Information Society Technologies

HiPEAC info 10

COMPILATION ARCHITECTURE

appears quarterly | April 2007

Network of Excellence on High Performance Embedded Architectures and Compilers

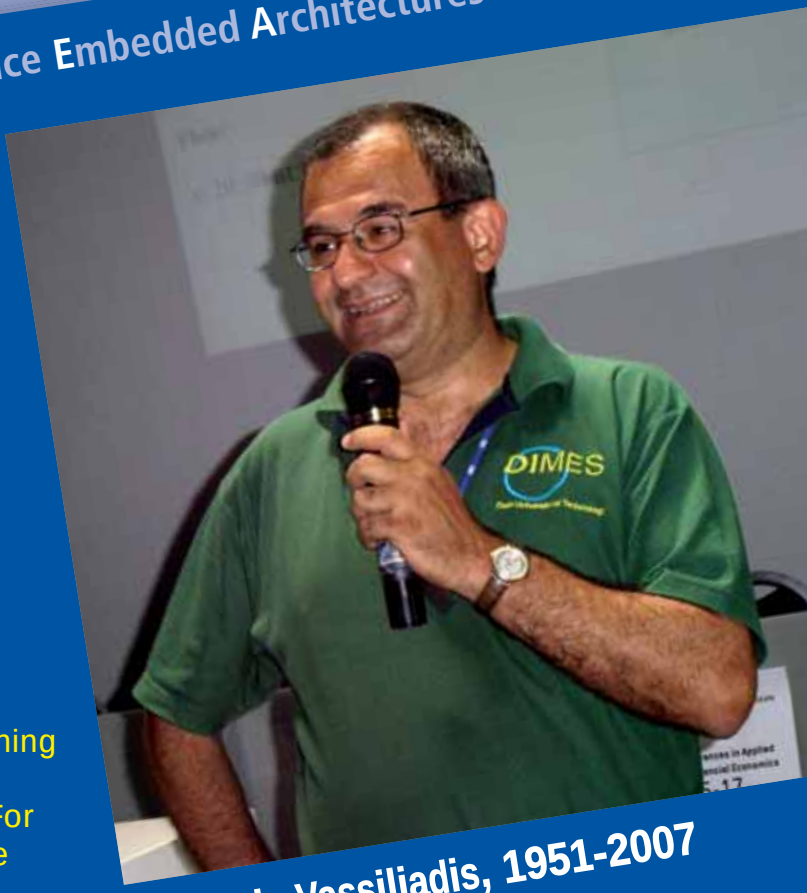

Stamatis Vassiliadis, 1951-2007

www.HiPEAC.net

HiPEAC 2008, Göteborg, Sweden

Deadline Call for Papers: June 10, 2007

Message from the HiPEAC coordinator

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

Dear colleagues,

This has been a hectic winter for HiPEAC. Apart from our ongoing research activity structured within the clusters, HiPEAC carried out its second Conference, published its first HiPEAC Roadmap, opened two new calls to offer new opportunities to our members and of course, more events are coming up.

The HiPEAC Winter Conference (January 28-30, 2007, Ghent, Belgium) welcomed more than 200 participants. 19 papers were accepted and 6 parallel workshops were organized, with 150 participants. The Conference proceedings were published by Springer LNCS and the best papers will be gathered in the Transactions on HiPEAC.

After its second edition, the HiPEAC Conference begins to consolidate as a major event for our community. The call for papers for the HiPEAC 2008 Conference, scheduled for January 27-29 in Göteborg, Sweden, has already been launched.

In January, HiPEAC published the HiPEAC Roadmap on high-performance embedded architecture and compilation. The objective was to point far in advance at the most fruitful and high-impact research directions relevant to the needs of the European industry. A total of 55 key challenges, organized in 10 themes are listed. The roadmap process has been a real community effort. The result is a straight to the point paper, available on our website, which will periodically be updated. I encourage you to read it carefully and to send your feedback, which as always will be much appreciated. I wish to thank again all HiPEAC members and people from outside our community, who contributed to this document.

The regular cluster meeting took place right after the Conference in Ghent. A new cluster call closed on February, 28. Forty proposals were received, that must be evaluated by April.

HiPEAC keeps on fostering industry-academia relationships on high-performance embedded systems, by means of internships and industrial workshops.

The second internship call, closing in March, will fund several company internships. This mechanism will allow PhD students to directly target company research groups on HiPEAC topics. HiPEAC member companies (ARM, IBM, Infineon, NXP and ST) have produced a list of the research topics for which they are seeking interns. Visit our website if you want to be updated about this and other enticing opportunities!

In April, a new HiPEAC Industrial Workshop has taken place in Haifa, Israel, in parallel with a general cluster meeting. A call for papers opened at the end of January. Papers were selected by industry researchers.

Finally, HiPEAC has already announced ACACES 2007, our Third International Summer School on Advanced Computer Architecture and Compilation, again in L'Aquila, Italy.

We are so enthusiastic about our achievements at HiPEAC and about the challenges we still have to meet, that we want it to be continued. The European Commission FP7 programme may give us again the chance of running an improved network from 2008 on. Thus, we presented our new plans at the InfoDay in Brussels, on March 7. That may be our future adventure. Our name? HiPEAC, too!

