

Commenced Publication in 1973

Founding and Former Series Editors:

Gerhard Goos, Juris Hartmanis, and Jan van Leeuwen

Editorial Board Members

David Hutchison

Lancaster University, Lancaster, UK

Takeo Kanade

Carnegie Mellon University, Pittsburgh, PA, USA

Josef Kittler

University of Surrey, Guildford, UK

Jon M. Kleinberg

Cornell University, Ithaca, NY, USA

Friedemann Mattern

ETH Zurich, Zurich, Switzerland

John C. Mitchell

Stanford University, Stanford, CA, USA

Moni Naor

Weizmann Institute of Science, Rehovot, Israel

C. Pandu Rangan

Indian Institute of Technology Madras, Chennai, India

Bernhard Steffen

TU Dortmund University, Dortmund, Germany

Demetri Terzopoulos

University of California, Los Angeles, CA, USA

Doug Tygar

University of California, Berkeley, CA, USA

More information about this series at <http://www.springer.com/series/7408>

Abhinav Bhatele · David Boehme ·
Joshua A. Levine · Allen D. Malony ·
Martin Schulz (Eds.)

Programming and Performance Visualization Tools

International Workshops, ESPT 2017 and VPA 2017
Denver, CO, USA, November 12 and 17, 2017
and ESPT 2018 and VPA 2018
Dallas, TX, USA, November 16 and 11, 2018
Revised Selected Papers

Editors

Abhinav Bhatele 
Lawrence Livermore National Laboratory
Livermore, CA, USA

Joshua A. Levine
The University of Arizona
Tucson, AZ, USA

Martin Schulz
Technical University of Munich
Munich, Germany

David Boehme
Lawrence Livermore National Laboratory
Livermore, CA, USA

Allen D. Malony
University of Oregon
Eugene, OR, USA

ISSN 0302-9743 ISSN 1611-3349 (electronic)
Lecture Notes in Computer Science
ISBN 978-3-030-17871-0 ISBN 978-3-030-17872-7 (eBook)
<https://doi.org/10.1007/978-3-030-17872-7>

LNCS Sublibrary: SL2 – Programming and Software Engineering

© Springer Nature Switzerland AG 2019

This work is subject to copyright. All rights are reserved by the Publisher, whether the whole or part of the material is concerned, specifically the rights of translation, reprinting, reuse of illustrations, recitation, broadcasting, reproduction on microfilms or in any other physical way, and transmission or information storage and retrieval, electronic adaptation, computer software, or by similar or dissimilar methodology now known or hereafter developed.

The use of general descriptive names, registered names, trademarks, service marks, etc. in this publication does not imply, even in the absence of a specific statement, that such names are exempt from the relevant protective laws and regulations and therefore free for general use.

The publisher, the authors and the editors are safe to assume that the advice and information in this book are believed to be true and accurate at the date of publication. Neither the publisher nor the authors or the editors give a warranty, expressed or implied, with respect to the material contained herein or for any errors or omissions that may have been made. The publisher remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.

This Springer imprint is published by the registered company Springer Nature Switzerland AG
The registered company address is: Gewerbestrasse 11, 6330 Cham, Switzerland

Preface

This volume contains the proceedings of two instances each of two workshops, held in conjunction with the International Conference on High Performance Computing, Networking, Storage and Analysis (SC). The workshops are – Workshop on Extreme-Scale Programming Tools (ESPT) and International Workshop on Visual Performance Analysis (VPA).

ESPT 2017

The 6th Workshop on Extreme-Scale Programming Tools (ESPT) was held in conjunction with the International Conference on High Performance Computing, Networking, Storage and Analysis (SC) in Denver, Colorado, USA, on Sunday, November 12, 2017. The workshop focused on how the path to exascale computing challenges HPC application developers in their quest to achieve the maximum potential that the machines have to offer. Factors such as limited power budgets, clock frequency variability, heterogeneous load imbalance, hierarchical memories, and shrinking I/O bandwidths make it increasingly difficult to create high performance applications. Tools for debugging, performance measurement and analysis, and tuning are needed to overcome the architectural, system, and programming complexities envisioned in exascale environments. At the same time, research and development progress for HPC tools faces equally difficult challenges from exascale factors. Increased emphasis on autotuning, dynamic monitoring and adaptation, heterogeneous analysis, and so on require new methodologies, techniques, and engagement with application teams.

