Tools for High Performance Computing 2013

Andreas Knüpfer • José Gracia •
Wolfgang E. Nagel • Michael M. Resch
Editors

Tools for High Performance Computing 2013

Proceedings of the 7th International Workshop on Parallel Tools for High Performance Computing, September 2013, ZIH, Dresden, Germany

Springer

Editors
Andreas Knüpfer
Zentrum für Informationsdienste und Hochleistungsrechnen (ZIH)
Technische Universität Dresden
Dresden
Germany

José Gracia
Höchstleistungsrechenzentrum Stuttgart (HLRS)
Universität Stuttgart
Stuttgart
Germany

Wolfgang E. Nagel
Zentrum für Informationsdienste und Hochleistungsrechnen (ZIH)
Technische Universität Dresden
Dresden
Germany

Michael M. Resch
Höchstleistungsrechenzentrum Stuttgart (HLRS)
Universität Stuttgart
Stuttgart
Germany

Front cover figure: Numerical simulation of a Kelvin–Helmholtz instability. Illustration by Lars Haupt, Center for Information Services and High Performance Computing (ZIH), 01062 Dresden, Germany

---

ISBN 978-3-319-08143-4 ISBN 978-3-319-08144-1 (eBook)
DOI 10.1007/978-3-319-08144-1
Springer Cham Heidelberg New York Dordrecht London

Library of Congress Control Number: 2014950999

Mathematics Subject Classification (2010): 68M20, 65Y20, 65Y05

Printed on acid-free paper

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

Preface

Ongoing research towards specifications and designs for prospective Exascale systems highlights daunting challenges: billion-way parallelism, more levels of parallelism, a very low memory per core ratio, potential for mean-time-between-failures of less than a day, high potential for silent bit errors even in the presence of error-correcting code memories, reduced I/O bandwidths per core, and energy considerations that could require application awareness. Without upcoming disruptive computer hardware, the operation and usage of such systems will require extreme diligence. Research in areas such as parallel programming paradigms, runtime systems, compilers, operating systems, system libraries, and fault-tolerant algorithms can provide answers to some of these challenges, but will not be able to hide all of these effects from application developers.

Such a situation motivates the use of specialized tools for high performance computing (HPC) more than ever. Performance optimizers can aid application programmers in the identification of additional potential for parallelism, identify inefficient use of specific levels of parallelism, and incorporate feedback from energy monitors to identify wasteful resource usage. Debugging and correctness tools could integrate feedback from the operating system and fault tolerant system libraries to highlight the presence and impact of system faults. Simulators and automatic optimizers could identify inefficient data movement patterns, optimize process/thread mappings, or even highlight potential errors from silent memory faults.

This year's International Parallel Tools Workshop in Dresden, the seventh in a series of workshops that started in 2007, highlights how tools must carefully consider these challenges. Presentations on the state of the art in HPC tools both include activities to incorporate advances in programming paradigms, system libraries, and novel information sources, as well as new concepts for data analysis, presentation, correlation, prediction, simulation, and visualization in order to ensure that application developers will not be overwhelmed by the massive amounts of data that tools can generate even for the Petascale era. The contributions of this year's workshop include the following topics:

- Performance analytics that maximize insight of performance data
- A workflow that captures the essential behavior of existing parallel applications and automatically generates benchmark codes with the ability to extrapolate said behavior at Exascale levels
- A simulator to analyze compute node performance and energy usage that overcomes the scalability limitations of cycle accurate simulators
- A performance prediction framework that both incorporates feedback from performance counters and energy measurements
- An approach to automatically detect and prioritize performance bottlenecks with differential analysis
- An evaluation of tool support for the task concept of the OpenMP shared memory programming paradigm
- Advanced functionality in an online performance analysis framework that includes an integration with automatic performance tuning
- An approach for holistic I/O analysis that combines application and system events
- An approach addressing a specific class of memory usage errors in GPGPU accelerated applications using CUDA

These topics highlight how tools face several of the challenges that increased hardware complexities impose. At the same time, the workshop presentations include updates and demonstrations of multiple widely available tools, as well as success stories on their use. Thus, the workshop serves both as a forum for application developers that want to apply tools, as well as for tool developers as a knowledge exchange. This audience combination enables fruitful and engaged discussions of current and future challenges in order to ensure that tools meet their requirements on current and future HPC systems.

Dresden, Germany
2014

Andreas Knüpfer
José Gracia
Wolfgang E. Nagel
Michael M. Resch

Contents

1 Performance Analytics: Understanding Parallel Applications Using Cluster and Sequence Analysis 1
Juan Gonzalez, Judit Gimenez, and Jesus Labarta

2 Tools for Simulation and Benchmark Generation at Exascale 19
Mahesh Lagadapati, Frank Mueller, and Christian Engelmann

3 Suitability of Performance Tools for OpenMP Task-Parallel Programs 25
Dirk Schmidl, Christian Terboven, Dieter an Mey, and Matthias S. Müller

4 Recent Advances in Periscope for Performance Analysis and Tuning 39
Yury Oleynik, Robert Mijaković, Isaías A. Comprés Ureña, Michael Firbach, and Michael Gerndt

5 MuMMI: Multiple Metrics Modeling Infrastructure 53
Xingfu Wu, Valerie Taylor, Charles Lively, Hung-Ching Chang, Bo Li, Kirk Cameron, Dan Terpstra, and Shirley Moore

6 Cudagrind: Memory-Usage Checking for CUDA 67
Thomas M. Baumann and José Gracia

7 Node Performance and Energy Analysis with the Sniper Multi-core Simulator 79
Trevor E. Carlson, Wim Heirman, Kenzo Van Craeynest, and Lieven Eeckhout

8 A Comparison of Trace Compression Methods for Massively Parallel Applications in Context of the SIOX Project ... 91
Alvaro Aguilera, Holger Mickler, Julian Kunkel, Michaela Zimmer, Marc Wiedemann, and Ralph Müller-Pfefferkorn

9 PAMDA: Performance Assessment Using MAQAO Toolset and Differential Analysis ... 107
Zakaria Bendifallah, William Jalby, José Noudohouenou, Emmanuel Oseret, Vincent Palomares, and Andres Charif Rubial