

Commenced Publication in 1973

Founding and Former Series Editors:

Gerhard Goos, Juris Hartmanis, and Jan van Leeuwen

Editorial Board

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, Zürich, 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, 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

Gerhard Weikum

Max Planck Institute for Informatics, Saarbrücken, Germany

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

Sascha Hunold · Alexandru Costan
Domingo Giménez · Alexandru Iosup
Laura Ricci · María Engracia Gómez Requena
Vittorio Scarano · Ana Lucia Varbanescu
Stephen L. Scott · Stefan Lankes
Josef Weidendorfer · Michael Alexander (Eds.)

Euro-Par 2015: Parallel Processing Workshops

Euro-Par 2015 International Workshops
Vienna, Austria, August 24–25, 2015
Revised Selected Papers

Editor
Sascha Hunold
Vienna University of Technology
Vienna
Austria

Workshop Editors *see next page*

ISSN 0302-9743 ISSN 1611-3349 (electronic)
Lecture Notes in Computer Science
ISBN 978-3-319-27307-5 ISBN 978-3-319-27308-2 (eBook)
DOI 10.1007/978-3-319-27308-2

Library of Congress Control Number: 2015955875

LNCS Sublibrary: SL1 – Theoretical Computer Science and General Issues

Springer Cham Heidelberg New York Dordrecht London
© Springer International Publishing Switzerland 2015

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, express or implied, with respect to the material contained herein or for any errors or omissions that may have been made.

Printed on acid-free paper

Springer International Publishing AG Switzerland is part of Springer Science+Business Media
(www.springer.com)

Workshop Editors

BigDataCloud

Alexandru Costan
IRISA/INSA Rennes
France
alexandru.costan@inria.fr

PELGA

Ana Lucia Varbanescu
University of Amsterdam
The Netherlands
a.l.varbanescu@uva.nl

Euro-EDUPAR

Domingo Giménez
University of Murcia
Spain
domingo@um.es

REPPAR

Sascha Hunold
Vienna University of Technology
Austria
sascha.hunold@tuwien.ac.at

HeteroPar

Alexandru Iosup
Delft University of Technology
The Netherlands
A.Iosup@tudelft.nl

Resilience

Stephen L. Scott
Tennessee Tech University and Oak
Ridge National Laboratory, USA
scottssl@ornl.gov

LSDVE

Laura Ricci
University of Pisa
Italy
laura.ricci@unipi.it

ROME

Stefan Lankes
RWTH Aachen University
Germany
slankes@eonerc.rwth-aachen.de

OMHI

María Engracia Gómez Requena
Universitat Politècnica de València
Spain
megomez@disca.upv.es

UCHPC

Josef Weidendorfer
Technische Universität München
Germany
Josef.Weidendorfer@tum.de

PADABS

Vittorio Scarano
Università di Salerno
Italy
vitsca@dia.unisa.it

VHPC

Michael Alexander
Vienna University of Technology
Austria
michael.alexander@tuwien.ac.at

Preface

Euro-Par is an annual, international conference on European ground, covering all aspects of parallel and distributed processing, ranging from theory to practice, from small to the largest parallel and distributed systems and infrastructures, from fundamental computational problems to full-fledged applications, from architecture, compiler, language and interface design and implementation to tools, support infrastructures, and application performance aspects. The Euro-Par conference itself is complemented by a workshop program, where workshops dedicated to more specialized themes, to cross-cutting issues, and to upcoming trends and paradigms can be easily and conveniently organized with little administrative overhead.

This year, 17 workshop proposals were submitted, and after a careful revision process, which was led by the workshop co-chairs, 13 workshops were accepted. One workshop had to be canceled later owing to a low number of submissions.

The workshops took place on the two days before the Euro-Par conference and the program included the following 12 workshops:

1. Big Data Management in Clouds (BIGDATA CLOUD)
2. Parallel and Distributed Computing Education for Undergraduate Students (EURO-EDUPAR)
3. Algorithms, Models, and Tools for Parallel Computing on Heterogeneous Platforms (HETERO PAR)
4. Large-Scale Distributed Virtual Environments (LSDVE)
5. On-Chip Memory Hierarchies and Interconnects: Organization, Management and Implementation (OMHI)
6. Parallel and Distributed Agent-Based Simulations (PADABS)
7. Performance Engineering for Large-Scale Graph Analytics (PELGA)
8. Reproducibility in Parallel Computing (REPPAR)
9. Resiliency in High-Performance Computing with Clouds, Grids, and Clusters (RESILIENCE)
10. Runtime and Operating Systems for the Many-Core Era (ROME)
11. UnConventional High Performance Computing (UCHPC)
12. Virtualization in High-Performance Cloud Computing (VHPC)

All workshops together received a total of 121 submissions from 34 different countries. Each workshop had an independent Program Committee, which was in charge of selecting the papers. The workshop papers received more than three reviews

per paper on average (403 reviews in total). Out of the 121 submissions, 67 papers were selected to be presented at the workshops.

The success of the Euro-Par workshops depends on the work of many individuals and organizations. We therefore thank all workshop organizers and reviewers for the time and effort that they invested. The Euro-Par vice-chair Luc Bougé provided guidance and support throughout the whole organizational process of the workshops. We would also like to express our sincere thanks to Springer for their help in publishing the proceedings.

Lastly, we thank all participants, panelists, and keynote speakers of the Euro-Par workshops for contributing to a productive meeting. It was a pleasure to organize and host the Euro-Par workshops 2015 in Vienna.



October 2015

Sascha Hunold

Organization

Euro-Par Steering Committee

Chair

Christian Lengauer University of Passau, Germany

Vice-Chair

Luc Bougé ENS Rennes, France

European Representatives

Marco Danelutto University of Pisa, Italy
Emmanuel Jeannot LaBRI-Inria, Bordeaux, France
Christos Kaklamanis Computer Technology Institute, Greece
Paul Kelly Imperial College, UK
Thomas Ludwig University of Hamburg, Germany
Emilio Luque Autonomous University of Barcelona, Spain
Tomàs Margalef Autonomous University of Barcelona, Spain
Wolfgang Nagel Dresden University of Technology, Germany
Rizos Sakellariou University of Manchester, UK
Fernando Silva University of Porto, Portugal
Henk Sips Delft University of Technology, The Netherlands
Domenico Talia University of Calabria, Italy
Felix Wolf Technische Universität Darmstadt, Germany

Honorary Members

Ron Perrott Oxford e-Research Centre, UK
Karl Dieter Reinartz University of Erlangen-Nuremberg, Germany

Observers

Jesper Larsson Träff Vienna University of Technology, Austria
Denis Trystram Grenoble Institute of Technology, France

Euro-Par 2015 Organization

Chair

Jesper Larsson Träff Vienna University of Technology, Austria

Proceedings

Francesco Versaci Vienna University of Technology, Austria

Workshops

Sascha Hunold

Vienna University of Technology, Austria

Local Organization

Alexandra Carpen-Amarie

Vienna University of Technology, Austria

Christine Kamper

Vienna University of Technology, Austria

Margret Steinbuch

Vienna University of Technology, Austria

Workshop Introduction and Organization

4th Workshop on Big Data Management in Clouds (BigDataCloud)

Workshop Description

The Workshop on Big Data Management in Clouds was created to provide a platform for the dissemination of recent research efforts that explicitly aim at addressing the challenges related to executing Big Data applications on the cloud. Initially designed for powerful and expensive supercomputers, such applications have seen an increasing adoption on clouds, exploiting their elasticity and economical model. While Map/Reduce covers a large fraction of the development space, there are still many applications that are better served by other models and systems. In such a context, we need to embrace new programming models, scheduling schemes, and hybrid infrastructures and scale out of single datacenters to geographically distributed deployments in order to cope with these new challenges effectively.