My only regret is not to be able to share this new path with my dear friend Stamatis. I have plenty of good memories of our meetings and dinners during these few years we worked together, and I can say he continuously surprised me. A very industrious and clever guy, he loved his job and would never stop working, but liked to combine business and pleasure, and he certainly achieved it, helping everybody to have a good time. He made of the yearly Samos Conference an opportunity for many students to enjoy his beloved Mediterranean sea. I personally cannot think of any occasion in which we were not joking while working. He liked people and people liked him. Sometimes he switched from a kind of "enfant terrible" attitude to the most serious, collaborative one in few seconds. I always suspected it was his particular way of dealing with long, tedious administrative issues that usually accompany large projects like HiPEAC or SARC. Stamatis was for many of our colleagues the "Happy Warrior" in our field. He was a most valuable colleague and friend, a very optimistic, positive kind of person who showed great courage until the end. Stamatis, tú sabes que siempre te llevaré en mi corazón.

Mateo Valero
HiPEAC Coordinator

Stamatis Vassiliadis, 1951-2007

On Saturday, April 7, 2007, after a long illness, Stamatis Vassiliadis peacefully passed away, surrounded by his family. The HiPEAC network lost one of its best members and its ultimate discussion partner. His family lost a loving father and a devoted husband, the world of computing a great scientist, and our world at large a very kind and generous person.

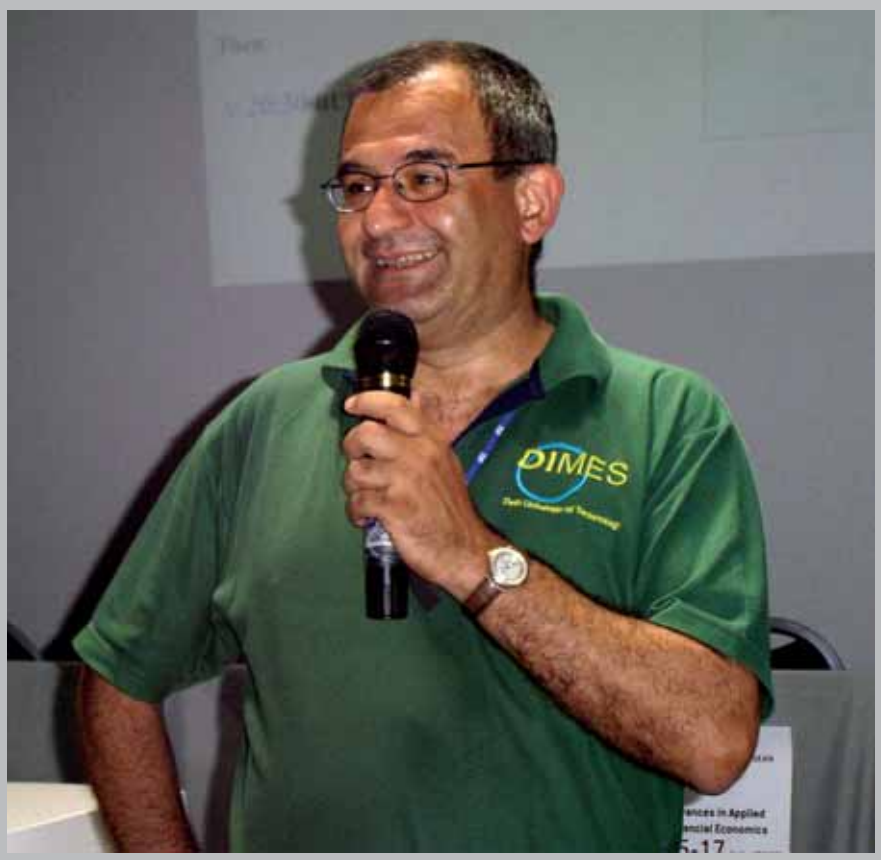

Stamatis was born in 1951 in the small village Manolates, on the island of Samos, in Greece. After his study at Politechnico di Milano, Stamatis moved to USA and worked for IBM at the Advanced Workstations and Systems laboratory in Austin, Texas, the Mid-Hudson Valley laboratory in Poughkeepsie, New York, and the Glendale laboratory in Endicott, New York. At IBM, he has been involved in a number of projects regarding computer design, organizations, and architectures and in the leadership of advanced research projects. In this period he was awarded 73 USA patents ranking him as the top all time IBM inventor. For his work he received numerous awards including 24 Publication Awards, 15 Invention Achievement Awards and an Outstanding Innovation Award for Engineering/Scientific Hardware Design in 1989. While working for IBM Stamatis also served in the ECE faculties of Cornell University, Ithaca, NY and State University of New York (S.U.N.Y.), Binghamton, NY. In 1995 he accepted a position at TU Delft and moved back to Europe. He made the Computer Engineering laboratory one of the strongest groups in the field with more that 50 PhD students from many different countries. Stamatis was IEEE and ACM fellow and a member of the Royal Dutch Academy of science (KNAW).

Stamatis was very proud of his island Samos, a small piece of Greek land that produced many great scientists. He loved this island very deeply and was returning there every summer. In addition, he made of the Samos conference an opportunity to enjoy his beloved Mediterranean sea with so many students and colleagues, who will always remember it as a great experience, so different from any other event.

Stamatis always inspired all of his friends and students with his enthusiasm, momentum, courage, audacity, warmth, friendliness, assistance and total support. Being a warm-hearted person, he always inspired friendly feelings in all those around him.

