Information Society Technologies

HiPEAC info 7

COMPILATION ARCHITECTURE

appears quarterly | July 2006

Network of Excellence on High Performance Embedded Architectures and Compilers

2 Message from the HiPEAC coordinator

2 Message from the project officer

4 Steering Committee News

4 ACACES 2006

5 Announce: A C++ data dependence graph library for optimizing compilation

6 HiPEAC Partner: Italy

8 ATMI : a microprocessor temperature model

9 Challenges in Embedded Computing

11 Transactions on High-performance Embedded Architecture and Compilers

11 PhD News

12 Upcoming events

Welcome to ACACES 2006

www.hipeac.net

HiPEAC 2007: Ghent, Belgium, January 28-30, 2007

Message from the HiPEAC coordinator

Mateo Valero
Coordinator
UPC Barcelona
mateo@ac.upc.edu

Dear colleagues,

An old proverb says that the best way to predict the future is to create it. At HiPEAC we aim to create a bright common future for our community; thus, we constantly review our procedures and tools. This quarter's novelties are the first edition of the industry workshops, conceived as a way to foster collaboration between the academy and industry, and the first issue of the HiPEAC Journal.

The purpose of the HiPEAC industry workshops is to let academic researchers in Europe present their work on topics of interest to companies, so that European industry can benefit from European academic research. In order to stimulate industry attendance, companies define the topics of the workshop, select the contributions to be presented, and host the workshop at a location close to one of their sites. Industry workshops are open to all European academic researchers, both HiPEAC and non HiPEAC members.

The first industry workshop took place on May 11th at INRIA in Grenoble, France, close to the STMicroelectronics main site. Philips will hold the next industry workshop in Eindhoven, in October 2006. Our periodic cluster meetings take place in the same week, in order to ease the attendance and participation of industry members.

Furthermore, as one of the Grenoble sessions, HiPEAC offered a tutorial on middle-end and back-end program manipulation in GCC. HiPEAC supports GCC as a compiler platform for research and development in compilation for high-performance and embedded systems. The tutorial aimed to bootstrap new research and development involving GCC and to help researchers already involved in GCC-related projects. The target audience was both the compilation/architecture researcher/engineer from industry or academia.

HiPEAC actively promotes industry-academy relationships by funding several company internships. This mechanism will allow PhDs to directly target company research groups in their domain. Each company lists and updates every year the research topics for which they are seeking interns. The first call took place on March 1st, for which 38 applications were received. As a result, seven students will be sent for periods of three to six months to ARM,

Message from the project officer

New projects in Embedded Reconfigurable Computing

The Fifth IST Call, which closed on September 2005, invited proposals for IPs (Integrated Projects), and STREPs (Specific Targeted Research Projects) aiming at research on design methods, programming models and compilation tools for reconfigurable architectures. As a result of the Call, four new projects are about to start: ACOTES, ANDRES, HARTES and MILEPOST. The total cost of these projects is estimated at 27 M€ and will be supported by the European Commission with about 17 M€.

HARTES: Holistic Approach for Real Time Embedded Systems
This IP will develop a methodology and a tool-chain to support the embedded system design flow, from high level to efficient modular implementation on a reconfigurable heterogeneous system. The approach is a general purpose tool chain that takes as input applications written in commercially available languages and produces - semi automatically - a "best" mapping onto heterogeneous reconfigurable hardware. Project results will be validated via audio and video applications in the automotive sector.

MILEPOST: MachIne Learning for Embedded PrOgramS opTimization
This STREP will research compiler technology that can automatically learn how to best optimise programs for reconfigurable heterogeneous multi-core embedded processors by employing machine learning approaches developed in artificial intelligence. If success-

Mercè Griera-I-Fisa (Merce.Griera-i-Fisa@cec.eu.int)

IBM, Infineon and STMicroelectronics.

The first issue of HiPEAC's new journal "Transactions on HiPEAC" has been put together. It features a leading article by Sir Maurice Wilkes along with five articles that are extended versions of the five papers that received the highest reviewing scores at the HiPEAC 2005 conference. The second issue, a special issue on compiler optimizations, is underway. I encourage you all to submit your best work to the journal. Directions can be found under the "journal link" at HiPEAC webpage.

Continuing with its activities, HiPEAC announced last May its fourth call for clusters. About 700.000 € has been distributed among members to date. The new Programme Committee is working on the winter HiPEAC Conference, scheduled for January, 2007 in Ghent, Belgium.

When you receive this newsletter, ACACES 2006, the HiPEAC Summer School, will be about to start in L'Aquila, Italy. As last year, ACACES 2006 has strived to offer a high quality programme, from which 200 students will benefit. HiPEAC provides 60 scholarships to selected students. I hope to have the pleasure of shaking hands with many of you there and of hearing your comments and suggestions so that ACACES can exceed its own expectations in the coming years.

As a child in my hometown, Alfamén, I used to daydream about how the future would be. Today, many years later, I like to think that we have contributed a bit to create this future, and that we continue to do so as HiPEAC members. As always, what really matters is not so much where we stand, but rather in what direction we are moving.

ful, it will dramatically reduce the time to market of reconfigurable systems.

