

HiPEAC^{info}11

COMPILATION ARCHITECTURE

appears quarterly | July 2007

Network of Excellence on High Performance Embedded Architectures and Compilers

- 3 Community News
- 4 In the Spotlight
Insight, Not (Random) Numbers:
A Perspective on Embedded Benchmarking
- 6 In the Spotlight
hArtes
- 6 HiPEAC Activity
ACACES 2007: Third international summer
school on Advanced Computer Architecture
and Compilation for Embedded Systems
- In the Spotlight
- 7 - New Book
- 8 - RWTH Aachen releases first version of
Virtual SHAPES Platform
- 8 Community News
- 9 HiPEAC Journal
Transactions on High-performance
Embedded Architectures and Compilers
- 10 PhD news
- 12 Upcoming Events



Mateo Valero
 Eckert-Mauchly Award winner 2007

www.HiPEAC.net

Message from the HiPEAC coordinator

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

Dear colleagues,

Time flies! When you read these lines, HiPEAC will be welcoming many of you to ACACES 2007, our Third International Summer School on Advanced Computer Architecture and Compilation, again in L'Aquila, Italy. As usual, courses will be taught by a prestigious group of professors. HiPEAC has for another year done its magic.



As already announced, call for papers on the HiPEAC Winter Conference 2008, scheduled for January 27-29, 2008 in Göteborg, has already been launched. Workshops and tutorials, which have proved very successful in the previous editions, will be run in parallel to the Conference.

The proposals received for the last call for collaborations were evaluated in April. This has been the largest HiPEAC call in terms of funds allocated for fellowships (around 225,000 ?). HiPEAC industry members participate in five of the new or extended clusters. Besides clusters, HiPEAC keeps on fostering industry-academia relationships on high-performance by means of internships and industrial workshops.

The second internship call closed in March. HiPEAC member companies (ARM, IBM, Infineon, NXP and ST) produced a list of the research topics to which HiPEAC members could apply. 64 applications were received and are

still under evaluation. Up to 10 internships could be granted this time. The internships are expected to start this summer, although companies are willing to host interns at other periods of the year.

In April, the Third HiPEAC Industrial Workshop took place. Representatives from companies and universities in more than a dozen European countries attended a week of events focusing on compiler and architecture technologies at the IBM Haifa Research Lab in Israel. Collocated with the HiPEAC Industrial Workshop were the IBM Haifa Compiler and Architecture seminar and the HiPEAC General Cluster Meeting. We are already planning our next industrial workshop for the last quarter of this year. Have a look at our website for future information on this event!

As planned, the HiPEAC community submitted a proposal for a new Network of Excellence to the first FP7

call for proposals, which closed last May. Our aim is to capitalise on the lessons learnt from HiPEAC in order to build an improved network, in which we are counting on many of you to join again as HiPEAC members; we also expect other excellent members to be able to join soon. Of course, we need the European Commission to keep on trusting us. I hope our good work so far, and our willingness and commitment to push forward our community research will help in convincing them. In the meantime, we will keep our spirits high, and continue to foster and undertake more activities of interest for our community. Moreover, at least 10 STREPs and one IP were submitted to this call, showing that HiPEAC is also an excellent breeding ground for new ideas. All the best, and see you in L'Aquila!

Mateo Valero
HiPEAC Coordinator

Message from the project officer

ARTEMIS-JTI: Industry, Member States and the EU pool resources in embedded computer systems

On the 15 May 2007 the European Commission adopted a proposal to launch the first ever Europe-wide, public-private R&D partnerships. The Commission presented two Joint Technology Initiatives

(JTIs) on Embedded Computing Systems and Innovative Medicines. These JTIs will pool industry, Member States and Commission resources into targeted research programmes



The ARTEMIS-JTI (Advanced Research and Technology for Embedded Intelligence

and Systems - Joint Technology Initiative) will be a public-private partnership between the Commission, Member States, industry and research institutes. The participation of industry and research organisations is orchestrated through ARTEMISIA (the ARTEMIS Industrial Association), which was established in January 2006 under Dutch law by Philips, ST Microelectronics, Thales, Nokia and

Mercè Grierà i Fisa Mercè.Grierà-i-Fisa@ec.europa.eu
Please take note of the European Commission's new e-mail addresses!