Stamatis was one of the earliest pioneers of HiPEAC, always trying to strengthen our community in Europe. We mourn his death, but are determined to continue the work on building the best scientific and engineering community on computer architecture and compilers in Europe, one of Stamatis' dreams.

Stamatis, we will always remember you.

Stamatis Vassiliadis
1951-2007

Integrity was his compass
Science his instrument
Advancement of humanity his final goal

The Netherlands

Scientists from the Netherlands continuously made significant contributions to the field of computing throughout the years. Edgser Dijkstra, Gerrit Blaauw and Willem van der Poel are three names that do not need any additional introduction. One very interesting fact is that the developments in this field were mainly driven by the needs of the industrial development and not by military demands.

The real boom in computing in the Netherlands happened in the 50's of the last century. There were three centers for building computers: the Dr Nether Laboratory at the Post Telephone and Telegraph (PTT) in Rotterdam (now KPN research), the Mathematic Centre (MC) in Amsterdam (nowadays CWI -Centrum voor Wiskunde en Informatica / Center for Mathematics and Computer Science) and Natuurkundig laboratorium (NATLAB) in Eindhoven (now Philips and NXP research).

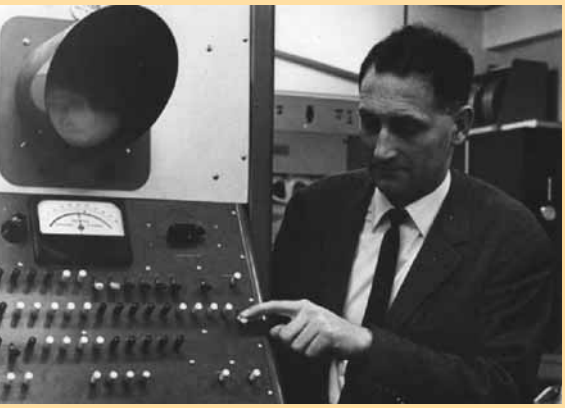

At Dr Nether Lab, W. van der Poel and his team created many successful computers in close collaboration with TU Delft where he built the Testudo (Turtle) in 1952 from parts donated by PTT. The PTT Testudo was used by TNO between 1952 and 1964, a long lifetime as the one of the animal whose name it borrows. The three Van der Poel's machines: ZERO (1953), PTERA (1953) and ZEBRA (1956) introduced many novel concepts. For example the ZEBRA (see the photo showing the computer and its inventor) was using microprogramming long before the term was even introduced. Willem van der Poel actually built his first relay-based computer already in 1944 as student in his hobby room.

The MC (see the photo) was founded in 1946 to promote pure and applied mathematics. The rekenafdeling (the calculation division) had two main tasks in 1947: performing advanced calculations and building a novel computing machine. The team involving Blaauw (also co-architect of IBM 360) Scholten and Loopstra developed ARRA I (1952), ARRA II (1954) and ARMAC (1956). Those machines were using proved methods and techniques and have been designed with maintainability and high productivity in mind. Many standard replaceable units where used (probably contribution of G. Blaauw). The programmer of those machines was E. Dijkstra.

At NATLAB in Eindhoven three different machines were built: PETER (1956), STEVIN (1960) and PASCAL (1960). They all were for internal Philips use only and were designed mainly to gain some experience with the circuits Philips was delivering to IBM.

Being very pragmatic, the programmers in the Netherlands in the 50's were also very confident in their models – two striking examples of that time are the following. In Delft wind tunnel forces were calculated using a technique later known as the finite elements method. As there was not enough processing power to calculate the whole structure, only small parts of it were computed. The scientists that did this job walked along the real structure under test when pressure above 3.5 Bar was applied to it for the first time (at the same time the builders took cover behind a concrete wall). Programmers at MC modeled the water movement produced when a big ship was let into water for the first time and stood very close on the pier watching the real situation, confident to keep their feet dry. The programmers in the Netherlands those days were living in the real world and not in any of the virtual ones,

which are very popular now.

The first programmer in the Netherlands, Edsger Wybe Dijkstra is considered as one of the most influential members of computing science's founding generation. His scientific contributions are fundamental in the domains of: algorithm design, programming languages, program design, operating systems, distributed processing, formal specification and verification, design of mathematical arguments. During his forty-plus years as a computer scientist, at both academia and industry, Dijkstra's contributions brought him many prizes and awards, including computing science's highest honor, the ACM Turing Award in 1972 (the "Nobel price" in computer science).

Many national and international companies are located in the Netherlands. The two most relevant to HiPEAC topics companies: NXP and ACE are members and active contributors to the development of high performance embedded architectures and compilers in the Netherlands and Europe.

HiPEAC Partners

The Computer Engineering (CE) laboratory, Microelectronics and Computer Engineering department at **TU Delft**, performs research and teaches the engineering discipline of how to determine, develop, and integrate software and hardware to build a computing system. The laboratory focuses on the definition of system requirements, from embedded to general purpose, their architecture and implementations, and the study and development of tools and software that allow improving the analysis and synthesis of computing systems. More precisely, the laboratory is actively involved in: computer architecture, machine organizations, and network processing, mapping of application and algorithm requirements to architectures of embedded systems (e.g. multimedia), compiler technology capable of directing system requirements to architectural definitions and improve implementations, architectural synthesis tools for semi-automatic implementation of architectures, computer arithmetic and logic design, algorithms and tools for testing memories, built in self-test of logic circuits, and automatic test pattern generation for combinational and sequential logic circuits, performance modelling and optimisation techniques and tools. CE is one of the founding HiPEAC partners and is a steering committee member. Currently 50 PhD students from many different countries are pursuing their research in various fields of Computer Engineering at CE. CE laboratory is the project coordinator of the Scalable ARChitecture **(SARC)** FP6 Integrated Project (contract number 027648) and the scientific coordinator of **hArtes** IP (035143).

