

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, 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, 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>

Julian M. Kunkel · Rio Yokota
Michela Taufer · John Shalf (Eds.)

High Performance Computing

ISC High Performance 2017 International Workshops
DRBSD, ExaComm, HCPM, HPC-IODC, IWOPH, IXPUG,
P³MA, VHPC, Visualization at Scale, WOPSSS
Frankfurt, Germany, June 18–22, 2017
Revised Selected Papers



Springer

Editors

Julian M. Kunkel

Deutsches Klimarechenzentrum (DKRZ)
Hamburg, Hamburg
Germany

Rio Yokota

TITECH
Tokyo
Japan

Michela Taufer

Department of Computer Science
University of Delaware
Newark, DE
USA

John Shalf

Lawrence Berkeley National Laboratory
Berkeley, CA
USA

ISSN 0302-9743

ISSN 1611-3349 (electronic)

Lecture Notes in Computer Science

ISBN 978-3-319-67629-6

ISBN 978-3-319-67630-2 (eBook)

<https://doi.org/10.1007/978-3-319-67630-2>

Library of Congress Control Number: 2017955780

LNCS Sublibrary: SL1 – Theoretical Computer Science and General Issues

© Springer International Publishing AG 2017, corrected publication 2017

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. The publisher remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.

Printed on acid-free paper

This Springer imprint is published by Springer Nature

The registered company is Springer International Publishing AG

The registered company address is: Gewerbestrasse 11, 6330 Cham, Switzerland

Preface

A separate workshop day attached to the ISC High Performance, formerly known as the International Supercomputing Conference, was first added to the Technical Program in 2015 under the leadership of Bernd Mohr (Forschungszentrum Jülich GmbH). ISC High Performance has renewed and further extended the workshop program this year welcoming 646 attendees to 21 workshops. This year Michela Taufer (University of Delaware, USA) served as workshop chair and led the workshop organization with workshop deputy chair John Shalf (Lawrence Berkeley National Laboratory). Julian Kunkel (German Climate Computing Center) served as the proceedings chair and managed the organization of proceedings for the workshops with Rio Yokota (Tokyo Institute of Technology) as the proceedings deputy chair.

The 21 workshops at ISC High Performance provided a focused, in-depth platform with presentations, discussion, and interaction on topics related to all aspects of research, development, and application of large-scale, high-performance experimental and commercial systems. Workshop topics included HPC computer architecture and hardware; programming models, system software, and applications; solutions for heterogeneity, reliability, power efficiency of systems; virtualization and containerized environments; big data and cloud computing; as well as international collaborations. Workshops were selected with a peer-review process by an international committee of 12 experts in the field from Europe, the United States, and Asia.

As in 2016, ISC High Performance provided a platform for workshops with their own call for papers and an individual peer-review process through an early deadline in December 2016. In all, 11 workshop proposals were submitted before this deadline from organizers all over the world; because of their high quality, all proposals were accepted by the committee (seven full-day and four half-day workshops) after a rigorous review process in which each proposal received at least three reviews. Additionally, each reviewer was given the possibility to discuss all the submissions.

Workshops without a call for papers were invited to submit their proposals in February 2017. For this second deadline, 15 workshop proposals were submitted and 10 workshops (one full-day and nine half-day workshops) were accepted by the committee with the same rigorous peer-review process as for workshops with proceedings.

The 21 workshops were held on Thursday, June 22, 2017, at the Frankfurt Marriott Hotel with 646 registered attendees, about 170 presentations, and over a dozen panel discussions. Workshop organizers were asked to collect the slides of all presentations at their workshops. PDF versions of the presentation slides were included in the ISC 2017 online proceedings, which were made available online to conference attendees a few days after the conference.

The workshop proceedings volume collects all accepted papers of the workshops received after the call for papers and a handful of invited papers. Each chapter of the

book contains the accepted and revised papers for one of the workshops. For some workshops, an additional preface describes the review process for the workshop and provides a summary of the outcome.

