# Lecture Notes in Computer Science 8551

*Commenced Publication in 1973*
Founding and Former Series Editors:
Gerhard Goos, Juris Hartmanis, and Jan van Leeuwen

## Editorial Board

More information about this series at http://www.springer.com/series/7407

Stephen Jarvis · Steven Wright
Simon Hammond (Eds.)

# High Performance Computing Systems

## Performance Modeling, Benchmarking and Simulation

4th International Workshop, PMBS 2013
Denver, CO, USA, November 18, 2013
Revised Selected Papers

Springer

*Editors*
Stephen Jarvis
Steven Wright
University of Warwick
Coventry, West Midlands
UK

Simon Hammond
Sandia National Laboratories CSRI
Albuquerque, NM
USA

ISSN 0302-9743 ISSN 1611-3349 (electronic)
ISBN 978-3-319-10213-9 ISBN 978-3-319-10214-6 (eBook)
DOI 10.1007/978-3-319-10214-6

Library of Congress Control Number: 2014952188

LNCS Sublibrary: SL1 – Theoretical Computer Science and General Issues

Springer Cham Heidelberg New York Dordrecht London



Printed on acid-free paper

Springer is part of Springer Science+Business Media (www.springer.com)

# 4th International Workshop on Performance Modeling, Benchmarking and Simulation of High-Performance Computing Systems (PMBS 2013)

This volume contains the 14 papers that were presented at the 4th International Workshop on Performance Modeling, Benchmarking and Simulation of High-Performance Computing Systems (PMBS 2013), which was held as part of the 25th ACM/IEEE International Conference for High-Performance Computing, Networking, Storage and Analysis (SC 2013) at the Colorado Convention Center in Denver during November 17–22, 2013.

The SC conference series is the premier international forum for high-performance computing, networking, storage, and analysis. The conference is unique in that it hosts a wide range of international participants from academia, national laboratories, and industry; this year's conference attracted over 10,000 attendees and featured over 350 exhibitors in the industry's largest HPC technology fair.

This year's conference was themed *HPC Everywhere, Everyday*, recognizing the impact that high-performance computing has on all aspects of society around the world, from weather forecasting, drug design, and finance, to aerospace, national security, and commerce.

SC offers a vibrant technical program, which includes technical papers, tutorials in advanced areas, birds of a feather sessions (BoFs), panel debates, a doctoral showcase, and a number of technical workshops in specialist areas (of which PMBS is one).

The focus of the PMBS 2013 workshop was comparing high-performance computing systems through performance modeling, benchmarking, or the use of tools such as simulators. We were particularly interested in receiving research papers that reported the ability to measure and make trade-offs in hardware/software co-design to improve sustained application performance. We were also keen to capture the assessment of future systems, for example, through work that ensured continued application scalability through peta- and exa-scale systems.

The aim of the PMBS 2013 workshop was to bring together researchers from industry, national labs, and academia, who were concerned with the qualitative and quantitative evaluation and modeling of high-performance computing systems. Authors were invited to submit novel research in all areas of performance modeling, benchmarking, and simulation, and we welcomed research that combined novel theory and practice. We also expressed an interest in submissions that included analysis of power consumption and reliability, and were receptive to performance modeling research that made use of analytical methods as well as those based on tracing tools and simulators.

Technical submissions were encouraged in areas including: performance modeling and analysis of applications and high-performance computing systems; novel techniques and tools for performance evaluation and prediction; advanced simulation techniques and tools; micro-benchmarking, application benchmarking, and tracing; performance-driven code optimization and scalability analysis; verification and

validation of performance models; benchmarking and performance analysis of novel hardware; performance concerns in software/hardware co-design; tuning and auto-tuning of HPC applications and algorithms; benchmark suites; performance visualization; real-world case studies; studies of novel hardware such as Intel Xeon Phi coprocessor technology, NVIDIA Kepler GPUs, and AMD Fusion APU.

# PMBS 2013

We received an excellent number of submissions for this year's workshop. This meant that we were able to be very selective in those papers that were chosen; the acceptance rate for full papers was approximately 30%. The resulting papers show worldwide programs of research committed to understanding application and architecture performance to enable peta-scale computational science.

Contributors to the workshop included Argonne National Laboratory, ETH Zurich, Georgia Institute of Technology, Inria, Indiana University, Lawrence Berkeley National Laboratory, NASA, Oak Ridge National Laboratory, San Diego Supercomputer Center, Sandia National Laboratories, Tokyo Institute of Technology, University of Grenoble, University of Tennessee, US National Energy Research Scientific Computing Center, amongst others.

Several of the papers are concerned with performance benchmarking and analysis (see Sect. A). The paper by Jeffrey Vetter et al. quantifies the architectural requirements of contemporary extreme-scale scientific applications. The paper by Subhash Saini et al. presents a performance evaluation of NASA's Pleiades, one of the world's most powerful supercomputers, using scientific and engineering applications. Matthew Cordery and colleagues document an analysis of the performance of the Cray XC30 using NERSC benchmarks, and contrast these results with those from the Cray XE6 and an IBM BG/Q. The paper by Pericàs et al. proposes a low-overhead methodology with which to benchmark data reuse in task-parallel runtimes, and in so doing correlates reuse distance with processor and memory configuration.

Section B of the proceedings collates papers concerned with performance modeling and simulation. Scott Levy and colleagues use simulation to identify system characteristics necessary for the exploration of resiliency in high-performance computing systems and applications. The paper by Collin McCurdy et al. introduces two analytic techniques with which to characterize and understand the impact of hardware and software data prefetching on scientific application performance. The performance modeling of the 3D gyrokinetic toroidal code GTC is examined by Matthew Anderson et al. in order to improve the efficiency and scalability of this code on many-tasking runtime systems. A novel flow-based hybrid network model for accurately simulating MPI applications on Ethernet/TCP networks is proposed by Paul Bédaride et al. The paper by Joo Hwan Lee et al. presents a model-driven co-design framework for high-performance computing architectures employing GPUs. Guillaume Aupy and colleagues provide a model and mathematical foundation for exploring the optimal checkpointing period for future exascale systems.

The final section of the proceedings, Sect. C, is concerned with performance optimization. The paper by Abhinav Sarje et al. presents the optimization and tuning of HipGISAXS, a parallel X-ray scattering simulation code, on general-purpose multi-core architectures as well as those with many-core accelerators. Prasanna Balaprakash and colleagues document a multi-objective optimization of HPC kernels for performance, power, and energy, with validation on three key architectures, an IBM BG/Q, hardware comprising Intel Xeon Phis, and a current generation Intel Xeon platform. Hongzhang Shan et al. present the performance tuning of fock matrix and two-electron

integral calculations for NWChem, a computational chemistry package, on the Cray XE6, the Cray XC30, and the IBM BG/Q. The final paper by Priyanka Ghosh et al. reports on a performance analysis of NWChem's Tensor Contraction Engine module, exploring alternative communications patterns through a new proxy application.

## Acknowledgments

The PMBS 2013 workshop was extremely well attended and we thank the participants for the lively discussion and positive feedback received throughout the workshop. We hope to be able to repeat this success in future years.

The SC conference series was sponsored by the IEEE Computer Society and the ACM (Association for Computing Machinery). We are extremely grateful for the support we received from the SC 2013 Steering Committee, and in particular from Barbara Chapman and Wolfgang Nagel, the SC 2013 workshop chairs.

The PMBS 2013 workshop was only possible thanks to significant input from AWE in the UK, and from Sandia National Laboratories and the Lawrence Livermore National Laboratory in the US. We acknowledge the support of the AWE Technical Outreach Programme (project CDK0724) and the Royal Society Industry Fellowship scheme (IF090020).

We are also grateful to LNCS for their support, and to Alfred Hofmann and Anna Kramer for assisting with the production of this volume.

May 2014

Stephen A. Jarvis
Steven A. Wright
Simon D. Hammond

# Organization

## Program Committee

### Workshop Chairs

| | |
|---|---|
| Stephen Jarvis | University of Warwick, UK |
| Simon Hammond | Sandia National Laboratories (NM), USA |

### Workshop Technical Program Committee

| | |
|---|---|
| Pavan Balaji | Argonne National Laboratory, USA |
| Patrick Carribault | CEA, France |
| Todd Gamblin | Lawrence Livermore National Laboratory, USA |
| Jeff Hammond | Argonne National Laboratory, USA |
| Andrew Jones | NAG Ltd., UK |
| Darren Kerbyson | Pacific Northwest National Laboratory, USA |
| David Lecomber | Allinea Software Ltd., UK |
| John Pennycook | Intel, UK |
| Karthik Raman | Intel Corporation, USA |
| Rolf Riesen | IBM Research, Dublin, Ireland |
| Arun Rodrigues | Sandia National Laboratory (NM), USA |
| Ali Saidi | ARM Research and Development, USA |
| Christian Trott | Sandia National Laboratories (NM), USA |
| Ash Vadgama | UK Atomic Weapons Establishment, UK |
| Meghan Wingate-McClelland | Xyratex, USA |
| Steven Wright | University of Warwick, UK |
| Yunquan Zhang | Chinese Academy of Sciences, China |

# Contents