Against this backdrop, the BigDataCloud workshop aims to provide a venue for researchers to present and discuss results on all aspects of data management in clouds, new developments, and deployment efforts in running data-intensive computing workloads. In particular, we are interested in how the use of cloud-based technologies can meet the data-intensive scientific challenges of HPC applications that are not well served by the current supercomputers or grids, and are being ported to cloud platforms. The goal of the workshop is to support the assessment of the current state, introduce future directions, and present architectures and services for future clouds supporting data-intensive computing.

BigDataCloud 2015 followed the successful previous editions held in conjunction with EuroPar. Its goal is to aggregate the data management and clouds/grids/P2P communities in order to complement the Big Data handling issues with a comprehensive system/infrastructure perspective. This year's edition was held on August 24 and gathered around 30 enthusiastic researchers from academia and industry. We received six papers, out of which three were selected for presentation. The Big Data theme was strongly reflected in the keynote given this year by Prof. Luc Bougé from École Normale Supérieure Rennes. The talk focused on the challenges of computing in distributed, very-large clouds from the execution and programming models perspective.

We wish to thank all the authors, the keynote speaker, the Program Committee members and the workshop chairs of EuroPar 2015 for their contribution to the success of this edition of BigDataCloud.

Program Chairs

Alexandru Costan
Frédéric Desprez

IRISA/INSA Rennes, France
Inria ENS Lyon, France

Program Committee

Gabriel Antoniu

Inria, France

Luc Bougé

ENS Rennes, France

Shadi Ibrahim

Inria, France

Olivier Nano

Microsoft Research ATLE, Germany

Bogdan Nicolae

IBM Research, Ireland

Maria S. Pérez

Universidad Politecnica de Madrid, Spain

Florin Pop

University Politehnica of Bucharest, Romania

Anna Queralt

Barcelona Supercomputing Center, Spain

Leonardo Querzoni

University of Rome La Sapienza, Italy

Balaji Subramaniam

Virginia Tech, USA

Domenico Talia

University of Calabria, Italy

Osamu Tatebe

University of Tsukuba, Japan

Radu Tudoran

Huawei European Research Center, Germany

First European Workshop on Parallel and Distributed Computing Education for Undergraduate Students (Euro-EDUPAR)

Workshop Description

Today, parallel and distributed computing (PDC) is omnipresent. It is encountered in all computational environments, from mobile devices, laptops, and desktops, to clusters of multicore nodes and supercomputers, usually comprising one or several coprocessors of different types (GPU, MIC, FPGA). This explains why it is vital to educate new generations of scientists and engineers about a range of PDC-related topics as we prepare them to effectively use modern computational systems. In a word, PDC-related topics must appear early and often in modern courses in computational science, computer science, and computer engineering.

In 2010, the IEEE Computer Society Technical Committee on Parallel Processing (TCPP) launched the Curriculum Initiative on Parallel and Distributed Computing, with Core Topics for Undergraduates. This led in 2011 to the EduPar workshop, which is dedicated to parallel and distributed computing education. Given the differences in educational environments in different parts of the world, the Euro-EDUPAR workshop starts with the aim of analyzing PDC education in a European context, i.e., within the structure and organization of European education.

Thus, Euro-EDUPAR is dedicated to analyzing where and how to include topics related to both PDC and HPC (high-performance computing) within the curricula of programs in computer science and engineering and computational science, while emphasizing European undergraduate teaching. The workshop especially seeks papers that report on experiences with incorporating PDC-related topics into undergraduate core courses taken by the majority of students on a degree course. Methods, pedagogical approaches, tools, and techniques that have potential for adoption across the European teaching community are of particular interest.

Topics of interest include: PDC teaching in the European space; pedagogical issues in PDC, educational methods, and learning mechanisms; novel ways of teaching PDC topics, including informal learning environments; curriculum design, models for incorporating PDC topics in core CS/CE curriculum; experience with incorporating PDC topics into core CS/CE courses; experience with incorporating PDC topics in the context of other applications learning; pedagogical tools, programming environments, and languages for PDC; e-Learning, e-Laboratory, Massive Open Online Courses (MOOC), Small Private Online Courses (SPOC); PDC experiences at non-university levels, secondary school, postgraduate, industry, diffusion of PDC.

Program Chairs

Domingo Giménez	University of Murcia, Spain
Sushil K. Prasad	Georgia State University, USA
Arnold L. Rosenberg	Northeastern University, Boston, USA

Program Committee

Marco Aldinucci	University of Turin, Italy
Paolo Bientinesi	RWTH Aachen University, Germany
Florina Ciorba	Technische Universität Dresden, Germany
Pierre-François Dutot	Université Pierre-Mendès Grenoble, France
Anshul Gupta	IBM Research, USA
Emmanuel Jeannot	Inria, France
Eleni Karatza	Aristotle University of Thessaloniki, Greece
Kishore Kothapalli	International Institute of Information Technology, Hyderabad, India
Friedhelm Meyer auf der Heide	University of Paderborn, Germany
Milan D. Mihajlovic	The University of Manchester, UK
Julio Ortega	University of Granada, Spain
Dana Petcu	West University of Timisoara, Romania
Cynthia Phillips	Sandia National Laboratories, USA
Martin Quinson	University of Lorraine/Inria Nancy, France
Noemi Rodriguez	PUC-Rio, Brazil
Julio Sahuquillo	Technical University of Valencia, Spain
Mitsuhsisa Sato	RIKEN Advanced Institute of Computational Science, Japan
Christian Scheideler	Universität Paderborn, Germany
Leonel Sousa	INESC-ID, IST, Universidade de Lisboa, Portugal
Peter Strazdins	Australian National University, Australia
Frédéric Vivien	Inria, France
Roman Wyrzykowski	Czestochowa University of Technology, Poland
Julius Žilinskas	Vilnius University, Lithuania

Additional Reviewers

Javier Cuenca	University of Murcia, Spain
Diego Fabregat-Traver	RWTH Aachen University, Germany
Daniel Ruprecht	University of Leeds, UK

13th International Workshop on Algorithms, Models, and Tools for Parallel Computing on Heterogeneous Platforms (HeteroPar)

Workshop Description

HeteroPar is a forum for researchers working on algorithms, programming languages, tools, and theoretical models aimed at efficiently solving problems on heterogeneous platforms. Heterogeneity is emerging as one of the most profound and challenging characteristics of today's parallel environments. From the macro level, where networks of distributed computers, composed by diverse node architectures, are interconnected with potentially heterogeneous networks, to the micro level, where deeper memory hierarchies and various accelerator architectures are increasingly common, the impact of heterogeneity on all computing tasks is increasing rapidly. Traditional parallel algorithms, programming environments and tools, designed for legacy homogeneous multiprocessors, will at best achieve a small fraction of the efficiency and the potential performance that we should expect from parallel computing in tomorrow's highly diversified and mixed environments. New ideas, innovative algorithms, and specialized programming environments and tools are needed to efficiently use these new and multifarious parallel architectures.

The 13th International Workshop on Algorithms, Models and Tools for Parallel Computing on Heterogeneous Platforms (HeteroPar 2015) was held in Vienna, Austria. For the seventh time, this workshop was organized in conjunction with the Euro-Par 2015 annual series of international conferences. The format of the workshop includes a keynote, followed by technical presentations, and ending with a panel. The workshop was well attended—there was 40 attendees.

