (cf. "Efficient Static Single-Assignment Form for Predication", by F. de Ferrière and A. Stouthinin, published at MICRO'34 2001), a specific form of SSA invented for predicated instruction sets;

Then in register coalescing, (cf. "On the Complexity of Register Coalescing, by F. Bouchez, A. Darte and F. Rastello, best paper at CGO 2007), which allowed better understanding of the algorithms for eliminating variables copies and for exhibiting their complexity;

In variable liveness checking, (cf. "Fast Liveness Checking for SSA-Form Programs", by B. Boissinot, S. Hack, D. Grund, B. Dupont de Dinechin, F. Rastello, best paper at CGO 2008), with an algorithm. Used for checking if a variable is in use in a given region of a program, of $O(1)$ complexity;

In out-of-SSA translation, (cf. "Revisiting Out-of-SSA For Correctness, Efficiency and Speed", by B. Boissinot, A. Darte, B. Dupont de Dinechin, C. Guillon, F. Rastello, best paper at CGO 2009), with a complete revisit of out-of-SSA algorithm, more accurate, more robust, twice as fast and consuming ten times less working memory.

With the current trend for virtual machines and processor virtualization (Java, Android, .NET, Javascript, etc) new challenges exist for SSA representation. SSA algorithms have not only to be correct and robust, but also very efficient because the compilation process runs on the target system - hence the idea of Fabrice Rastello (INRIA Comsys) to organize a 3-day seminar dedicated to SSA.

The way a program is represented by a compiler has a large influence on the efficiency and effectiveness of the compiler. Static Single-Assignment form is widely used in modern compilers even at the code-generation level, since it allows for simple, yet efficient optimizations and analyses. Nowadays, we see compilers emerging that are completely based on SSA. Thus, SSA will play an even more important role in the future of compilation.



Seminar attendees

(slides available at: <http://www.prog.uni-saarland.de/ssasem/>)

HiPEAC Computing Systems Week

June 2-4 2009, Infineon, Munich, Germany



The seminar brought 55 compiler researchers and practitioners together from all over the world: Austria, Canada, Germany, France, United Kingdom, and United States were the most represented countries. People such as Keith Cooper, Vivek Sarkar, Keshav Pingali, or Kenneth Zadeck were present. Some of the main contributors of compilers such as Firm, GCC, HotSpot, LAO, LLVM, Mono, or Open64 were also present.

This is the first time a workshop exclusively focussed on SSA has been held. 28 presentations were given, covering a large spectrum of topics on SSA. The seven sessions that structured the workshop dealt with: (1) history, properties, motivations; (2) semantics, construc-

tion, and destruction; (3) programming languages, such as scripting languages; (4) memory, especially alias analysis; (5) SSA-based optimizations and code generation such as register allocation and instruction selection; (6) analyses and optimizations such as liveness checking, code plagiarism, constant propagation, if-conversion and handling of predication; (7) SSA-based compilers with Firm, GCC, and HotSpot.

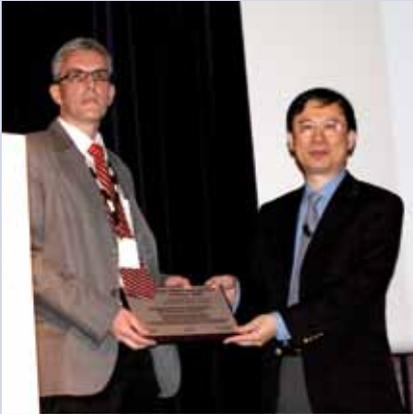
The primary goal was to exchange ideas and foster the development of SSA. This goal has been fully achieved: the program was very comprehensive, and the participants usually interacted until late in the evening. It was difficult to find time for the so-called "book session":

in fact, the plan is to write a textbook dealing with SSA-based compilers. 16 participants decided to collectively write the first book fully related to compilers under SSA form. It will contain four main parts: (1) Vanilla SSA; (2) Extensions; (3) Analysis; (4) Machine-dependent optimizations and code generation. The plan is to complete the authoring by mid-2010.

Fabrice Rastello
Laboratoire d'informatique du parallélisme (LIP), Lyon, France
Fabrice.Rastello@ens-lyon.fr

Christian Bertin
STMicroelectronics, Grenoble, France
Christian.Bertin@st.com

Best Paper Award: Design Automation and Test in Europe (DATE)



Prof. Rainer Leupers at award ceremony

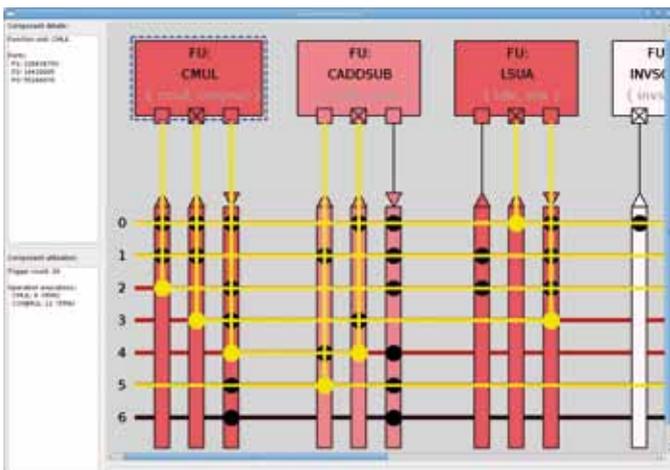
Manuel Hohenauer and Rainer Leupers from RWTH Aachen University, Balpreet Singh from NXP and Gerrit Bette from ACE received a Best Paper Award for their contribution to DATE 2008 entitled "Retargetable Code Optimization for Predicated Execution".