The ACM/IEEE Eckert-Mauchly Award

The ACM/IEEE Eckert-Mauchly Award was institutionalized in 1979 and is presented annually to an individual for contributions to computer and digital systems architecture. It is the most prestigious award in computer architecture and is often nicknamed as the "The Nobel Prize of Computer Architecture." The winner of this prestigious award this year is **Professor Mateo Valero**, UPC, with the following citation:



"For extraordinary leadership in building a world-class computer architecture research center, for seminal contributions in the areas of vector computing and multithreading, and for pioneering basic new approaches to instruction-level parallelism"

Among the 29 individuals that have received the award so far, Mateo is the third European. He is in great company with Maurice Wilkes and Tom Kilburn who received it in 1980 and 1983, respectively, for outstanding contributions to our field. Following



the tradition, the award was presented to him at the award luncheon at the 34th International Symposium on Computer Architecture in San Diego on June 12.

As reflected in the citation, Mateo's contributions to our field have an outstanding breadth. While many of us only manage to focus on a few topics in computer architecture over our professional career, Mateo has made contributions to such diverse areas of our field as vector architectures, superscalar, and, more recently, multithreaded architectures to name a few.

As also reflected in the citation, his organizational and human qualities have been instrumental to form a world-class computer architecture center in Barcelona, as well as joining computer architects and compiler designers under our so successful European HiPEAC umbrella. It is surprising that the technical and organizational excellence of his caliber can be found in a single human being. It is with pride that we congratulate him on a well-deserved award of this level.

On behalf of the HiPEAC Community
Per Stenström

DaimlerChrysler. It currently has over 100 members and new applications for membership are in the pipeline from industry, SMEs and research organisations. Members of ARTEMISIA can vote in elections, participate in key decisions, and shape the policies and evolution of ARTEMIS' "strategic research agenda". Membership also provides access to an extensive network of respected research partners. ARTEMIS is to be set up as a joint undertaking based in Brussels. Public research funding will be allocat-

ed, following open calls for proposals. The first of these is expected to be published in 2008.

The overall budget of the ARTEMIS initiative is €2.7 billion over seven years with around 60% coming from industry. In short, each Euro contributed by the Commission will leverage 7 euros of research effort, with €1.8 coming from Member States and €4.2 from Industry. Overall, the Commission is expecting to contribute €420 million during the

seven years, starting with €42.5 million in 2008.

Details on the launch of the ARTEMIS-JTI can be found at:

http://ec.europa.eu/information_society/newsroom/cf/itemlongdetail.cfm?item_id=3413 and people interested in entering into ARTEMISIA can apply at: <http://www.artemisoffice.org/DotNetNuke/ARTEMISIA/tabid/114/Default.aspx>

Mercè Griera-I-Fisa

Insight, Not (Random) Numbers: A Perspective on Embedded Benchmarking

The good example of the challenges of designing embedded systems is today's feature-rich smartphones. Cellphone consumers are increasingly receiving the majority of their computing needs from these handheld, multi-standard communicators. This demands requirements of yesterday's supercomputers in the form factor of a handheld. But whereas supercomputers were designed with assumptions of near limitless power, unconstrained form factor, and dedicated teams of programmers, cellphones have none of these luxuries. For this reason, embedded systems might be better thought of as highly efficient systems.

Solutions to the problems of embedded system design has often been less than scientific. One chief approach is to let the market decide what to build. This has obvious drawbacks of limiting vendors to evolutionary instead of revolutionary designs. The second popular approach is to let the industrial design of the phone sell it—form over function. Clearly that is not scientific. As engineering scientists, we approach the problem by measurement and modeling. In 1962, computer and communications pioneer Richard Hamming cast a sour note on the use of modeling, stating that “the purpose of computing [modeling] is insight, not numbers.” Hamming's caution still holds weight to us today in embedded systems design.

Hamming was not saying that we should model, get numbers, and then stop. What he said was that we could not go directly from modeling to insight. Numbers are not bad in and of themselves. They are an integral part of insight. Parsing Hamming's advice further, we can conclude that we need to develop a method to gain insight from numbers and guarantee the quality of the numbers so that you have a hope of gaining that insight!

