

Workshop on Automatic Performance Analysis (WAPA 2005)

Dagstuhl Seminar 05501

Michael Gerndt
Technische Universität München

Allen D. Malony
University of Oregon

Barton Miller
University of Wisconsin, Madison

Wolfgang Nagel
Technische Universität Dresden

The Workshop on Automatic Performance Analysis (WAPA 2005, Dagstuhl Seminar 05501), held December 13-16, 2005, brought together performance researchers, developers, and practitioners with the goal of better understanding the methods, techniques, and tools that are needed for the automation of performance analysis for high performance computing.

High Performance Computing is an crucial component of current and future advances in science and engineering. These advances depend on simulations that run on massively parallel systems to obtain results within a reasonable response time. HPC are becoming commonplace, but exist in a wide variety of forms, and the forms dictate what kind of computation will running efficiently. Most current HPC systems are clustered SMP architectures with a private address space for each node. A few large-scale shared memory systems exist with a NUMA characteristic. Architectural aspects, like the memory hierarchy within a node, the latency and bandwidth of the communication network, and the connection of the node to the network are exposed to the programmer. Their individual complexity and mutual dependencies make the process of iterative program tuning, combining performance measurement, analysis, and optimization, an essential technique for high performance.

The workshop approached this problem by addressing techniques for automating the performance analysis process. Automatic performance analysis will: enable analysis of large application runs with a different behavior than scaled-down versions, allow triggering on-the-fly optimizations, and free the application expert from technical details of the analysis process. New advances in performance analysis automation are timely and achievable, and are the first step towards a fully automated tuning process.

WAPA 2005 builds on the successful Dagstuhl Seminar 02342 on Performance Analysis and Distributed Computing held in August 2002. The workshop is also a continuation of the European Working Group on Automatic Performance Analysis which ended in July, 2005. The experiences from those two activities played an important role in shaping the WAPA focus.

The WAPA program allowed a variety of opportunities for interaction among the workshop participants. Many leaders in the HPC performance community were able to attend WAPA and contribute to the technical content and discussion. As outlined in the attached agenda, the workshop was loosely organized around four themes:

- Performance analysis for large-scale parallel systems
- Automated performance analysis and diagnosis
- Automated performance tuning and performance prediction

- Performance tools and technology

Each theme session consisted of a set of technical talks. In addition, we were fortunate to have two perspective talks by John Levesque (Cray Inc.) and Phil Roth (Oak Ridge National Laboratory, on behalf of Jeff Vetter). Complementing the technical presentations, two panels provided an interesting perspective on the state of performance analysis:

- What works? What's missing? What's all this about productivity?
- Wizard Tales and Santa's Wish List

To complete the program, there was the opportunity for performance tool demonstrations. Eight tools were presented covering the techniques for HPC distributed memory analysis, scalable tracing, open performance technology, performance databases, and performance data mining.

The goals of WAPA were to increase the exchange of ideas among the tool developers, to transfer knowledge on existing and planned automatic techniques, to engage people supporting application development, and to start a dialog between researchers, developers, and users of automatic performance analysis methods and tools. By all accounts, the WAPA meeting was a tremendous success. Many participants commented on the high quality of the technical talks and the fruitful discussions they had during the week. And of course, everyone enjoyed our excursion and dinner in the historic city of Trier.

It is the energy and meaningful involvement of each participant that is most responsible for the positive WAPA experience. The organizers would like to thank everyone who came to Schloss Dagstuhl and we look forward to when our paths cross again in the future.

Best wishes,
Michael, Allen, Bart, Wolfgang



Figure 1: WAPA 2005 Group Photo

AGENDA

Tuesday, December 13, 2005

07:30 - 08:30 Breakfast

08:30 - 09:00 Opening remarks and workshop introduction
Michael, Wolfgang, Bart, Allen

THEME: Performance analysis for large-scale parallel systems

09:00 - 09:45 *Perspective:* How Tool Usage Has Progressed in HPC Over the Past 30 Years
John M. Levesque, Director, Cray's Supercomputing Center of Excellence, Cray Inc.

09:45 - 10:00 break

10:00 - 10:30 *Talk 1:* On-line Automated Performance Diagnosis on Thousands of Processes
Phil Roth, Oak Ridge National Laboratory

10:30 - 11:00 *Talk 2:* Zero Cost Reliability for Scalable Tree-based Overlay Networks
Dorian Arnold, University of Wisconsin

11:00 - 11:30 *Talk 3:* Program Trace Analysis based on Compressed Complete Call Graphs
Andreas Knuepfer, ZHR, Technische Universitaet Dresden

11:30 - 12:00 *Talk 4:* Scalable Automatic Trace Analysis
Felix Wolf, Forschungszentrum Juelich