DATE is the leading conference on Design Automation in Europe. In combination with the parallel technical exhibition a total of more than 2500 people attended the 2009 event in Nice on the beautiful Côte d'Azur. The spectrum of topics was very broad. The DATE parallel tracks ranged from VLSI design over architecture design to software tools and novel applications. The core conference program was accompanied by a variety of panel sessions, tutorials the day before the main conference and a number of workshops afterwards. This unique combination of events makes DATE a very attractive event for researchers and industry members alike.

The paper honoured at DATE 2009 presents a fully retargetable predicated execution support integrated into ACE's

CoSy compiler framework. Predicated execution (PE) is a processor feature often found in modern embedded processors, where the execution of instructions is dependent on a Boolean predicate. This allows control dependencies to be efficiently converted to data dependencies, thereby increasing the ILP available in the instruction stream. The main challenges the paper addresses are to concisely describe the peculiarities of a specific PE implementation and to efficiently determine the regions of the code where to optimally perform if-conversion. The net effect is a significant increase in code quality for many embedded applications. The results have subsequently been produced by ACE and have been released as part of the CoSy retargetable compiler generation framework. ■

Announcing the Open Source Release of TTA-Based Codesign Environment 1.0



A toolkit for customizable low power processors design

Energy efficiency plays an increasingly important role in embedded and mobile devices. TTA-based Codesign Environment (TCE) is a software toolkit for customizable low power processors being developed at Tampere University

of Technology. The toolset provides a complete codesign flow from C programs down to synthesizable VHDL and parallel program binaries.

The processor template used in TCE is based on the Transport Triggered Architecture (TTA), which originated at Delft University of Technology in the 1990s. TTA is a modular, statically scheduled processor paradigm resembling VLIWs, and it provides a high degree of flexibility for the designer in choosing the boundary between hardware and software. The architecture places very few restrictions

on functional units, so it is easy to integrate custom units with multiple inputs and outputs, long latencies, etc in a TTA processor.

All the tools in TCE, such as the compiler and the processor simulator are runtime retargetable by means of an architecture definition file. Custom operations/special function units can be used from C code with minimal effort. Other processor customization points include register files, function units, the set of supported operations by each function unit, and the interconnection network.

The compiler of TCE is based on LLVM 2.5 and supports compiling ANSI C/ISO C99 to a wide range of TTA architectures. The register allocator, instruction selector and other target-dependent

HiPEACinfo Staff Changes

The HiPEAC Newsletter, one of HiPEAC's channels for spreading excellence, has been produced at RWTH Aachen University for the past 1.5 years. As most of the readers and all contributors know, Prof. Rainer Leupers, assisted by Jeronimo Castrillon, has been taking care of editing the past five issues of the Newsletter. As of this issue, Jeronimo has transferred his duties to a fresh HiPEAC member, Anastasia Stulova. The HiPEAC community thanks Jeronimo for his excellent work and is pleased to intro-



duce Anastasia to the future contributors. From now on, it will be Anastasia who will be interacting with contributors and getting the best out of them in order to maintain the high quality and friendly style of the Newsletter. ■

HiPEAC Mini-sabbatical Program

HiPEAC has a program of mini-sabbatical visits for HiPEAC faculty members as part of its overall efforts to maintain a high quality research program and to facilitate research collaborations. We want to stimulate short sabbatical leaves for senior researchers and professors. Mini-sabbatical visits are typical stays of one month or longer at another member company or academic site, or even at a non-HiPEAC institution that is looking to join or collaborate somehow with HiPEAC. The goal of the mini-sabbatical visits is to stimulate collaboration with the aim of coordinating or refocusing the research portfolio of the two institutions involved. The candidate for a sabbatical leave and his/her host have to prepare a common sabbatical project, including a budget, which has to be approved by the Steering

Committee. BSC is in charge of promoting and managing this program. Mini-sabbatical application calls are open all-year long. The goal is to have five to ten mini-sabbatical visits per year. Reimbursement is done through a simple per diem rate plus transportation costs.

Two mini-sabbaticals have already been granted in 2009.

Prof. Rainer Leupers, from RWTH Aachen University, has spent some weeks at ACE Associated Compiler Experts by, Amsterdam, The Netherlands, working on parallel compilation for MPSoC, including e.g. reconstruction of C code from the compiler IR. The visit has also enabled intensive consulting on future developments and market perspectives

in the Electronic Design Automation context as well as joint software demonstrators on how to employ CoSy for efficient MPSoC software design.

As already reported in the April Newsletter, Enrique Torres, Assistant Professor of the Computer Science and Systems Engineering Department at the University of Zaragoza, Spain, spent his sabbatical leave for study and research at the University of California at Berkeley, at the International Computer Science Institute (ICSI), collaborating with Professor Krste Asanovic, which was partially funded by this program. He is working on scalable cache coherence protocols in the Parlab and RAMP projects.