The current state of affairs is to bombard




Figure 1: Function unit usage requirements to meet 85% of any cycle's issue needs in a superscalar processor. Results are for each of the seven EEMBC benchmark suites.

customers with numbers—a particular car can go from 0-100 kph in 4.9 sec, with 500 horsepower and 383 lb.-ft. of torque, and rev with a redline of 8,250, total top speed of 300 kph. This doesn't correspond to the needs of every car buyer. Surely the family car does not need these kinds of sportscar statistics, regardless of the desires of the family's father! What you plan to do with a system should dictate which aspect of performance matters. In buying a car, the decisions are clearer than in buying an embedded computer system, however.

Benchmarking has as its chief goal to serve as the stand-in, or the proxy, for the users. Benchmarks, if chosen carefully, represent what real users do with their systems. Or, that is the message that benchmark vendors would have you believe. But how can a user tell whether or not a benchmark matches what he or she does with a cellphone? The traditional answer to this problem is to give meaningful descriptions of the benchmarks. But does saying that “176.gcc is the GNU CC compiler” help a cellphone user decide that “176.gcc” is a good proxy for what they play to do with their phone? Can we be more scientific about this problem? The answer is benchmark characterization [1].

EEMBC is the Embedded Microprocessor Benchmark Consortium. It produces a suite of benchmarks for embedded systems. Member companies include the likes of ARM, Ltd., IBM, AMD, Freescale, Mips, AMD, Texas Instruments, etc. Recently, my research group looked at the

problem of benchmark characterization for EEMBC benchmarks. Figure 1 shows the results of characterizing the processor resource needs of the various suites of the EEMBC benchmark set. From this figure, you can see that automotive applications require a lot of memory access bandwidth (“LSU”), but not much floating point hardware. Telecomm benchmarks, however, do not require much memory access bandwidth (perhaps due to their serial nature). You can perform this exercise for the cache memory requirements, or the branch prediction capacity needs, etc.

A way to gain insight from benchmark results is to allow the consumers to characterize their own applications. We call this “WAID,” which is short for “what am I doing?” A WAID tool can take any application and develop the same kind of characteristics as we presented above for EEMBC. Then the user can match up their application's characteristics with the benchmarks. Perhaps the matching would result in the consumer deciding to purchase a Smart ForOne instead of a Porsche.

The other aspect of Hamming's advice is that we need high quality numbers. Even if we model the benchmarks, if the model was flawed we will not have a chance for insight. Today, it is past time to adopt aspects of engineering science that other fields of engineering have been using for decades (or even centuries).

SAMPLING

There are two ingredients to producing

high-quality numbers: high quality simulators and accurate representations of benchmarks. I will only address the latter in this article, as the former is well addressed elsewhere.

On the surface of it, the accurate representations of benchmarks given an accurate simulator should be an easy task: one runs the benchmark on the simulator, collects the results, and reports them. Unfortunately, the effort in simulation is not decreasing over time. I modestly call this “Conte’s Law”, namely that benchmark sizes scale to the computing power present for simulation, so that architectural simulators always take too long! SPECcpu 2006 is an example of this. We and others have found it to be two orders of magnitude more simulation-resource intensive than SPECcpu 2000.

How does one speed up simulation? There are many approaches today, many with impressive nicknames implying precision and intelligence. I will present a rather older concept: statistical sampling.

Statistical sampling is a cornerstone of statistical science. The goal of sampling is to summarize within a bounded precision an aspect of an overall population. In brief, sampling summarizes the whole from a small subset. Sampling for computer simulation is non-trivial because of several factors, like non-Normal distributions and indirect sampling.

It is easy to imagine that the underlying events in a computer are non-Normally distributed. The way that statistical sampling theory deals with such populations is to sample them randomly, and to oversample these populations. Although this sounds rather intuitive, the two major approaches today to reducing simulation time both use systematic (i.e., non-random) sampling regimens. One in fact uses a singular systematic technique that is not statistical at all. It is valid to poll every 10th person that walks through the door of a party, for example, if you are sure there is a random order to the line of people as they enter the party. But it is not valid if

there is a structure to it. With these techniques, precision is difficult to calculate. Precision—or in the parlance of sampling theory, confidence intervals—is the key to Hamming’s suggestion for high quality numbers. Why do we not sample randomly? This author is at a loss as to why this is the case today.