ACOTES: Advanced COmpiler Technologies for Embedded Streaming

This STREP will address the complexity and cost of programming future embedded architectures by implementing program transformation techniques in compilation tools. This will enable efficient programming of applications that process massive streams of data on low-cost, highly parallel embedded architectures. The target application areas are video processing and advanced radio communications for consumer systems.

ANDRES: ANalysis and Design of run-time REconfigurable, heterogeneous Systems

Currently no methodology exists to seamlessly specify, simulate, synthesize and verify heterogeneous embedded hardware/software systems, because each domain comes with its own computational models, languages and design tools. The prime outcome of this STREP will be a seamless design flow that will provide a holistic system approach. Verifications will be possible at design time and therefore long, costly and time consuming system integration reiterations will be avoided.

Many of the partners in these four new projects are part of the HiPEAC community, therefore you may have direct links to some or all of them. Details of each project will be soon available at http://cordis.europa.eu/ist/embedded/index.html

Steering committee

- **Industry Internships**
  Seven Ph.D. students from HiPEAC member institutions were selected for a three to six month internship at HiPEAC industry members:

  - **IBM: Increasing ease-of-use in programming of Cell Broadband Engines**
    Mauricio Alvarez, UPC, Spain
    Friman Sánchez Castaño, UPC, Spain
  - **INFINEON:**
    Florian Kluge, University of Augsburg, Germany
  - **ARM:**
    Anton Lokhmotov, University of Cambridge, United Kingdom
    Marios Kleanthous, University of Cyprus, Cyprus
  - **STMicro: GCC4 front-end for.NET**
    Ricardo Fernández Pascual, University Politecnica de Valencia, Spain
  - **STMicro: Static Single Assignment (SSA) form in an Embedded Compiler**
    Nikola Puzovic, University of Siena, Italy

- **HiPEAC Summer School**
  - **60 scholarships provided to attend ACACES 2006**
  - **ACACES 2007 will take place in L'Aquila, July 15-21, 2007**

- **HiPEAC Conference**
  - **The list of workshops/tutorials for HiPEAC 2007 has been approved:**
    1. Workshop on Statistical and Machine Learning approaches applied to Architectures and Compilation (SMART)
    2. Workshop on Interconnection Network Architectures: On-Chip, Multi-Chip
    3. Workshop on Reconfigurable Computing
    4. Tutorial on UNISIM Simulation Environment
    5. Tutorial on GCC
    6. Tutorial on Microthreading
  - **HiPEAC 2008 will take place in Sweden. General Chair: Per Stenström**

- **Cluster meeting**
  - **Next industry Workshop in Eindhoven, October 16-17, 2006**
  - **Funding for participation at cluster meetings of non-members available** ■

HiPEAC Activity

ACACES 2006: Second International summer school on advanced computer architecture and compilation for embedded systems

July 23 to July 29, 2006, L'Aquila, Italy

This year we received 273 applications from 38 countries (88 applicants also attended ACACES 2005). Since the steering committee decided after ACACES 2005 to limit the size of the summer school to 200 participants, we could not admit all applicants. This was a tough decision, but important to guarantee a good summer school experience. The steering committee wanted to have as many nationalities as possible in order to make it a truly international event, and gave special attention to female participants, industry and senior researchers. The remaining places were distributed among the HiPEAC countries according to their number of remaining applicants. The 60 HiPEAC grants were distributed in proportion to the participation from the different HiPEAC countries.

About 70 particpants will present their work at the summer school during the poster session. The abstracts will again be published by Academia Press in the poster proceedings.

The summer school will as before be a great networking event where several HiPEAC clusters and affiliated projects will meet. Furthermore the program committee of the HiPEAC conference and the HiPEAC pre-review will also take place in the same week.

The massive interest in the ACACES summer school clearly shows that there is a real need for this type of event. The HiPEAC network of excellence is strongly committed to further develop the ACACES summer school into a major event in our community.

Koen De Bosschere
Summer school organizer ■

In the spotlight

A C++ data dependence graph library for optimizing compilation

We are pleased to announce the release of an object oriented data dependence graph library (built on top of the LEDA library). Our graph library is for people wanting to make quick, robust and modular implementations of code optimization techniques for basic blocks and simple loops (modelled by regular mono-dimensional data dependences). We are willing to extend it to model non-perfectly nested loops (multi-dimensional data dependences using algebraic polyhedrons).

At present, we manage directed acyclic graphs (DAGs) for basic blocks and cyclic graphs for innermost loops. The user is able to take advantage of many standard algorithms for graphs and general abstract data structures. Also, some well known algorithms on data dependence graphs are implemented (such as computing critical cycles and so on). It is also possible to configure the library for different instruction set architectures.