This year, we received 26 articles, from 15 countries for review. After a thorough peer-reviewing process, we selected eight articles for presentation at the workshop. The review process focused on innovation and on proven applicability to heterogeneous settings. As a consequence, the quality and the relevance of the selected articles are very high. The acceptance ratio of 31 % is a result of the reviewers' discussion and not of cut-off selection; none of the articles submitted to HeteroPar 2015 was rejected because of the acceptance of other articles. The accepted articles represent an interesting mix of topics, techniques, applications, and scales, exhibiting nicely the diversity and growth of the heterogeneous computing field.

The Panel on Next Generation Heterogeneous Computing was led by, in alphabetical order, Prof. Dr. Henri Bal (VU Amsterdam, The Netherlands), Dr. Guojing Cong (IBM T.J. Watson Research Center, NY, USA), Prof. Dr. Miriam Leeser (Northeastern University in Boston, MA, USA), Dr. Martin Schultz (Lawrence Livermore National Lab, CA, USA), and Christian Iwainsky (TU Darmstadt, Germany).

Last, but certainly not least, I would like to thank the HeteroPar Steering Committee and the HeteroPar 2015 Program Committee, who made the workshop possible. I would also like to thank Euro-Par for hosting our community, and the Euro-Par workshops chair Dr. Sascha Hunold for his timely help.

Steering Committee

Domingo Giménez	University of Murcia, Spain
Alexey Kalinov	Cadence Design Systems, Russia
Alexey Lastovetsky	University College Dublin, Ireland
Yves Robert	Ecole Normale Supérieure de Lyon, France
Leonel Sousa	INESC-ID/IST, TU Lisbon, Portugal
Denis Trystram	LIG, Grenoble, France

Program Chair

Alexandru Iosup	Delft University of Technology, The Netherlands
-----------------	---

Program Committee

Rosa M. Badia	Barcelona Supercomputing Center, Spain
Jorge Barbosa	Faculdade de Engenharia do Porto, Portugal
Olivier Beaumont	Inria Futurs Bordeaux, LABRI, France
Cristina Boeres	Universidade Federal Fluminense, Brazil
George Bosilca	ICL, UTK, USA
Louis-Claude Canon	University of Franche-Comte, France
Alexandre Denis	Inria, France
Toshio Endo	Tokyo Institute of Technology, Japan
Jianbin Fang	NUDT, China
Edgar Gabriel	University of Houston, USA
Shuichi Ichikawa	Toyohashi University of Technology, Japan
Emmanuel Jeannot	Inria, France
Helen Karatza	Aristotle University of Thessaloniki, Greece
Hatem Ltaief	KAUST, Saudi Arabia
Pierre Manneback	University of Mons, Belgium
Satoshi Matsuoka	Tokyo Institute of Technology, Japan
Rafael Mayo	Universidad Jaume I, Spain
Wahid Nasri	ESST de Tunis, Tunisia
Nacho Navarro	Barcelona Supercomputing Center, Spain
Dana Petcu	University of Timisoara, Romania
Antonio J. Plaza	University of Extremadura, Spain
Thomas Rauber	University of Bayreuth, Germany
Matei Ripeanu	University of British Columbia, Canada
Vladimir Rychkov	University College Dublin, Ireland
Mitsuhsa Sato	University of Tsukuba, Japan
Erik Saule	University of North Carolina at Charlotte, USA

Tom Scogland

Henk Sips

Ana Lucia Varbanescu

Antonio M. Vidal

Frederic Vivien

Jon Weissman

Lawrence Livermore National Laboratory, USA

Delft University of Technology, The Netherlands

University of Amsterdam, The Netherlands

Universidad Politecnica de Valencia, Spain

Inria, France

University of Minnesota, USA

Third Workshop on Large-Scale Distributed Virtual Environments (LSDVE)

Workshop Description

The focus of the workshop has been the investigation of different aspects of distributed cooperative applications. Several novel applications have emerged in this area in the last few years. These include distributed social networks, distributed social games, collaborative recommender systems, collaborative learning systems, large-scale crowd-based applications, supported collaborative work (CSCW), and massively multi-player games.

The realization of these applications requires affording several challenges, such as the definition of user interfaces, coordination protocols, and proper middleware and architectures supporting distributed cooperation. Collaborative applications may greatly benefit from the support of different kinds of platforms, both cloud and peer to peer and also platforms recently proposed for the Internet of Things (IoT), such as fog computing. The integration of different platforms, for instance, mobile and cloud environments, is currently a challenge.

Some important challenges in the area of large-scale virtual environments are collaborative protocols design, latency reduction/hiding techniques for guaranteeing real-time constraints, large-scale processing of user information, privacy and security issues, state consistency/persistence.

The workshop investigated open challenges in this area, related to the design of new applications and to the definition of proper environments and frameworks for their development. LSDVE 2015 was a venue for researchers to present and discuss important aspects of large-scale collaborative applications and of the platforms supporting them.

The workshop opened with the keynote “Distributed Virtual Environments: From Client Server to Cloud and P2P Architectures: A Tutorial” given by Prof. Laura Ricci, University of Pisa, Italy.

The workshop organizers want to thank the authors of the papers for joining us in Vienna, the Program Committee and all the referees for doing the hard work of reviewing all the submissions, the conference organizers for proving a great support, and the researchers attending the workshop in Vienna.

Program Chairs

Laura Ricci	University of Pisa, Italy
Alexandru Iosup	TU Delft, Delft, The Netherlands
Radu Prodan	University of Innsbruck, Austria

Program Committee

Michele Amoretti	University of Parma, Italy
Emanuele Carlini	ISTI CNR, Pisa, Italy
Massimo Coppola	ISTI CNR, Pisa, Italy
Patrizio Dazzi	ISTI CNR, Pisa, Italy
Juan J. Durillo	University of Innsbruck, Austria
Kalman Graffi	University of Dusseldorf, Germany
Barbara Guidi	University of Pisa, Italy
Alexandru Iosup	TU Delft, The Netherlands
Pedro Garcia Lopez	University Rovira i Virgili, Spain
Pietro Michiardi	EURECOM, France
Dana Petcu	West University of Timisoara, Romania
Radu Prodan	University of Innsbruck, Austria
Laura Ricci	University of Pisa, Pisa, Italy
Alexey Vinel	Tampere University of Technology, Finland

Additional Reviewers

Valerio Arnaboldi	IIT, CNR, Pisa, Italy
Jean Botev	Université du Luxembourg, Luxembourg
Hanna Kavalionak	ISTI, CNR, Pisa, Italy

4th International Workshop on On-Chip Memory Hierarchies and Interconnects (OMHI)

Workshop Description

Current chip multiprocessors (CMPs) include several levels of on-chip caches to avoid the huge latencies of accessing the off-chip DRAM main memory modules. These caches must be efficiently interconnected to avoid performance penalties. On-chip networks are used to interconnect the memory hierarchy inside the processor chip. Latencies can be significantly affected by the devised on-chip memory hierarchy and the interconnect design, whose impact on the overall latency strongly depends on the core count. Consequently, this problem aggravates with the increasing core counts, which is the current commercial trend. By contrast, the main concern in GPUs is on memory bandwidth instead of latencies. Current GPUs are designed to hide memory latencies through fine-grained multithreading. The main goal of on-chip memories in current GPUs is to reduce off-chip memory traffic. In this context, the programmer plays a key role in improving cache access locality. Hence we can conclude that CPUs and GPUs require memory organizations with different characteristics. Thus, as current heterogeneous CPU-GPU systems are proliferating in the market, the memory system must be designed to efficiently support both types of memory organizations: latency-oriented and bandwidth-oriented.