The second factor that makes computer simulation difficult for sampling is due to indirect sampling. Indirect sampling is the case where the population being sampled is different from the population being measured. An example of this would be asking people in Europe who their friends in the US will vote for for president. Clearly, it makes more sense to ask the citizens in the US directly! But in computing, we do not have this luxury. What we are sampling are events before they are simulated, and what we are measuring are cycle-by-cycle metrics about the system’s performance. Sampling the system’s performance directly would require an entire run of the benchmark, and save no effort nor time at all!

The problem with sampling events instead of cycle-by-cycle system performance is that at each sampling point, the state in the machine—the contents of the registers, the caches, etc—is entirely unknown. There are several solutions to this “state recovery” problem. One is to simulate the caches for the entire trace (i.e., not sampling cache behavior). But this approach, of course, results in a much slower simulation. There are other approaches to state recovery that use some more sophistication. Recently, we presented an example where the state is recovered by simulating events backwards until a known state is reached, and then simulating forwards to measure the performance of a sampling point [2].

If sampling is properly applied, then the precision of the results can be known using statistical confidence tests. The best known of these is the Student-t test, found in any undergraduate statistics textbook. An example of how to apply this to simulation, complete with a crib sheet of

the appropriate statistical equations, can be found in [3].

In summary, I believe that when Hamming said “the purpose of computing [modeling] is insight, not numbers” today we should hear that we need accurate and precise modeling to gain insight. Accuracy comes from using good benchmarks as proxies for the end users. Precision comes from applying statistics correctly when simulating so that we know the confidence in our results. I think Hamming would be quite happy with that!

[1] Conte, T.M.; Hwu, W.W.; “Benchmark Characterization,” Proceedings of the Twenty-Fourth Annual Hawaii International Conference on System Sciences, Jan. 1991.

[2] P. D. Bryan and T. M. Conte, “Reverse State Reconstruction for Sampled Microarchitectural Simulation,” Proceedings of the 2007 IEEE International Symposium on Performance Analysis of Systems and Software (San Jose, CA), April 2007.

[3] T. M. Conte, M. A. Hirsch, and K. N. Menezes, “Reducing state loss for effective trace sampling of superscalar processors,” in Proceedings of the 1996 International Conference, on Computer Design, (Austin, TX), Oct. 1996.

About Tom Conte

Tom Conte’s research is in the areas of microprocessor architecture, compiler code generation/optimization, performance evaluation and embedded computer systems. Conte is the



past chair of the ACM Special Interest Group on Microarchitecture (SIGMICRO), the past chair of the IEEE Computer Society Technical Committee on Microprogramming and Microarchitecture (TC-uARCH), and also a fellow of the IEEE. He was the editor-in-chief of the Journal of Instruction-Level Parallelism from 1997?2001 and 2002?2005. He is an associate editor of ACM Transactions on Embedded Computer Systems, ACM Transactions on Architecture and Compiler Optimization and IEEE Micro magazine.

hArtes

hArtes is an European-founded IP project (CONTRACT NUMBER 035143) in FP6. hArtes is an acronym for "Holistic Approach to heterogeneous Reconfigurable real Time Embedded Systems".



The complexity of future real-time embedded systems for consumer and professional products is becoming too big to design monolithic processing platforms.

The adoption of heterogeneous processing platform including combination of general-purpose processors, digital signal processors (or multi-cores devices that integrate both like Diopsis or OMAP) and reconfigurable hardware is the solution proposed by hArtes. To efficiently exploit such platforms, a new holistic approach to complex embedded system design is required.

The objective of hArtes is to simplify the development of embedded products based on heterogeneous reconfigurable embedded platforms. This objective will be achieved by:

- developing a toolchain and a methodology to deal with the different challenges of effective embedded system designs.
- designing a scalable heterogeneous and reconfigurable hardware platform that can be re-targeted to produce optimized real-time embedded systems.