A great amount of usual algorithms are implemented by LEDA that we can use inside our DDG library. The DDG library allows the direct use of basic graph algorithms (sorts algorithms, strongly connected components, bipartite components, transitive closure and reduction), shortest path algorithms, flow algorithms (max flow, min cut max flow, min cut), maximum cardinality matchings in bipartite and general graphs, stable matching, minimum spanning trees, Euler tours, algorithms for planar graphs, graph drawing algorithms and graph morphism algorithms. There are also many implemented algorithms usually used for instruction scheduling in basic blocks and loops. For DAGs, the DDG library implements shortest and longest paths between any pair of nodes, Dilworth decomposition and maximal antichain. It is also possible to compute the register saturation of a DAG, which is the exact maximal register pressure of any instructions schedule (see [1]). For

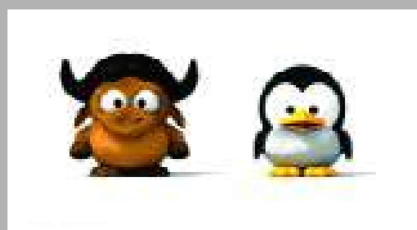

simple loops, the DDG library allows the computation of the critical cycle (even null cycles are detected) using the method with the best known alorithmic complexity.

For more details, please consult the online documentation and download URL at:
http://www.prism.uvsq.fr/~touati/sw/DDG

References

[1] Sid-Ahmed-Ali Touati. Register Saturation in Instruction Level Parallelism. International Journal of Parallel Programming, Springer-Verlag, Volume 33, Issue 4, August 2005. ■

HiPEAC Partner

Italy

The story of Italian computer science started in Pisa, Rome and Milan, in an academic environment. While Rome and Milan acquired foreign computers, Pisa was developing its own systems. In 1955, the CSCE (Centro Studi Calcolatrici Elettroniche) was founded by the University of Pisa, on the suggestion of the Nobel Prize winner Enrico Fermi. In such a context, the CEP (Calcolatrice Elettronica Pisana) was designed. It was the first wholly Italian "stored-program computer" project, and was centred on a microprogrammed machine working with the ForTran language. Also in Pisa, in 1959, a research group at Olivetti labs designed and built the ELEA computer, the first fully transistorized computer. The transistor era saw the foundation of the firm SGS, which later evolved in ST Microelectronics; Federico Faggin started his career at SGS, and later went on to direct the project that built the 1st microprocessor at Intel (Intel 4004, 1971).

As for the present, research capabilities in embedded systems are provided by professorships at computer science and computer and information engineering Departments in more than 40 Universities and research centers.

In addition, many national and international companies have settled in Italy over the years, including research and development centers operating in sectors involving the utilization of computer science and embedded systems. A few examples of such companies and their relevant sectors include:

- Automotive, mechanical and process engineering: Fiat, CRF, Ferrari, Piaggio, Magneti Marelli, Finmeccanica, Pirelli, Siemens Automation;
- TLC and microwave systems: Marconi Communications, Siemens Communications, Siemens Microwave Division, WAS, IDS, Telecom Italia, T-Lab, Ericsson, Motorola;
- Manufacturers of computer components and devices: Texas Instruments, STMicroelectronics, Atmel, IBM, Eurotech;
- Aerospace: Agusta-Westland, AerMacchi, Alenia, Galileo Avionics.

*Prof. Alfio Andronico (on the left) programming the CEP.**

Besides, networked SMEs, working in automation, Robotics and medicine, collaborate in product development, contributing to the 40% of exportations in the embedded systems market.
All these companies actively collaborate with Universities and research centers, as shown by the many joint research projects.

The University of Pisa, one of the most ancient universities in the world (established in 1343), has a long tradition of studies related to electronics and computer science and engineering. It established its MS degree in Electronic Engineering in 1961, and the MS degree in Computer Science in 1969 (the first in Italy). Since 1994 it has offered an MS Degree in Computer Engineering. Thanks to the University, many start-ups and SMEs actively utilizing computers and embedded systems exist all around Pisa, most operating in the fields of process control, mechatronics, robotics and medicine.

The Hipeac Group in Pisa.

The Department of Information Engineering of the University of Pisa (Computer, Electronics and Telecommunications Engineering), **DIIEIT**, consists of 97 professors and 60 PhD students and post-doctoral fellows. Among them, 19 professors and 9 PhD students are involved in research and development of embedded systems.

A view of Pisa from the Arno River

The department was founded in the early '60s. Industrial cooperation includes: STMicroelectronics, Marconi Communications, Telecom Italia, Cisco, Nokia, Alenia, Vodafone, Piaggio, Siemens, Alcatel, Philips, Starcore DSP, Microsoft, and Atmel. Other research institutions cooperating with the Department are the CNR, ASI, ESA, NATO, INFN,CERN, and the Italian Navy. Inside the Department, together with STMicroelectronics, there exists the ST LAB, involved in radio-frequency, monolithic microwave integrated circuits and MEMS research.

Inside the Department, **the Computer Architecture Research Group** has been active for 20 years, and is involved in a variety of European, Italian and local research projects, funded by public and private sources. The research interests of the group concern architectures, optimization techniques and development environments for embedded systems; development environments for parallel and distributed systems; high performance multi-core systems; coherence protocols for cache memories and non conventional cache architectures; design tools for hardware-software codesign; and methodologies and tools for ensuring usability.