June 2017

Julian M. Kunkel
Michaela Taufer
Rio Yokota
John Shalf

Organization

ISC High Performance Workshops Chair

Michela Taufer University of Delaware, USA

ISC High Performance Workshops Deputy Chair

John Shalf Lawrence Berkeley National Laboratory, USA

ISC High Performance Workshops Committee

Rosa M. Badia	Barcelona Supercomputing Center, Spain
François Bodin	IRISA, France
Bronis R. de Supinski	Lawrence Livermore National Laboratory, USA
Jay Lofstead	Sandia National Laboratories, USA
Naoya Maruyama	Lawrence Livermore National Laboratory, USA
Simon McIntosh-Smith	University of Bristol, UK
Bernd Mohr	Jülich Supercomputing Centre, Germany
Marie-Christine Sawley	Intel, France
Seetharami Seelam	IBM T.J. Watson Research Center, USA
John Shalf	Lawrence Berkeley National Laboratory, USA
Michela Taufer	University of Delaware, USA
Carsten Trinitis	Technische Universität München, Germany
Antonino Tumeo	Pacific Northwest National Laboratory, USA
Didem Unat	Koç Universitesi, Turkey
Rio Yokota	Tokyo Institute of Technology, Japan

ISC High Performance Workshops Proceedings Chair

Julian Kunkel DKRZ, Germany

ISC High Performance Workshops Proceedings Deputy Chair

Rio Yokota Tokyo Institute of Technology, Japan

Experiences on Intel Knights Landing at the One-Year Mark

Organizing Committee

Estela Suarez Jülich Supercomputing Centre, Germany
Michael A. Lysaght ICHEC, Ireland

Simon J. Pennycook Intel, USA
Richard Gerber NERSC, USA

Program Committee

Damian Alvarez	Jülich Supercomputing Centre, Germany
Carlo Cavazzoni	CINECA, Italy
Gilles Civario	DELL, USA
Doug Doerfler	Lawrence Berkeley National Laboratory, USA
Richard Gerber	Lawrence Berkeley National Laboratory/NERSC, USA
Clayton Hughes	Sandia National Laboratories, USA
Balint Joo	Thomas Jefferson National Accelerator Facility, USA
Rakesh Krishnaiyer	Intel, USA
Michael A. Lysaght	ICHEC, Ireland
Simon McIntosh-Smith	University of Bristol, UK
Andrew Mallinson	Intel, UK
David E. Martin	Argonne National Laboratory, USA
Hideki Saito	Intel, USA
Thomas Steinke	Zuse Institute Berlin, Germany
Estela Suarez	Jülich Supercomputing Centre, Germany
Zhengji Zhao	Lawrence Berkeley National Laboratory, USA

HPC I/O in the Data Center

Organizing Committee

Julian Kunkel	DKRZ, Germany
Jay Lofstead	Sandia National Laboratory, USA
Colin McMurtrie	CSCS, Switzerland

Program Committee

Wolfgang Frings	Jülich Supercomputing Center, Germany
Javier Garcia Blas	University Carlos III of Madrid, Spain
Rob Ross	Argonne National Laboratory, USA
Carlos Maltzahn	University of California, Santa Cruz, USA
Thomas Bönisch	HLRS, Germany
Sai Narasimhamurthy	Seagate, UK
Jean-Thomas Acquaviva	DDN, France
Julian Kunkel	DKRZ, Germany
Jay Lofstead	Sandia National Laboratory, USA
Colin McMurtrie	CSCS, Switzerland

Workshop on Performance and Scalability of Storage Systems (WOPSSS)

Organizing Committee

Jean-Thomas Acquaviva	DDN, France
Jalil Boukhobza	Université de Bretagne Occidentale, France
Philippe Deniel	CEA/DIF, France
Massimo Lamanna	CERN, Switzerland
Pedro Javier García	University of Castilla-La Mancha, Spain
Allen D. Malony	University of Oregon, USA

