

2018 IEEE/ACM Performance Modeling, Benchmarking and Simulation of High Performance Computer Systems (PMBS 2018)

**Dallas, Texas, USA
12 November 2018**



**IEEE Catalog Number: CFP18J43-POD
ISBN: 978-1-7281-0183-5**

**Copyright © 2018 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:	CFP18J43-POD
ISBN (Print-On-Demand):	978-1-7281-0183-5
ISBN (Online):	978-1-7281-0182-8

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

CURRAN ASSOCIATES INC.
proceedings
.com

2018 IEEE/ACM Performance Modeling, Benchmarking and Simulation of High Performance Computer Systems (PMBS) **PMBS 2018**

Table of Contents

Session 1: HPC Interconnects

Improving MPI Reduction Performance for Manycore Architectures with OpenMP and Data
Compression .1.....
*Hongzhang Shan (Lawrence Berkeley National Laboratory, United States),
Samuel Williams (Lawrence Berkeley National Laboratory, United
States), and Calvin W. Johnson (San Diego State University, United
States)*

Exploring and Quantifying How Communication Behaviors in Proxies Relate to Real
Applications .12.....
*Omar Aaziz (Sandia National Labs), Jeanine Cook (Sandia National
Labs), Jonathan Cook (New Mexico State University), and Courtenay
Vaughan (Sandia National Labs)*

Session 2: Machine/Deep Learning

Deep Learning at Scale on NVIDIA V100 Accelerators .23.....
Rengan Xu (Dell EMC), Frank Han (Dell EMC), and Quy Ta (Dell EMC)

Benchmarking Machine Learning Methods for Performance Modeling of Scientific Applications .33.....
*Preeti Malakar (Argonne National Laboratory), Prasanna Balaprakash
(Argonne National Laboratory), Venkatram Vishwanath (Argonne National
Laboratory), Vitali Morozov (Argonne National Laboratory), and Kalyan
Kumaran (Argonne National Laboratory)*

Session 3: Late-breaking Research and Preliminary Techniques

Algorithm Selection of MPI Collectives Using Machine Learning Techniques .45.....
Sascha Hunold (TU Wien) and Alexandra Carpen-Amarie (Fraunhofer ITWM)

miniVite: A Graph Analytics Benchmarking Tool for Massively Parallel Systems .51.....	
	<i>Sayan Ghosh (Washington State University), Mahantesh Halappanavar (Pacific Northwest National Laboratory), Antonino Tumeo (Pacific Northwest National Laboratory), Ananth Kalyanaraman (Washington State University), and Assefaw H. Gebremedhin (Washington State University)</i>
Unified Cross-Platform Profiling of Parallel C++ Applications .57.....	
	<i>Vladyslav Kucher (University of Muenster), Florian Fey (University of Muenster), and Sergej Gorlatch (University of Muenster)</i>

Session 4: System Evaluation

A Metric for Evaluating Supercomputer Performance in the Era of Extreme Heterogeneity .63.....	
	<i>Brian Austin (Lawrence Berkeley National Laboratory), Chris Daley (Lawrence Berkeley National Laboratory), Douglas Doerfler (Lawrence Berkeley National Laboratory), Jack Deslippe (Lawrence Berkeley National Laboratory), Brandon Cook (Lawrence Berkeley National Laboratory), Brian Friesen (Lawrence Berkeley National Laboratory), Thorsten Kurth (Lawrence Berkeley National Laboratory), Charlene Yang (Lawrence Berkeley National Laboratory), and Nicholas J. Wright (Lawrence Berkeley National Laboratory)</i>
Evaluating SLURM Simulator with Real-Machine SLURM and Vice Versa .72.....	
	<i>Ana Jokanovic (Barcelona Supercomputing Center (BSC), Spain), Marco D'Amico (Barcelona Supercomputing Center (BSC), Spain), and Julita Corbalan (Universitat Politecnica de Catalunya (UPC), Spain)</i>

Session 5: Performance

Is Data Placement Optimization Still Relevant on Newer GPUs? .83.....	
	<i>Md Abdullah Shahneous Bari (Stony Brook University, Stony Brook, NY, USA), Larisa Stoltzfus (University of Edinburgh, Edinburgh, UK), Pei-Hung Lin (Lawrence Livermore National Laboratory, Livermore, CA , USA), Chunhua Liao (Lawrence Livermore National Laboratory, Livermore, CA , USA), Murali Emani (Lawrence Livermore National Laboratory, Livermore, CA , USA), and Barbara Chapman (Stony Brook University, Stony Brook, NY, USA; Brookhaven National Laboratory, Upton, NY , USA)</i>
Approximating a Multi-Grid Solver .97.....	
	<i>Valentin Le Fèvre (ENS de Lyon), Leonardo Bautista-Gomez (Barcelona Supercomputing Center), Osman Unsal (Barcelona Supercomputing Center), and Marc Casas (Barcelona Supercomputing Center)</i>

Session 6: Simulation and Prediction

Evaluating the Impact of Spiking Neural Network Traffic on Extreme-Scale Hybrid Systems .108.....	
	<i>Noah Wolfe (Rensselaer Polytechnic Institute), Mark Plagge (Rensselaer Polytechnic Institute), Christopher D. Carothers (Rensselaer Polytechnic Institute), Misbah Mubarak (Argonne National Laboratory), and Robert B. Ross (Argonne National Laboratory)</i>

Automated Instruction Stream Throughput Prediction for Intel and AMD Microarchitectures .121.....

*Jan Laukemann (Friedrich-Alexander-Universität Erlangen-Nürnberg),
Julian Hammer (Friedrich-Alexander-Universität Erlangen-Nürnberg),
Johannes Hofmann (Friedrich-Alexander-Universität Erlangen-Nürnberg),
Georg Hager (Friedrich-Alexander-Universität Erlangen-Nürnberg), and
Gerhard Wellein (Friedrich-Alexander-Universität Erlangen-Nürnberg)*

Author Index .133