As for the activities, the group has worked on multiprocessor systems, participating in the design and realization of a multiprocessor (MUTEAM, supported by CNR), the development of a kernel and a hardware and software debugger for the MARA multiprocessors (built by Selenia – Rome, now Finmeccanica), and introduced cache memories and a signaling subsystem in the M100 multiprocessor developed for Olivetti. More recently, the group has designed a Cartographic, ARM based, single chip multiprocessor in the framework of the SPP/ESPRIT project.
Starting from this multiprocessor experience, the group has focused its interests on coherence protocols and process migration in multiprocessor architectures devoted to the execution of general purpose and commercial workloads, developing a coherence protocol (PSCR) able to reduce the effects of process migration on coherence traffic. As part of its activities, the group built the simulation environment needed to perform the evaluations; the resulting simulation platform,

*Archivio fotografico dell'Università di Pisa

Trace Factory, based on a hybrid approach for generating workload, has been utilized in many universities for research and teaching activities. Currently the group is working on wire delay effects on cache memories, developing ad-hoc architectures and also considering power consumption issues.

The activities in embedded systems have focused on development environments, cache optimization techniques and simulators. The group's simulator for ARM processors has been included in the VLSI Technology, San Josè, CA (now Philips) ChARM/Jumpstart product; the code rearrangement technique minimizing cache misses, developed by the group has been utilized in the VLSI Technology, Sophia Antipolis, CCO product.

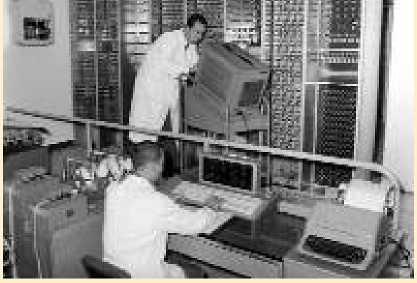

*The CEP**

Since 2005, the Computer Architecture Research Group has been collaborating with IMT. IMT Lucca is an International PhD school with interdisciplinary research areas. The Computer Science and Engineering area performs research on networking systems, global computing and embedded technologies also ensuring usability for final users.

People: Cosimo Antonio Prete, Alessandro Bardine, Alessio Bechini, Luca Fanucci, Pierfrancesco Foglia, Francesco Panicucci, Marco Solinas, Michele Zanda.

Website: http://www.iet.unipi.it
Contact: hipeac@garga.iet.unipi.it

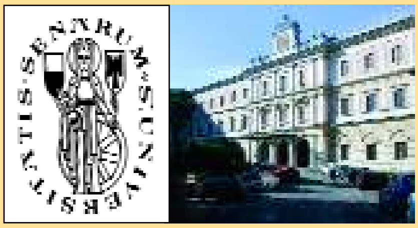

The faculty

University of Siena

Although founded only 12 years ago, the **Faculty of Engineering** of Siena ranks first in Italy in the area of Industrial and Information Engineering, as reported by an independent national evaluation conducted by our Ministry. Our Faculty relies on a single Department: **the Department of Information Engineering**, which consists of 45 professors and 42 PhD students and post-doctoral fellows. Among them, 8 professors and 6 PhD students are involved in research and development of embedded systems and related issues from the Computer Architecture and Engineering perspective. Prof. Alfio Andronico is among the founders of our Faculty: he gave his contribution for the first diagnostic software and machine simulator of the CEP.

In 1992, the University of Siena recognized a Honoris Causa Laurea to Konrad Zuse for his design of the first electronic computer in 1938.

The Department of Information Engineering has many research collaborations with Italian and foreign high-tech companies such as: STMicroelectronics, Telecom Italia, Marconi Communications, Siemens. Other Italian and European institutions cooperating with us on research projects are: MIUR, CNR, CNIT, ASI, start-up companies, founded by ex-students and faculty members.

Our current research activities address the processor and co-processor architecture, the design of memory hierarchies for both constrained and high-performance systems, the optimization of application and system-level software for the exploitation of architectural resources via compile and runtime approaches. More in details, our main research activities focus on the following topics:

- performance evaluation of wide-range of applications including embedded system targets,
- implementation of security mechanisms based on system-wide cryptographic hardware: circuits resistant to side-channel attacks, certification of hardware and software components of whole system,
- accelerators, CPU cores, and DSPs for special-purpose systems,
- high-level design of memory hierarchy and cache memory for embedded and general-purpose DSPs,
- system-level design and optimizations for embedded systems,
- feedback-directed compiler optimizations,
- low-power design of pervasive and mobile devices and applications.
- analysis of problems related to the interconnection of multi-core systems: coherence protocols, consistency protocols.

Our research group uses to research in these domains in a "vertical" way, from the implementation level of on-chip architectural elements (e.g. studying the interaction between power consumption and management, and cell design in cache systems), to the architectural and ISA level (e.g. studying instruction-set and architectural extensions for specific applications), to the system level design and the related system-application interface (e.g. chip tiling and multi-core interconnection for exploitation of application parallelism), up to the application-level aiming to enable the application to match the features of the underlying system and to exploit its capabilities (e.g. compiler approaches for memory hierarchy exploitation).