Program Committee

Julien Bigot	CEA, France
Jason Chun Xue	City University of Hong Kong, Hong Kong, SAR China
Stefano Cozzini	CNR, Italy
Jesus Escudero-Sahuquillo	University of Castilla-La Mancha, Spain
Maria E. Gomez	Polytechnic University of Valencia, Spain
Pilar Gonzalez Ferez	Universidad de Murcia, Spain
Denis Gutfreund	ATOS, France
Julian Kunkel	DKRZ, Germany
Duo Liu	Chongqing University, China
Manolis Marazakis	Forth, Greece
Lars Nagel	Johannes Gutenberg-Universität Mainz, Germany
Ramon Nou	BSC, Spain
Juan Piernas Cánovas	Universidad de Murcia, Spain
Rekha Singhal	Tata Consultancy Services, India
Josef Weidendorfer	TUM, Germany
Soraya Zertal	University of Versailles, France

ExaComm: Third International Workshop on Communication Architectures for HPC, Big Data, Deep Learning and Clouds at Extreme Scale

Organizing Committee

Hari Subramoni	The Ohio State University, USA
Dhabaleswar	The Ohio State University, USA
K. (DK) Panda	

Program Committee

Taisuke Boku	University of Tsukuba, Japan
Ron Brightwell	Sandia National Laboratories, USA
Hans Eberle	NVIDIA, USA
Jesus Escudero-Sahuquillo	University of Castilla-La Mancha, Spain
Ada Gavrilovska	Georgia Institute of Technology, USA
Brice Goglin	Inria, France
Dror Goldenberg	Mellanox Technologies, Israel
R. Govindarajan	Indian Institute of Science, India
Ryan Grant	Sandia National Laboratories, USA
Hai Jin	Huazhong University of Science and Technology, China
Sven Karlsson	Technical University of Denmark, Denmark
Nectarios Koziris	National Technical University of Athens, Greece
Takeshi Nanri	University of Kyushu, Japan
Dimitrios Nikolopoulos	Queen's University of Belfast, UK
Antonio Pena	Barcelona Supercomputing Center, Spain
Sebastien Rumley	Columbia University, USA
Smruti Ranjan Sarangi	Indian Institute of Technology, India
Martin Schulz	Lawrence Livermore National Laboratory, USA
John M. Shalf	Lawrence Berkeley National Laboratory, USA
Tor Skeie	Simula Research Laboratory, Norway
Sayantan Sur	Intel, USA
Xin Yuan	Florida State University, USA
Jidong Zhai	Tsinghua University, China

12th Workshop on Virtualization in High-Performance Cloud Computing (VHPC'17)

Organizing Committee

Michael Alexander	scaledinfra technologies, Austria
Anastassios Nanos	OnApp, UK
Balazs Gerofi	RIKEN, Japan

Program Committee

Stergios Anastasiadis	University of Ioannina, Greece
Jakob Blomer	CERN, Europe
Ron Brightwell	Sandia National Laboratories, USA
Eduardo César	Universidad Autonoma de Barcelona, Spain
Julian Chesterfield	OnApp, UK
Stephen Crago	USC ISI, USA
Christoffer Dall	Columbia University, USA
Patrick Dreher	MIT, USA

Robert Futrick	Cycle Computing, USA
Maria Girone	CERN, Europe
Kyle Hale	Northwestern University, USA
Romeo Kinzler	IBM, Switzerland
Brian Kocoloski	University of Pittsburgh, USA
Nectarios Koziris	National Technical University of Athens, Greece
John Lange	University of Pittsburgh, USA
Che-Rung Lee	National Tsing Hua University, Taiwan
Giuseppe Lettieri	University of Pisa, Italy
Qing Liu	Oak Ridge National Laboratory, USA
Nikos Parlantzas	IRISA, France
Kevin Pedretti	Sandia National Laboratories, USA
Amer Qouneh	University of Florida, USA
Carlos Reaño	Technical University of Valencia, Spain
Thomas Ryd	CFEngine, Norway
Na Zhang	VMWare, USA
Borja Sotomayor	University of Chicago, USA
Craig Stewart	Indiana University, USA
Anata Tiwari	San Diego Supercomputer Center, USA
Kurt Tutschku	Blekinge Institute of Technology, Sweden
Yasuhiro Watashiba	Osaka University, Japan
Nicholas Wright	Lawrence Berkeley National Laboratory, USA
Chao-Tung Yang	Tunghai University, Taiwan