Mateo Valero, BSC ■

compiler phases adapt to the target architecture dynamically, i.e. without requiring recompilation of the tool-chain. The TCE compiler backend implements an instruction scheduler and TTA-specific optimizations such as software bypassing.

The TCE processor simulator supports compiled simulation for fast design space exploration, and a slower interpretive-style engine with debugging features and graphical user interface. The simulator is accurate at instruction cycle level and is designed to be used in automated test benches by means of its

Tcl script interpreter console. In addition to simulation, processors generated by the TCE tools can also be mapped to FPGA, for example the DE2 evaluation board from Altera.

TTAs are well-suited for applications with plenty of instruction-level parallelism and relatively static control flow, e.g. multimedia and DSP. TCE has so far been applied to video decoding (variable length code decoding, inverse transform, motion compensation), baseband processing (QR decomposition, FIR filtering, Viterbi, and others), and graphics rendering.

Version 1.0 of the toolkit is now available under the MIT open source license and can be downloaded at <http://tce.cs.tut.fi>. You are encouraged to download TCE and experiment with it. We welcome feedback on the tools and are happy to help in case there are any problems.

Links: <http://tce.cs.tut.fi>
<http://llvm.org>

Contact Information:
pertti.kellomaki@tut.fi ■

ACE Associated Compiler Experts bv



Going Dutch – Coffee & Compilation in Amsterdam

The summer is here. Aficionados of compiler technology are invited to visit ACE to compare notes on multi-core and other challenges over coffee in Amsterdam.

Adjacent to the Central Station in A'dam you will find the collection of computer scientists and compiler engineers that comprise ACE Associated Compiler Experts. For 35 years, we have been in the systems software business having taken on everything from the complete control system for a particle accelerator through to multiprocessor operating systems. Today ACE is synonymous with CoSy, its compiler development system.

Next Generation Compilation System

CoSy emerged from an ESPRIT project where the design brief was a no-holds-barred next generation compilation system with the high performance computing market in mind. The focus of the project was to produce a highly scalable and robust compiler engineering environment rather than specific analyses and optimisations.

Parallel Culture

Conceived to run analysis & optimisation engines concurrently on massively parallel architectures, primarily to generate optimising compilers for highly parallel systems, CoSy was well-ahead of its time. Even today, we hardly make use of its ability to run engines in parallel and of the fine-grain locking mechanisms that were originally built onto the central IR access functions to support this. It is only in the last couple of years that users have experimented with more elaborate compilation flows incorporating speculative, iterative, profile



feedback and machine-learning elements. These are features which are designed into the heart of CoSy.

Similarly, CoSy's extensible IR (using a SDL specification that can be distributed across different engines) took users some time to get to grips with but is now proving invaluable in the multi-core domain as engineering starts to catch up with computer science.

Change of Tack

Just as the technology was being productised, the HPC market went into hibernation. Fortunately, Philips stepped into the role of Dutch uncle with an ambitious VLIW/DSP design that was a natural target for CoSy. Subsequently, the DSP market and unusual architectures proved fertile soil, with CoSy being used for numerous processors - over 100 have been targeted. Over the last decade, ACE has established a position as the leading supplier of compiler technology to the silicon & design industry.

ESL

In search of the next challenge, ACE forged a partnership with ESL pioneer CoWare. Following on from research conducted at RWTH, CoWare formulated a new engineering paradigm enabling system designers to generate the core software tools, including ISSs required for hardware-software co-design. The joint brief

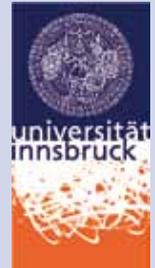
was to add a compiler to the tool-chain to support high-level application code. Within a few months, the teams had a flow generating validating, optimising C compilers from an augmented LISA-based architecture specification. The resulting product, C Compiler Designer, is used extensively in both industry and academia where it is proving invaluable for design, research and teaching. It has recently been released via Europractice to the research community.

To MultiCore, Reconfigurable Systems and beyond

ACE continues to raise the bar, helping the community apply compiler technology to multi-core and reconfigurable systems working not only with large operations but also innovative start-ups and academia. If you have some particularly thorny issues you need solving, or want a bouncing board, or are interested in a demanding HiPEAC internship, then A'dam is as good a place as any.

Joseph van Vlijmen
ACE Associated Compiler Experts bv,
De Ruyterkade 113, 1011AB
Amsterdam, The Netherlands,
tel: +31 20 6646416,
email: info@ace.nl

Professor Thomas Fahringer, University of Innsbruck, Austria



Since the end of the nineties, only a few efforts in the area of software for high-performance computing have been funded on a European scale. It was mainly national projects and initiatives that tried to keep pace with the many research groups in the USA and Japan that greatly flourished as a result of continuous funding on a large scale. Only recently, through the introduction of multi-core processors as a mainstream processor technology, is high performance computing finally gaining considerable momentum in Europe again, with HiPEAC as the primary network of excellence for high-performance computing research funded by the European Commission.

In the last ten years, Austria has managed to maintain two important national parallel processing initiatives. Firstly, the Austrian Center for Parallel Processing (ACPC), which involves most research groups working on algorithms, models and software for parallel computing, is comparable to a national centre of excellence. Secondly, the AURORA priority research program on advanced models, applications and software systems for high performance computing supported by the Austrian Science Fund. This Fund consolidates all major research groups in Vienna with a special focus on high-productivity languages, compiler and tool development for both cluster computing and distributed high performance computing.

