

# **2022 IEEE/ACM International Workshop on Performance Modeling, Benchmarking and Simulation of High Performance Computer Systems (PMBS 2022)**

**Dallas, Texas, USA  
13 – 18 November 2022**



**IEEE Catalog Number: CFP22J43-POD  
ISBN: 978-1-6654-5186-4**

**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:	CFP22J43-POD
ISBN (Print-On-Demand):	978-1-6654-5186-4
ISBN (Online):	978-1-6654-5185-7

**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](mailto:curran@proceedings.com)  
Web: [www.proceedings.com](http://www.proceedings.com)

CURRAN ASSOCIATES INC.  
**proceedings**  
.com

# 2022 IEEE/ACM International Workshop on Performance Modeling, Benchmarking and Simulation of High Performance Computer Systems (PMBS) **PMBS 2022**

## Table of Contents

Message from the Workshop Chairs .....	vi
Workshop Organization .....	viii

### Machine Learning and AI

ML-based Performance Portability for Time-Dependent Density Functional Theory in HPC environments .....	1
<i>Adrian Perez Dieguez (Lawrence Berkeley National Laboratory, USA), Min Choi (University of California Riverside, USA), Xinran Zhu (Cornell University, USA), Bryan M. Wong (University of California Riverside, USA), and Khaled Z. Ibrahim (Lawrence Berkeley National Laboratory, USA)</i>	
A Comprehensive Evaluation of Novel AI Accelerators for Deep Learning Workloads .....	13
<i>Murali Emani (Argonne National Laboratory, USA), Zhen Xie (Argonne National Laboratory, USA), Siddhisanket Raskar (Argonne National Laboratory, USA), Varuni Sastry (Argonne National Laboratory, USA), William Arnold (Argonne National Laboratory, USA), Bruce Wilson (Argonne National Laboratory, USA), Rajeev Thakur (Argonne National Laboratory, USA), Venkatram Vishwanath (Argonne National Laboratory, USA), Zhengchun Liu (Argonne National Laboratory, USA), Michael E. Papka (Argonne National Laboratory, USA; University of Illinois, USA), Cindy Orozco Bohorquez (Cerebras Systems, USA), Rick Weisner (SambaNova Systems Inc., USA), Karen Li (SambaNova Systems Inc., USA), Yongning Sheng (SambaNova Systems Inc., USA), Yun Du (SambaNova Systems Inc., USA), Jian Zhang (SambaNova Systems Inc., USA), Alexander Tsyplikhin (Graphcore Inc., USA), Gurdaman Khaira (Graphcore Inc., USA), Jeremy Fowers (Groq Inc., USA), Ramakrishnan Sivakumar (Groq Inc., USA), Victoria Godsoe (Groq Inc., USA), Adrian Macias (Groq Inc., USA), Chetan Tekur (Groq Inc., USA), and Matthew Boyd (Groq Inc., USA)</i>	

## Heterogeneous Systems

Frontier vs the Exascale Report: Why So Long? And Are We Really There Yet? .....	26
<i>Peter M. Kogge (Univ. of Notre Dame, USA) and William J. Dally (NVIDIA Corp., USA)</i>	
Evaluating ISO C++ Parallel Algorithms on Heterogeneous HPC Systems .....	36
<i>Wei-Chen Lin (University of Bristol, UK), Tom Deakin (University of Bristol, UK), and Simon McIntosh-Smith (University of Bristol, UK)</i>	
Going Green: Optimizing GPUs for Energy Efficiency through Model-Steered Auto-Tuning .....	48
<i>Richard Schoonhoven (Computational Imaging Group, Centrum Wiskunde &amp; Informatica, Netherlands; Leiden Institute of Advanced Computer Science, Netherlands), Bram Veenboer (Netherlands Institute for Radio Astronom(ASTRON), Netherlands), Ben van Werkhoven (Computational Imaging Group, Centrum Wiskunde &amp; Informatica, Netherlands; Netherlands eScience Center, Netherlands), and Kees Joost Batenburg (Computational Imaging Group, Centrum Wiskunde &amp; Informatica, Netherlands; Leiden Institute of Advanced Computer Science, Netherlands)</i>	