Key personnel: Massimo Alioto, Alfio Andronico, Sandro Bartolini, Roberto Giorgi, Enrico Martinelli.
Website: http://www.dii.unisi.it ■

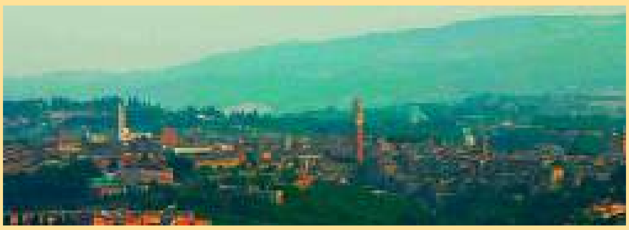

View of the city Siena

ATMI : a microprocessor temperature model

During the last decade, electric power consumption has become one of the major concerns of microarchitects. However in many cases, power consumption per se is not the real problem

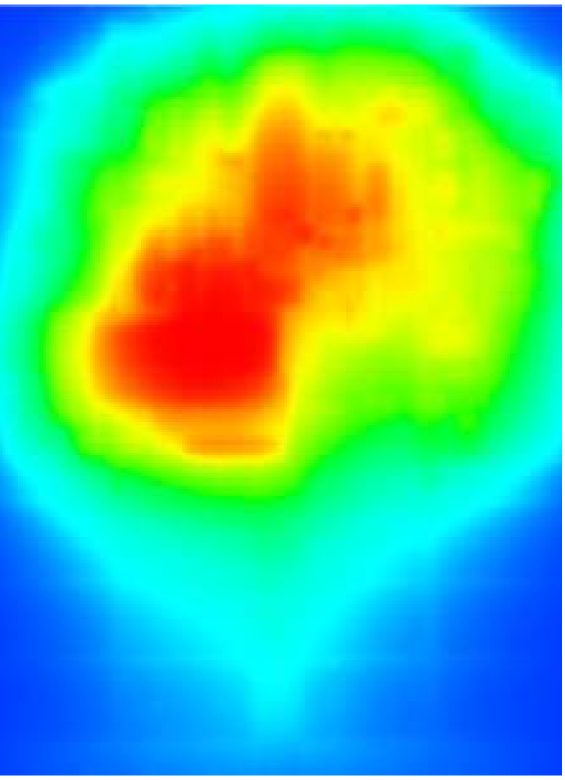

Example : steady-state temperature map

but just a convenient quantity. In most cases, quantities with more direct links to the core problem are energy and temperature. Energy is the price we pay (battery life, electric bill,...) for doing a given work. If one has a power consumption model, energy consumption is easily obtained as an integral over time.

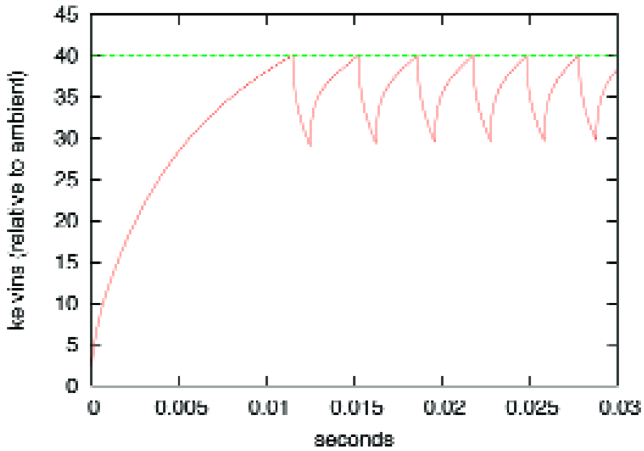

Example : time-varying temperature

Temperature is a different problem. Circuits are designed to work below a certain temperature limit, because temperature impacts circuit delays and mean-time to failure. Power is indeed related to temperature, but the relation between power (or power density) and temperature is not as obvious as the one between power and energy. Technology miniaturization has many advantages (reduced transistor delay, reduced capacitances, ...), not to mention Moore's law. But miniaturizing a circuit without decreasing its power consumption leads to a temperature increase, unless the chip packaging is improved. In that sense, temperature is a stronger constraint than power in the long term.

The CAPS project from INRIA/IRISA has released a microprocessor temperature model, ATMI, for temperature-aware microarchitecture research. The ATMI software is written in C. It provides a mathematical solution to a three-dimensional boundary-value problem of heat conduction. ATMI takes as input power numbers and packaging characteristics and returns a steady-state or transient temperature. The ATMI model is based on some physical idealization. Nevertheless, it produces qualitative behaviors that are consistent with three-dimensional heat conduction. Compared with numerical methods like finite differences or finite elements, the main advantages of ATMI are its speed and ease of use. The ATMI software is released under GNU GPL and is downloadable at http://www.irisa.fr/caps/projects/ATMI.
Contact : pmichaud@irisa.fr.

About CAPS
CAPS (Compiler and Architecture for Superscalar and embedded Processors) is as common project from INRIA, CNRS, University of Rennes 1 and INSA Rennes, located at IRISA in Rennes. The CAPS team studies both hardware and software issues for the design of high performance computer systems. http://www.irisa.fr/caps.