12:00 - 13:30 Lunch

13:30 - 15:00 *Panel:* What works? What's missing? What's all this about productivity?
- Mary Zosel, Lawrence Livermore National Laboratory
- Wolfgang Nagel, ZHR, Technische Universitaet Dresden
- Dieter en Mey, High Performance Computing, RWTH Aachen University

15:00 - 15:30 break

15:30 - 17:00 *Demonstrations:*

* Distributed Automated Performance Analysis with Performance Properties
Karl Fuerlinger, Technische Universitaet Munchen

* Open—SpeedShop - Open Source Performance Analysis for MPI Clusters
Martin Schulz, Lawrence Livermore National Laboratory

* The PerfTrack Tool for Performance Data Management
Kathryn Mohror, Portland State University

* CallGrind and Kcachegrind
Josef Weidendorfer, Technische Universitaet Munchen

17:30 - 19:00 Dinner

19:00 - late Informal discussions

Wednesday, December 14, 2005

07:30 - 08:30 Breakfast

THEME: Automated performance analysis and diagnosis

09:00 - 09:45 *Perspective:* Performance Tools for Future High-End Production Platforms
Jeff Vetter, Future Technologies Group Leader, Computer Science and Mathematics
Division, Oak Ridge National Laboratory (Talk given by: Phil Roth)

09:45 - 10:00 break

10:00 - 10:30 *Talk 1:* Remote Performance Monitoring (RPM) for Resource-Constrained Devices
Chandra Krintz, University of California, Santa Barbara

10:30 - 11:00 *Talk 2:* Automatic Performance Analysis for Cache Architectures in SMPs
Edmund Kereku, Technische Universitaet Munchen

11:00 - 11:30 *Talk 3:* Model-based Automatic Performance Diagnosis: the Hercule Approach
Allen D. Malony, University of Oregon

11:30 - 12:00 *Talk 4:* Integrative Concepts for Scalable Distributed Performance Analysis and
Visualization of Parallel Programs
Holger Brunst, Technische Universitaet Dresden

12:00 - 13:30 Lunch

13:30 - 15:00 *Panel:* Wizard Tales and Santa's Wish List
- Sameer Shende, University of Oregon
- Bernd Mohr, Forschungszentrum Juelich
- Holger Brunst, Technische Universitaet Dresden

15:00 - 15:30 break

15:30 - 17:30 *Demonstrations:*
* Cray Apprentice2 and CrayPat
Luiz DeRose, Cray
* TAU Parallel Performance System
Sameer Shende, University of Oregon
* Kojak
Bernd Mohr, Forschungszentrum Juelich
* Vampir-NG
Holger Brunst, Technische Universitaet Dresden

17:30 - 19:00 Dinner

19:00 - ??? Informal discussions

Thursday, December 15, 2005

07:30 - 08:30 Breakfast

THEME: Automated performance tuning and performance prediction

09:00 - 09:30 *Talk 1: Active Harmony: Parallel Automated Tuning of Parallel Programs*
Jeff Hollingsworth, University of Maryland

09:30 - 10:00 *Talk 2: Development and Tuning Framework of M/W Applications*
Paola Caymes-Scutari, Universitat Autònoma de Barcelona

10:00 - 10:30 break

10:30 - 11:00 *Talk 3: Modeling and Predicting Resource Availability in Federated Distributed Computing Environments*
Rich Wolski, University of California, Santa Barbara

11:00 - 11:30 *Talk 4: Automatic Model Generation for Performance Prediction*
Martin Schulz, Lawrence Livermore National Laboratory

11:30 - 12:00 *Talk 5: Holistic Hardware Counter Performance Analysis of Parallel Programs*
Brian Wylie, Forschungszentrum Juelich

12:00 - 13:30 Lunch

15:00 - 22:00 Excursion to Trier

22:00 - ??? Informal discussions

Friday, December 16, 2005

07:30 - 08:30 Breakfast

THEME: Performance tools and technology

09:00 - 09:30 *Talk 1:* Performance analysis of Grid workflows in ASKALON and K-WfGrid
Hong-Linh Truong, University of Innsbruck

09:30 - 10:00 *Talk 2:* Vertical Performance Monitoring
Calin Cascaval, IBM TJ Watson Research Center

10:00 - 10:15 break

10:15 - 10:45 *Talk 3:* Comparative Performance Analysis using PerfTrack
Karen Karavanic, Portland State University

10:45 - 11:15 *Talk 4:* Observing and Controlling: The Case for Organic Computing Systems
Wolfgang Karl, Universitaet Karlsruhe

11:15 - 11:45 *Talk 5:* Data Profiling for Cache Optimizations on SMP Clusters
Thomas Brandes, German National Research Center for Information Technology
(GMD)

11:45 - 12:00 *Closing remarks*
Michael, Wolfgang, Bart, Allen

12:00 - 13:30 Lunch