The results will be evaluated using audio and video systems that support next-generation communication and entertainment. The benefit of hArtes for the industry is the possibility of moving the focus of embedded system design from the implementation details to the algorithm design allowing faster product innovation and differentiation.

The expected outcomes of the hArtes project are:

- 1) The hArtes toolchain supporting the heterogeneous reconfigurable platforms;
- 2) The hardware platforms, including heterogeneous reconfigurable devices;
- 3) A set of innovative and legacy applications in the audio and video domain validating the hArtes approach.

Innovations of the hArtes toolchain include:

- a framework that allows implementations of novel algorithms for design space exploration, supporting design partitioning automation, task transformation, choice of data representation, and metric evaluation for HW and SW components;
- a system synthesis tool producing near-optimal implementations that best exploits the capability of each type of processing element; dynamic HW reconfigurability can be exploited to support system upgrade or adaptation to operating conditions;
- diagrammatic and textual formats in algorithm description and exploration.

Innovative and legacy applications will be considered. The innovative ones will be used to validate the hArtes concept: keep the focus of embedded system design to higher added value activities such as algorithm development. The porting of legacy applications allows having clear metrics to assess the benefits of the hArtes approach.

More information can be found at: <http://www.hartes.org>

HiPEAC Activity

ACACES 2007: Third international summer school on Advanced Computer Architecture and Compilation for Embedded Systems



July 15-21, 2007, L'Aquila, Italy

This year, we received 289 applications from 33 countries (about 50% of the applicants already attended a previous edition, for the other 50% it was their first attempt to attend the summer school). Since the attendance is limited to 200, we again had to make some tough decisions. The invitation criteria were: geographical distribution, gender balance, balance between academics and people from industry, between senior and junior people.

The HiPEAC 60 grants have been distributed among HiPEAC members and attendees from new member states.

On Wednesday afternoon, 78 participants will present their work during a huge poster session. The abstracts are again to be published by Academia Press in a nice poster proceedings.

The summer school will, as before, be a great networking event where several HiPEAC clusters and affiliated projects will meet. The program commit-

New Book


Manish Verma, Peter Marwedel:
Advanced Memory Optimization
Techniques for Low-Power
Embedded Processors, Springer,
2007, ISBN 13 978 1 420 5896-7

Abstract: Memories are known to be a very crucial component of current and future embedded systems. For some of today's systems, they consume more than 50 % of the energy required for information processing. Furthermore, the gap between the speeds of processors and memories is increasing. There is the very realistic thread that the speed of future systems will be limited by the speed of the memories, less so by the speed of processing. This has led to the term "memory wall", designating the insurmountable barrier of memory speed limits. Moreover, classical techniques for solving the memory wall problem (such as caches) are aiming at improving the average case access times of memories. For these techniques, the predictability of their timing behavior

is rather low, contrasting with the fact that many embedded systems are real-time systems.

Scratch pad memories (SPMs) are potential means of reducing the severity of the problem. SPMs are small memories mapped into the address space of the system. They can be used successfully if frequently used memory objects are mapped into the corresponding address range. If this is guaranteed, SPMs provide fast, energy-efficient and timing-predictable access to memories. However, this approach requires the use of tools mapping hot spots of applications to the address range. This book extends the knowledge about mapping algorithms by a major step.

The book starts with an introduction. The problem to be tackled is clearly described. The second chapter presents some of the related work. Some fundamental facts concerning power and energy dissipation as well as their reduction are explained. The remaining chapters describe various



approaches for mapping data and instructions to SPMs. This book on memory-architecture aware compilation is based on the thesis of Manish Verma, a former PhD student at the University of Dortmund.

tee of the HiPEAC conference will again take place at the end of the same week.

The massive and still growing 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 as one of the flagship events of our community.

Koen De Bosschere
ACACES 2007 Coordinator



RWTH AACHEN
UNIVERSITY

Within SHAPES, the responsibility of ISS is to provide other partners with a fast and accurate simulation platform for



This version has been built with CoWare's new Virtual Platform Designer product in order to meet high speed requirements.