People:
Stamatis Vassiliadis, Georgi Gaydadjiev, Koen Bertels, Mladen Berekovic, Ben Juurlink and Sorin Cotofana
URL: http://ce.et.tudelft.nl

The Computer Systems Architecture (CSA) group from the **University of Amsterdam** performs research in the fields of scalable, instruction-level parallel architectures as well as system-level computer architecture modelling and simulation. In our MicroGrid project, we are developing a novel approach to micro-architecture that supports massive on-chip concurrency, which is scalable, flexible and amenable to analysis. It has the potential to provide for the management of on-chip resources (processors etc.) so as to autonomously configure a system for performance, power dissipation or fault tolerance. To this end, we have introduced the concept of microthreads. Microthreading is an execution model that breaks code down into fragments that can execute simultaneously, on a single microthreaded microprocessor or distributed over different processors. Within the Aether project, we extend the above concepts to develop computer systems that support self-adaptation in their software, architecture and implementation.

In addition, our group also investigates more generic methods and techniques for system-level design and analysis of (future) system-on-chip based computer architectures. This work focuses on architectural design space exploration (DSE) during the very early stages of design, where design decisions have great impact on (the success of) the final product. To this end, we study both analytical modeling methods as well as simulation methods for system-level (performance) analysis. These methods and techniques are incorporated in our Sesame simulation framework for system-level DSE.

People:
Chris Jesshope, Andy Pimentel and Peter Knijnenburg
URL:
http://www.science.uva.nl/research/csa/

NXP is a new independent semiconductor company (founded by Philips) with a fifty-year history of providing engineers and designers with semiconductors and software that deliver better sensory experiences for mobile communications, consumer electronics, security applications, contactless payment and connectivity, and in-car entertainment and networking.

Building on its heritage in consumer research, significant R&D investment and world-class industry partners, NXP's "vibrant media technologies" allow consumers to enjoy better sensory experiences - brilliant images, crisp clear sound and easy sharing of information in homes, cars and mobile devices.
NXP was established in 2006 (from some formal Philips divisions) and inherits more than 50 years of experience in semiconductors. The company headquarters are in Eindhoven, the Netherlands. The net sales in 2005 amounted € 4.77 billion. The company employs approximately 37,000 people in more than 20 countries and has more than 24 R&D centers world-wide. NXP has 10 wafer fabs and 8 test and assembly sites spread in different countries. Company Customers include Apple, Bosch, Dell, Ericsson, Flextronics, FoxConn, Nokia, Philips, Samsung, Siemens and Sony.

NXP holds Worldwide **number one** positions in: 5-V CMOS logic products for the automotive industry, Automotive In-Vehicle Networks, Car radio Digital Signal Processors, Contactless identification for e-passports, Digital cordless chips, FM radio ICs for mobile, GSM/GPRS/EDGE system solutions, Interface products, Mobile speaker systems, Near Field Communication, PC TV chips, RF products for CATV and satellite tuners, RFID for electronic ticketing in public transport, System solutions for automotive immobilizers and keyless entry/go, TV chips and USB.

NXP current focus markets are: Mobile & Personal, Home, Identification, Automotive, Multimarket Semiconductors and Software.

People:
Marc Duranton, Jan Hoogerbrugge
URL: http://www.nxp.com

HiPEAC Partners

The Parallel and Distributed Systems group (PDS) in the Department of Software Technology of **Delft University of Technology** (TU Delft) focuses on the following research topics: grid computing, parallel programming models, peer-to-peer systems, and sensor networks.

The research in *grid computing* focuses on the problem of scheduling in multi-cluster systems and grids. The subsystems making up a grid are to a large extent autonomous, since grids are often heterogeneous and also may exhibit many failures. Therefore, scheduling in grids is highly non-trivial. In addition, when jobs only employ resources in a single grid subsystem, grid schedulers are not much more than load-balancing devices. For this reason, we focus on the problem of co-allocation, i.e., on allocating resources (processors) in multiple clusters or subsystems to single jobs. We study grid scheduling by actually designing, implementing, and deploying a working grid scheduler (called KOALA) in the Dutch National Research Grid system (the DAS).

High performance computing has as focus research in parallel languages and programming environments, more specifically in languages and compilation techniques for distributed memory architectures such as multi-core systems. We focus on HPC-extensions to Java (SPAR) and compilers that semi-automatically generates code for distributed memory systems. Another area of focus is stream processing, which stems from consumer electronics applications and scientific applications where streams are generated by sensors (e.g. radio telescopes).

The research in *peer-to-peer networks* focuses on adding social features such as friends and taste buddies, and adding support for TV distribution (both live and recorded) to BitTorrent, which is one of the most popular p2p systems. Among the most important features we concentrate on are IP support for the notion of friends, an efficient gossip-based algorithm for doing recommendations for content, support for improved download performance and streaming videos across p2p systems, and decentralization of a number of mechanisms.

