2022 IEEE/ACM International Workshop on Performance, Portability and Productivity in HPC (P3HPC 2022)

Dallas, Texas, USA
13 – 18 November 2022

IEEE Catalog Number: CFP22S71-POD
ISBN: 978-1-6654-6022-4

Copyright © 2022 by the Institute of Electrical and Electronics Engineers, Inc.
All Rights Reserved

Copyright and Reprint Permissions: Abstracting is permitted with credit to the source. Libraries are permitted to photocopy beyond the limit of U.S. copyright law for private use of patrons those articles in this volume that carry a code at the bottom of the first page, provided the per-copy fee indicated in the code is paid through Copyright Clearance Center, 222 Rosewood Drive, Danvers, MA 01923.

For other copying, reprint or republication permission, write to IEEE Copyrights Manager, IEEE Service Center, 445 Hoes Lane, Piscataway, NJ 08854. All rights reserved.

****** This is a print representation of what appears in the IEEE Digital Library. Some format issues inherent in the e-media version may also appear in this print version.***

IEEE Catalog Number: CFP22S71-POD
ISBN (Print-On-Demand): 978-1-6654-6022-4
ISBN (Online): 978-1-6654-6021-7

Additional Copies of This Publication Are Available From:

Curran Associates, Inc
57 Morehouse Lane
Red Hook, NY 12571 USA
Phone: (845) 758-0400
Fax: (845) 758-2633
E-mail: curran@proceedings.com
Web: www.proceedings.com

2022 IEEE/ACM International Workshop on Performance, Portability and Productivity in HPC (P3HPC)

P3HPC 2022

Table of Contents

Message from the Workshop Chairs ..... **vi**
Workshop Organization ..... **vii**

Session 1

Heterogeneous Programming for the Homogeneous Majority 1
Tom Deakin (University of Bristol, UK), James Cownie (University of Bristol, UK), Wei-Chen Lin (University of Bristol, UK), and Simon McIntosh-Smith (University of Bristol, UK)

Leveraging Compiler-Based Translation to Evaluate a Diversity of Exascale Platforms 14
Jacob Lambert (Advanced Micro Devices, USA), Mohammad Alaul Haque Monil (Oak Ridge National Laboratory, USA), Seyong Lee (Oak Ridge National Laboratory, USA), Allen D. Malony (University of Oregon, USA), and Jeffrey S. Vetter (Oak Ridge National Laboratory, USA)

Session 2

Understanding Strong Scaling on GPUs Using Empirical Performance Saturation Size 26
David Eberius (Oak Ridge National Laboratory, USA), Philip Roth (Oak Ridge National Laboratory, USA), and David Rogers (Oak Ridge National Laboratory, USA)

Portable and Efficient Dense Linear Algebra in the Beginning of the Exascale Era 36
Mark Gates (University of Tennessee, USA), Asim YarKhan (University of Tennessee, USA), Dalal Sukkari (University of Tennessee, USA), Kadir Akbudak (University of Tennessee, USA), Sebastien Cayrols (University of Tennessee, USA), Daniel Bielich (University of Tennessee, USA), Ahmad Abdelfattah (University of Tennessee, USA), Mohammed Al Farhan (KAUST, Saudi Arabia), and Jack Dongarra (University of Tennessee, USA)

Exploiting Dynamic Sparse Matrices for Performance Portable Linear Algebra Operations 47
Christodoulos Stylianou (EPCC, The University of Edinburgh, United Kingdom) and Michele Weiland (EPCC, The University of Edinburgh, United Kingdom)

Performance Portability of Sparse Block Diagonal Matrix Multiple Vector Multiplications on GPUs 58
Khaled Ibrahim (Lawrence Berkeley National Laboratory), Chao Yang (Lawrence Berkeley National Laboratory), and Pieter Maris (Iowa State University)

Session 3

Performance Portable Vlasov Code with C++ Parallel Algorithm 68
Yuuichi Asahi (Japan Atomic Energy Agency, Japan), Thomas Padioleau (Maison de la Simulation, France), Guillaume Latu (DES/IRESNE/DEC, CEA, France), Julien Bigot (Maison de la Simulation, France), Virginie Grandgirard (IRFM, CEA, France), and Kevin Obrejan (IRFM, CEA, France)

Towards Cross-Platform Portability of Coupled-Cluster Methods with Perturbative Triples using SYCL 81
Abhishek Bagusetty (Argonne National Laboratory, USA), Ajay Panyala (Pacific Northwest National Laboratory, USA), Gordon Brown (Codeplay Software Ltd, UK), and Jack Kirk (Codeplay Software Ltd, UK)

From Task-Based GPU Work Aggregation to Stellar Mergers: Turning Fine-Grained CPU Tasks into Portable GPU Kernels 89
Gregor Daiß (University of Stuttgart, Germany), Patrick Diehl (Louisiana State University, United States of America), Dominic Marcello (Louisiana State University, United States of America), Alireza Kheirkhahan (Louisiana State University, United States of America), Hartmut Kaiser (Louisiana State University, United States of America), and Dirk Pflüger (University of Stuttgart, Germany)

Piper: Pipelining OpenMP Offloading Execution through Compiler Optimization for Performance 100
Konstantinos Parasyris (Lawrence Livermore National Laboratory), Giorgis Georgakoudis (Lawrence Livermore National Laboratory), Johannes Doerfert (Argonne National Laboratory), Ignacio Laguna (Lawrence Livermore National Laboratory), and Thomas R.W. Scogland (Lawrence Livermore National Laboratory)

Session 4

Towards Performance Portability of AI Graphs Using SYCL 111
Kumudha Narasimhan (Codeplay Software Ltd, UK), Ouadie El Farouki (Codeplay Software Ltd, UK), Mehdi Goli (Codeplay Software Ltd, UK), Muhammad Tanvir (Codeplay Software Ltd, UK), Svetlozar Georgiev (Codeplay Software Ltd, UK), and Isaac Ault (Codeplay Software Ltd, UK)

ECP SOLLVE: Validation and Verification Testsuite Status Update and Compiler Insight for OpenMP 123
Thomas Huber (University of Delaware), Swaroop Pophale (Oak Ridge National Laboratory), Nolan Baker (University of Delaware), Michael Carr (University of Delaware), Nikhil Rao (University of Delaware), Jaydon Reap (University of Delaware), Kristina Holsapple (University of Delaware), Joshua Hoke Davis (University of Maryland), Tobias Burnus (Siemens), Seyong Lee (Oak Ridge National Laboratory), David E. Bernholdt (Oak Ridge National Laboratory), and Sunita Chandrasekaran (University of Delaware)

Author Index **137**