Introducing ACM SIGPLAN International Workshop on Adaptive Self-Tuning Computing Systems for the Exaflop Era http://exadapt.org

co-located with PLDI 2011 at FCRC 2011
San Jose, CA, USA, June 5th, 2011

Grigori Fursin
UVSQ, France
grigori.fursin@exascale-computing.eu

Robert Hundt
Google, CA, USA
rhundt@google.com

Jason Mars
University of Virginia, USA
jom5x@cs.virginia.edu

Yuriy Kashnikov
UVSQ, France
yuriy.kashnikov@exascale-computing.eu

"It is not the strongest of the species that survives, or the most intelligent; it is the one most capable of change"

attributed to Charles Darwin

Modern large scale computing systems are rapidly evolving and may soon feature millions of cores with exaflop performance. However, this leads to a tremendous complexity with an unprecedented number of available design and optimization choices for architectures, applications, compilers and run-time systems. Using outdated, non-adaptive technology results in an enormous waste of expensive computing resources and energy, while slowing down time to market.

This workshop is intended to become a regular inter-disciplinary forum for researchers, practitioners, developers and application writers to discuss ideas, experience, methodology, applications, practical techniques and tools to improve or change current and future computing systems using self-tuning technology. Such systems should be able to automatically adjust their behavior to multi-objective usage scenarios at all levels (hardware and software) based on empirical, dynamic, iterative, statistical, collective, bio-inspired, machine learning and alternative techniques while fully utilizing available resources.

All full and position papers have been peer-reviewed and had to include unpublished ideas on how to simplify, automate and standardize the design, programming, optimization and adaptation of large-scale computing systems for multiple objectives to improve performance, power consumption, utilization, reliability and scalability including the following topics:

- whole system parameterization and modularization to enable self-tuning across the whole hardware and software stack
- transformation space of static, JIT and source-to-source compilers
- run-time resource management/scheduling

- task/process/thread/data migration
- design space of architectures including heterogeneous multi-cores, accelerators, memory hierarchy and IO
- propagation and usage of the feedback between various system layers
- static and dynamic code and data partitioning/modification for self-tuning
- application conversion to support multi-level, hybrid parallelization
- modification of existing tools and applications to enable auto-tuning
- resource and contention aware scheduling
- performance, power and reliability evaluation methodologies
- scalable performance evaluation tools
- detection, classification, and mitigation of resource contentions
- collaborative optimization repositories and benchmarks
- characterization of static program constructs
- characterization of dynamic program behavior under various system load scenarios
- software/hardware co-design and co-optimization
- analysis of interactions between different parts of a large application
- prediction of optimizations and architectural designs based on prior knowledge
- scalable system and processor simulation
- hardware support for self-tuning and scheduling
- virtualization
- fault-tolerance

We would like to thank all our Program Committee members for providing very deep and detailed reviews on time that allowed us to select 8 high quality papers to appear at this workshop and in ACM Digital Library (International Conference Proceedings Series, ISBN 978-1-4503-0708-6):

- *Erik R. Altman*, IBM TJ Watson, USA
- *David H. Bailey*, Lawrence Berkeley National Laboratory, USA
- *Steve Blackburn*, Australian National University, Australia
- *Wenguang Chen*, Tsinghua University, China
- *Keith Cooper*, Rice University, USA
- *Lieven Eeckhout*, Ghent University, Belgium
- *Julia Fedorova*, Intel, Russia
- *Rajiv Gupta*, University of California, Riverside, USA
- *William Jalby*, UVSQ, France
- *Geoff Lowney*, Intel, USA
- *Bernd Mohr*, Julich Supercomputing Centre, Germany
- *Tipp Moseley*, Google, USA
- *Toshio Nakatani*, IBM Tokyo Research Lab, Japan
- *Michael O'Boyle*, University of Edinburgh, UK
- *Kunle Olukotun*, Stanford University, USA
- *David Padua*, UIUC, USA
- *Keshav Pingali*, University of Texas at Austin, USA
- *Markus Puschel*, ETH Zurich, Switzerland
- *Mary Lou Soffa*, University of Virginia, USA
- *Richard Vuduc*, Georgia Tech, USA
- *Ben Zorn*, Microsoft, USA

We would also like to thank *Prof. Katherine Yelick* (LBNL and UC Berkeley, USA) who kindly agreed to give a keynote on Autotuning in the Exascale Era.

Finally, we are grateful to *Prof. Mary Hall* (PLDI general chair, University of Utah, USA) for the guidelines when submitting this workshop proposal for PLDI 2011.