It is a great honour to have been invited by the HiPEAC steering committee to become a member of HiPEAC. My name is Thomas Fahringer. I have been working for more than ten years at the University of Vienna participating in the above-mentioned national initiatives. During this time I was leading a research group that developed tools to support the development of mostly

High Performance Fortran (HPF), MPI and OpenMP parallel programs. We have crafted performance prediction tools that examined the performance behaviour of parallel programs without actually running them on a target architecture. A whole range of performance instrumentation, measurement and analysis tools have been created to measure performance data at the binary code level which was related back to the input source code. A special strength of our approach was interpretation of performance data. In contrast to many existing tools which drowned a user or compiler in vast amounts of performance data, we provided technology that can interpret performance data and refer a user to those code sections that require his attention. We also created an HPF debugger, which at that time was a brave undertaking as high-level languages such as HPF had not been widely accepted. Instead most programmers preferred to code a low level of abstraction in order to avoid losing control over vital strategic optimization decisions. Later on, we extended some of our work towards parallel Java codes with an interesting side product that lead to the first language to describe performance problems for a wide range of languages and parallel programming paradigms, which was done as part of EU-funded APART working.

In 2003 I joined the Institute of Computer Science at the University of Innsbruck where I founded the Distributed and Parallel Systems group. We are currently working on a programming, analysis, and optimization environment for scalable parallel programs with increased performance and power-efficiency for homogeneous and heterogeneous many-core on-chip computing systems. For this environment, we work on an extension of OpenMP to unleash the

potential of parallelism for wide classes of applications in the domain of industry, business and science. We explore novel static and dynamic program and system analysis to support optimization and to enable scalable performance even for computing systems with a large number of cores. The underlying runtime environment incorporates dynamic information of the system behaviour to detect parallelism and to optimize performance and power consumption. Of particular interest are trade-offs between performance and energy consumption as part of an optimization system which has not yet been sufficiently resolved and which has an impact far beyond multi-core parallel processing

My group is very interested in contributing to the HiPEAC research clusters "programming models and operating system", "design methodology and tools" as well as "adaptive compilation". We look forward to attending many HiPEAC activities and events to exchange new ideas on programming models, compiler technologies and tools for high-performance computing. Moreover, we are seeking new collaborations with HiPEAC members from industry and academia.



Thomas Fahringer
Institut für Informatik
Gruppe Verteilte und Parallele Systeme
Universität Innsbruck
<http://www.dps.uibk.ac.at>

Sigasi Enters EDA Market with Public Beta Program

Sigasi, an early stage EDA company, recently announced the launch of a Public Beta Program for Sigasi HDT, an Intelligent Development Environment (IDE) for VHDL. Sigasi HDT offers a smart and comprehensive design platform for VHDL. Building upon the widely accepted Eclipse platform, it contains an ultra-fast VHDL parser and compiler which run transparently in the background. The tool fully understands the complete design in terms of VHDL concepts, allowing the digital designer to make modifications faster and smarter.

"This technology enables a wide range of innovative features, such as intelligent navigation, instant error reporting, code completion and code refactoring," said Sigasi's CTO Hendrik Eeckhaut.

Public Beta Program

Sigasi HDT has already been extensively tested by early adopters. Through the Public Beta program, Sigasi will build a community of hardware designers, who will take advantage of modern development techniques. To participate in the program, users should visit <http://www.sigasi.com/signup>.

Business Benefits

Sigasi HDT significantly increases design productivity for both occasional and experienced VHDL designers by helping them to write, inspect and reuse their designs in an intuitive way. "Our technology is inspired by the most advanced software development environments, but it is unique in the hardware design world," said Philippe

Faes, Sigasi's CEO. "Operations that may take hours when done by hand can be done in minutes with our tool."

About Sigasi

Sigasi was founded by two former PhD students from HiPEAC member Ghent University and is an early stage Electronic Design Automation (EDA) company. Sigasi focuses on the creation of an Intelligent Development Environment (IDE) for Hardware Description Languages (HDLs). Sigasi HDT drastically increases the productivity of hardware designers by helping them to write, inspect and reuse their designs in an intuitive way. The company is headquartered in Ghent, Belgium. For further information, please visit <http://www.sigasi.com>. ■

Collaboration Grant Report - Ricardo Quisiant



My name is Ricardo Quisiant and I am a PhD student at the University of Malaga, Spain. I am reaching the end of my second year of research in the topic of Transactional Memory

(TM). Specifically, I am focusing on signatures in Hardware Transactional Memory (HTM). Recording the read and write sets of transactions in a time-and-space-efficient way is crucial for conflict detection in HTM systems. Therefore, signatures are implemented as Bloom filters which exhibit such efficient features, although they introduce false conflicts that may affect the performance of transactional applications. Hence, enhancement of HTM signatures is a hot research topic.

Last winter I was awarded a three-month collaboration grant funded by the HiPEAC Network of Excellence to collaborate on Transactional Memory

with host Per Stenstrom at Chalmers University of Technology, Goteborg, Sweden.

At Chalmers, we focused on enhancing conflict detection in TM systems from the coherence protocol point of view.

I was working on the GEMS simulator at the University of Wisconsin, Madison. Such a simulator implements a version of LogTM-SE, a Transactional Memory proposal of the Multifacet group. We used STAMP benchmark suite as workloads on simulations since this suite is designed for Transactional Memory research and includes a wide range of

Collaboration Grant Report - Francesco Regazzoni



Security is a fundamental requirement for modern embedded systems. Mathematically strong cryptographic algorithms are insufficient due to the advent of side channel attacks, which exploit weaknesses in the underlying hardware platform rather than directly attacking the algorithm itself. To improve the security of embedded devices, protected logic styles have been proposed as an alternative to

CMOS; however, the area and power consumption of protected logic styles are both significantly larger than for CMOS. As a consequence, they should only be used sparingly.

In this HiPEAC collaboration grant work, conducted at EPF Lausanne, ALaRI and UCL Louvain-la-Neuve, we have enabled the automatic partition of cryptographic applications into protected and non-protected regions. The design flow we have developed is based on standard CAD tools and allows for the synthesis and place-and-route of such hybrid designs. The flow is integrated into a simulation and evaluation environment to quantify the security achieved on a sound basis. Furthermore, we have exploited our design flow to augment an embedded processor, realized in CMOS, with custom instruction set extensions. Such extensions are realized in a protected logic and are designed with security and performance as the primary objectives.

Using MCML logic as a case study, we have explored different partitions of cryptographic algorithms between protected and unprotected logic. Our experiments illustrate the trade-off between the type and amount of application-level functionality implemented in protected logic and the level of security achieved by the design, paving the way for automatic design optimizations based on security.

A complete description of our approach can be found in the reference paper: Francesco Regazzoni, Alessandro Cevrero, François-Xavier Standaert, Stephane Badel, Theo Kluter, Philip Brisk, Yusuf Leblebici, Paolo Ienne, A Design Flow and Evaluation Framework for DPA-resistant Instruction Set Extensions. Accepted for publication at Workshop on Cryptographic Hardware and Embedded Systems 2009 - CHES 2009.

applications with emphasis on those with long-running transactions and large read and write sets.

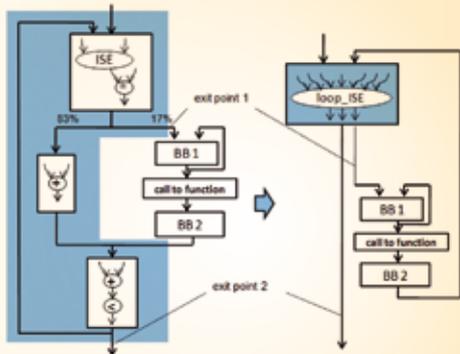
Our simulations showed that some applications spend a significant part of the time aborting or stalling transactions due to conflicts, specially those with large transactions, so there were opportunities for optimization. Therefore, in order to remove the time of transactions due to abortions and stalls, we proposed a way to release isolation of conflicting data so that other transactions are prevented from stalling or aborting. However, we must ensure atomicity and isolation in order to follow the transactional memory correctness cri-

teria. Hence, the proposed isolation release may be allowed only in some well-defined cases to make it correct and invisible for the user.

These kind of transactional memory system changes involve modifications in the cache coherence protocol and other critical parts of the transactional memory system. Cache coherence protocols are complex finite state machines. GEMS provides a MESI filter directory protocol with huge transition tables for L1 and L2 caches, with lots of transition states and actions, so modifications are to be made carefully. The host's experience in this topic is essential for the research.

Several activities were performed during my stay at Chalmers University. The host set up weekly meetings in which we discussed the work done during the week. I also prepared a talk about my research at the University of Malaga. Finally, we set up a special meeting which was attended by the host, his PhD students and my advisor Oscar Plata to reinforce the collaboration and to discuss future plans. Punctual meetings will be arranged in the near future and we hope to establish a long term and prolific collaboration effort beyond the scope of this research.

Collaboration Grant Report – Marcela Zuluaga



multi-cycle ISEs

I am a PhD student at the University of Edinburgh in the Compiler and Architecture Design (CARd) group. I am part of the PASTA project, which aims to automate the design and optimization of customizable embedded processors. The core funding of the PASTA project is from EPSRC (<http://groups.inf.ed.ac.uk/pasta>). My research activities within the project are oriented towards automating the micro-architecture synthesis of processors that exploit instruction-level parallelism through instruction set extensions (ISEs). Last year, I received a HiPEAC collaboration grant that gave me the wonderful opportunity to visit Paolo lenne and Philip Brisk at EPFL (Switzerland) in order to develop a shared research plan. It was an invaluable experience to be able to work with a different group, in a different environment and to exchange ideas and resources. The work carried out in this collaboration focused on studying the viability of control-flow inclusion to support pipelining in ISEs and led to a paper that will appear in the IEEE Symposium on Application Specific Processors (SASP'09) in July.

