

Executive Summary
of the
Dagstuhl Seminar
on

**Code Instrumentation and Modeling
for Parallel Performance Analysis
CIMPPA 2007**

20. to 24. August 2007

<http://www.dagstuhl.de/07341/>

[Adolfy Hoisie](#) (Los Alamos National Lab., USA)

[Barton P. Miller](#) (University of Wisconsin, USA)

[Bernd Mohr](#) (Forschungszentrum Jülich, DE)

Given the exponential increase in the complexity of modern parallel systems, parallel applications often fail to exploit the full power of the underlying hardware. At scale, it is not uncommon for applications to run at parallel efficiencies in the low single digits. Moreover, their optimization is extremely difficult due to the inherent complexity of the systems and the applications themselves. Therefore, a variety of projects have been developing tools and techniques for the measurement, analysis, and visualization of parallel program performance in order to help guide users in the optimization process.

This meeting was the third in a series of seminars related to the topic “Performance Analysis of Parallel and Distributed Programs”, with previous meetings being the Dagstuhl Seminar 02341 on “Performance Analysis and Distributed Computing” held in August 2002 and Seminar 05501 on “Automatic Performance Analysis” in December 2005. While these seminars concentrated on the “analysis” part of performance analysis, at the most recent seminar the focus was on the building blocks of program instrumentation and modeling that are prerequisites for the analysis phase. As a result, the presentations of the participants concentrated on several fundamental issues related to instrumentation for generating high-quality performance data, methodologies for performance modeling leading to accurate predictions for the performance, and on the ways in which these techniques are combined for the performance analysis of applications and systems.

The program consisted of 28 presentations and practical tool demonstrations as well as two “open mic” sessions where time was set aside for spontaneous discussions and “brain storming”. The seminar brought together a total of 48 researchers and developers working in the area of performance from universities, national research laboratories and, especially important, from three major computer vendors. The goals were to increase the exchange of ideas, knowledge transfer, foster a multidisciplinary approach to attacking this very important research problem with direct impact on the way in which we design and utilize parallel systems to achieve high application performance. The presentations can be grouped thematically as follows:

Session “Performance Analysis in General”

- *Application Performance Analysis on High Performance Computers*
Sudip Dosanjh, SNL (USA)
- *System Architectures and Application Code Performance Bottlenecks*
Jim Tomkins, SNL (USA)
- *Using Communication Patterns for Parallel Program Optimization*
Dieter Kranzlmüller, Univ. Linz (A)
- *Performance Analysis and Tuning of Parallel/Distributed Applications*
Anna Morajko, UAB (E)
- *The IBM HPC Toolkit*
David Klepacki, IBM (USA)

Session “Instrumentation”

- *The Impact of Scalable Tool Infrastructure on Systems with 3D Mesh and Torus Interconnection Networks*
Phil Roth, ORNL (USA)
- *Peering Inside Black Boxes: Strategies for Uncovering Scaling Issues In Lower Software Stack Levels*
Terry Jones, LLNL (USA)
- *Deep Binary Analysis & The Need for Cooperating Analyses*
Jeff Hollingsworth, Univ. Maryland (USA)
- *The Deconstruction of Dyninst*
Andrew Bernat, Univ. Wisconsin (USA)
DEMO: SymtabAPI + Stackwalker

Session “Modeling”

- *Constructing Application Performance Models using Neural Networks*
Sally McKee, Cornell Univ. (USA)
Martin Schulz, LLNL (USA)
DEMO
- *Characterization of Parallel Applications*
Jesus Labarta, UPC (E)
DEMO: Paraver and Dimemas analysis of WRF
- *Application Performance Modeling: Predictive Accuracy in the Presence of Simplifying Assumptions*
Kevin Barker, LANL (USA)
- *Performance Analysis and Modeling of High Performance Networks*
Darren Kerbyson, LANL (USA)

Session “Scalability”

- *Chasing Scalability Bottlenecks: Recent Activities*
Pat Worley, ORNL (USA)
- *(Are you sure that) Tracing is not scalable?*
Judith Gimenez, UPC (E)
- *Scalable Compression and Replay of Communication Traces*
Frank Mueller, Univ. NC (USA)
- *Large-scale Application Measurement and Analysis with Scalasca*
Brian Wylie, FZJ (DE)
DEMO: Scalasca

Session “Tools”

- *Knowledge Support for Parallel Performance Data Mining*
Kevin Huck, Univ. Oregon (USA)
DEMO: Perfexplorer 2
- *Self-consistent MPI Performance Requirements*
Jesper Larsson Träff, NEC (DE)
- *The Cray Performance Tools*
Luiz De Rose, Cray Inc. (USA)
DEMO: Cray Apprentice 2
- *Sampling-based performance measurement and analysis*
John Mellor-Crummey, Rice Univ. (USA)
DEMO: hpcviewer
- *PⁿMPI: A Dynamic MPI Tool Infrastructure*
Martin Schulz, LLNL (USA)
DEMO: PⁿMPI
- *Exploiting the MPI Profiling Interface*
Rusty Lusk, ANL (USA)
DEMO: Jumpshot
- *Timestamp Synchronization for Event Traces of Large-Scale Message-Passing Applications*
Daniel Becker, FZJ (DE)
- *Applying Performance Tools to Real World Applications*
Matthias Müller, Michael Kluge, TU Dresden (DE)
DEMO: Vampir
- *Random Access to Event Traces with OTF*
Heike Jagode, TU Dresden (DE)
- *Inclusion of Co-Processor Usage in Program Traces*
Guido Juckeland, TU Dresden (DE)
- *Compiler Support for Scalable Program Instrumentation*
Barbara Chapman, Univ. Houston (USA)

“Open Mic” Sessions

- *Instrumentation of Heterogeneous Systems and New Languages*
Moderator: Jeffrey K. Hollingsworth, Univ. of Maryland (USA)
- *Performance Modeling and Analysis on Multi-Paradigm Computing Systems - Challenges and Opportunities*
Moderator: Jeffrey Vetter, Oak Ridge National Laboratory (USA)

Session “Short Announcements”

- *Implementation and Usage of the PERUSE-Interface in Open MPI*
Rainer Keller, Universität Stuttgart (D)
- *Summary of CScADS Workshop on Performance Tools for Petascale Systems*
John Mellor-Crummey, Rice Univ. (USA)
- *The Virtual Institute High-Productivity Supercomputing*
Felix Wolf , Forschungszentrum Jülich (D)
- *Summary of SDTPC Workshop on Software Development Tools for PetaScale Computing*
Barton Miller, University of Wisconsin (USA)
- *Preview of the SC 2007 Technical Program*
Harvey Wasserman, Lawrence Berkeley National Laboratory (USA)

Despite the larger than normal number participants, the seminar was very successful due to the dedicated professionalism and discipline of the participants on one side and the very helpful and professional staff of Dagstuhl on the other side. Lively discussions and spontaneous computer demonstrations continued every day well beyond midnight. It is important to note that the group meeting and residential aspects of Dagstuhl and the five-day format provide a continuity of thought and discussion unavailable in other conference, workshop, or meeting settings. At Dagstuhl, we have time for considered (and reconsidered!) dialogs whose impact last well beyond the meeting week.

A half-day excursion including a guided tour of “Burg Eltz” and a wine-tasting and dinner at the lovely down of Beilstein on the Mosel river completed the program.