

Organizers:

Grigori Fursin

INRIA Saclay, France

Robert Hundt

Google, CA, USA

Jason Mars

University of Virginia, USA

Yuriy Kashnikov

Exascale Computing Research, France

Program Committee:

Erik R. Altman

IBM TJ Watson, USA

Jose Nelson Amaral

University of Alberta, Canada

David H. Bailey

Lawrence Berkeley National Laboratory, USA

Steve Blackburn

Australian National University, Australia

Francois Bodin

CAPS Entreprise, France

Franck Capello

INRIA, France / UIUC, USA

Koen DeBoschere

Ghent University, Belgium

Rajiv Gupta

University of California, Riverside, USA

Anton Likhomotov

ARM, UK

Geoff Lowney

Intel, USA

Allen Malony

University of Oregon, USA

Bernd Mohr

Jülich Supercomputing Centre, Germany

Tipp Moseley

Google, USA

Toshio Nakatani

IBM Tokyo Research Lab, Japan

Nacho Navarro

UPC, Spain

Boyana Norris

Argonne National Laboratory, USA

David Padua

UIUC, USA

Keshav Pingali

University of Texas at Austin, USA

Markus Püschel

ETH Zurich, Switzerland

Xipeng Shen

College of William & Mary, USA

Allan Snively

San Diego Supercomputer Center, USA

Felix Wolf

Aachen University, Germany

Chengyong Wu

ICT, China

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 Using outdated, non-adaptive technology results in an enormous waste of expensive computing resources and energy, while slowing down time to market.

The 2nd International Workshop on Self-tuning, Large Scale Computing Systems for Exaflop Era (EXADAPT) is an interdisciplinary 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All papers will be peer-reviewed including short position papers and should include unpublished ideas on how to simplify, automate and standardize the design, programming, optimization and adaptation of large-scale computing systems for multiple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.

Important Dates:

Final submissions due: **January 6, 2012 (23:59:59 submitter's time zone)**

Author notification: **January 20, 2012**

Final paper: **February 10, 2012**

Paper Submission Guidelines:

We invite papers in two categories:

Full papers should be at most 12 pages long including bibliography and appendices. Papers in this category are expected to have relatively mature content. Full paper presentations will be 25 minutes each.

Position papers should be at most 6 pages long including bibliography and appendices. Preliminary and exploratory work are welcome in this category, including wild & crazy ideas. Position paper presentations will be 10 minutes each. Authors submitting papers in this category must prepend the phrase Position Paper: to the title of the submitted paper.

Submissions should be PDF documents typeset in the ACM proceedings format using 10pt fonts. Both full and position papers must describe work not published in other refereed venues (see the SIGPLAN republication policy for more details).

The proceedings of this workshop will be published in the ACM International Conference Proceedings Series (ISBN 978-1-4503-1147-2) and ACM Digital Library.

Prior workshops and proceedings:

1st ACM SIGPLAN EXADAPT 2011 at PLDI/FCRC 2011, San Jose, USA

<http://exadapt.org/2011>

Proceedings:



Sponsors:

