Lecture Notes in Computer Science 13289

More information about this series at https://link.springer.com/bookseries/558

Ana-Lucia Varbanescu · Abhinav Bhatele ·
Piotr Luszczek · Baboulin Marc (Eds.)

High Performance Computing

37th International Conference, ISC High Performance 2022
Hamburg, Germany, May 29 – June 2, 2022
Proceedings

Springer

Editors
Ana-Lucia Varbanescu
University of Twente
Enschede, The Netherlands

Abhinav Bhatele
University of Maryland
College Park, MD, USA

Piotr Luszczek
University of Tennessee
Knoxville, TN, USA

Baboulin Marc
Université Paris-Saclay
Orsay, France

ISSN 0302-9743 ISSN 1611-3349 (electronic)
Lecture Notes in Computer Science
ISBN 978-3-031-07311-3 ISBN 978-3-031-07312-0 (eBook)
https://doi.org/10.1007/978-3-031-07312-0

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

Preface

ISC High Performance Computing—formerly known as the International Supercomputing Conference—was founded in 1986 as the Supercomputer Seminar. Originally organized by Hans Meuer, Professor of Computer Science at the University of Mannheim, and former director of its computer center, the 1986 edition of the seminar brought together a group of 81 scientists and industrial partners who shared an interest in high-performance computing (HPC). Since then, the annual conference has become a major international event within the HPC community, growing beyond its humble beginnings, and moving out of Mannheim into other cities throughout the years: Frankfurt, Heidelberg, Dresden, Hamburg, Leipzig, and this year back to Hamburg. Prior to the coronavirus pandemic, the conference had seen a steady increase in the number of submissions of high-quality research papers and corresponding growth in the number of conference attendees. Benefiting from the improving health indicators, ISC-HPC 2022 was again held in person in Hamburg.

The call for papers to ISC-HPC 2022 was issued in Fall 2021, inviting the researchers, developers, and practitioners alike to submit their latest results for consideration for one of the five conference tracks: (1) Architecture, Networks, and Storage; (2) HPC Algorithms and Applications; (3) Machine Learning, AI, and Emerging Technologies; (4) Performance Modeling, Evaluation, and Analysis; and (5) Programming Environments and Systems Software. In all, 53 full submissions were received from authors all over the world. The Research Papers Program Committee consisted of 71 members from 18 countries. After initial reviews were completed, a rebuttal process offered authors an opportunity to respond to reviewers' questions and help clarify issues the reviewers might have had. A virtual Program Committee meeting was held to discuss all the papers and to finalize consensus on the papers. Finally, the committee selected 18 papers for publication.

For the past several years, the ISC-HPC conference has presented an ISC-sponsored award to encourage outstanding research in HPC and to honor the overall best research paper submitted to the conference. Four years ago, this annual award was renamed in memory of the late Dr. Hans Meuer, who was general chair of the ISC-HPC conference from 1986 through 2014, and a co-founder of the TOP500 project. This year, from all research papers submitted, the Best Paper Committee selected the best paper based on its technical merit, its novelty, and impact on the HPC community. During a live ceremony, the following paper was awarded the Hans Meuer Award: *Remote OpenMP Offloading* by Atmn Patel from University of Waterloo, Canada, and Johannes Doerfert from Argonne National Laboratory, USA. The paper extended the canonical scope of OpenMP, which is traditionally confined to a shared memory domain, by utilizing the standard's modern features to offload the computational workload to GPU accelerators housed in remote cluster nodes. The Best Paper Committee appreciated the paper's unique combination of OpenMP familiarity and ease of use with accelerated and distributed computing that were accompanied by analysis of scaling capabilities. The winning paper also showed the

solution's versatility and how it could use multiple transport layers, each of which offers a different set of trade-offs between performance, portability, and scalability potential.

As the chairs of the Research Papers Committee, we would like to express our gratitude to our colleagues for submitting high-quality papers to all five ISC-HPC scientific tracks. Also, we wish to extend our thanks to the track, area, and conflict chairs, as well as the members of the Best Paper Committee, and finally to the Research Papers Committee that provided the reviews and manuscript evaluation throughout the submission stages. We hope to express our thanks in person during this year's meeting and upcoming ISC-HPC 2023.