Visualization at Scale: Deployment Case Studies and Experience Reports

Organizing Committee

Glendon Holst	KAUST, Saudi Arabia
Thomas Theussl	KAUST, Saudi Arabia
Julien Jomier	Kitware, France
Joachim Pouderoux	Kitware, France

Program Committee

Second International Workshop on Performance Portable Programming Models for Accelerators (P^3 MA)

Organizing Committee

Sunita Chandrasekaran	University of Delaware, USA
Graham Lopez	ORNL, USA

Program Committee

Samuel Thibault	Inria, University of Bordeaux, France
James Beyer	NVIDIA, USA
Wei Ding	AMD, USA
Saber Feki	KAUST, Saudi Arabia
Robert Henschel	Indiana University, USA
Eric Stotzer	Texas Instruments, USA
Amit Amritkar	University of Houston, USA
Guido Juckeland	Helmholtz-Zentrum Dresden-Rossendorf, Germany
Will Sawyer	ETH, Zurich
Sameer Shende	University of Oregon, USA
Costas Bekas	IBM, Zurich
Toni Collis	University of Edinburgh, UK
Adrian Jackson	University of Edinburgh, UK
Henri Jin	NASA, USA
Andreas Knuepfer	TU Dresden, Germany
Steven Olivier	Sandia National Laboratory, USA
Suraj Prabhakaran	TU Darmstadt, Germany
Bora Ucar	ENS De Lyon, France
Veronica Vergara Larrea	ORNL, USA
Manisha Gajbe	Intel, USA
Daniel Tian	PGI, USA

Second International Workshop on OpenPOWER for HPC (IWOPH'17)

Organizing Committee

Dirk Pleiter	Jülich Supercomputing Centre, Germany
Jack Wells	Oak Ridge National Laboratory, USA

Program Committee

Nishant Agrawal	TCS, India
Carlo Cavazzoni	CINECA, Italy
Norbert Eicker	Jülich Supercomputing Centre, Germany
Holger Fröning	University of Heidelberg, Germany
Christoph Hagleitner	IBM Research, Switzerland
Oscar Hernandez	Oak Ridge National Laboratory, USA
Guido Juckeland	Helmholtz-Zentrum Dresden-Rossendorf, Germany
M. Graham Lopez	Oak Ridge National Laboratory, USA
Lena Oden	Jülich Supercomputing Centre, Germany
Dirk Pleiter	Jülich Supercomputing Centre, Germany
Swaroop Pophale	Oak Ridge National Laboratory, USA

Tiago Quintino	ECMWF, UK
Sebastiano F. Schifano	University and INFN Ferrara, Italy
Sameer Shende	University of Oregon, USA
Tjerk Straatsma	Oak Ridge National Laboratory, USA
Xiaonan Tian	NVIDIA, USA
Piero Vicini	INFN, University of Rome Sapienza, Italy
Jack Wells	Oak Ridge National Laboratory, USA
Michael Wolfe	PGI, USA
Bronis de Supinski	Lawrence Livermore National Laboratory, USA

First International Workshop on Data Reduction for Big Scientific Data (DRBSD-1)

Organizing Committee

Ian Foster	National Laboratory/University of Chicago, USA
Scott Klasky	Oak Ridge National Laboratory, USA
Gary Liu	New Jersey Institute of Technology, USA
Mark Ainsworth	Brown University/Oak Ridge National Laboratory, USA

Program Committee

Frank Cappello	Argonne National Laboratory, USA
Peter Lindstrom	Lawrence Livermore National Laboratory, USA
Tamara Kolda	Sandia National Laboratory, USA
Todd Munson	Argonne National Laboratory, USA
George Ostroumov	Oak Ridge National Laboratory, USA
Scott Klasky	Oak Ridge National Laboratory, USA
Mark Ainsworth	Brown University/Oak Ridge National Laboratory, USA
John Wu	Lawrence Berkeley National Laboratory, USA
Todd Munson	Argonne National Laboratory, USA
Eric Suchyta	Oak Ridge National Laboratory, USA
Martin Burtscher	Texas State University, USA