HW prototype is available, the VSP gives the programmer more detailed observation capabilities by providing full access to platform resources such as processor

Community news

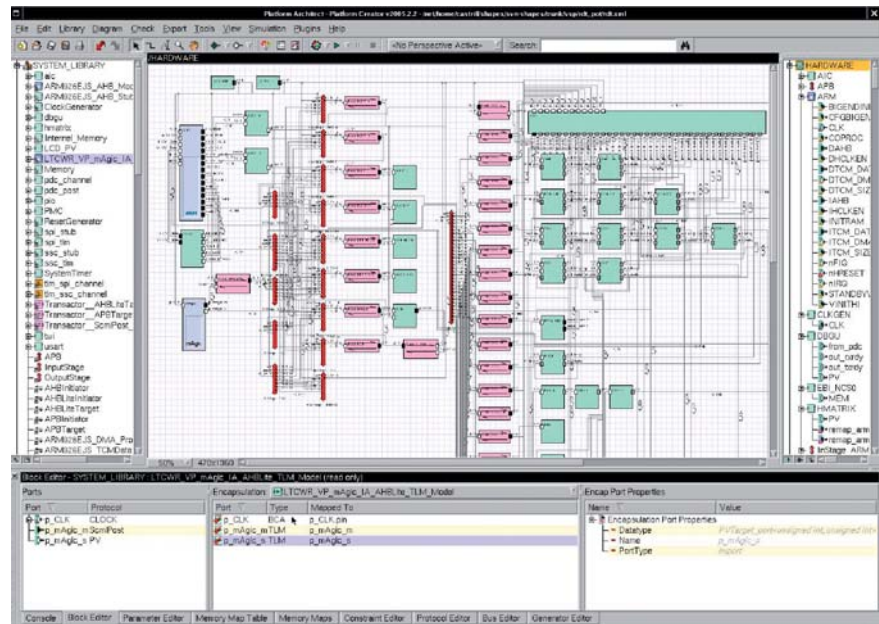
Per Stenström was nominated IEEE Fellow for contributions to design of high-performance memory systems.





About ISS

The Institute for Integrated Signal Processing Systems at RWTH Aachen University of Technology focuses on the design of wireless communication systems. A number of successful design automation tools for application specific systems have been developed at ISS, with LISATek (now available from CoWare Inc.) being the most recent example. Current research activities concentrate on multi-processor SoC design tools as well as compilers for embedded processors. Further information can be found at www.iss.rwth-aachen.de.



Virtual SHAPES platform (HW view)

registers, registers of memory mapped peripherals and the general system memory. VSP permits software partners in charge of hardware dependent software (HdS) to develop drivers and communication primitives concurrently with the development of the HW. In addition to this, VSP will empower high level software exploration (task and communication allocation and scheduling) by providing accurate performance esti-

mates through a well-defined query interface.

Taking the initial VSP as a starting point, ISS activities will now move towards a multi-tile virtual platform. This will be available long before the actual silicon and thus will also allow system architects to perform high-level architecture exploration. In order to cope with the increasing demands for higher simulation speed and the complexity of multi

processor system simulation, ISS is focusing its research on new simulation techniques that use higher abstraction levels and take advantage of multi processor host machines. More information about SHAPES is available at www.shapes-p.org.

HiPEAC Journal

Transactions on High-performance Embedded Architectures and Compilers

The following papers were accepted for the second issue of Volume 2:

W. Choi, S-J Park, and M. Dubois.

Accurate Instruction Pre-scheduling in Dynamically Scheduled Processors.

D. Chanet, J. Cabezas, E. Morancho, N. Navarro, K. De Bosschere.

Linux Kernel Compaction through Cold Code Swapping.

H. Vandierendonck and A. Sez nec.

Fetch Gating Control through Speculative Instruction Window Weighting.



Advanced Link-Time Program Analysis

By Ludo Van Put
(ludo.vanput@elis.ugent.be)
Prof. Koen De Bosschere, Ghent University
May 24, 2007

This work discusses three whole-program analysis that efficiently compute important information for a link-time program rewriter. The interprocedural dominator analysis extends the dominator relation to a complete program. We show that in the interprocedural case, there is no unique immediate domina-

tor for every node in the graph. Using a new efficient algorithm, the dominator relation of control flow graphs with several hundreds of thousand of basic blocks can be computed within tens of seconds.