May 2022

Ana Lucia Varbanescu
Abhinav Bhatele

Organization

Program Chair

Keren Bergman	Columbia University, USA

Program Deputy Chair

John Shalf	Lawrence Berkeley National Laboratory, USA

Research Papers Program Committee

Research Papers Chairs

Ana Lucia Varbanescu (Chair)	University of Amsterdam, Netherlands
Abhinav Bhatele (Deputy Chair)	University of Maryland, USA

Architecture, Networks, and Storage

Jay Lofstead (Chair)	Sandia National Laboratories, USA
Edson Borin	University of Campinas, Brazil
Elsa Gonsiorowski	Lawrence Livermore National Laboratory, USA
Mozhgan Kabiri Chimeh	NVIDIA, UK
Nectarios Koziris	National Technical University of Athens, Greece
Michael Kuhn	Otto von Guericke University Magdeburg, Germany
Jay Lofstead	Sandia National Laboratories, USA
Preeti Malakar	Indian Institute of Technology Kanpur, India
Dhabaleswar Panda	Ohio State University, USA
Guangming Tan	Institute of Computing Technology (ICT), China
Osamu Tatebe	University of Tsukuba, Japan
Carsten Trinitis	Technical University of Munich, Germany
Venkatram Vishwanath	Argonne National Laboratory, USA

HPC Algorithms and Applications

Didem Unat (Chair)	Koç University, Turkey
Sameh Abdulah	KAUST, Saudi Arabia
Mehmet Belviranli	Colorado School of Mines, USA
Xing Cai	Simula Research Laboratory and University of Oslo, Norway

Lin Gan	Tsinghua University and National Supercomputing Center in Wuxi, China
Clemens Grelck	University of Amsterdam, Netherlands
Fuerlinger Karl	Ludwig Maximilian University Munich (LMU), Germany
Kamer Kaya	Sabancı University, Turkey
Simon McIntosh-Smith	University of Bristol, UK
Gabriel Noaje	NVIDIA, Singapore
Lena Oden	Fernuniversität in Hagen and Forschungszentrum Jülich GMBH, Germany
Johann Rudi	Argonne National Laboratory, USA
Tuğba Torun	Koç University, Turkey
Miwako Tsuji	RIKEN, Japan

Machine Learning, AI, and Emerging Technologies

Theodore L. Willke (Chair)	Intel Corporation, USA
Nikoli Dryden	ETH Zurich, Switzerland
Gurbinder Gill	Katana Graph Inc., USA
Jiajia Li	William and Mary College, USA
Maryam Mehri Dehnavi	University of Toronto, Canada
Bogdan Nicolae	Argonne National Laboratory, USA
Mostofa Patwary	NVIDIA, USA
Shaden Smith	Microsoft, USA
Edgar Solomonik	University of Illinois at Urbana-Champaign, USA
Sofia Vallecorsa	CERN, Switzerland
Abhinav Vishnu	AMD, USA
Yang You	National University of Singapore, Singapore

Performance Modeling, Evaluation, and Analysis

Nathan Tallent (Chair)	Pacific Northwest National Laboratory, USA
Ivy B. Peng	Lawrence Livermore National Laboratory, USA
Alexandru Calotoiu	ETH Zürich, Germany
Marc Casas	Barcelona Supercomputing Center, Spain
Tom Deakin	University of Bristol, UK
Seyong Lee	ORNL, USA
Simon McIntosh-Smith	University of Bristol, UK
Xiaozhu Meng	Rice University, USA
Bernd Mohr	Juelich Supercomputing Centre, Germany
Scott Pakin	Los Alamos National Laboratory, USA
Xian-He Sun	Illinois Institute of Technology, USA
Jidong Zhai	Tsinghua University, China
Tianwei Zhang	Nanyang Technological University, Singapore