In the *wireless sensor networks* area, we focus on the development of new protocols and algorithms for the efficient management of the resource- and energy-limited sensor nodes making up such networks. Prototype deployments have to establish the feasibility of such application domains.

People:
Henk Sips
URL: http://www.pds.ewi.tudelft.nl

The mission of the section Electronic Systems at (TU/e) is to provide a scientific basis for design trajectories of digital electronic circuits and systems 'from (generalized) algorithm to realization'. To identify the key problems, and verify the validity, robustness and completeness of our results, we develop, implement and maintain consistent and complete flows, and use them for realizing innovative multimedia hardware with emphasis on video processing and embedded architectures.
Implied in the mission statement is the question of how to convert the "art" of designing electronic systems into methodology. This is an absolute necessity because:
- the complexity of modern integrated circuits continues to increase,
- new physical phenomena at submicron feature dimensions are having more and more impact, not only on performance, but even on the functionality
- and the heavy demand pull from signal processing applications, in particular multimedia and telecommunications, requires rigorous and robust answers.
Algorithms play a key role here, and with a dual nature:
- they still form the basis for effectively using a computer in design assistance, so in the first instance we want to support or develop algorithms for synthesis and verification of complex integrated systems where we do not stop at the level of point-wise solutions to specific problems, but integrate them into complete design environments: this is the methodological challenge.
- they are the core of signal processing, and since video processing is our major application area, we aim at discovering effective algorithms to treat video signals in multimedia systems.
We like to emphasize our generalized view on algorithms and more specifically our view towards computational networks, that is graphs with computations at the nodes and transfer protocols on the arcs. Of course, groups at other universities and industries are facing these challenges, and we gladly adopt their results and tools to come to full trajectories and innovative processors. Our contributions come from tackling the fundamental problems and filling the essential gaps, revealed by careful analysis of the methodological scene, using our insight in video processing and our experience with design environments.

The chair's expertise is firmly rooted in industrial research of video processing, multiprocessor architectures and design automation for deep submicron VLSI, and complemented with a solid theoretical basis in combinatorics and process algebra. This makes a distinctively design-oriented group, which aims to push ideas so far that they are implemented and used in design programs and in major applications.

People:
Ralph Otten, Henk Corporaal, Twan Basten, Bart Mesman, Jeroen Voeten
URL: http://www.es.ele.tue.nl/

Message from the project officer

Mercè Griera i Fisa Merce.
Griera-i-Fisa@ec.europa.eu

Computing Systems Call closes on 8 May at 17:00

In the first call of FP7, there are 25 M€ reserved for research on Computing Systems. Proposals in this area must address one of the following domains:

- Novel architectures, the corresponding system-level software and programming environments for **on-chip multi-core** computing systems.
- Reference designs /architectures for generic **embedded** platforms cutting across application domains, accompanied by appropriate tools and component libraries.

There are 5 M€ earmarked for a NoE covering the first domain and the rest is for STREPs.

The challenge in this area is to strengthen Europe's position as a leading supplier of computing systems in order to support the competitiveness of industrial strongholds such as consumer electronics, telecoms or medical systems. As a consequence:

- Appropriate industry participation is to be proposed by all consortia
- All STREPs are to be build with a systems approach in mind and this is to be reflected in both the consortium composition and the workplan structure.
- Hardware and software are to be considered together in all proposals.

Details on the call are at http://cordis.europa.eu/fp7/ict/
In addition to the Work Programme, I would like to recommend that you read the "guide for applicants" in detail, because being it the first call of FP7, there is a need for understanding of the new rules and procedures. To facilitate this work we organized on the 7th of

March an Information Day in Brussels at the Commission premises. A considerable number of people from the HiPEAC community attended the meeting. You will find the slides of the different presentations at: http://cordis.europa.eu/ist/embedded/infoday-070307.htm

If you have any questions related to the Call content and on the new rules and procedures do not hesitate to contact me.
Mercè Griera i Fisa (Merce.Griera-i-Fisa@ec.europa.eu)

HiPEAC Conference Ghent January 28-30 2007

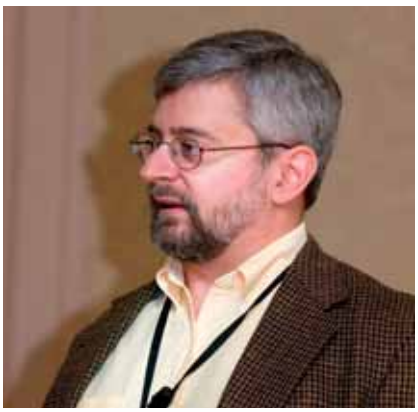

Modifying GCC to enable automatic tuning of optimization heuristic

Current innovations in science and industry demand ever-increasing computing resources while placing strict requirements on system performance, power consumption, size, response, reliability, portability and design time. However compilers often fail to deliver satisfactory levels of performance on modern processors, due to rapidly evolving hardware, lack/cost of expert resources, fixed and black-box optimization heuristics, simplistic hardware models, inability to fine-tune the application of transformations, and highly dynamic behavior of the system.