State-of-the-art ISEs are multi-cycle and may read their data from the memory subsystem in addition to the register file. These ISEs may be pipelined in order to increase their throughput. Pipelining divides the circuit into several execution stages,

allowing it to operate concurrently on different inputs. To exploit a hardware pipeline, several ISEs must be emitted in consecutive cycles; however, typical program streams rarely contain consecutive invocations of the same ISE that would permit this type of execution. This requires the use of ISEs that cover entire loop bodies. In the cases where an ISE can effectively cover all of the computations of the loop body, there is still one limiting factor - several operations relating to the loop itself must be done in software. In particular, the loop counter must be incremented compared with the maximum loop count, and a conditional branch that determines whether the loop continues must all execute in software.

This will require that the processor issues, in most cases, a few instructions per iteration to facilitate the loop; this, in effect, may cancel most or all of the benefit from the pipelined Customized Unit (CU). To address this concern, the software operations described above are integrated into a single complex instruction that controls the iteration of the loop.

Furthermore, we have identified another limiting factor - loops often contain structures that cannot be included in a single ISE without introducing control dependencies. These structures include multiple control flow paths, multiple exits, inner loops and calls to functions that cannot be inlined.

In these cases, unimportant paths with high-resource usage can prohibit the optimization of the execution of more important paths. To mitigate this problem and further expose instruction-level parallelism, we propose ISEs that support loops whose bodies form hyperblocks. To facilitate correct execution in the presence of multiple exit points, the CU returns the address of the correct exit point



to the processor. A loop ISE has itself as the default destination, which is achieved by implicitly issuing the same instruction within the CU. The remaining destinations are the hyperblock exits, including the eventual completion of the loop.

In short, we propose an innovative method to create multi-cycle ISEs that are executed as hardware pipelines that comprise complete loops. This approach is justifiable because most programs spend the bulk of their runtime in a few deeply nested loops, and this is particularly true in embedded applications. These ISEs broaden the scope of instruction-level parallelism and obtain higher speed ups compared to traditional ISEs, primarily through pipelining, the exploitation of spatial parallelism, and reducing the overhead of control flow statements and branches.

As a follow up to this work, we will continue our collaboration over the next few months in order to automate the process of identifying the proposed type of ISEs. ■

Methodologies and CAD Tools Targeting 2D/3D Reconfigurable Architectures

By Kostas Siozios
(ksiop@ee.duth.gr)
Prof. Dimitrios Soudris
(currently in National Technical
Univ. of Athens, Greece)
Democritus University of Thrace,
Greece
November 2008

This PhD deals with the development of novel methodologies for hardware/software co-design, targeting to more efficient reconfigurable architectures. More specifically, during my thesis a novel reconfigurable architecture was designed in full-custom approach. This device exhibits remarkable power/energy savings compared to existing approaches, without impact on delay and silicon area. An equivalent important problem is the development of supporting algorithms and CAD tools that fully employ the features introduced by the target FPGA. Due to this, a design framework targeting to FPGAs (named MEANDER – public available at <http://proteas.microlab.ntua.gr>) was also developed. The framework full-filling both the needs

of experienced designers by providing practical answers to state-of-the-art problems (e.g. logic synthesis, P&R, configuration), as well as to novice designers by providing a simple and consistent set of tools.

Since the application's delay and energy consumption in FPGAs are interconnection driven, the design of an advanced routing network was also studied. Moreover, as the interconnection resources are not fully utilized over FPGAs, we conclude that a more careful design is required. In order to resolve this problem, a novel methodology for designing general-purpose heterogeneous interconnection structures was proposed.

The power/energy consumption of FPGAs, among others affects the on-chip temperature. This problem is even more crucial than the corresponding one from ASICs, due to the increased power density of FPGAs. In order to face this problem, during this thesis, a novel methodology that manages the on-chip temperature is

proposed. This methodology guarantees to reduce the maximal temperature values and to distribute uniformly the power sources over the whole architecture.

Another potential solution to interconnection problem is the usage of three dimensional (3D) integration. However, up to now there are only a few available CAD tools (with limited features) that quantify the potential gains of employing such a new design approach. Due to this, the last issue that is discussed in this PhD thesis affects the development of a novel design methodology, as well as the supporting CAD tools, targeting to 3D FPGAs.

The results of my PhD work are presented on: 4 book chapters, 8 journals, 32 conference papers and the 6 PhD Forums. Moreover, two design contest awards for hardware platform (at ASP-DAC 2005) and software framework (VLSI 2005) are won.

High-Performance Decimal Floating Point Units

By Alvaro Vazquez
(alvaro.vazquez@usc.es)
Prof. Elisardo Antelo
University of Santiago de
Compostela, Spain
March 2009

Current financial, e-commerce and user-oriented applications make intensive use of decimal data. This provides an attractive opportunity for the microprocessor manufacturers to include dedicated DFUs (decimal floating-point units) in their new high-end processors. This interest is supported by the recent revision of

the IEEE 754-2008 standard for floating point, which incorporates a specification for decimal arithmetic.

In this context, this PhD thesis presents the research and design of new methods and high-performance architectures for decimal fixed and floating point hardware units to improve the evaluation of the basic decimal arithmetic operations; i.e., addition, subtraction, multiplication, fused multiply-addition and division. We have studied the use of new algorithms, decimal codings and methods to improve the performance and

efficiency of the resultant architectures. With the application of these concepts, the proposed designs are very competitive with other proposals from both academia and industry.

A further contribution is focused on improving the reliability of computations. A new method for sum error checking is applied to BCD addition/subtraction, which presents a reduced hardware complexity when compared with other solutions used in current microprocessors.