The ESPT 2017 workshop served as a forum for HPC application developers, system designers, and tools researchers to discuss the requirements for exascale-enabled tools and the roadblocks that need to be addressed. It was the sixth instantiation of successful SC conference workshops organized by the Virtual Institute – High Productivity Supercomputing (VI-HPS), an international initiative of HPC researchers and developers focused on parallel programming and performance tools for large-scale systems. The workshop topics of interest included:

- Programming tools (e.g., performance analysis, tuning, debuggers, IDEs)
- Methodologies for performance engineering
- Tool technologies for extreme-scale challenges (e.g., scalability, resilience, power)
- Tool infrastructures and environments
- Evolving/future application requirements for programming tools and technologies
- Application developer experiences with programming and performance tools

More information can be found at: <http://www.vi-hps.org/symposia/espt/espt-sc17.html>.

ESPT 2017 was a full-day workshop consisting of a keynote address in the morning followed by research paper presentations and an open debate in the afternoon. All submitted research papers underwent a rigorous review process. A total of 13 papers were submitted, with three to four reviews provided for each paper. The ESPT organizers made final decisions on paper selection. A total of eight papers were accepted and each paper was allotted 30 minutes for presentation. This volume contains seven of the eight accepted papers from the ESPT 2017 proceedings.

Organizing Committee

William Jalby (Chair)	Université de Versailles St-Quentin-en-Yvelines, France
Allen D. Malony (Vice-chair)	University of Oregon, Eugene, USA
Martin Schulz	Technische Universität München, Germany
Judit Gimenez	Barcelona Supercomputing Center, Spain

Program Committee

Jean-Baptiste Besnard	ParaTools, SAS, France
Michael Gerndt	Universität München, Germany
Kevin Huck	University of Oregon, Eugene, USA
Andreas Knüpfer	Technische Universität Dresden, Germany
Heike McCraw	University of Tennessee, Knoxville, USA
Barton P. Miller	University of Wisconsin, Madison, USA
Pablo Oliveira	Université de Versailles St-Quentin-en-Yvelines, France
Sameer Shende	University of Oregon, Eugene, USA
Jan Treibig	Friedrich-Alexander-Universität Erlangen-Nürnberg, Germany
Felix Wolf	Technische Universität Darmstadt, Germany
Brian Wylie	Jülich Supercomputing Centre, Germany

Special thanks: Marc-Andre Hermanns, Jülich Supercomputing Centre, Germany

ESPT 2018

The 2018 Workshop on Extreme-Scale Programming Tools (ESPT) was held in conjunction with the International Conference on High Performance Computing, Networking, Storage and Analysis (SC) in Dallas, Texas, USA, on Friday, November 16, 2018. It was the seventh instantiation of successful SC conference workshops organized by the Virtual Institute – High Productivity Supercomputing (VI-HPS), an international initiative of HPC researchers and developers focused on parallel programming and performance tools for large-scale systems.

The path to extreme computing keeps broadening: Large-scale systems toward exascale and beyond, growing many-core systems with deep memory hierarchies and massively parallel accelerators are just a few of the platforms we can expect. This trend will challenge HPC application developers in their quest to achieve the maximum potential that their systems have to offer, both on and across nodes. Factors such as limited power budgets, heterogeneity, hierarchical memories, shrinking I/O bandwidths, and performance variability will make it increasingly difficult to create productive applications on future platforms. To address these challenges, we need tools for debugging, performance measurement and analysis, and tuning to overcome the architectural, system, and programming complexities expected in these environments. At the same time, research and development progress for HPC tools faces equally difficult challenges from exascale factors. Increased emphasis on autotuning, dynamic monitoring and adaptation, heterogeneous analysis, and so on require new methodologies, techniques, and engagement with application teams.

