

# **2021 International Workshop on Performance Modeling, Benchmarking and Simulation of High Performance Computer Systems (PMBS 2021)**

**St. Louis, Missouri, USA  
15 November 2021**



**IEEE Catalog Number: CFP21J43-POD  
ISBN: 978-1-6654-1119-6**

**Copyright © 2021 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:	CFP21J43-POD
ISBN (Print-On-Demand):	978-1-6654-1119-6
ISBN (Online):	978-1-6654-1118-9

**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

# 2021 International Workshop on Performance Modeling, Benchmarking and Simulation of High Performance Computer Systems (PMBS) **PMBS 2021**

## Table of Contents

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

### Best Papers

Memory Demands in Disaggregated HPC: How Accurate Do We Need to Be? .....	1
<i>Felippe Vieira Zacarias (Universitat Politècnica de Catalunya, Barcelona Supercomputing Center, Spain), Paul Carpenter (Barcelona Supercomputing Center, Spain), and Vinicius Petrucci (University of Pittsburgh, United States)</i>	
Architectural Requirements for Deep Learning Workloads in HPC Environments .....	7
<i>Khaled Ibrahim (Lawrence Berkeley National Laboratory), Tan Nguyen (Lawrence Berkeley National Laboratory), Hai Ah Nam (Lawrence Berkeley National Laboratory), Wahid Bhimji (Lawrence Berkeley National Laboratory), Steven Farrell (Lawrence Berkeley National Laboratory), Leonid Oliker (Lawrence Berkeley National Laboratory), Michael Rowan (Lawrence Berkeley National Laboratory), Nicholas J. Wright (Lawrence Berkeley National Laboratory), and Samuel Williams (Lawrence Berkeley National Laboratory)</i>	

## Modeling and Simulation

Multilevel Simulation-based Co-design of Next Generation HPC Microprocessors .....	18
<i>Lilia Zaourar (Université Paris-Saclay, France), Mohamed Benazouz (Université Paris-Saclay, France), Ayoub Mouhagir (Université Paris-Saclay, France), Fatma Jebali (Université Paris-Saclay, France), Tanguy Sassolas (Université Paris-Saclay, France), Jean-Christophe Weill (CEA, DAM, France), Carlos Falquez (Institute for Advanced Simulation, Germany), Nam Ho (Institute for Advanced Simulation, Germany), Dirk Pleiter (Institute for Advanced Simulation, Germany), Antoni Portero (Institute for Advanced Simulation, Germany), Estela Suarez (Institute for Advanced Simulation, Germany), Polydoros Petrakis (Foundation for Research and Technology, Greece), Vassilis Papaefstathiou (Foundation for Research and Technology, Greece), Manolis Marazakis (Foundation for Research and Technology, Greece), Milan Radulovic (Barcelona Supercomputing Center, Spain), Francesc Martinez (Barcelona Supercomputing Center, Spain), Adrià Armejach (Barcelona Supercomputing Center, Spain), Marc Casas (Barcelona Supercomputing Center, Spain), Alejandro Nocua (ATOS, France), and Romain Dolbeau (Sipearl, France)</i>	
An Extended Roofline Performance Model with PCI-E and Network Ceilings .....	30
<i>Amanda S. Dufek (Lawrence Berkeley National Laboratory, USA), Jack R. Deslippe (Lawrence Berkeley National Laboratory, USA), Paul T. Lin (Lawrence Berkeley National Laboratory, USA), Charlene J. Yang (NVIDIA, USA), Brandon G. Cook (Lawrence Berkeley National Laboratory, USA), and Jonathan Madsen (AMD, USA)</i>	
Exploration of Congestion Control Techniques on Dragonfly-class HPC Networks Through Simulation .....	40
<i>Neil McGlohon (Rensselaer Polytechnic Institute, USA), K. Scott Hemmert (Sandia National Laboratories, USA), Kevin A. Brown (Argonne National Laboratory, USA), Michael Levenhagen (Sandia National Laboratories, USA), Sudheer Chunduri (Argonne National Laboratory, USA), Robert B. Ross (Argonne National Laboratory, USA), and Christopher D. Carothers (Rensselaer Polytechnic Institute, USA)</i>	
Understanding Power Variation and its Implications on Performance Optimization on the Cori Supercomputer .....	51
<i>Sridutt Bhalachandra (Lawrence Berkeley National Laboratory, USA), Brian Austin (Lawrence Berkeley National Laboratory, USA), and Nicholas J. Wright (Lawrence Berkeley National Laboratory, USA)</i>	

## Short Papers

Using the Semi-Stencil Algorithm to Accelerate High-Order Stencils on GPUs .....	63
<i>Ryuichi Sai (Rice University, USA), John Mellor-Crummey (Rice University, USA), Xiaozhu Meng (Rice University, USA), Mauricio Araya-Polo (TotalEnergies EP Research &amp; Technology USA, USA), and Jie Meng (TotalEnergies EP Research &amp; Technology USA, USA)</i>	
MicroBench Maker: Reproduce, Reuse, Improve .....	69
<i>Sascha Hunold (TU Wien, Austria), Jordy I. Ajanohoun (TU Wien, Austria), and Alexandra Carpen-Amarie (Fraunhofer ITWM, Kaiserslautern, Germany)</i>	

Enabling Cache Aware Roofline Analysis with Portable Hardware Counter Metrics .....	75
<i>Brian J Gravelle (Los Alamos National Laboratory; University of Oregon), William David Nystrom (Los Alamos National Laboratory), Dewi Yokelson (University of Oregon), and Boyana Norris (University of Oregon)</i>	

## Performance Portability and Optimization

Customized Monte Carlo Tree Search for LLVM/Polly's Composable Loop Optimization Transformations .....	82
<i>Jaehoon Koo (Argonne National Laboratory, USA), Prasanna Balaprakash (Argonne National Laboratory, USA), Michael Kruse (Argonne National Laboratory, USA), Xingfu Wu (Argonne National Laboratory, USA), Paul Hovland (Argonne National Laboratory, USA), and Mary Hall (University of Utah, USA)</i>	
Comparing Julia to Performance Portable Parallel Programming Models for HPC .....	94
<i>Wei-Chen Lin (University of Bristol, UK) and Simon McIntosh-Smith (University of Bristol, UK)</i>	
Bayesian Optimization for Auto-Tuning GPU Kernels .....	106
<i>Floris-Jan Willemsen (University of Amsterdam, The Netherlands), Rob van Nieuwpoort (University of Amsterdam, The Netherlands), and Ben van Werkhoven (The Netherlands eScience Center, The Netherlands)</i>	
Narrowing the Search Space of Applications Mapping on Hierarchical Topologies .....	118
<i>Nicolas Denoyelle (Argonne National Laboratory), Emmanuel Jeannot (Inria Bordeaux Sud-Ouest), Swann Perarnau (Argonne National Laboratory), Brice Videau (Argonne National Laboratory), and Pete Beckman (Argonne National Laboratory)</i>	

<b>Author Index</b> .....	<b>129</b>
---------------------------	------------