Contents

The 1st International Workshop on Data Reduction for Big Scientific Data (DRBSD-1)

Toward Decoupling the Selection of Compression Algorithms from Quality Constraints	3
<i>Julian Kunkel, Anastasiia Novikova, Eugen Betke, and Armin Schaare</i>	
On the Scalability of Data Reduction Techniques in Current and Upcoming HPC Systems from an Application Perspective	15
<i>Axel Huebl, René Widera, Felix Schmitt, Alexander Matthes, Norbert Podhorszki, Jong Youl Choi, Scott Klasky, and Michael Bussmann</i>	
Toward a Multi-method Approach: Lossy Data Compression for Climate Simulation Data	30
<i>Allison H. Baker, Haiying Xu, Dorit M. Hammerling, Shaomeng Li, and John P. Clyne</i>	
Exploration of Pattern-Matching Techniques for Lossy Compression on Cosmology Simulation Data Sets	43
<i>Dingwen Tao, Sheng Di, Zizhong Chen, and Franck Cappello</i>	
Third International Workshop on Communication Architectures for HPC, Big Data, Deep Learning and Clouds at Extreme Scale (ExaComm)	
Design Space Exploration of the Dragonfly Topology	57
<i>Min Yee Teh, Jeremiah J. Wilke, Keren Bergman, and Sébastien Rumley</i>	
High-Throughput Sockets over RDMA for the Intel Xeon Phi Coprocessor	75
<i>Aram Santogidis and Spyros Lalis</i>	
Workshop on HPC Computing in a Post Moore's Law World (HCPM)	
Reconfigurable Silicon Photonic Interconnect for Many-Core Architecture . . .	89
<i>Hang Guan, Sébastien Rumley, Ke Wen, David Donofrio, John Shalf, and Keren Bergman</i>	
Instruction Set Architectures for Quantum Processing Units	98
<i>Keith A. Britt and Travis S. Humble</i>	

Eliminating Dark Bandwidth: A Data-Centric View of Scalable, Efficient Performance, Post-Moore	106
<i>Jonathan C. Beard and Joshua Randall</i>	
Towards an Integrated Strategy to Preserve Digital Computing Performance Scaling Using Emerging Technologies	115
<i>Dilip Vasudevan, Anastasiia Butko, George Michelogiannakis, David Donofrio, and John Shalf</i>	
HPC I/O in the Data Center (HPC-IODC)	
HPC I/O in the Data Center Workshop (HPC-IODC)	127
<i>Julian Kunkel, Jay Lofstead, and Colin McMurtrie</i>	
Simulation of Hierarchical Storage Systems for TCO and QoS	132
<i>Jakob Luettgau and Julian Kunkel</i>	
GPU Erasure Coding for Campaign Storage	145
<i>Walker Haddock, Matthew L. Curry, Purushotham V. Bangalore, and Anthony Skjellum</i>	
PIOM-PX: A Framework for Modeling the I/O Behavior of Parallel Scientific Applications	160
<i>Pilar Gomez-Sanchez, Sandra Mendez, Dolores Rexachs, and Emilio Luque</i>	
Real-Time I/O-Monitoring of HPC Applications with SIOX, Elasticsearch, Grafana and FUSE	174
<i>Eugen Betke and Julian Kunkel</i>	
Output Performance Study on a Production Petascale Filesystem	187
<i>Bing Xie, Jeffrey S. Chase, David Dillow, Scott Klasky, Jay Lofstead, Sarp Oral, and Norbert Podhorszki</i>	
Second International Workshop on OpenPOWER for HPC (IWOPH'17)	
GPU-Accelerated Particle-in-Cell Code on Minsky	205
<i>Andreas Herten, Dirk Brömmel, and Dirk Pleiter</i>	
Pushing Big Data into Accelerators: Can the JVM Saturate Our Hardware?	220
<i>Johan Peltenburg, Ahmad Hesam, and Zaid Al-Ars</i>	
The Technological Roadmap of Parallware and Its Alignment with the OpenPOWER Ecosystem	237
<i>Manuel Arenaz, Oscar Hernandez, and Dirk Pleiter</i>	