Like its predecessors, the ESPT 2018 workshop served as a forum for HPC application developers, system designers, and tools researchers to discuss the requirements for exascale-enabled tools and the roadblocks that need to be addressed. The workshop topics of interest included:

- Programming tools (e.g., performance analysis, tuning, debuggers, IDEs)
- Methodologies for performance engineering
- Tool technologies for extreme-scale challenges (e.g., scalability, resilience, power)
- Tool support for accelerated architectures and large-scale multi-cores
- Tool infrastructures and environments
- Evolving/future application requirements for programming tools and technologies
- Application developer experiences with programming and performance tools

More information can be found at: <http://www.vi-hps.org/symposia/espt/espt-sc18.html>.

ESPT 2018 was a half-day workshop consisting of two keynote addresses and a session of research paper presentations. All submitted research papers underwent a rigorous review process. A total of five papers were submitted, with three to five reviews provided for each paper. The ESPT organizers made final decisions on paper selection. A total of four papers were accepted and each paper was allotted 30 minutes for presentation. This volume contains the ESPT 2018 proceedings.

Organizing Committee

Martin Schulz

David Boehme

Marc-André Hermanns

Felix Wolf

Technische Universität München, Germany

Lawrence Livermore National Laboratory, USA

Jülich Supercomputing Centre, Germany

Technische Universität Darmstadt, Germany

Program Committee

Dorian C. Arnold	Emory University, USA
Jean-Baptiste Besnard	ParaTools, France
Karl F�rlinger	Ludwig Maximilian University of Munich, Germany
Michael Gerndt	Technical University of Munich, Germany
Judit Gimenez	Barcelona Supercomputing Center, Spain
Marc-Andr� Hermanns	Forschungszentrum J�lich, Germany
Kevin Huck	University of Oregon, USA
William Jalby	Universit� de Versailles St-Quentin-en-Yvelines, France
Andreas Kn�pfer	Technical University of Dresden, Germany
John Linfood	ARM, USA
Allen D. Malony	University of Oregon, USA
John Mellor-Crummey	Rice University, USA
Bart Miller	University of Wisconsin Madison, USA
Heidi Poxon	Cray Inc., USA
Nathan Tallent	Pacific Northwestern National Laboratory, USA
Christian Terboven	RWTH Aachen University, Germany
Josef Weidendorfer	Leibniz Supercomputing Centre, Germany
Gerhard Wellein	University of Erlangen-N�rnberg, Germany
Felix Wolf	Technical University of Darmstadt, Germany
Brian J. N. Wylie	Forschungszentrum J�lich, Germany

VPA 2017

The Fourth International Workshop on Visual Performance Analysis (VPA 17) was held in conjunction with the International Conference on High Performance Computing, Networking, Storage and Analysis (SC 17) in Denver, Colorado, USA, on Friday, November 17, 2017 and in cooperation with TCHPC: The IEEE Computer Society Technical Consortium on High Performance Computing.

Over the past decades an incredible amount of resources has been devoted to building ever more powerful supercomputers. However, exploiting the full capabilities of these machines is becoming exponentially more difficult with each new generation of hardware. To help understand and optimize the behavior of massively parallel simulations, the performance analysis community has created a wide range of tools and APIs to collect performance data, such as flop counts, network traffic, or cache behavior at the largest scale. However, this success has created a new challenge, as the resulting data are far too large and too complex to be analyzed in a straightforward manner. Therefore, new automatic analysis and visualization approaches must be developed to allow application developers to intuitively understand the multiple, interdependent effects that their algorithmic choices have on the final performance.

This workshop brought together researchers from the fields of performance analysis and visualization to discuss new approaches of applying visualization and visual analytics techniques to large-scale applications. The workshop topics of interest included:

- Scalable displays of performance data
- Data models to enable scalable visualization
- Graph representation of unstructured performance data
- Presentation of high-dimensional data
- Visual correlations between multiple data source
- Human–computer interfaces for exploring performance data
- Multiscale representations of performance data for visual exploration

More information can be found here: <https://vpa17.github.io>.

VPA 17 was a half-day workshop consisting of a keynote address by Dr. Lucy Nowell of the U.S. Department of Energy titled “Visual Performance Analysis for Extremely Heterogeneous Systems” as well as a panel discussion on “Challenges and the Future of HPC Performance Visualization.” Research paper presentations were mixed into these two sessions. All submitted research papers underwent a rigorous review process. A total of six papers were submitted, with five reviews provided for each paper. The VPA workshop chairs made final decisions on paper selection. A total of three papers were accepted and each paper was allotted 25 minutes for presentation. This volume contains two of three accepted papers from the VPA 2017 proceedings.