The on-chip memory hierarchy occupies two thirds of the processor area and consumes a significant fraction of the overall system power. To deal with processor scalability issues, new technologies have emerged to implement the on-chip hierarchy. Regarding on-chip memory technologies, current SRAM technologies deployed in on-chip caches present important design challenges in terms of density and leakage currents. Instead, alternative technologies addressing leakage and density, such as eDRAM or MRAM, are being implemented and explored in large CMPs. Also the current electronic technology used in on-chip networks has important performance and power scalability limitations and designs using alternative technologies such as photonics or wireless are being proposed.

To efficiently leverage any on-chip memory hierarchy design, efforts must be focused on the management of shared resources, especially in the context of multicore systems where multiple threads contend while accessing these resources. This management involves, among others, thread allocation policies, cache management strategies, and NoC design. In this context, the synergy between the research on memory organization and management, interconnection networks, as well as novel technologies becomes a key strategy for fostering further developments. With this aim, the International Workshop on On-chip Memory Hierarchy and Interconnects (OMHI) started in 2012 and continued with its fourth edition in 2015, which was held in Vienna, Austria. This workshop is organized in conjunction with the Euro-Par annual

series of international conferences dedicated to the promotion and advancement of all aspects of parallel computing.

The goal of the OMHI workshop is to provide a forum for engineers and scientists to address the aforementioned challenges, and to present new ideas for future on-chip memory hierarchies and interconnects focusing on organization, management, and implementation. The specific topics covered by the OMHI workshop have been kept up to date according to technology advances and industrial and academia interests.

The chairs of OMHI were proud to present Prof. Sandro Bartolini as keynote speaker, who gave an interesting talk focusing on the key topics of the workshop entitled “Illuminating Processors: How Photonics Will Help Computing,” which together with the paper session resulted in an interesting and very exciting one-day program.

Finally, the chairs would like to thank the members of the Program Committee for their reviews, the Euro-Par organizers, Sandro Bartolini, and all of the attendees. Based on the positive feedback from all of them, we plan to continue the OMHI workshop in conjunction with Euro-Par 2016.

Program Chairs

Julio Sahuquillo	Universitat Politècnica de València, Spain
María Engracia Gómez	Universitat Politècnica de València, Spain
Salvador Petit	Universitat Politècnica de València, Spain

Program Committee

Manuel Acacio	Universidad de Murcia, Spain
Sandro Bartolini	Università di Siena, Italy
João M.P. Cardoso	University of Porto, Portugal
Marcello Coppola	STMicroelectronics, France
Giorgos Dimitrakopoulos	Democritus University of Thrace, Greece
Pierfrancesco Foglia	Università di Pisa, Italy
Holger Fröning	University of Heidelberg, Germany
Crispín Gómez	Universitat Politècnica de València, Spain
Kees Goossens	NXP Semiconductors and Delft University of Technology, The Netherlands
David Kaeli	Northeastern University, USA
Sonia López	Rochester Institute of Technology, USA
Iakovos Mavroidis	Foundation for Research and Technology - Hellas (FORTH), Greece
Pierre Michaud	Inria, France
Tor Skei	Simula Research Laboratory, Norway
Rafael Ubal	Northeastern University, USA

Third Workshop on Parallel and Distributed Agent-Based Simulations (PADABS)

Workshop Description

Agent-based simulation models are an increasingly popular tool for research and management in many fields such as ecology, economics, sociology, etc. In some fields, such as social sciences, these models are seen as a key instrument to the generative approach, essential for understanding complex social phenomena. But also in policy-making, biology, military simulations, control of mobile robots and economics, the relevance and effectiveness of agent-based simulation models has been recently recognized. The computer science community has responded to the need for platforms that can help the development and testing of new models in each specific field by providing tools, libraries, and frameworks that speed up and make massive simulations. The key objective of this workshop is to bring together researchers who are interested in getting more performance from their simulations, by using synchronized, many-core simulations (e.g., GPUs), strongly coupled, parallel simulations (e.g., MPI) and loosely coupled, distributed simulations (distributed heterogeneous setting).

Several frameworks have been recently developed and are active in this field. They range from the GPU-manycore approach, to parallel, to distributed simulation environments. In the first category, you can find FLAME GPU, which also allows non-GPU specialists to harness the GPUs performance for real-time simulation and visualization. For tightly-coupled, large computing clusters and supercomputers a very popular framework is Repast for High-Performance Computing (REPAST-HPC), a C++-based modeling system. On the distributed side, recent work on Distributed Mason, allows non-specialists to use heterogeneous hardware and software in local area networks for enlarging the size and speeding up the simulation of complex agent-based models.

Therefore, our focus and positioning is on the applied side of parallel computing, with a particular emphasis on performance but also on the expressivity of the frameworks, since the field that is the target of our research is multidisciplinary and does not include only “hard-science” scientists.

Program Chairs

Vittorio Scarano (Chair)	Università di Salerno, Italy
Gennaro Cordasco	Seconda Università di Napoli, Italy
Ugo Erra	Università della Basilicata, Italy
Carmine Spagnuolo (Publicity Chair)	Università di Salerno, Italy

Program Committee

Maria Chli	Aston University, UK
Claudio Cioffi-Revilla	George Mason University, USA
Biagio Cosenza	University of Innsbruck, Austria
Nick Collier	Argonne National Laboratory, USA
Rosaria Conte	CNR, Italy
Andrew Evans	University of Leeds, UK
Bernardino Frola	The MathWorks, Cambridge, UK
Joanna Kolodziej	Cracow University of Technology and AGH University of Science and Technology, Cracow, Poland
Nicola Lettieri	Università del Sannio and ISFOL, Italy
Sean Luke	George Mason University, USA
Michael North	Argonne National Laboratory, USA
Mario Paolucci	CNR, Italy
Paul Richmond	The University of Sheffield, UK
Arnold Rosenberg	Northeastern University, USA
Flaminio Squazzoni	Università di Brescia, Italy
Michela Taufer	University of Delaware, USA

Additional Reviewers

Carmine Spagnuolo	Università di Salerno, Italy
Luca Vicidomini	Università di Salerno, Italy

First Workshop on Performance Engineering for Large-Scale Graph Analytics (PELGA)

Workshop Description

The knowledge economy is based on data, of which graphs represent an increasing part, in advanced marketing, in social networking, in life sciences, in health and bioinformatics services, in academic networks, in hiring of professionals, etc. As a consequence, graph analytics is fast becoming a significant consumer of computing resources, due to the ever larger graphs of hundreds of millions up to hundreds of billions of edges, and to the increased complexity of analysis tasks. To enable existing algorithms to fit modern architectures and scale with these new requirements, there is a growing need for performance engineering.

PELGA is a venue that aims to address this need. Its goal is to bring together specialists from both industry and academia to discuss the state of the art of graph processing systems, with a special focus on performance. Hosting PELGA with EuroPar allows the largest community of parallel and distributed systems in Europe and elsewhere to participate in the discussion and acknowledge the new research opportunities that large-scale graph processing presents.