Cache Architecture for Wire-Delay Dominated CMP Systems

By Marco Solinas
 (marco.solinas@iet.unipi.it)
Prof. Cosimo Antonio Prete,
Ing. Pierfrancesco Foglia
Università di Pisa, Italy
May 2009

Increasing on-chip wire delay and growing off-chip miss latency are two key challenges in designing large Last Level Cache (LLC) CMP caches. Currently, some CMPs use a shared LLC cache to maximize cache capacity and minimize off-chip misses. Others use private LLC caches, replicating data to limit the delay from slow on-chip wires and minimize cache access time. Ideally, to improve performance for a wide variety of workloads, CMPs prefer both the capacity of a shared cache and the access latency of private caches. In this context, NUCA caches have proven to be effective in tolerating wire delay effects while maintaining a huge on-chip storage capacity, thanks to their sub-banked organization. When adopted in a CMP system, a NUCA usually represents the LLC shared among all cores, while the

other levels of the cache hierarchy are private and have to be kept coherent among themselves. As the communication infrastructure of NUCAs is a Network-on-chip (NoC), all the nodes in the system communicate via message passing: consequently, a directory-based protocol is the most suitable solution for managing the coherence of private cache levels.

The first part of this thesis focuses on the design of directory-based coherence protocols that are particularly suitable for NUCA caches, by proposing to adopt write-invalidate protocols based on a non-blocking and distributed directory. In particular, it investigates the choice of the coherence strategy (MESI and MOESI) with respect to the whole system topology (i.e. the relative position of cores and cache banks) in S-NUCA based CMP systems. Results discussed in this thesis demonstrate that choosing between MESI and MOESI has not had a significant impact on the overall performance, as a consequence of i) the reduced directory latencies and ii) the low number of block transfers

between two different private caches. Instead, topology variation has a greater impact on system behaviours due to latency variations, which depend on the access pattern to NUCA banks of the running application.

The second part of this thesis proposes a novel block migration mechanism for D-NUCA based CMP systems, designed to avoid particular race conditions that can arise due to multiple traffic sources. The proposed protocol minimizes the number of off-chip accesses while avoiding the need of a centralized directory structure. Results show that the adoption of migration strongly reduces the NUCA response time, with respect to S-NUCA. However, overall performance still depends on topology and application characteristics. This thesis demonstrates that a CMP configuration in which all the cores are plugged to the same side of the D-NUCA is the topology that always takes advantage from block migration, with respect to a configuration in which cpus are plugged to two different NUCA sides.

A Technique for Reducing Power Consumption of Wire Delay Tolerant Cache Memories

By Alessandro Bardine
 (alessandro.bardine@iet.unipi.it)
Prof. C.A. Prete, Prof P. Foglia,
Prof. P. Stenstrom
Università di Pisa, Italy
May 2009

NUCA caches (Non-Uniform Cache Architectures) are large on-chip cache memories that are designed to hide wire delay effects typical of current and future generation nanoscale processors. Thanks to the high number of independently accessible banks of which they are composed, they exhibit high hit rates while keeping the access latency low. Proposed designs for such caches are Static NUCA (S-NUCA), in which data are statically allocated to the cache

banks, and Dynamic NUCA (D-NUCA), in which data may reside in different banks, and a migration mechanism is introduced to better tolerate wire delay effects. The two architectures allow different performances to be achieved by acting on architectural parameters and data management policies, at the cost of different balances between static and dynamic power consumption and energy dissipation.

In this thesis, we develop an energy model for NUCA caches and we characterize such balances by presenting an evaluation of performance and energy consumption of conventional UCA (Uniform Cache Architectures) and of Static and Dynamic NUCA caches. Results indicate that the migration of

data contributes to increased dynamic energy consumption in D-NUCA caches with respect to the other considered cache architectures. However, the higher IPC achieved by D-NUCA caches and the consequent reduced execution time enables static energy to be saved, which, similar to the other considered designs, dominates their power/energy balance.

In the second part of the thesis, we propose the "Way Adaptable D-NUCA cache": a micro architectural technique to reduce the static power consumption of a D-NUCA cache by dynamically adapting the number of active (i.e. powered-on) ways to the need of the running application. The proposed technique leverages the data

migration mechanism and the fact that the distribution of hits across the ways of a D-NUCA cache varies across applications as well as across different execution phases within the same application. On a regular basis, a prediction algorithm measures the cache

ways usage and adapts the number of the active ways to the current needs. The experimental evaluation of the proposed technique shows that it is possible to reduce the average energy consumption of a D-NUCA cache more than 30% while introducing a slight

3% performance degradation. The work is completed by a proposal for a methodology to calculate the parameters on which the prediction algorithm relies and by an evaluation of the low sensitiveness of the technique to those parameters.

Heterogeneity-Awareness in Multithreaded Multicore Processors

By Carmelo Acosta
(cacosta@ac.upc.edu)

**Prof. Mateo Valero, Alex Ramirez,
Francisco J. Cazorla**
**Universitat Politècnica de
Catalunya, Barcelona, Spain**
June 2009

In this PhD dissertation we analyze in depth the inherent heterogeneity present in software behaviour. We identify the main issues and sources of this heterogeneity, which hamper most of the state-of-the-art processor designs from obtaining their maximum potential. Hence, the heterogeneity in software turns most of the state-of-the-art processors, which may be found both in our desktop workstations and laptops – com-

monly called general-purpose processors – into overdesigned ones. This means they have much more hardware resources than really needed to execute the software running on them. This fact would not represent a main problem if we were not concerned about the additional power consumption involved in software computation.