Experiences Evaluating Functionality and Performance of IBM POWER8+ Systems	254
<i>Verónica G. Vergara Larrea, Wayne Joubert, Mark Berrill, Swen Boehm, Arnold Tharrington, Wael R. Elwasif, and Don Maxwell</i>	
Power/Performance Controlling Techniques in OpenPOWER	275
<i>Todd Rosedahl, Martha Broyles, Charles Lefurgy, Bjorn Christensen, and Wu Feng</i>	
Performance Evaluation of Container-Based High Performance Computing Ecosystem Using OpenPOWER	290
<i>Animesh Kuity and Sateesh Kumar Peddoju</i>	
Pre-exascale Architectures: OpenPOWER Performance and Usability Assessment for French Scientific Community	309
<i>Gabriel Hautreux, Alfredo Buttari, Arnaud Beck, Victor Cameo, Dimitri Lecas, Dominique Aubert, Emeric Brun, Eric Boyer, Fausto Malvagi, Gabriel Staffelbach, Isabelle d'Ast, Joeffrey Legaux, Ghislain Lartigue, Gilles Grasseau, Guillaume Latu, Juan Escobar, Julien Bigot, Julien Derouillat, Matthieu Haefele, Nicolas Renon, Philippe Parnaudeau, Philippe Wautelaet, Pierre-Francois Lavallee, Pierre Kestener, Remi Lacroix, Stephane Requena, Anthony Scemama, Vincent Moureau, Jean-Matthieu Etancelin, and Yann Meurdesoif</i>	
Experiences on Intel Knights Landing at the One-Year Mark (IXPUG)	
IXPUG: Experiences on Intel Knights Landing at the One Year Mark	327
<i>Estela Suarez, Michael Lysaght, Simon J. Pennycook, and Richard A. Gerber</i>	
Analyzing Performance of Selected NESAP Applications on the Cori HPC System	334
<i>Thorsten Kurth, William Arndt, Taylor Barnes, Brandon Cook, Jack Deslippe, Doug Doerfler, Brian Friesen, Yun (Helen) He, Tuomas Koskela, Mathieu Lobet, Tareq Malas, Leonid Oliker, Andrey Ovsyannikov, Samuel Williams, Woo-Sun Yang, and Zhengji Zhao</i>	
On the Mitigation of Cache Hostile Memory Access Patterns on Many-Core CPU Architectures	348
<i>Tom Deakin, Wayne Gaudin, and Simon McIntosh-Smith</i>	
From Knights Corner to Landing: A Case Study Based on a Hodgkin-Huxley Neuron Simulator	363
<i>George Chatzikonstantis, Diego Jiménez, Esteban Meneses, Christos Strydis, Harry Sidiropoulos, and Dimitrios Soudris</i>	

Porting Tissue-Scale Cardiac Simulations to the Knights Landing Platform	376
<i>Johannes Langguth, Chad Jarvis, and Xing Cai</i>	
KART – A Runtime Compilation Library for Improving HPC Application Performance	389
<i>Matthias Noack, Florian Wende, Georg Zitzlsberger, Michael Klemm, and Thomas Steinke</i>	
Performance Evaluation of NWChem Ab-Initio Molecular Dynamics (AIMD) Simulations on the Intel® Xeon Phi™ Processor.	404
<i>Eric J. Bylaska, Mathias Jacquelin, Wibe A. de Jong, Jeff R. Hammond, and Michael Klemm</i>	
Performance Variability on Xeon Phi.	419
<i>Brandon Cook, Thorsten Kurth, Brian Austin, Samuel Williams, and Jack Deslippe</i>	
Optimizing Fusion PIC Code Performance at Scale on Cori Phase Two	430
<i>Tuomas Koskela and Jack Deslippe</i>	
amask: A Tool for Evaluating Affinity Masks in Many-Core Processors	441
<i>Kent Milfeld</i>	
Second International Workshop on Performance Portable Programming Models for Accelerators (P^3MA)	
Analyzing Offloading Inefficiencies in Scalable Heterogeneous Applications	457
<i>Robert Dietrich, Ronny Tschüter, Guido Juckeland, and Andreas Knüpfer</i>	
Performance Portability Analysis for Real-Time Simulations of Smoke Propagation Using OpenACC	477
<i>Anne Küsters, Sandra Wienke, and Lukas Arnold</i>	
Tuning and Optimization for a Variety of Many-Core Architectures Without Changing a Single Line of Implementation Code Using the Alpaka Library	496
<i>Alexander Matthes, René Widera, Erik Zenker, Benjamin Worpitz, Axel Huebl, and Michael Bussmann</i>	
An Embedded Domain Specific Language for General Purpose Vectorization	515
<i>Przemysław Karpiński and John McDonald</i>	