PELGA is a venue that welcomes contributions focusing on graph-centric performance engineering tools and methods, workload characterization, new algorithms and new graph processing systems, and performance modeling. Less conventional workshop topics such as surveys, performance studies, comparative analyses are also encouraged, given the young age of the large-scale graph processing community. We strive to cover the specifics of three large classes of topics.

1. Systems invites contributions focusing on new graph processing systems focused on high-performance analytics, performance studies of existing systems to be used for graph processing, and comparative and/or in-depth analysis of graph processing systems.
2. Algorithms, Applications, and Architectures is the largest topic cluster, including work focusing on new high-performance graph processing algorithms, new performance-aware applications for graph processing algorithms, platform-specific algorithms and their performance optimization (e.g., GPUs, Xeon Phi, heterogeneous platforms) for graph analytics, algorithms and/or architectures for large-scale graph analytics, and partitioning methods for large-scale or otherwise challenging graphs.
3. Characterization, Modeling, and Engineering is the core of the workshop. We encourage novel contributions focusing on graph models for performance tuning and/or prediction of analytics workloads, performance models for prediction or ranking of graph processing platforms, performance analysis and engineering of

existing graph processing algorithms, and tools and benchmarks for graph-centric performance engineering.

In summary, large-scale graph processing is a high-impact field in full development, driven by both data owners and the analytics world. As we recognize the need to adapt traditional performance evaluation, analysis, and modeling to the needs of this dynamic new topic, PELGA is a workshop with a strong community focus, aiming to bring the challenges of large-scale graph processing to the attention of the EuroPar community as an unconventional, yet very relevant topic for parallel and distributed computing.

Program Chairs

Ana Lucia Varbanescu
Alexandru Iosup

University of Amsterdam, The Netherlands
Delft University of Technology, The Netherlands

Program Committee

Arnau Prat-Perez
Claudio Martella
Hannes Muhleisen
Hassan Chafi
Sungpack Hong
Jan Hidders
Josep Lluís Larriba Pey
Matei Ripeanu
Mihai Capota
Ted Willke

UPC, Spain
VU University Amsterdam, The Netherlands
CWI Amsterdam, The Netherlands
Oracle Labs, USA
Oracle Labs, USA
Delft University of Technology, The Netherlands
UPC, Spain
The University of British Columbia, Canada
Intel Labs, USA
Intel Labs, USA

Second International Workshop on Reproducibility in Parallel Computing (REPPAR)

Workshop Description

Conducting sound and reproducible experiments in parallel computing is not easy, as hardware and software architectures of current parallel computers are most often very complex. This high complexity makes it difficult—and often impossible—for computer scientists to model such systems mathematically. For that reason, scientists rely on experiments to study new parallel algorithms, different software solutions (e.g., operating systems), or novel hardware architectures. The situation in parallel computing is made even more difficult than it would be otherwise, as parallel systems are in a constant state of flux, e.g., the total core count is rapidly growing and many programming paradigms for parallel machines have emerged and are actively being used in a hybrid fashion, e.g., MPI, OpenMP, or PGAS.

For these reasons, the workshop is concerned with experimental practices in parallel computing research. We solicit research papers and experience reports on a number of relevant topics, particularly: methods for analysis and visualization of experimental data, best-practice recommendations, results of attempts to replicate previously published experiments, and tools for experimental computational sciences. Some examples of the latter include workflow management systems, experimental testbeds, and systems for archiving and querying large data files.

Program Chairs

Sascha Hunold	Vienna University of Technology, Austria
Arnaud Legrand	CNRS, LIG, Grenoble, France
Lucas Nussbaum	Université de Lorraine, LORIA, France
Mark Stillwell	Imperial College London, UK

Program Committee

Andrew Davison	CNRS, UNIC, Gif-sur-Yvette, France
Georg Hager	University of Erlangen-Nuremberg, Germany
Sascha Hunold	Vienna University of Technology, Austria
Arnaud Legrand	CNRS, LIG, Grenoble, France
Lucas Nussbaum	Université de Lorraine, LORIA, France
Olivier Richard	Université Joseph Fourier, LIG, Grenoble, France
Lucas M. Schnorr	Universidade Federal do Rio Grande do Sul, Porto Alegre, Brazil
Mark Stillwell	Imperial College London, UK
Jesper Larsson Träff	Vienna University of Technology, Austria

8th Workshop on Resiliency in High-Performance Computing in Clusters, Clouds, and Grids (Resilience)

Workshop Description

Clouds, grids, and clusters are three different computational paradigms with the potential to support high-performance computing (HPC) and enterprise IT infrastructure. Currently, they consist of hardware, management, and usage models particular to different computational regimes [e.g., high-performance cluster systems designed to support tightly coupled scientific simulation codes typically utilize high-speed interconnects and commercial cloud systems designed to support software as a service (SAS) typically do not]. However, in order to support HPC, all must at least utilize large numbers of resources and hence effective HPC in any of these paradigms must address the same issue of resiliency at a very large scale.

Recent trends in HPC systems have clearly indicated that future increases in performance, in excess of those resulting from improvements in single-processor performance, will be achieved through corresponding increases in system scale, i.e., using a significantly larger component count. As the raw computational performance of the world's fastest HPC systems increases from today's current multi-petascale to next-generation exascale capability and beyond, their number of computational, networking, and storage components will grow from the ten to one hundred thousand compute nodes of today's systems to several hundreds of thousands of compute nodes in the foreseeable future. This substantial growth in system scale, and the resulting component count, poses a challenge for HPC system and application software with respect to reliability, availability, and serviceability (RAS).

The goal of this workshop is to bring together experts in the area of fault tolerance and resilience for HPC to present the latest achievements and to discuss the challenges ahead. The program of the Resilience 2015 workshop included one keynote and five high-quality papers. The keynote was given by Christian Engelmann from Oak Ridge National Laboratory with the title "Toward A Fault Model and Resilience Design Patterns for Extreme Scale Systems."

Third Workshop on Runtime and Operating Systems for the Many-Core Era (ROME)

Workshop Description

Since the beginning of the multicore era, parallel processing has become prevalent across the board. However, in order to continue a performance increase according to Moore's law, the next step needs to be taken: away from common multicores toward innovative many-core architectures. Such systems, equipped with a significantly higher amount of cores per chip than multicores, pose challenges in both hardware and software design. On the hardware side, complex on-chip networks, scratchpads, hybrid memory cubes, non-volatile memory and stacked memory as well as deep cache hierarchies and novel cache-coherence strategies will enrich the current research areas in the future.

However, the ROME workshop (Runtime and Operating Systems for the Many-Core Era) focuses on the software side because without complying system software, runtime and operating system support, all these new hardware facilities cannot be exploited. Hence, the new challenges in hardware/software co-design are to step beyond traditional approaches and to create new programming models and operating system designs in order to exploit the theoretically available performance of future hardware as effectively and as power-aware as possible.

This focus of the ROME workshop stands in the tradition of a successful series of events originally hosted by the Many-Core Applications Research Community (MARC). Prior MARC symposia took place at ONERA Research Center in Toulouse, at the Hasso Plattner Institute in Potsdam, and at the RWTH Aachen University. Starting in 2013, the organizers continued this series by establishing ROME as one of the co-located workshops of Euro-Par, the prime European conference for parallel and distributed computing.

While the first ROME workshop, which was hosted at Euro-Par 2013 in Aachen, was still a MARC-related follow-up event but for a broader audience, the second ROME workshop, held in conjunction with Euro-Par 2014 in Porto, already expanded its focus to research questions arising from the upcoming generation of heterogeneous and/or massive parallel systems stepping toward a many-core-dominated exascale era.