The final goal of this PhD dissertation consists of assigning to each portion of software the exact amount of hardware resources really needed to fully exploit its maximal potential without consuming more energy than is strictly needed; i.e., obtaining complexity-effective executions using the inherent heterogeneity in software behaviour as a steering indi-

cator. Thus, we start to analyze in depth the heterogeneous behaviour of the software run on top of general-purpose processors and then we match it on top of a heterogeneously distributed hardware, which explicitly exploits heterogeneous hardware requirements. Only by being heterogeneity-aware in software, and by appropriately matching this software heterogeneity on top of hardware heterogeneity, may we effectively obtain better processor designs. The results of such a heterogeneity-awareness are both lower energy consumption and processor designs able to exploit the peak potential of the software running on them.

Predictable Embedded Multiprocessor Architecture for Streaming Applications

By Arno Moonen
(A.J.M.Moonen@tue.nl)
**Prof.dr.ir. R.H.J.M. Otten,
Prof.dr. H. Corporaal**
**Eindhoven University of
Technology, Netherlands**
June 2009

The focus of this thesis is on embedded media systems that execute applications from the application domain car infotainment. These applications, which we refer to as jobs, typically fall into the class of streaming, i.e. they process on a stream of data. The jobs are executed on heterogeneous multiprocessor platforms, for performance and power efficiency

reasons. Most of these jobs have firm real-time requirements, like throughput and end-to-end latency. Car-infotainment systems become increasingly more complex due to an increase in the supported number of jobs and an increase in resource sharing. Therefore, it is hard to verify for each job that the real-time requirements are satisfied. To reduce the verification effort, we elaborate on an architecture for a predictable system from which we can verify at design time that the job's throughput and end-to-end latency requirements are satisfied.

This thesis introduces a network-

based multiprocessor system that is predictable.

This is achieved by starting with an architecture where processors have private local memories and execute tasks in a static order, so that the uncertainty in the temporal behaviour is minimised. As an interconnect, we use a network that supports guaranteed communication services so that it is guaranteed that data is delivered in time. The architecture is extended with shared local memories, run-time scheduling of tasks, and a memory hierarchy.

Dataflow modelling and analysis techniques are used for verification because

they allow for cyclic data dependencies that influence the job's performance. We show how to construct a dataflow model from a job that is mapped onto our predictable multiprocessor platforms. This dataflow model takes into account the computation of tasks,

communication between tasks, buffer capacities, and scheduling of shared resources. The job's throughput and end-to-end latency bounds are derived from a self-timed execution of the dataflow graph, by making use of existing dataflow-analysis techniques. It is shown that the derived bounds

are tight, e.g. for our channel equaliser job, the accuracy of the derived throughput bound is within 10.1%. Furthermore, it is shown that the dataflow modelling and analysis techniques can be used despite the use of shared memories, run-time scheduling of tasks, and caches.

Upcoming Events

International Conference on Parallel Computing (ParCo2009)

September 1-4 2009, Lyon, France, <http://www.parco.org/>



A Mini-Symposium on Parallel Computing with FPGAs (ParaFPGA2009)

September 1-4 2009, Lyon, France, www.elis.ugent.be/parafpga



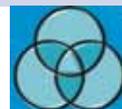
ENISA-FORTH Summer School on Network & Information Security (NIS'09)

September 14-18 2009, Crete, Greece <http://www.nis-summer-school.eu/>



XXVII International Conference on Computer Design 2009 (ICCD 2009)

October 4-7 2009, Resort at Squaw Creek, Lake Tahoe, California, <http://www.iccd-conference.org>



The International Symposium on System-on-Chip 2009 (SOC 2009)

October 5-7, 2009, Tampere, Finland, <http://soc.cs.tut.fi/>



IEEE Workshop on Signal Processing Systems (SiPS 2009)

October 7-9, 2009, Tampere, Finland <http://www.sips09.org/>



22nd ACM Symposium on Operating Systems Principles (SOSP 09)

October 11-14, 2009, Big Sky Resort, Big Sky, Montana, <http://www.sigops.org/sosp/sosp09/>



The IASTED International Conference on Modelling, Simulation and Identification (MSI 2009)

October 12 -14, 2009, Beijing, China, <http://www.iasted.org/conferences/home-659.html>

The 5th ACM/IEEE Symposium on Architectures for Networking and Communications Systems (ANCS 2009)

October 19-20, 2009, Princeton, New Jersey, USA, <http://www.ancsconf.org>



The 5th International Conference on High Performance and Embedded Architectures and Compilers (HiPEAC 2010)

January 25-27, 2010, Pisa, Italy, <http://www.hipeac.net/conference>



23rd International Conference on Architecture of Computing Systems (ARCS 2010)

February 22 - 25, 2010, Hannover, Germany, <http://www.arcs2010.de>



Contributions

If you are a HiPEAC member and would like to contribute to future HiPEAC newsletters, please contact Rainer Leupers at leupers@iss.rwth-aachen.de