Workshop Chairs

Fabian Beck	University of Duisburg-Essen, Germany
Abhinav Bhatele	Lawrence Livermore National Laboratory, USA
Judit Gimenez	Barcelona Supercomputing Center, Spain
Joshua A. Levine	University of Arizona, USA

Steering Committee

Peer-Timo Bremer	Lawrence Livermore National Laboratory, USA
Bernd Mohr	Jülich Supercomputing Center, Germany
Valerio Pascucci	University of Utah, USA
Martin Schulz	Lawrence Livermore National Laboratory, USA

Program Committee

Harsh Bhatia	Lawrence Livermore National Laboratory, USA
Holger Brunst	TU Dresden, Germany
Alexandru Calotoiu	Technical University Darmstadt, Germany
Todd Gamblin	Lawrence Livermore National Laboratory, USA
Marc-Andre Hermanns	Jülich Supercomputing Center, Germany
Kevin Huck	University of Oregon, USA
Katherine Isaacs	University of Arizona, USA
Yarden Livnat	University of Utah, USA
Naoya Maruyama	Lawrence Livermore National Laboratory, USA

Bernd Mohr	Jülich Supercomputing Center, Germany
Ananya Muddukrishna	KTH Royal Institute of Technology, Sweden
Matthias Mueller	RWTH Aachen University, Germany
Valerio Pascucci	University of Utah, USA
Paul Rosen	University of South Florida, USA
Carlos Scheidegger	University of Arizona, USA
Chad Steed	Oak Ridge National Laboratory, USA

VPA 2018

The Fifth International Workshop on Visual Performance Analysis (VPA 2018) was held in conjunction with the International Conference on High Performance Computing, Networking, Storage and Analysis (SC 2018) in Dallas, Texas, USA, on Sunday, November 11, 2018.

Over the past decades an incredible amount of resources has been devoted to building ever more powerful supercomputers. However, exploiting the full capabilities of these machines is becoming exponentially more difficult with each new generation of hardware. To help understand and optimize the behavior of massively parallel simulations, the performance analysis community has created a wide range of tools and APIs to collect performance data, such as flop counts, network traffic, or cache behavior at the largest scale. However, this success has created a new challenge, as the resulting data are far too large and too complex to be analyzed in a straightforward manner. Therefore, new automatic analysis and visualization approaches must be developed to allow application developers to intuitively understand the multiple, interdependent effects that their algorithmic choices have on the final performance.

This workshop brought together researchers from the fields of performance analysis and visualization to discuss new approaches of applying visualization and visual analytics techniques to large-scale applications. The workshop topics of interest included:

- Scalable displays of performance data
- Case studies demonstrating the use of performance visualization in practice
- Data models to enable scalable visualization
- Graph representation of unstructured performance data
- Presentation of high-dimensional data
- Visual correlations between multiple data source
- Human–computer interfaces for exploring performance data
- Multi-scale representations of performance data for visual exploration

More information can be found here: <https://vpa18.github.io>.

VPA 18 was a half-day workshop consisting of a keynote address by Dr. Allen Malony of the University of Oregon, USA titled “Not Your Mama’s Angry Fruit Salad: Ruminations on 30 Years of Performance Visualization and Visual Performance Analysis” and research paper presentations. All submitted research papers underwent a rigorous review process. A total of five papers were submitted, with three to five reviews provided for each paper. The VPA workshop chairs made final decisions on

paper selection. A total of four papers were accepted and each paper was allotted 25 minutes for presentation. This volume contains the VPA 2018 proceedings.