In 2015, this broader focus was essentially retained for the third ROME workshop, which was held in conjunction with Euro-Par 2015 in Vienna, but the relevance of runtime and operating system aspects was stressed once again as being the primary scope of the ROME workshop series. In this spirit, the organizers were very happy that Dr. Carsten Weinhold from the Operating Systems Group of TU Dresden, Germany, volunteered to give an invited keynote for this third ROME workshop with the title "A Microkernel-Based Operating System for Exascale Computing."

Program Chairs

Stefan Lankes RWTH Aachen University, Germany
Carsten Clauss ParTec Cluster Competence Center GmbH

Program Committee

Jens Breitbart TU München, Germany
André Brinkmann Johannes Gutenberg-Universität Mainz, Germany
Carsten Clauss ParTec Cluster Competence Center GmbH
Christos Kartsaklis Oak Ridge National Laboratory, USA
Stefan Lankes RWTH Aachen University, Germany
Timothy Mattson Intel Labs, USA
Jörg Nolte BTU Cottbus, Germany
Michael Riepen IAV GmbH, Germany
Bettina Schnor University of Potsdam, Germany
Christian Terboven RWTH Aachen University, Germany
Theo Ungerer Universität Augsburg, Germany
Josef Weidendorfer TU München, Germany

Additional Reviewers

Steffen Christgau University of Potsdam, Germany
Sonja Kolen RWTH Aachen University, Germany
Stefan Petri University of Potsdam, Germany
Simom Pickartz RWTH Aachen University, Germany
Lukas Razik RWTH Aachen University, Germany

8th Workshop on UnConventional High-Performance Computing 2015 (UCHPC)

Workshop Description

Recent issues regarding the power consumption of conventional HPC hardware has resulted both in new interest in accelerator hardware and in usage of mass-market hardware originally not designed for HPC. The most prominent examples are GPUs but FPGAs, DSPs, and embedded designs are also possible candidates to provide higher-power efficiency, as they are used in energy-restricted environments, such as smartphones or tablets. The so-called dark silicon forecast, i.e., not all transistors may be active at the same time, may lead to even more specialized hardware in future mass-market products. Exploiting this hardware for HPC can be a worthwhile challenge.

As the word “UnConventional” in the title suggests, the workshop focuses on usage of hardware or platforms for HPC, which are not (yet) conventionally used today, and may not have been designed for HPC in the first place. Reasons for its use can be raw computing power, good performance per watt, or low cost in general. To address this unconventional hardware, often new programming approaches and paradigms are required to make best use of it. Another focus of the workshop is on innovative, (yet) unconventional, new programming models and algorithms (e.g., Big Data) exploiting unconventional HPC hardware or software.

To this end, UCHPC tries to capture solutions for HPC that are unconventional today but could become conventional and significant tomorrow, and thus provide a glimpse into the future of HPC.

This year was the eighth time the UCHPC workshop took place, and it was the sixth time in a row it was co-located with Euro-Par (each year since 2010). Before that, it was held in conjunction with the International Conference on Computational Science and Its Applications 2008 and with the ACM International Conference on Computing Frontiers 2009. However, UCHPC is a perfect addition to the scientific fields of Euro-Par, and this is confirmed by the continuous interest we see among Euro-Par attendees for this workshop.

While the general focus of the workshop is fixed, the topic is actually a moving target. GPUs were quite unconventional for HPC a few years ago, but today a notable portion of the machines in the Top500 list are making use of them. Currently, the exploitation of mobile processors for HPC – including on-chip GPU and DSPs – is a hot topic. A recent technological breakthrough is mass-market production of 3D stacking technology, which allows us to put memory and logic nearer together. This may result in a revival of the processing-in-memory idea, which is quite unconventional from a programmer’s point of view and seems to be a good fit for UCHPC. To this end, we invited Zehra Sura from the IBM T.J. Watson Center to give a keynote

about IBM's recent research on "The Active Memory Cube: A Processing-in-Memory System for High-Performance Computing."

These proceedings include the final versions of the papers presented at UCHPC and accepted for publication. They take the feedback from the reviewers and workshop audience into account.

The workshop organizers/program chairs want to thank the authors of the papers for joining us in Vienna, the Program Committee for doing the hard work of reviewing all submissions, the conference organizers for providing such a nice venue, and last but not least the large number of attendees this year.

Program Chairs

Jens Breitbart	Technische Universität München, Germany
Josef Weidendorfer	Technische Universität München, Germany

Program Committee

Michael Bader	Technische Universität München, Germany
Denis Barthou	University of Bordeaux, France
Alex Bartzas	National Technical University of Athens, Greece
Lars Bengtsson	Chalmers University of Technology, Sweden
James Beyer	Cray Inc., USA
Jens Breitbart	Technische Universität München, Germany
Georgios Dimitrakopoulos	Democritus University of Thrace, Greece
Karl Furlinger	LMU München, Germany
Frank Hannig	University of Erlangen-Nuremberg, Germany
Anders Hast	Uppsala University, Sweden
Paul Keir	University of the West of Scotland, UK
Rainer Keller	Hochschule für Technik Stuttgart, Germany
Gaurav Khanna	University of Massachusetts Dartmouth, USA
Harald Köstler	University of Erlangen-Nuremberg, Germany
Stefan Lankes	RWTH Aachen, Germany
Dimitar Lukarski	Paralution Labs, Germany
Manfred Mücke	Materials Center Leoben, Austria
Yannis Papaefstathiou	Technical University of Crete, Greece
Bertil Schmidt	University of Mainz, Germany
Ioannis Sourdis	Chalmers University of Technology, Sweden
Dylan Stark	Sandia National Laboratories, USA
Robert Strzodka	Universität Heidelberg, Germany
Carsten Trinitis	Technische Universität München, Germany
Josef Weidendorfer	Technische Universität München, Germany
Jan-Philipp Weiss	COMSOL, Sweden
Gerhard Wellein	University of Erlangen-Nuremberg, Germany
Ren Wu	Baidu, China
Peter Zinterhof jun	University of Salzburg, Austria

10th Workshop on Virtualization in High-Performance Cloud Computing (VHPC)

Workshop Description

Virtualization technologies constitute a key enabling factor for flexible resource management in modern data centers, cloud environments, and increasingly in HPC as well. Providers need to dynamically manage complex infrastructures in a seamless fashion for varying workloads and hosted applications, independently of the customers deploying software or users submitting highly dynamic and heterogeneous workloads. Thanks to virtualization, we have the ability to manage vast computing and networking resources dynamically and close to the marginal cost of providing the services, which is unprecedented in the history of scientific and commercial computing.

OS-level virtualization, such as provided by Docker, allows for multiple isolated user-space environments within the same OS kernel. It promises to provide many of the advantages of machine virtualization with high levels of responsiveness and performance; coupled with lightweight OSs it forms a potent architecture with the potential of becoming a mainstream environment for HPC workloads.

Machine virtualization, with its capability to enable consolidation of multiple under-utilized servers with heterogeneous software and operating systems (OSs), and its capability to live-migrate a fully operating virtual machine (VM) with a very short downtime, enables novel and dynamic ways to manage physical servers.

I/O virtualization allows physical network adapters to take traffic from multiple VMs; network virtualization, with its capability to create logical network overlays that are independent of the underlying physical topology and IP addressing, provides the fundamental ground on top of which evolved network services can be realized with an unprecedented level of dynamicity and flexibility. These technologies have to be inter-mixed and integrated in an intelligent way, to support workloads that are increasingly demanding in terms of absolute performance, responsiveness, and interactivity, and have to respect well-specified service-level agreements (SLAs), as needed for industrial-grade provided services.