Recently, we started developing an Interactive Compilation Interface (ICI) to connect external optimization drivers to the GCC. This interface is meant to facilitate the prototyping and evaluation of iterative optimization, fine-grain customization and design-space exploration strategies. An early design, able to provide non-intrusive feature extraction and meddling with heuristic's decisions, was presented at the SMART'07 workshop. Currently, we are working on a more advanced design, incrementally modifying Tree-SSA to support dynamic pass reordering, structured split of analysis and optimization code, and a component model for passes to enable dynamic linking of external optimization plugins. We believe these modifications will simplify the tuning process of new optimization heuristics and will eventually simplify the whole compiler design, so that compiler heuristics will be learned automatically, continuously and transparently, aiding users using statistical and machine learning techniques.

More information will be available at the development site:
http://gcc-ici.sourceforge.net

Grigori Fursin, INRIA Futurs, France

MiDataSets: Creating The Conditions For A More Realistic Evaluation of Iterative Optimization

Iterative optimization has become a popular technique to obtain improvements over the default settings in a compiler for performance-critical applications, such as embedded applications. An implicit assumption, however, is that the best configuration found for any arbitrary data set will work well with other data sets that a program uses.

We created 20 different datasets per program for MiBench benchmark to evaluate this assumption and analyze the behavior of various programs with multiple datasets. This work has been presented at HiPEAC'07. After resolving some copyright issues, we plan to make our datasets publicly available to enable more realistic benchmarking and practical iterative optimization research.

More information will be available at the development site:
http://midatasets.sourceforge.net

Grigori Fursin, INRIA Futurs, France

Community news

Jose Duato was awarded the research prize "Premio Rey Jaime I a las Nuevas Tecnologías 2006"

For his discoveries of high international impact regarding traffic optimization in interconnection networks, with special impact in the domain of supercomputing where they have been applied to the BlueGene/L supercomputer.

Four HiPEAC members (UPV, Simula, FORTH ICS and University of Murcia) have joined the HyperTransport consortium

HyperTransport is the system area network communications standard that delivers the highest bandwidth and lowest latency in the market.
(see http://www.hypertransport.org/consortium/cons_members.cfm?m=3)

Software/Hardware Techniques for Mesh Compression in Computer Graphics

By Paula Novío (paulanm@dec.usc.es), Prof. Javier Bruguera and Prof. Montserrat Bóo, University: Santiago de Compostela, Spain
June 1, 2006

The work presented in this dissertation is focused on hardware-oriented compression and decompression algorithms for primitives commonly used in computer graphics applications, such as triangles, points and tetrahedra. Hardware units for decompression are also proposed. The size of the objects employed in computer graphics is continuously growing, reaching billions of primitives. This produces bottlenecks caused by the lack of storage space and limited bandwidth between the CPU and the GPU. Therefore, hardware compression algorithms are currently of great interest.

In particular, we present a compression algorithm for triangle meshes based on concentric strips. This algorithm achieves a high compression ratio. Firstly, the hardware unit for decompression is presented. The algorithm is then extended to tetrahedral meshes. Finally, point meshes are considered, proposing two algorithms for the compression of the geometry, obtaining good compression ratios.

Locality optimization techniques for irregular codes on multiprocessor and multithreading architectures

By Juan C. Pichel (jcpichel@dec.usc.es), Prof. Dora Blanco, Prof. Jose C. Cabaleiro, University: Santiago de Compostela, Spain
September 9, 2006

In this work, several techniques for optimizing data locality in irregular codes of sparse matrix algebra are proposed. These approaches were applied to different parallel architectures. The sparse matrix algebra codes are present in a lot of problems from engineering. Their main characteristics are their poor data locality and the fact that their memory accesses are unpredictable. These issues explain why for these applications, the memory hierarchy performs poorly.

The proposed techniques are based on the reordering of the data structures (sparse matrices) that determine the locality of the code under study. These reorderings are guided by a locality model previously developed by our group. In this model four distance functions are established. These functions are then evaluated between pairs of rows (or columns) of the matrix and measure the locality in the irregular accesses that these rows (or columns) address. The technique is general enough since, as is shown in this work, it can be successfully applied to any sparse matrix (without limitation in its pattern) on different multiprocessor systems (both share memory and distributed memory) and multithreaded architectures.

Request-Grant Scheduling for Congestion Elimination in Multistage Networks

By Nikolaos Chrysos (nchrysos@ics.forth.gr), Prof. Manolis Katevenis, University of Crete and FORTH,
December 2006

Interconnection networks suffer from congestion any time multiple inputs, unaware of each other's decisions, inject into the network traffic in excess of some output or link capacity(ies); this results to delays not only for the responsible packets, but for other unrelated flows as well. This thesis introduces a request-grant scheduling scheme that allows packet injection only after reserving all necessary buffer space, thus eliminating head-of-line (HOL) blocking. The scheme works as a hybrid between traditional scheduling of bufferless fabrics, which is too complex, and providing per-flow queues, which requires too much buffer space. We have tested the new architecture by simulating a specific design that sustains robust operation under any number of congested outputs in a 1024-port, 10 Tb/s, three-stage Clos/Benes network, built using just 96 buffered crossbar chips and 1 control chip.