Exploiting Auto-tuning to Analyze and Improve Performance Portability on Many-Core Architectures	538
<i>James Price and Simon McIntosh-Smith</i>	
OpenACC 2.5 Validation Testsuite Targeting Multiple Architectures	557
<i>Kyle Friedline, Sunita Chandrasekaran, M. Graham Lopez, and Oscar Hernandez</i>	
12th Workshop on Virtualization in High-Performance Cloud Computing (VHPC'17)	
A Survey of Fast Packet I/O Technologies for Network Function Virtualization	579
<i>Giuseppe Lettieri, Vincenzo Maffione, and Luigi Rizzo</i>	
Machine Learning Using Virtualized GPUs in Cloud Environments.	591
<i>Uday Kurkure, Hari Sivaraman, and Lan Vu</i>	
A Locality-Aware Communication Layer for Virtualized Clusters	605
<i>Simon Pickartz, Jonas Baude, Stefan Lankes, and Antonello Monti</i>	
YASMIN: Efficient Intra-node Communication Using Generic Sockets	617
<i>Michalis Rozis, Stefanos Gerangelos, and Nectarios Koziris</i>	
Dynamic Paging Method Switching - An Implementation for KVM	629
<i>Yu Zhang, Peter Tröger, and Matthias Werner</i>	
Aggregating and Managing Memory Across Computing Nodes in Cloud Environments	642
<i>Luis A. Garrido and Paul Carpenter</i>	
Visualization at Scale: Deployment Case Studies and Experience Reports	
In-situ Visualization for Computation Workflows	655
<i>Alejandro Ribes, Ovidiu Mircescu, Anthony Geay, and Yvan Fournier</i>	
From Big Data to Big Displays High-Performance Visualization at Blue Brain	662
<i>Stefan Eilemann, Marwan Abdellah, Nicolas Antille, Ahmet Bilgili, Grigory Chevtchenko, Raphael Dumusc, Cyrille Favreau, Juan Hernando, Daniel Nachbaur, Pawel Podhajski, Jafet Villafranca, and Felix Schürmann</i>	

Workshop on Performance and Scalability of Storage Systems (WOPSSS)

An MPI-IO In-Memory Driver for Non-volatile Pooled Memory of the Kove XPD	679
<i>Julian Kunkel and Eugen Betke</i>	
HetFS: A Heterogeneous File System for Everyone.	691
<i>Georgios Koloventzos, Ramon Nou, Alberto Miranda, and Toni Cortes</i>	
Scientific Applications Performance Evaluation on Burst Buffer	701
<i>George S. Markomanolis, Bilel Hadri, Rooh Khurram, and Saber Feki</i>	
JULEA: A Flexible Storage Framework for HPC	712
<i>Michael Kuhn</i>	
Delivering LHC Software to HPC Compute Elements with CernVM-FS	724
<i>Jakob Blomer, Gerardo Ganis, Nikola Hardi, and Radu Popescu</i>	
Scaling the EOS Namespace.	731
<i>Andreas J. Peters, Elvin A. Sindrilaru, and Georgios Bitzes</i>	
Erratum to: Performance Portability Analysis for Real-Time Simulations of Smoke Propagation Using OpenACC.	E1
<i>Anne Küsters, Sandra Wienke, and Lukas Arnold</i>	
Author Index	741