Programming Environments and Systems Software

Michele Weiland (Chair)	EPCC, University of Edinburgh, UK
Bilel Hadri	KAUST Supercomputing Laboratory, Saudi Arabia
Guido Juckeland	HZDR, Germany
Michael Klemm	AMD and OpenMP ARB, Germany
Pouya Kousha	Ohio State University, USA
John L. inford	Arm, USA
István Z. Reguly	Pázmány Péter Catholic University, Hungary
Harvey Richardson	Hewlett Packard Enterprise, UK
Martin Ruefenacht	Leibniz Supercomputing Centre, Germany
Roxana Rusitoru	Arm, UK
Thomas R. W. Scogland	Lawrence Livermore National Laboratory, USA
Simon Smart	ECMWF, UK
Hiroyuki Takizawa	Tohoku University, Japan
Christian Terboven	RWTH Aachen University, Germany
Justs Zarins	EPCC, University of Edinburgh, UK

Birds of a Feather Committee

Roman Wyrzykowski (Chair)	Czestochowa University of Technology, Poland
Iosif Meyerov (Deputy Chair)	Lobachevsky State University of Nizhni Novogorod, Russia
Michael Bader	Technical University of Munich, Germany
Claudia Blaas-Schenner	TU Wien, VSC Research Center, Austria
Dominik Göddeke	University of Stuttgart, Germany
Aleksandar Ilic	INESC-ID and Universidade de Lisboa, Portugal
Jacek Kitowski	AGH University of Science and Technology, Poland
Dieter Kranzlmueller	Ludwig Maximilian University Munich (LMU) and Leibniz Rechenzentrum, Germany
Carola Kruse	Centre Européen de Recherche et de Formation Avancée en Calcul Scientifique (CERFACS), France
Krzysztof Kurowski	Poznań Supercomputing and Networking Center, Poland
Marco Lapegna	University of Naples Federico II, Italy
Simon McIntosh-Smith	University of Bristol, UK
Iosif Meyerov	Lobachevsky State University of Nizhni Novogorod, Russia
Koji Nakano	Hiroshima University, Japan
Gabriel Oksa	Slovak Academy of Sciences, Slovakia
Dana Petcu	West University of Timisoara, Romania

Antonio J. Peña	Barcelona Supercomputing Center, Spain
Thomas Rauber	University of Bayreuth, Germany
Lubomir Riha	IT4Innovations National Supercomputing Center and Technical University of Ostrava, Czech Republic
Masha Sosonkina	Old Dominion University, USA
Vladimir Stegailov	Higher School of Economics and JIHT RAS, Russia
Dave Turner	Kansas State University, USA
Bora Ucar	CNRS and ENS-Lyon, France

Project Posters Committee

Christian Perez (Chair)	Inria, France
Are Magnus Bruaset (Deputy Chair)	Simula Research Laboratory, Norway
Marco Aldinucci	University of Torino, Italy
Bartosz Bosak	Poznań Supercomputing and Networking Center, Poland
Nick Brown	EPCC, University of Edinburgh, UK
Theodoros Christoudias	The Cyprus Institute, Cyprus
Andrew Ensor	Auckland University of Technology, New Zealand
Ana Gainaru	Oak Ridge National Laboratory, USA
Andra Hugo	Apple, France
Kamer Kaya	Sabancı University, Turkey
Francesc Lordan Gomis	Barcelona Supercomputing Center, Spain
Maciej Malawski	Sano Centre for Computational Medicine, Institute of Computer Science AGH, Poland
Kengo Nakajima	University of Tokyo and RIKEN, Japan
Bogdan Nicolae	Argonne National Laboratory, USA
Eric Petit	Intel, France
Phil Ridley	Arm, UK
Jonathan Rouzaud-Cornabas	Inria and INSA de Lyon, France
Kentaro Sano	RIKEN, Japan
Francieli Zanon Boito	Inria, France
Ameli Chi Zhou	Shenzhen University, China

Research Posters Committee

Aparna Chandramowlishwaran (Chair)	UCI, USA
Hartwig Anzt	Karlsruhe Institute of Technology, Germany and University of Tennessee, USA