On-chip traffic statistical analysis

By Antoine Scherrer (Antoine.Scherrer@insa-lyon.fr), Dr. Tanguy Risset and Dr. Antoine Fraboulet, Inria Compsys
December 11, 2006

This PhD deals with the analysis and synthesis of on-chip traffic, i.e. the traffic occurring in a system-on-chip (SoC). In these systems, the introduction of networks on chip (NoC) has brought up the interconnection system as a major issue of the design flow. In order to prototype these NoC rapidly, fast simulations need to be done, and replacing the components by traffic generators is a good way to achieve this purpose. We have set up and developed a complete and flexible on-chip traffic generation environment that is able to replay a previously recorded trace, to generate a random load on the network, to produce a stochastic traffic fitted to a reference trace and to take into account traffic phases.

Optimization of a parallel 3D simulator applied to the study of intrinsic parameters on HEMT devices

By Natalia Seoane (natalia@dec.usc.es), Prof. Antonio Garcia Loureiro, University: Santiago de Compostela, Spain
January 9, 2007

Three-dimensional numerical simulation of semiconductor devices is extremely demanding in terms of computational time because it involves complex numerical schemes. The large amount of memory and floating point operations needed necessitate the use of parallel machines and appropriate algorithms in order to obtain the maximum performance and reduce simulation time. The simulation study of the intrinsic parameter fluctuations in nanometer devices requires full scale 3D device simulations on a statistical scale and this is computationally expensive. Therefore, the simulation technique used to study intrinsic parameter fluctuations has to be fast and efficient, allowing the simulation of a large ensemble of devices in a relatively short period of time.
A 3D parallel device simulator, based on the drift-diffusion (D-D) approach to the semiconductor transport, was implemented to study HEMT (High Electron Mobility Transistors). This approach constitutes a system of coupled, nonlinear partial differential equations that have been discretized using finite element methods. Domain decomposition methods, implemented by the PSPARSLIB library, were used to solve the linear systems arising from the linearisation of these equations. The 3D simulator has been developed for multicomputers using a Multiple Instruction Multiple Data strategy (MIMD) under the Single Program Multiple Data paradigm (SPMD) and the Message Passing Interface (MPI) standard library.

Clustered VLIW Architectures: a Quantitative Approach

By Andrei Terechko (andrei.terechko@nxp.com), Prof. H. Corporaal, Prof. R.H.J.M. Otten, Dr. P. Stravers
Technical University of Eindhoven, The Netherlands
February 6, 2007

Achieving the best characteristics from clustering a VLIW processor requires a thorough selection of an Inter-Cluster Communication (ICC) model, which is the way clustering is exposed to the ISA. Our VLSI layouts and instruction scheduling show that performance of the popular copy operation model is severely limited by copies hampering scheduling of regular operations. The dedicated issue slots model, deals with this limitation by dedicating extra issue slots for ICC, reaching at most a 1.74 speedup relative to the unicluster. Lowering the area and energy consumption by 55% and 57% relative to the unicluster, respectively, is achieved by the extended operands model.
Site: http://www.terechko.net/cgi-bin/moin.cgi/PhD_thesis

Processor Architecture Design for Smart Cameras

By Hamed Fatemi (H.Fatemi@tue.nl), Prof. Henk Corporaal, Prof. Twan Basten and Dr. Bart Mesman, Technische Universiteit Eindhoven, March 2007

In many networked embedded systems, sensing with cameras is combined with processing to achieve certain communication, measurement or control goals. The advent and subsequent popularity of low cost, low power CMOS vision sensors enables us to integrate processing logic (in a single package or board) on the camera, thereby creating the so-called smart sensors. They have a SIMD data processing array driven by a controller. On top of this, a separate powerful instruction level parallelism (ILP) or general-purpose processor (GPP), is usually needed in embedded applications for feature and object processing and control tasks. Integrating all this functionality (possible in a single package) will have a positive effect on the cost, power consumption, latency and inter-processor bandwidth. The result is a low-cost Smart Camera (so-called SmartCam) solution. In this thesis, we investigate new opportunities and contribute to a better and more quantitatively guided design trajectory for an efficient SmartCam template by considering constraints such as power, performance, and cost.

Contributions to the design of reliable and programmable, high-performance systems: principles, interfaces, algorithms and tools

By Prof. Albert Cohen, INRIA, France
March 23, 2007
Moore's law on semiconductors is coming to an end. Scaling the von Neumann architecture over the 40 years of the microprocessor has led to unsustainable circuit complexity, very low compute-density, and high power consumption. On the other hand, parallel computing practices are nowhere close to the portability, accessibility, productivity and reliability levels of single-threaded software engineering. This dangerous gap translates into exciting challenges for compilation and programming language research in high-performance, general purpose and embedded computing. This habilitation thesis motivates our approach to these challenges, presents our main directions and results, and draws some research perspectives.

Transactions on High-performance Embedded Architectures and Compilers

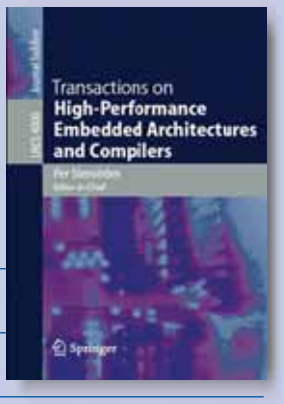

The first issue of Volume 2 contains the following papers:

Introduction to Part 1 by Per Stenstrom and David Whalley

G. Keramidas, P. Xekalakis, S. Kaxiras, **Recruiting Decay for Dynamic Power Reduction in Set-Associative Caches.**

V. Nagarajan, R. Gupta, A. Krishnaswamy, C**ompiler-Assisted Memory Encryption for Embedded Processors.**

S. Kluyskens, L. Eeckhout, **Branch Predictor Warmup for Sampled Simulation through Branch History Matching.**

M. Bhadauria, S. A. McKee, K. Singh, G.S. Tyson, **Data Cache Techniques to Save Power and Deliver High Performance in Embedded Systems.**