Linear-constant propagation computes linear relations between two registers at every point in the program. The resulting information is used to determine the stack layout of procedures and to eliminate unnecessary spill code from the program, hereby reducing the number of store operations by up to 7%.

A third analysis revisits the interprocedural register liveness analysis by incorporating the available debug information. From the debug information, function signatures are derived which provide additional information on the register usage in a program. On average, more than one extra dead register is found throughout the program. This result serves as a proof of concept for using compiler-generated high-level information in a link-time program rewriter.

Low Power Design of Block-Based Video Codecs

By Kristof Denolf
(denolf@imec.be), **Prof. Henk Corporaal, Dr. Diederik Verkest, Technische Universiteit Eindhoven**
June 7, 2007

The improving display resolution of new mobile video appliances increases the throughput requirements of video codecs and further complicates the challenges encountered during their cost-efficient design. In contrast, their energy and heat dissipation limitations create the demand for low-power implementations.

This PhD proposes a memory and

communication centric design methodology for dedicated implementations. Its high-level steps combine memory and algorithmic optimizations on a sequential executable description. Then, a partitioning exploration introduces parallelism using a cyclostatic dataflow model. To maintain the effect of the high-level optimizations, also implementation specific aspects of communication channels, like using shared buffers, are expressed without extending the model of computation. Consequently, all analysis potential at design time is preserved. Aiming at dedicated hardware, these channels

are implemented as a restricted, but sufficient set of communication primitives. They enable an automated RTL test and development strategy for rigorous functional testing. In this way, the design time is reduced.

The introduced methodology is applied to the design of a high-performance MPEG-4 video encoder sustaining 30 4CIF frames per second. The core consumes only 71 mW in a 180 nm, 1.62V UMC technology. This energy efficiency is equivalent to the state of the art for high resolution video encoders.

The Molen Compiler for Reconfigurable Architectures

By Elena Moscu Panainte
(E.Panainte@ewi.tudelft.nl), **Prof. Koen Bertels, TUDelft, The Netherlands**
June 20, 2007

In this dissertation, we present the Molen compiler framework that targets reconfigurable architectures under the Molen Programming Paradigm. More specifically, we introduce a set of compiler optimizations that address one of the main shortcomings of the reconfigurable archi-

tectures, namely the reconfiguration overhead. The proposed optimizations are based on data flow analyses both at the intraprocedural and the interprocedural level and take into account the competition for reconfigurable hardware resources and the spatio-temporal mapping. The hardware configuration instructions are scheduled in advance of hardware execution instructions, in order to exploit the available parallelism between the hardware configuration phase and the sequential execution on the core

processor. The intraprocedural optimization uses the min s-t cut graph algorithm to reduce the number of executed hardware configurations by identifying the redundant hardware configurations. We also introduce two allocation algorithms for the reconfigurable hardware resources that aim to minimize the total reconfigured area and to maximize the overall performance gain. Based on profiling results and software/hardware estimations, the compiler optimization and allocation algorithms generate optimized

code for the spatio-temporal constraints of the target reconfigurable architecture and input application. Additionally, they guide the selection of the hardware/software execution of the operations candidate for reconfigurable hardware execution. In order to evaluate the Molen compiler, we first present an experiment with a multimedia benchmark application com-

plied by the Molen compiler and executed on the Molen polymorphic media processor with an overall speedup of 2.5 compared to the pure software execution. Subsequently, we estimate that the intraprocedural compiler optimization contributes to up to 94 % performance improvement compared to the pure software execution, while the intraprocedural compiler

optimization and the allocation algorithms significantly reduce the number of executed reconfigurations for the considered benchmarks. Finally, we determine that the important performance impact of our compiler optimizations and allocation algorithms increases for the future faster FPGAs.

Characterization and Reduction of Memory Usage in 64-Bit Java Virtual Machines

**By Kris Venstermans (kris.venstermans@elis.ugent.be), Prof. Koen De Bosschere, Prof. Lieven Eeckhout, Ghent University, Belgium
June 25, 2007**

Modern general purpose computer systems are more and more equipped with 64-bit processors, while a few years ago these computer systems almost exclusively contained 32-bit processors. The use of 64-bit computer systems has both advantages as well as disadvantages over the use of

32-bit computer systems.