Workshop Chairs

Abhinav Bhatele	Lawrence Livermore National Laboratory, USA
Katherine Isaacs	University of Arizona, USA
Kevin Huck	University of Oregon, USA

Steering Committee

Peer-Timo Bremer	Lawrence Livermore National Laboratory, USA
Bernd Mohr	Jülich Supercomputing Center, Germany
Valerio Pascucci	University of Utah, USA
Martin Schulz	Lawrence Livermore National Laboratory, USA

Program Committee

Harsh Bhatia	Lawrence Livermore National Laboratory, USA
Holger Brunst	TU Dresden, Germany
Alexandru Calotoiu	Technical University Darmstadt, Germany
Todd Gamblin	Lawrence Livermore National Laboratory, USA
Judit Gimenez	Barcelona Supercomputing Center, Spain
Marc-Andre Hermanns	Jülich Supercomputing Center, Germany
Aaditya Landge	Twitter, Inc., USA
Joshua A. Levine	University of Arizona, USA
Yarden Livnat	University of Utah, USA
Naoya Maruyama	Lawrence Livermore National Laboratory, USA
Matthias Mueller	RWTH Aachen University, Germany
Paul Rosen	University of South Florida, USA

Contents

ESPT 2017

Enhancing PAPI with Low-Overhead rdpmc Reads	3
<i>Yan Liu and Vincent M. Weaver</i>	
Generic Library Interception for Improved Performance Measurement and Insight.	21
<i>Ronny Brendel, Bert Wesarg, Ronny Tschüter, Matthias Weber, Thomas Ilsche, and Sebastian Oeste</i>	
Improved Accuracy for Automated Communication Pattern Characterization Using Communication Graphs and Aggressive Search Space Pruning	38
<i>Philip C. Roth</i>	
Moya—A JIT Compiler for HPC	56
<i>Tarun Prabhu and William Gropp</i>	
Polyhedral Optimization of TensorFlow Computation Graphs	74
<i>Benoît Pradelle, Benoît Meister, Muthu Baskaran, Jonathan Springer, and Richard Lethin</i>	
CAASCADE: A System for Static Analysis of HPC Software Application Portfolios	90
<i>M. Graham Lopez, Oscar Hernandez, Reuben D. Budiardja, and Jack C. Wells</i>	
Visual Comparison of Trace Files in Vampir	105
<i>Matthias Weber, Ronny Brendel, Michael Wagner, Robert Dietrich, Ronny Tschüter, and Holger Brunst</i>	

ESPT 2018

Understanding the Scalability of Molecular Simulation Using Empirical Performance Modeling.	125
<i>Sergei Shudler, Jadran Vrabec, and Felix Wolf</i>	
Advanced Event-Sampling Support for PAPI	144
<i>Forrest Smith and Vincent M. Weaver</i>	

PARLOT: Efficient Whole-Program Call Tracing for HPC Applications.	162
<i>Saeed Taheri, Sindhu Devale, Ganesh Gopalakrishnan, and Martin Burtscher</i>	
Gotcha: An Function-Wrapping Interface for HPC Tools	185
<i>David Poliakoff and Matt LeGendre</i>	
VPA 2017	
Projecting Performance Data over Simulation Geometry Using SOSflow and ALPINE	201
<i>Chad Wood, Matthew Larsen, Alfredo Gimenez, Kevin Huck, Cyrus Harrison, Todd Gamblin, and Allen Malony</i>	
Visualizing, Measuring, and Tuning Adaptive MPI Parameters	219
<i>Matthias Diener, Sam White, and Laxmikant V. Kale</i>	
VPA 2018	
Visual Analytics Challenges in Analyzing Calling Context Trees	233
<i>Alexandre Bergel, Abhinav Bhatele, David Boehme, Patrick Gralka, Kevin Griffin, Marc-André Hermanns, Dušan Okanović, Olga Pearce, and Tom Vierjahn</i>	
PaScal Viewer: A Tool for the Visualization of Parallel Scalability Trends. . .	250
<i>Anderson B. N. da Silva, Daniel A. M. Cunha, Vitor R. G. Silva, Alex F. de A. Furtunato, and Samuel Xavier-de-Souza</i>	
Using Deep Learning for Automated Communication Pattern Characterization: Little Steps and Big Challenges	265
<i>Philip C. Roth, Kevin Huck, Ganesh Gopalakrishnan, and Felix Wolf</i>	
Visualizing Multidimensional Health Status of Data Centers.	273
<i>Tommy Dang</i>	
Author Index	285