## Performance Analysis

Performance Analysis with Unified Hardware Counter Metrics .....	60
<i>Brian J Gravelle (Los Alamos National Laboratory), William David Nystrom (Los Alamos National Laboratory), and Boyana Norris (University of Oregon)</i>	
A Methodology for Evaluating Tightly-integrated and Disaggregated Accelerated Architectures .....	71
<i>Taylor Groves (Berkeley Lab, USA), Chris Daley (Berkeley Lab, USA), Rahul Kumar Gayatri (Berkeley Lab, USA), Hai Ah Nam (Berkeley Lab, USA), Nan Ding (Berkeley Lab, USA), Lenny Olikier (Berkeley Lab, USA), Nicholas J. Wright (Berkeley Lab, USA), and Sam Williams (Berkeley Lab, USA)</i>	

## Benchmarking

Benchmarking Fortran DO CONCURRENT on CPUs and GPUs Using BabelStream .....	82
<i>Jeff R. Hammond (NVIDIA Helsinki Oy, Finland), Tom Deakin (HPC Research Group, University of Bristol, UK), James Cownie (HPC Research Group, University of Bristol), and Simon McIntosh-Smith (HPC Research Group, University of Bristol, UK)</i>	
WfBench: Automated Generation of Scientific Workflow Benchmarks .....	100
<i>Tainā Coleman (University of Southern California, USA), Henri Casanova (University of Hawaii, USA), Ketan Maheshwari (Oak Ridge National Laboratory, USA), Loïc Pottier (University of Southern California, USA), Sean R. Wilkinson (Oak Ridge National Laboratory, USA), Justin Wozniak (Argonne National Laboratory, USA), Frédéric Suter (Oak Ridge National Laboratory, USA), Mallikarjun Shankar (Oak Ridge National Laboratory, USA), and Rafael Ferreira da Silva (Oak Ridge National Laboratory, USA)</i>	

High-Performance GMRES Multi-Precision Benchmark: Design, Performance, and Challenges .....	112
<i>Ichitaro Yamazaki (Sandia National Laboratories, USA), Christian Glusa (Sandia National Laboratories, USA), Jennifer Loe (Sandia National Laboratories, USA), Piotr Luszczek (University of Tennessee, USA), Sivasankaran Rajamanickam (Sandia National Laboratories, USA), and Jack Dongarra (University of Tennessee, USA)</i>	

## Short Papers

OMPICollTune: Autotuning MPI Collectives by Incremental Online Learning .....	123
<i>Hunold Sascha (TU Wien, Austria) and Steiner Sebastian (TU Wien, Austria)</i>	
AppEKG: A Simple Unifying View of HPC Applications in Production .....	129
<i>Mohammad Al-Tahat (New Mexico State University, USA), Strahinja Trecakov (New Mexico State University, USA), and Jonathan Cook (New Mexico State University, USA)</i>	
An Initial Evaluation of Arm’s Scalable Matrix Extension .....	135
<i>Finn Wilkinson (University of Bristol, United Kingdom) and Simon McIntosh-Smith (University of Bristol, United Kingdom)</i>	
Time-Series ML-regression on Graphcore IPU-M2000 and Nvidia A100 .....	141
<i>Jan Balewski (NERSC, Lawrence Berkeley National Laboratory, USA), Zhenying Liu (Graphcore, USA), Alexander Tsyplikhin (Graphcore, USA), Manuel Lopez Roland (Graphcore, USA), and Kristofer Bouchard (Scientific Data Division and Biological Systems &amp; Engineering Division, LBNL, Helen Wills Neuroscience Institute &amp; Redwood Center for Theoretical Neuroscience, USA)</i>	
<b>Author Index .....</b>	<b>147</b>