The Workshop on Virtualization in High-Performance Cloud Computing (VHPC) aims to bring together researchers and industrial practitioners facing the challenges posed by virtualization in order to foster discussion, collaboration, and mutual exchange of knowledge and experience, thereby enabling research to ultimately provide novel solutions for virtualized computing systems of tomorrow.

Program Chairs

Michael Alexander	TU Wien, Austria
Anastassios Nanos	NTUA, Greece
Balazs Gerofi	RIKEN, Japan

Program Committee

Stergios Anastasiadis	University of Ioannina, Greece
Costas Bekas	IBM, Switzerland
Jakob Blomer	CERN, Switzerland
Ron Brightwell	Sandia National Laboratories, USA
Roberto Canonico	University of Napoli Federico II, Italy
Julian Chesterfield	OnApp, UK
Piero Castoldi	Sant'Anna School of Advanced Studies, Italy
Patrick Dreher	MIT, USA
William Gardner	University of Guelph, Canada
Kyle Hale	Northwestern University, USA
Marcus Hardt	Karlsruhe Institute of Technology, Germany
Iftekhhar Hussain	Infinera, USA
Krishna Kant	Temple University, USA
Eiji Kawai	National Institute of Information and Communications Technology, Japan
Romeo Kinzler	IBM, Switzerland
Kornilios Kourtis	ETH, Switzerland
Nectarios Koziris	National Technical University of Athens, Greece
Massimo Lamanna	CERN, Switzerland
Che-Rung Roger Lee	National Tsing Hua University, Taiwan
William Magato	University of Cincinnati, USA
Helge Meinhard	CERN, Switzerland
Jean-Marc Menaud	Ecole des Mines de Nantes, France
Christine Morin	Inria, France
Amer Qouneh	University of Florida, USA
Seetharami Seelam	IBM T.J. Watson Research Center, USA
Josh Simons	VMWare, USA
Borja Sotomayor	University of Chicago, USA
Kurt Tutschku	Blekinge Institute of Technology, Sweden
Yasuhiro Watashiba	Osaka University, Japan
Chao-Tung Yang	Tunghai University, Taiwan

Contents

BigDataCloud - Big Data Management in Clouds

Distributed Range-Based Meta-Data Management for an In-Memory Storage	3
<i>Florian Klein, Kevin Beineke, and Michael Schöttner</i>	
Network-Based Data Processing Architecture for Reliable and High-Performance Distributed Storage System	16
<i>Hiroki Ohtsuji and Osamu Tatebe</i>	
File-Less Approach to Large Scale Data Management	27
<i>Bartosz Kryza and Jacek Kitowski</i>	

Euro-EDUPAR - Parallel and Distributed Computing Education for Undergraduate Students

Parallel Computing vs. Distributed Computing: A Great Confusion? (Position Paper)	41
<i>Michel Raynal</i>	
SAUCE: A Web-Based Automated Assessment Tool for Teaching Parallel Programming	54
<i>Moritz Schlarb, Christian Hundt, and Bertil Schmidt</i>	
Teaching Parallel Programming in Interdisciplinary Studies	66
<i>Eduardo Cesar, Ana Cortés, Antonio Espinosa, Tomàs Margalef, Juan Carlos Moure, Anna Sikora, and Remo Suppi</i>	
On-line Service for Teaching Parallel Programming	78
<i>Marek Nowicki, Maciej Marchwiany, Maciej Szpindler, and Piotr Bala</i>	
Challenges of a Systematic Approach to Parallel Computing and Supercomputing Education	90
<i>Vladimir Voevodin, Victor Gergel, and Nina Popova</i>	
Teaching Heart Modeling and Simulation on Parallel Computing Systems . . .	102
<i>Andrey Sozykin, Mikhail Chernoskutov, Anton Koshelev, Vladimir Zverev, Konstantin Ushenin, and Olga Solovyova</i>	
Integration of ICT in Concurrent and Parallel Programming Lectures.	114
<i>Antonio J. Tomeu-Hardasmal, Alberto G. Salguero, and Manuel I. Capel</i>	

Teamwork Across Disciplines: High-Performance Computing Meets Engineering	125
<i>Philipp Neumann, Christoph Kowitz, Felix Schraner, and Dmitrii Azarnykh</i>	
An Educational Module Illustrating How Sparse Matrix-Vector Multiplication on Parallel Processors Connects to Graph Partitioning	135
<i>M. Ali Rostami and H. Martin Buecker</i>	
FERBJMON Tools - Visualizing Thread Access on Java Objects using Lightweight Runtime Monitoring	147
<i>Marvin Ferber</i>	
Interdisciplinary Practical Course on Parallel Finite Element Method Using HiFlow ³	160
<i>Markus Hoffmann, Simon Gawlok, Eva Treiber, Wolfgang Karl, and Vincent Heuveline</i>	
HeteroPar - Algorithms, Models, and Tools for Parallel Computing on Heterogeneous Platforms	
A Randomized LU-based Solver Using GPU and Intel Xeon Phi Accelerators	175
<i>Marc Baboulin, Amal Khabou, and Adrien Rémy</i>	
Identifying Optimization Opportunities Within Kernel Execution in GPU Codes	185
<i>Robert Lim, Allen Malony, Boyana Norris, and Nick Chaimov</i>	
Modeling Contention and Mapping Effects in Multi-core Clusters	197
<i>Juan-Antonio Rico-Gallego, Juan-Carlos Díaz-Martín, and Alexey L. Lastovetsky</i>	
Towards Community Detection on Heterogeneous Platforms	209
<i>Stijn Heldens, Ana Lucia Varbanescu, Arnau Prat-Pérez, and Josep-Lluis Larriba-Pey</i>	
A Design Proposal for a Next Generation Scientific Software Framework . . .	221
<i>Anshu Dubey and Daniel T. Graves</i>	
Accelerating Direction-Optimized Breadth First Search on Hybrid Architectures	233
<i>Scott Sallinen, Abdullah Gharaibeh, and Matei Ripeanu</i>	
FiNS: A Framework for Accelerating Nested Simulations on Heterogeneous Platforms	246
<i>Joris Cramwinckel, Stefan Singor, and Ana Lucia Varbanescu</i>	

Communication Models Insights Meet Simulations 258
Pierre-François Dutot, Millian Poquet, and Denis Trystram

LSDVE - Large Scale Distributed Virtual Environments

Community Discovery for Interest Management in DVEs: A Case Study 273
*Emanuele Carlini, Patrizio Dazzi, Matteo Mordacchini,
Alessandro Lulli, and Laura Ricci*

Continuation Complexity: A Callback Hell for Distributed Systems. 286
Edgar Zamora-Gómez, Pedro García-López, and Rubén Mondéjar

Offloading Service Provisioning on Mobile Devices in Mobile Cloud
Computing Environments. 299
Marco Conti, Davide Mascitti, and Andrea Passarella

A Systematic Quality Analysis of Virtual Desktop Infrastructure
Technologies 311
Arman Sheikholeslami and Kalman Graffi

A Trustworthy Distributed Social Carpool Method 324
*Francisco Martín-Fernández, Cándido Caballero-Gil,
and Pino Caballero-Gil*

**OMHI - On-Chip Memory Hierarchies and Interconnects:
Organization, Management and Implementation**

Efficient DVFS Operation in NoCs Through a Proper Congestion
Management Strategy 339
José V. Escamilla, José Flich, and Pedro Javier García