Maryam Mehri Dehnavi	University of Toronto, Canada
Jee Choi	University of Oregon, USA
Ana Gainaru	Oak Ridge National Laboratory, USA
Lin Gan	Tsinghua University and National Supercomputing Center in Wuxi, China
Amal Khabou	Université Paris-Saclay, France
Mariam Kiran	Lawrence Berkeley National Laboratory, USA
Penporn Koanantakool	Google LLC, USA
Ronald Kriemann	MPI for Math. i.t.S., Germany
Ang Li	Pacific Northwest National Laboratory, USA
Jiajia Li	William and Mary College, USA
Piyush Sao	Oak Ridge National Laboratory, USA
Christian Terboven	RWTH Aachen University, Germany
Bo Wu	Colorado School of Mines, USA
Rio Yokota	Tokyo Institute of Technology, Japan
Rohit Zambre	AMD Research, USA

Tutorials Committee

Kathryn Mohror (Chair)	Lawrence Livermore National Laboratory, USA
Suren Byna (Deputy Chair)	Lawrence Berkeley National Laboratory, USA
Ritu Arora	University of Texas at San Antonio, USA
Rosa M. Badia	Barcelona Supercomputing Center, Spain
Wahid Bhimji	Lawrence Berkeley National Laboratory, USA
Philip Carns	Argonne National Laboratory, USA
James Dinan	NVIDIA, USA
Ann Gentile	Sandia National Laboratories, USA
Tanzima Islam	Texas State University, USA
Simon McIntosh-Smith	University of Bristol, UK
Diana Moise	Cray, and HPE, Switzerland
Sarah Neuwirth	Goethe-University Frankfurt, Germany
C. J. Newburn	NVIDIA, USA
Dhabaleswar Panda	Ohio State University, USA
Raghunath Raja Chandrasekar	Frau, USA
Michela Taufer	The University of Tennessee, USA
Michele Weiland	EPCC – University of Edinburgh, UK

Workshops Committee

Hartwig Anzt (Chair)	Karlsruhe Institute of Technology, Germany and University of Tennessee, USA
Amanda Bienz (Deputy Chair)	University of New Mexico, USA
Cody Balos	Lawrence Livermore National Laboratory, USA

Harun Bayraktar	NVIDIA, USA
Natalie Beams	University of Tennessee, USA
Luc Berger-Vergiat	Sandia National Laboratories, USA
George Bosilca	University of Tennessee, USA
Lisa Claus	LBNL, USA
Terry Cojean	Karlsruhe Institute of Technology, Germany
Anthony Danalis	University of Tennessee Knoxville, USA
Edoardo Di Napoli	Juelich Supercomputing Centre, Germany
Markus Goetz	Karlsruhe Institute of Technology, Germany
Aditya Kashi	Karlsruhe Institute of Technology, Germany
Sarah Knepper	Intel, USA
Andreas Knuepfer	Technische Universität Dresden, Germany
Martin Kronbichler	Technical University of Munich, Germany
Weifeng Liu	China University of Petroleum, China
Simone Pezzuto	Università della Svizzera italiana, Switzerland
Enrique S. Quintana-Orti	Universitat Politècnica de València, Spain
Estela Suarez	Jülich Supercomputing Centre, Germany
Nico Trost	AMD, Germany
Markus Wittmann	Friedrich-Alexander-Universität Erlangen-Nürnberg, Germany

HPC in Asia Committee

Kento Sato (Chair)	RIKEN, Japan
James Lin (Deputy Chair)	Shanghai Jiao Tong University, China

Inclusivity Committee

Laura Schulz (Chair)	Frau, Germany

Publicity Committee

Carsten Trinitis (Chair)	Technical University of Munich, Germany

Proceedings Chairs

Piotr Luszczek (Chair)	University of Tennessee, USA
Marc Baboulin (Deputy Chair)	Université Paris-Saclay, France

Contents

HPC Algorithms and Applications

Performance Modeling, Evaluation, and Analysis

Programming Environments and System Software