The most visible advantage of 64-bit computer systems is that they have a much larger memory addressability and so they are able to execute programs that need lots of memory. However, the most important disadvantage of 64-bit computer systems is that the executed programs also use more memory than when executed on a 32-bit computer system.

This dissertation first characterizes the memory usage and overall performance impact of the transition from 32-bit to 64-bit computing for Java appli-

cations. We observe that the average object size increases by 45.3%. Next we propose two techniques that improve the memory usage of 64-bit applications in the context of a Java Virtual Machine. The first technique, Object-Relative Addressing, reduces the memory usage by compressing object pointers and the second technique, Selective Typed Virtual Addressing, reduces the size of the object header by e.g., allocating objects of the same type in typed memory segments.

Instrumentation Techniques for Layered Execution Environments

**By Jonas Maebe (jmaebe@elis.ugent.be), Prof. Koen De Bosschere, Ghent University,
June 29, 2007**

By delegating responsibilities to different components, layers and virtual machines, modern computer programs are able to solve very complex problems without the need for a single programmer to understand the details of the whole execution environment. However, analyzing such execution environments is hard. Firstly, all executed code must be analyzable. Secondly,

all observed facts have to relate back to higher level concepts in the observed program. Finally, there is the problem of specifying what should be instrumented. We describe a combination of three techniques to tackle these problems. First of all, to be able to instrument all executed code, we base our instrumentation infrastructure on dynamic binary instrumentation (DBI). The second part of our solution is vertical instrumentation, which means that multiple execution layers assist in making the instrumentation possibilities as rich as possible.

The third part concludes our work with

the introduction of Aspect-Oriented Instrumentation (AOI). It is derived from Aspect-Oriented Programming (AOP), a technique introduced to more easily express concerns which crosscut entire programs in a modular fashion.

In summary, in our work we introduce a vertically integrated approach to instrumentation, based on using DBI which is assisted by information coming from other execution layers. The end result is that it becomes easier to accurately instrument complex execution environments.


Upcoming events

Euro-Par 2007

Rennes, France, 28-31 August 2007, <http://europar2007.irisa.fr/>

HPPC: Workshop on Highly Parallel Processing on a Chip

IRISA, Rennes, France, August 28, 2007, <http://www.hppc-workshop.org/>



PATMOS'07

Göteborg, Sweden, September 3-5, 2007, <http://www.ce.chalmers.se/research/conference/patmos07/>



Parallel Computing with FPGA's (with ParCo 2007)

September 4-7, 2007, Aachen, Germany, <http://www.elis.ugent.be/parafpga/>

PACT 2007: The Sixteenth International Conference on Parallel Architectures and Compilation Techniques

Brasov, Romania, September 15-19, 2007, <http://parasol.tamu.edu/pact07/>



Call for Paper for the MEDEA Workshop 2007 (MEemory performance: DEaling with Applications, systems and architecture), held with IEEE/ACM PACT, September 15-19, 2007, Brasov, Romania, <http://garga.iet.unipi.it/medea07>

CellSim: a Modular Simulator for Heterogeneous Chip Multiprocessors (held with PACT)


September 15, 2007, <http://parasol.tamu.edu/pact07/tutorials-workshops.php#T2>

GREPS 2007: International Workshop on GCC for Research in Embedded and Parallel Systems

(held with PACT) September 16, 2007, <http://sysrun.haifa.il.ibm.com/hrl/greps2007/>

IISWC-2007 : 2007 IEEE International Symposium on Workload Characterization

Boston, MA, USA, September 27-29 2007, <http://csl.cse.psu.edu/iiswc2007/index.htm>



CASES'2007

Salzburg, Austria, September 30 - October 5, 2007, <http://www.casesconference.org/>

Conference on Nanotechnologies

Braga, Portugal, November 20-21, 2007

HiPEAC Conference 2008

Göteborg, Sweden, January 27-29, 2008, <http://www.hipeac.net/hipeac2008>



CFP: EuroSys 2008

Glasgow, April 2-4, 2008, <http://www.dcs.gla.ac.uk/conferences/eurosys2008/>

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