Superoptimizing Memory Subsystems for Multiple Objectives 352
Joseph G. Wingbermuehle, Ron K. Cytron, and Roger D. Chamberlain

PADABS - Parallel and Distributed Agent-Based Simulations

On Evaluating Graph Partitioning Algorithms for Distributed Agent
Based Models on Networks 367
*Alessia Antelmi, Gennaro Cordasco, Carmine Spagnuolo,
and Luca Vicidomini*

Distributed Agent-Based Simulation and GIS: An Experiment
with the Dynamics of Social Norms 379
Nicola Lettieri, Carmine Spagnuolo, and Luca Vicidomini

Behavioral Spherical Harmonics for Long-Range Agents' Interaction. 392
Biagio Cosenza

Graph-Based Automatic Dynamic Load Balancing for HPC	
Agent-Based Simulations	405
<i>Claudio Márquez, Eduardo César, and Joan Sorribes</i>	
Preliminary Evaluation of a Parallel Trace Replay Tool for HPC	
Network Simulations	417
<i>Bilge Acun, Nikhil Jain, Abhinav Bhatele, Misbah Mubarak, Christopher D. Carothers, and Laxmikant V. Kale</i>	
Road Network Simulation Using FLAME GPU	430
<i>Peter Heywood, Paul Richmond, and Steve Maddock</i>	
A Communication Schema for Parallel and Distributed Multi-agent Systems Based on MPI	442
<i>Alban Rousset, Bénédicte Herrmann, Christophe Lang, and Laurent Philippe</i>	
Large-Scale Agent-Based Modeling with Repast HPC: A Case Study in Parallelizing an Agent-Based Model	454
<i>Nicholson Collier, Jonathan Ozik, and Charles M. Macal</i>	
RAMSES: Reversibility-Based Agent Modeling and Simulation Environment with Speculation-Support	466
<i>Davide Cingolani, Alessandro Pellegrini, and Francesco Quaglia</i>	
PELGA - Performance Engineering for Large-Scale Graph Analytics	
Can Embedding Solve Scalability Issues for Mixed-Data Graph Clustering? . . .	481
<i>Nadezhda Fedorova, Josep Blat, and David F. Nettleton</i>	
Using the Marshall-Olkin Extended Zipf Distribution in Graph Generation	493
<i>Ariel Duarte-López, Arnau Prat-Pérez, and Marta Pérez-Casany</i>	
Highspeed Graph Processing Exploiting Main-Memory Column Stores	503
<i>Matthias Hauck, Marcus Paradies, Holger Fröning, Wolfgang Lehner, and Hannes Rauhe</i>	
A Multi-layer Framework for Graph Processing via Overlay Composition . . .	515
<i>Alessandro Lulli, Patrizio Dazzi, Laura Ricci, and Emanuele Carlini</i>	
Quantifying the Performance Impact of Graph Structure on Neighbour Iteration Strategies for PageRank.	528
<i>Merijn Verstraaten, Ana Lucia Varbanescu, and Cees de Laat</i>	
Accelerating Minimum Spanning Forest Computations on Multicore Platforms	541
<i>Guojing Cong, Ilie Tanase, and Yinglong Xia</i>	

Importance of Runtime Considerations in Performance Engineering
of Large-Scale Distributed Graph Algorithms 553
*Jesun Sahariar Firoz, Thejaka Amila Kanewala, Marcin Zalewski,
Martina Barnas, and Andrew Lumsdaine*

Characterizing Communication Patterns of Parallel Programs
Through Graph Visualization and Analysis. 565
Denise Stringhini and Alvaro Fazenda

REPPAR - Reproducibility in Parallel Computing

Reproducible and User-Controlled Software Environments in HPC
with Guix 579
Ludovic Courtès and Ricardo Wurmus

Reproducibility in Practice: Lessons Learned from Research
and Teaching Experiments 592
Antonio Maffia, Helmar Burkhart, and Danilo Guerrero

Towards Complete Tracking of Provenance in Experimental
Distributed Systems Research 604
Tomasz Buchert, Lucas Nussbaum, and Jens Gustedt

**Resilience - Resiliency in High Performance Computing
with Clouds, Grids, and Clusters**

A Case Study of Application Structure Aware Resilience Through
Differentiated State Saving and Recovery. 619
*Anshu Dubey, Hajime Fujita, Zachary Rubenstein, Brian Van Straalen,
and Andrew A. Chien*

A Holistic Approach to Log Data Analysis in High-Performance
Computing Systems: The Case of IBM Blue Gene/Q. 631
Alina Sîrbu and Ozalp Babaoglu

Addressing the Last Roadblock for Message Logging in HPC:
Alleviating the Memory Requirement Using Dedicated Resources. 644
Tatiana Martsinkevich, Thomas Ropars, and Franck Cappello

Towards Understanding Post-recovery Efficiency for Shrinking
and Non-shrinking Recovery 656
Aiman Fang, Hajime Fujita, and Andrew A. Chien

Canaries in a Coal Mine: Using Application-Level Checkpoints
to Detect Memory Failures. 669
Patrick M. Widener, Kurt B. Ferreira, Scott Levy, and Nathan Fabian

ROME - Runtime and Operating Systems for the Many-Core Era

Energy Characterization and Optimization of Parallel Prefix-Sums Kernels . . . 685
Angelos Papatriantafyllou

An OS-Oriented Performance Monitoring Tool for Multicore Systems 697
*Juan Carlos Saez, Jorge Casas, Abel Serrano,
 Roberto Rodríguez-Rodríguez, Fernando Castro, Daniel Chaver,
 and Manuel Prieto-Matias*

A Topology-Aware Performance Monitoring Tool for Shared Resource
 Management in Multicore Systems 710
Nicolas Denoyelle, Brice Goglin, and Emmanuel Jeannot

Diamond Rings: Acknowledged Event Propagation in Many-Core
 Processors 722
*Stefan Nürnberger, Randolph Rotta, Gabor Drescher, Daniel Danner,
 and Jörg Nolte*

UCHPC - UnConventional High Performance Computing

Energy-Performance Tradeoffs for HPC Applications on Low
 Power Processors 737
Enrico Calore, Sebastiano Fabio Schifano, and Raffaele Tripiccione

A Cache-Aware Performance Prediction Framework
 for GPGPU Computations 749
Alexander Pöppel and Alexander Herz

Towards Application Variability Handling with Component Models:
 3D-FFT Use Case Study 761
Vincent Lanore, Christian Perez, and Jérôme Richard

Optimized Force Calculation in Molecular Dynamics Simulations
 for the Intel Xeon Phi 774
*Nikola Tchipev, Amer Wafai, Colin W. Glass, Wolfgang Eckhardt,
 Alexander Heinecke, Hans-Joachim Bungartz, and Philipp Neumann*

VHPC - Virtualization in High-Performance Cloud Computing

A Simplified TDP with Large Tables. 789
Yu Zhang

GPGPU Virtualisation with Multi-API Support Using Containers 802
John Walsh and Jonathan Dukes

Performance Evaluation of Containers for HPC. 813
Cristian Ruiz, Emmanuel Jeanvoine, and Lucas Nussbaum

The Virtual Puppet Master: Adaptive Streaming on Top of an SDN-Enabled
Virtual Infrastructure 825
*Roberto Canonico, Enrico De Maio, Pasquale Di Rienzo,
and Simon Pietro Romano*

Author Index 837