



PEPPHER

Performance Portability and Programmability
for Heterogeneous Many-core Architectures (PEPPHER)

FP7-248481

2011 PEPPHER Workshop

Saturday, January the 22nd, 2011
Galaxy Hotel, Heraklion, Crete, Greece
in conjunction with the HiPEAC'11 conference
<<http://www.peppher.eu/hipeac11>>





Context

The emergence of multi-core processors with often significantly different architectures and capabilities not only reflects different application needs (graphics, signal processing, numerical acceleration, power constraints) but also a potentially critical lack of convergence toward a common base architecture upon which parallel applications can be built. Ensuring portability of real applications across such differing (heterogeneous) multi-core platforms, including means of reasoning about performance (processor utilization, time-to-solution, power consumption), is evidently a major issue for the software-intensive industry. Such methods must also extend to Systems on Chip (SoC processors) -- hybrids -- that themselves are composed of different types of parallel processing elements. The European FP7 project PEPPHER (for "PERformance Portability and Programmability of Heterogeneous many-core aRchitectures") is an ambitious research effort that addresses some of these issues.

The aim of workshop is to provide a forum for discussion and exchange of opinions and experience on PEPPHER related issues. Furthermore, relevant PEPPHER recent achievements will be presented and feedback will be sought.

Main topics

Workshop topics include, but are not limited to,

- 1) Methods and tools for component-based software;
- 2) Portable compilation techniques;
- 3) Adaptive, auto-tuned algorithms and data structures;
- 4) Efficient, flexible run-time systems;
- 5) Hardware support mechanisms for auto-tuning, synchronization and scheduling.

Workshop format

The workshop program includes a keynote, invited talks, and a panel.

Venue

This workshop is organized in conjunction with the 6th International Conference on High-Performance and Embedded Architectures and Compilers (HiPEAC'11) Heraklion, Crete, Greece, January 22, 2011.

Organizers and their affiliations

Sabri Pllana and Jesper Larsson Träff
University of Vienna
Faculty of Computer Science
Research Group Scientific Computing
Nordbergstrasse 15/C/3
1090 Vienna, AUSTRIA
T +43-1-4277-39411 | +43-1-4277-39432
F +43-1-4277-9394
E pllana AT par.univie.ac.at | traff AT par.univie.ac.at



WORKSHOP PROGRAM

09:00-10:30 SESSION 1

Chair: Sabri Pllana, University of Vienna

09:00-10:00 Keynote Address: Markus Püschel (ETH Zürich) "Automatic Performance Tuning and Machine Learning"

10:00-10:30 Siegfried Benkner (University of Vienna) "PEPPHER Vision & Overview"

10:30-11:00 Coffee Break**11:00-12:30 SESSION 2**

Chair: Bev Bachmayer, Intel

11:00-11:30 Martin Wimmer (University of Vienna) "Work-stealing for mixed-mode parallelism by deterministic team-building"

11:30-12:00 Daniel Cederman (Chalmers University) "Data Structures in Work-Stealing"

12:00-12:30 Cedric Augonnet (INRIA) "StarPU: A Unified Runtime System for Heterogeneous Multicore Architectures"

12:30-14:00 Lunch Break**14:00-15:30 SESSION 3**

Chair: Koen De Bosschere, Ghent University

14:00-14:30 Kunle Olukotun (Stanford University) "Taming Heterogeneous Parallelism with Domain Specific Languages"

14:30-15:00 François Bodin (CAPS entreprise) "Incremental Migration of C and Fortran Applications to GPGPU using HMPP"

15:00-15:30 Ben Juurlink (TU Berlin) "The ENCORE Project - Enabling Technologies for a Programmable Many-core"

15:30-16:00 Coffee Break**16:00-17:00 SESSION 4**

Chair: Jesper Larsson Träff, University of Vienna

Panelists: François Bodin (CAPS entreprise), Ben Juurlink (TU Berlin), Christoph Kessler (Linköping University), and Kunle Olukotun (Stanford University).

Panel discussion "on the convergence/standardization/future of directive/annotation-based languages for programming heterogeneous multi/many-core systems."



KEYNOTE ADDRESS

Markus Püschel (ETH Zürich) "Automatic Performance Tuning and Machine Learning"**Abstract**

Automatic performance tuning has emerged as a paradigm complementing traditional compilers to port software and performance between platforms. Several techniques have proven useful including adaptive libraries, program generation, domain-specific languages, and architecture models. However, one technique is shared by almost all approaches: search for the fastest among a set of alternative implementations. Typically the search space is huge and hence the search is costly. This may be bearable in offline tuning (e.g., ATLAS) that is performed during installation but becomes cumbersome in online tuning (e.g., FFTW) that is performed at runtime since the input size is required. We argue that machine learning, which is already studied in the compiler community, can solve this problem and should be added to the portfolio of performance tuning tools. As example we show a successful approach to automatically convert Spiral-generated online-tunable transform libraries into offline-tunable ones.

About the Keynote Speaker

Markus Püschel is a Professor of Computer Science at ETH Zurich, Switzerland since 2010. Before, he was a Professor of Electrical and Computer Engineering at Carnegie Mellon University, where he still has an adjunct status. He received his Diploma (M.Sc.) in Mathematics and his Doctorate (Ph.D.) in Computer Science, in 1995 and 1998, respectively, both from the University of Karlsruhe, Germany. He served on the editorial boards of the IEEE Transactions on Signal Processing and the IEEE Signal Processing Letters, was a Guest Editor of the Proceedings of the IEEE and the Journal of Symbolic Computation, and served on various program committees of conferences in computing, compilers, and programming languages. He is a recipient of the Outstanding Research Award of the College of Engineering at Carnegie Mellon and the Eta Kappa Nu Award for Outstanding Teaching. He also holds the title of Privatdozent at the University of Technology, Vienna, Austria. In 2009 he cofounded SpiralGen Inc.



PANEL DISCUSSION

"On the convergence/standardization/future of directive/annotation-based languages for programming heterogeneous multi/many-core systems."

Moderator

Jesper Larsson Träff, University of Vienna

Panelists

François Bodin (CAPS entreprise), Ben Juurlink (TU Berlin),
Christoph Kessler (Linköping University) and Kunle Olukotun (Stanford University).

Description

There is currently a number of proposed and already partly successful directives and annotation based language extensions "on the market" for making programming of heterogeneous (accelerator-based) multi/many-core systems feasible/productive and efficient/performant. The panelists will briefly summarize and contrast some of the proposals with a view to larger questions like,

- will this approach eventually be feasible and efficient? Is this the "right" approach? What can it and what can it not provide?
- why is there currently little convergence between such proposals?
- or are the underlying assumptions already similar enough that a convergence on common standards can be expected to happen?

The panelists will have 5-8 minutes to explain their view. This will hopefully provoke a lively discussion involving also the workshop participants.