Community news

The public school of Alfamén (Spain) has been renamed 'Mateo Valero' in honor of its most famous alumnus.

Challenges in Embedded Computing

With the advent of logic and RTL synthesis tools in the 90's, coupled with the relative stability of the underlying layout techniques, electronic design has experienced a significant rise of design abstraction and productivity. The remarkable success of logic synthesis justified optimism in research for the next envisioned levels of abstraction, notably behavioral and system-level synthesis, promising to raise design productivity further through the use of generally applicable programming models and tools.

Unfortunately, as we all know from hard experience, such optimistic predictions didn't materialize because of two main forces: complexity, from the top of the abstraction pyramid causing mismatches in the stack of abstraction layers, and process technology changes from the bottom, destabilizing the whole thing from the basis, with the recent acceleration of this latter factor caused by the new challenges of deep submicron design.

While EDA tools have been struggling with these issues, process technology has actually kept the pace of Moore's law, causing the exponential growth of the well known productivity gap between design productivity and process integration capabilities. Such a gap is actually filled in through IP reuse and memory, but the lack of proper programming models and tools has lead to highly unstructured designs, whose complexity has now reached unmanageable levels. It is indeed common for mainstream products, such as the next generation application processors for the mobile market, or for high-definition set-top boxes, to incur NRE costs in the range of 100's of M$, with teams of several 100's of people, software content of several millions of lines of code, and 2-3 years of elapsed time for the introduction of a new chip that must sell in the range of 10's of $'s – definitively not a sustainable trend.

To make things worse, the current transition to deep submicron design is not going to bring new integration capability for free anymore. Indeed, the concurrent push of most process technologies to their physical limits starts to introduce systematic failures in the fabrication process, which can be avoided only by introducing design for manufacturability (DFM) practices.

DFM is going to have a profound impact in the whole design flow, from the layout of standard cells to subsystem architectures, enforcing new design styles strongly based on regularity at all levels. At layout level, lithographic issues, such as the limits of optical proximity correction, are going to force the definition of very few, highly regular cells, whose combinations into larger designs are also optically pre-verified. Those few regular cells will therefore constitute the new basis upon which the whole design flow will be based.

Because of their already high regularity, FPGA-style layouts will be favored in the future, in the sense that their area overhead, compared to future standard cells, will be less of a penalty. For similar reasons, other forms of regular programmable or reconfigurable designs will emerge through the use of coarser grain regular cells, up to the level of complete functional units interconnected by regular networks on chip. Indeed, the DFM-friendly regular design style of these structures will complement well the opportunity to reduce verification offered by these predefined reconfigurable regular designs. In addition, by leveraging other important additional factors such as flexibility, product longevity, the productivity gap and time to market, we can extrapolate this push towards regularity at even higher levels of granularity, up to regularly interconnected dedicated microprocessors.

Software tools will be the real enablers making it possible to fully exploit this convergence towards regular coarse-grain reconfigurable architectures. A lot of R&D is therefore needed in this area, given the traditional delay of several years between the advent of new architectures and the availability of mature and stable tools. But regularity is a good basis for tools, so this trend is actually good news!

A key question though is up to which level regularity applies. Can we extrapolate to

Marco Cornero, is Director of Compilers, Operating Systems and Applications within the Advanced Systems Technology division of STMicroelectronics. He holds a Ph.D. from the University of Genova, Italy, on CAD technologies for high-level synthesis, in cooperation with IMEC, Belgium. He has been working in industry, at STMicroelectronics, for over ten years, contributing to the development of several tools and compilers for advanced microprocessors, including the ST200 VLIW family, in cooperation with HP Labs Cambridge. He is currently managing three main activities, namely compiler development for advanced architectures, operating systems and applications for embedded parallel processors.

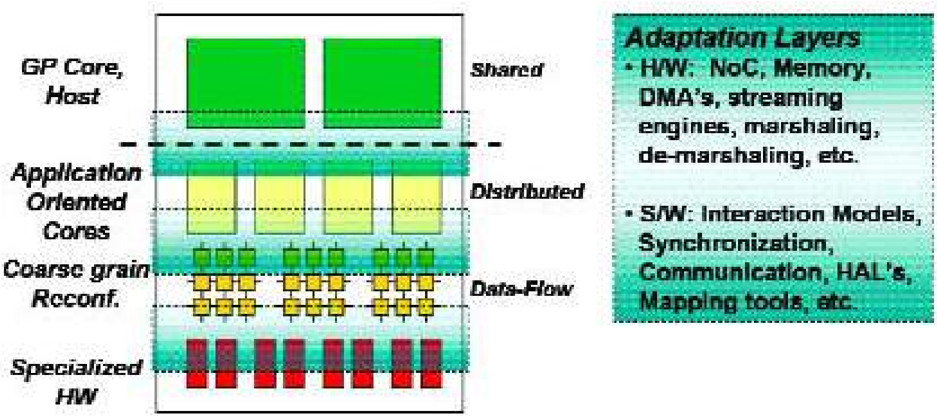

future systems composed entirely of uniform regular parallel structures? Unfortunately, from long experience in supercomputing we know that the answer is no, for well-known reasons. Parallel processing tends to be application-specific, i.e. there is not a single parallel structure that works well for all applications, but rather specific applications require specific parallel architectures. This leads to future systems composed of multiple different parallel subsystems. In addition, a lot of control is present in most modern and future applications, so Amdhal's law comes into the picture, imposing heavy limits on the scalability that we can obtain, and forcing the presence of powerful traditional microcontrollers.