C. Hu, D.A. Jiménez, U. Kremer, **Combining Edge Vector and Event Counter for Time-dependent Power Behavior Characterization.**

HiPEAC 2008 Conference

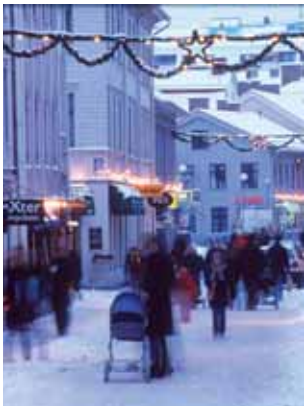

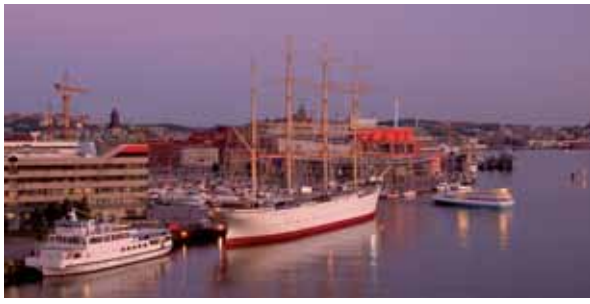

The HiPEAC 2008 conference will take place in Göteborg on the west coast of Sweden on January 27-29, 2008. Göteborg is the second largest city in Sweden and it is an important academic as well as industrial center hosting Chalmers University of Technology and Göteborg University as well as Volvo, SKF, and Ericsson.

The general co-chairs of the conference are Per Stenström (Chalmers) and Michel Dubois (University of Southern California). The program co-chairs are Manolis Katevenis (University of Crete/FORTH) and Rajiv Gupta (University of Arizona). Deadline for paper submissions is June 10, 2007.

For more information:
http://www.hipeac.net/conference
Topics of interest:

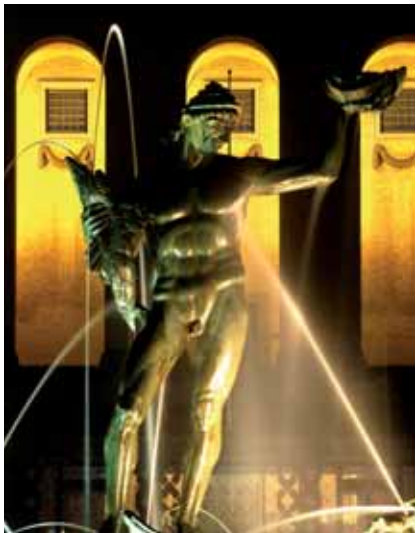

- Processor architectures
- Memory system optimization
- Power, performance and implementation efficient designs
- Interconnection networks, networks-on-chip, network interfaces and processors
- Security, dependability, and predictability support
- Application specific processors and accelerators
- Reconfigurable architectures
- Simulation and methodology
- Compiler techniques for embedded processors
- Feedback-directed optimization
- Program characterization and analysis techniques
- Dynamic compilation, adaptive execution, and continuous profiling/optimization
- Back-end code generation
- Binary translation/optimization
- Code size/memory footprint optimizations

Upcoming events

ISPASS-2007, 2007 IEEE International Symposium on Performance Analysis of Systems and Software,
San Jose, California, USA, April 25-27, 2007, http://ispass.org/

ACM International Conference on Computing Frontiers,
Ischia, Italy, May 7-9, 2007, http://www.computingfrontiers.org/

DAC'44: 44th Design Automation Conference
San Diego, California, June 4-8, 2007, http://www.dac.com/44th/index.html

ISCA: The 34th International Symposium on Computer Architecture
San Diego, CA, USA, June 9-13, 2007, http://www.cse.ucsd.edu/isca2007/

PLDI 2007, Programming Language Design and Implementation,
San Diego, CA, USA, June 10-13, 2007, http://ties.ucsd.edu/PLDI/

ICS07: 21st ACM International Conference on Supercomputing
Crowne Plaza Seattle, Seattle, WA, USA, June 16-20, 2007, http://ics07.ac.upc.edu/

7th Int'l Workshop on Worst-Case Execution Time Analysis (WCET'07)
Pisa, Italy, July 3, 2007, http://www.irit.fr/wcet2007

SIES'2007: IEEE Seconda Symposium on Industrial Embedded Systems
Hotel Costa da Caparica, Lisbon, Portugal, July 4-6, 2007, http://www.uninova.pt/sies2007/

ACACES 2007, Third HiPEAC Summer School,
L'Aquila, Italy, July 15-20, 2007, http://www.hipeac.net/summerschool

SAMOS VII: International Symposium on Systems, Architectures, MOdeling and Simulation
Samos, Greece, July 16-19, 2007, http://samos.et.tudelft.nl/samos_vii/

Euro-Par 2007
Rennes, France, 28-31 August 2007, http://europar2007.irisa.fr/

HiPEAC 2008 Conference
Göteborg, Sweden, 27-29 January 2008, http://www.hipeac.net/conference

Contributions

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

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