The push to regularity at subsystem level brings therefore relevant local benefits, but at the system level it actually makes things worse, at least for issues that concern software mapping, since from the current heterogeneous multi-processors systems we will move to heterogeneous multi-multi-processors – hardly an encouraging development, especially considering that software complexity is already unmanageable and growing worse; in our opinion this beats even Moore's law.

The only way to avoid crashing into the otherwise inevitable complexity wall is to actually step back and re-think the way we conceive both hardware and software. The next years are going to be crucial for the establishment of innovative architectures and programming models, simply because current practices have reached their limits, and their associated costs are becoming unbearable.

In our view, heterogeneity is here to stay in the embedded domain for the reasons already discussed. This trend will be exacerbated by the ever increasing performance requirements of next generation devices, such as high definition video and software radios just to mention two major ones, combined with the remaining huge cost pressure. So this is a fact of life, an unavoidable complexity whose management is indeed the next challenge in embedded computing, as it determines both hardware and software productivity, and therefore a large portion of the cost and time to market.

In our opinion, in order to address this challenge, two main evolutions need to occur:
1) On the system architecture side, heterogeneity needs to evolve from the current unstructured approach to a much more structured decomposition of the different computation domains. In so doing, verification can be greatly improved, but most importantly it is going to be the only way to establish appropriate programming models for the appropriate identifiable subsystems. As an example, the likely computational domains could be, starting form the general purpose computing side, shared memory symmetric multi-processors (SMPs) running the host operating system, interfacing with the user, and managing the system peripherals. We will then find a number of dedicated media processors, such as VLIWs or DSPs, most likely organized in a distributed memory fashion for scalability and predictability purposes, running dedicated operating systems and performing programmable real-time media functions. The next level is composed of the highly parallel coarse-grain reconfigurable/programmable subsystems that we covered previously, most likely programmed through data-flow based programming models, and finally the inevitable layer of non-programmable but highly efficient hardware accelerators. Each domain will come with its dedicated interconnect and memory hierarchy, while hardware/software adaptation layers will manage the interaction between different domains.
2) On the software side, known programming models can be applied to known separated domains, as in the case of shared and distributed memory subsystems, while recent or new ones will need to be developed for the reconfigurable parallel subsystems. But most important, proper models need to be established for managing the interaction between the different domains. This is going to be the toughest part in our opinion, and therefore provides the most scope for innovation.

The whole picture is clearly way too ambitious to be achieved in a single step though; there is no way such a transition can occur in our industry all at once. Realistically speaking we need a transition path, a kind of system-level assembly language, allowing us to express software the way we need to, without necessarily requiring sophisticated tools that take several years and several generations before being usable. For this intermediate step, and most likely also in the future, software components have the right characteristics. THINK [1], the C incarnation of the Fractal model [2] initially developed by INRIA and France Telecom, gives an example of the required features: 1) data and code encapsulation, which allows smooth and efficient adaptation to both shared and distributed memory models, 2) a highly flexible "binding" model, were links between components can actually be implemented using dedicated software components as well. This characteristic is particularly useful in our heterogeneous systems for abstracting the body of the application from the physical implementation of the communication links, which may differ in nature depending on the domains involved. And finally 3) flexible control interfaces, enabling observability and introspection within the hierarchical (fractal) composition of the components, and even their dynamic adaptation or substitution, for example for maintenance purposes.

Thanks to the above characteristics, software components can already be used today as part of good software engineering practice, but most importantly as a unifying system-level assembly model within and throughout the highly heterogeneous

[1] J.-P. Fassino, J.-B. Stefani, J. Lawall and G. Muller. THINK: A Software Framework for Component-based Operating System Kernels. In USENIX Annual Technical Conference, 2002.
[2] E. Bruneton, T. Coupaye, M. Leclerq, V. Quema, and J.-B. Stefani. An open component model and its support in Java. In CBSE, Volume 3054 of LNCS. Springer, 2004.

computational domains. In the future we envision more sophisticated tools to assist software developers in appropriately decomposing applications into components and in effectively mapping them on the target architectures.

In conclusion, big challenges are definitively in front of us, arising from the evolution of the process technologies, DFM issues in particular, and from the overwhelming complexity of the fast growing software content to be mapped onto the highly heterogeneous architectures that characterize, and will continue to characterize embedded systems. A deep restructuring of both hardware and software architectures is needed, together with a transition path to get there. Solutions exist of course, as for example the ones highlighted in this summary. These will bring fundamental evolutions, and hopefully a decisive step ahead, in the years to come, towards much higher productivity levels in designing and programming our increasingly complex systems on chip. ■

HiPEAC Journal

Transactions on High-performance Embedded Architectures and Compilers

The first issue of the Transactions on HiPEAC contains the following papers:

Introduction by Per Stenström, editor-in-chief

M. Wilkes, **High Performance Processor Chips.**

G. Fursin, A. Cohen, M. O'Boyle, O. Temam, **Quick and Practical Run-time Evaluation of Multiple Program Optimizations.**

M. Geiger, S.A. McKee, G.S. Tyson, **Specializing Cache Architectures for High-Performance and Energy Conservation in Embedded Systems.**

D. Buytaert, K. Venstermans, L. Eeckhout, K. De Bosschere, GCH: **Hints for Triggering Garbage Collection.**

W. Shi, C. Lu, H.S. Lee, **Memory-centric Security Architecture.**

K. Ning, D. Kaeli, **Power-Aware External Bus Arbitration for System- on-Chip Embedded Systems.**

A pdf version of this issue can be found at:
http://www.hipeac.net/transactions1.1 ■

PhD news

On the Systematic Design of Cost-Effective Branch Prediction

Veerle Desmet, veerle.desmet@elis.ugent.be, Ghent University, June 2006

This thesis explores machine learning techniques to gain insights into the fundamental properties of branch prediction. As starting point, it studies the importance of a large set of prediction attributes that can be used by a branch predictor. The predictive quality of these attributes is quantified by how much information they carry with respect to the branch prediction problem. The most interesting prediction attributes are used in a semi-automated way to design and dimension static and dynamic branch predictors. In addition, this research has lead to a branch predictor design that was qualified for the final of the Championship Branch Prediction 2004.

Techniques to Reduce Thread-Level Speculation Overhead

Fredrik Warg, warg@ce.chalmers.se, Department of Computer Science and Engineering, Chalmers University of Technology, June 2006.

With thread-level speculation (TLS), programs are aggressively parallelized at run-time -- threads are later squashed in case of dependence violations. Unfortunately, overhead is a problem.
This thesis quantifies the impact of overheads, and proposes remedies in the form of run-time techniques reducing the amount of violations and too short threads. The interaction between instruction and thread-level parallelism, and TLS performance for simultaneous multi-threaded processors, are other topics covered.

Upcoming events

ACACES 2006 Summer School + 6th HiPEAC cluster meeting
L'Aquila, Italy, July 23-29, http://www.hipeac.net/hipeac/summerschool/

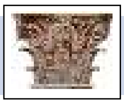

ISSPIT 2006, The 6th IEEE International Symposium on Signal Processing and Information Technology
Vancouver, Canada. August 27-30, 2006, http://web.unbc.ca/~zhoul/isspit/

Euro-Par 2006
Dresden, Germany, August 29 September 1, 2006, http://www.europar2006.de/

PATMOS 2006, Power and Timing Modeling, Optimization and Simulation
Montpellier, France, September 13-15, 2006, http://www.lirmm.fr/patmos06/

International Conference on Parallel Architectures and Compilation Techniques (PACT-2006)
Seattle, Washington, September 16-20, 2006, http://www.pactconf.org/

MEDEA Workshop, MEmory performance: DEaling with Applications, systems and architecture, held in conjunction with PACT 2006 Conference.
Seattle, Washington, September 16-20, 2006, http://garga.iet.unipi.it/medea06/

Twelfth International Conference on Architectural Support for Programming Languages and Operating Systems (ASPLOS XII)
San Jose, CA, October 21-25, 2006, http://www.princeton.edu/~asplos06/

CODES+ISSS 2006, International Conference on Hardware/Software Codesign and System Synthesis
Seoul, Korea, October 23-25, 2006, http://www.ida.liu.se/conferences/codes/

CASES 06, International Conference on Compilers, Architecture and Synthesis for Embedded Systems
Seoul, Korea, October 23-25, 2006

IISWC-2006, IEEE International Symposium on Workload Characterization
Hilton Hotel, San Jose, California, USA. October 25-28, 2006, http://www.iiswc.org/iiswc2006/

MICRO-39, The 39th Annual IEEE/ACM International Symposium on Microarchitecture
Orlando, Florida, USA, December, 9-13, 2006, http://www.microarch.org/micro39/

HiPEAC 2007, High-Performance Embedded Architecture and Compilation
Ghent, Belgium, January 28-30, 2007, http://www.hipeac.net/conference/

HiPEAC 2007 Workshops and Tutorials
Ghent, Belgium, January 28, 2007

1. Workshop on Statistical and Machine Learning approaches applied to Architectures and Compilation (SMART)
2. Workshop on Interconnection Network Architectures: On-Chip, Multi-Chip
3. Workshop on Reconfigurable Computing
4. Tutorial on UNISIM Simulation Environment
5. Tutorial on GCC
6. Tutorial on Microthreading

Contributions

If you are a HiPEAC member and you want to contribute to this newsletter, please contact Thomas Van Parys at **Thomas.VanParys@elis.UGent.be**

HiPEAC Info is a quarterly newsletter published by the HiPEAC network of excellence.
Funded by the 6th European Framework Programme (FP6), under contract no. IST-004408.
Website : http://www.hipeac.net.
Subscriptions: http://www.hipeac.net/newsletter

In the next issue:
A presentation of HiPEAC Belgium