

2021 IEEE International Parallel and Distributed Processing Symposium Workshops (IPDPSW 2021)

**Portland, Oregon, USA
17 – 21 June 2021**

Pages 1-519



**IEEE Catalog Number: CFP2151J-POD
ISBN: 978-1-6654-1192-9**

**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:	CFP2151J-POD
ISBN (Print-On-Demand):	978-1-6654-1192-9
ISBN (Online):	978-1-6654-3577-2

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

2021 IEEE International Parallel and Distributed Processing Symposium Workshops (IPDPSW) IPDPSW 2021

Table of Contents

Message from the 2021 General Co-Chairs	xxiv
IPDPS 2021 Message from the Workshops Chair and Vice-chair	xxvi

HCW: Heterogeneity in Computing Workshop

Introduction to HCW 2021	1
<i>Florina M. Ciorba (University of Basel, Switzerland) and Ryan D. Friese (Pacific Northwest National Laboratory, USA)</i>	
Message from the HCW Steering Committee Chair	2
<i>Behrooz Shirazi (Washington State University, USA)</i>	
Message from the HCW General Chair	3
<i>Florina M. Ciorba (University of Basel, Switzerland)</i>	
Message from the HCW Technical Program Committee Chair	N/A
<i>Ryan D. Friese (Pacific Northwest National Laboratory, USA)</i>	
HCW 2021 Keynote Speaker	N/A
<i>Roberto Gioiosa (Pacific Northwest National Laboratory, USA)</i>	
Adaptive Stochastic Gradient Descent for Deep Learning on Heterogeneous CPU+GPU Architectures	6
<i>Yujing Ma (University of California Merced, USA), Florin Rusu (University of California Merced, USA), Kesheng Wu (Lawrence Berkeley National Lab, USA), and Alexander Sim (Lawrence Berkeley National Lab, USA)</i>	
Providing In-Depth Performance Analysis for Heterogeneous Task-Based Applications with StarVZ	16
<i>Vinicius Garcia Pinto (Federal University of Rio Grande do Sul, Brazil), Lucas Leandro Nesi (Federal University of Rio Grande do Sul, Brazil), Marcelo Cogo Miletto (Federal University of Rio Grande do Sul, Brazil), and Lucas Mello Schnorr (Federal University of Rio Grande do Sul, Brazil)</i>	
A Streaming Accelerator for Heterogeneous CPU-FPGA Processing of Graph Applications	26
<i>Francis O'Brien (University of Toronto, Canada), Matthew Agostini (University of Toronto, Canada), and Tarek S. Abdelrahman (University of Toronto, Canada)</i>	

A New Double Rank-Based Multi-workflow Scheduling with Multi-objective Optimization in Cloud Environments	36
<i>Feng Li (Nanyang Technological University, Singapore), Moon Gi Seok (Nanyang Technological University, Singapore), and Wentong Cai (Nanyang Technological University, Singapore)</i>	
Pooling Acceleration in the DaVinci Architecture Using Im2col and Col2im Instructions	46
<i>Caio S. Rohwedder (Institute of Computing, University of Campinas (UNICAMP), Brazil; University of Alberta, Canada), João P. L. de Carvalho (Institute of Computing, University of Campinas (UNICAMP), Brazil; University of Alberta, Canada), José Nelson Amaral (University of Alberta, Canada), Guido Araújo (Institute of Computing, University of Campinas (UNICAMP), Brazil), Giancarlo Colmenares (Heterogeneous Compiler Lab., Huawei Canada Research Centre, Canada), and Kai-Ting Amy Wang (Heterogeneous Compiler Lab., Huawei Canada Research Centre, Canada)</i>	
Scheduling HPC Workflows with Intel Optane Persistent Memory	56
<i>Ranjan Sarpankala Venkatesh (Georgia Institute of Technology, USA), Tony Mason (University of British Columbia, Canada; Georgia Institute of Technology, USA), Pradeep Fernando (Georgia Institute of Technology, USA), Greg Einsenhauer (Georgia Institute of Technology, USA), and Ada Gavrilovska (Georgia Institute of Technology, USA)</i>	
Coding the Computing Continuum: Fluid Function Execution in Heterogeneous Computing Environments	66
<i>Rohan Kumar (University of Chicago, USA), Matt Baughman (University of Chicago, USA), Ryan Chard (Data Science and Learning Division, Argonne National Laboratory, USA), Zhuozhao Li (University of Chicago, USA; Data Science and Learning Division, Argonne National Laboratory, USA), Yadu Babuji (University of Chicago, USA; Data Science and Learning Division, Argonne National Laboratory, USA), Ian Foster (University of Chicago, USA; Data Science and Learning Division, Argonne National Laboratory, USA), and Kyle Chard (University of Chicago, USA; Data Science and Learning Division, Argonne National Laboratory, USA)</i>	
Practice and Experience in using Parallel and Scalable Machine Learning with Heterogenous Modular Supercomputing Architectures	76
<i>Morris Riedel (University of Iceland, Iceland), Rocco Sedona (Juelich Supercomputing Centre, Germany), Chadi Barakat (Juelich Supercomputing Centre, Germany), Petur Einarsson (University of Iceland, Iceland), Reza Hassanian (University of Iceland, Iceland), Gabriele Cavallaro (Juelich Supercomputing Centre, Germany), Matthias Book (University of Iceland, Iceland), Helmut Neukirchen (University of Iceland, Iceland), and Andreas Lintermann (Juelich Supercomputing Centre, Germany)</i>	

RAW: Reconfigurable Architectures Workshop

Introduction to RAW 2021	86
<i>Marco D. Santambrogio (Politecnico di Milano, Italy)</i>	
Accelerating ODE-Based Neural Networks on Low-Cost FPGAs	88
<i>Hirohisa Watanabe (Keio University, Japan) and Hiroki Matsutani (Keio University, Japan)</i>	

An FPGA-Based On-Device Reinforcement Learning Approach using Online Sequential Learning ..	96
<i>Hirohisa Watanabe (Keio University, Japan), Mineto Tsukada (Keio University, Japan), and Hiroki Matsutani (Keio University, Japan)</i>	
Plaster: An Embedded FPGA-Based Cluster Orchestrator for Accelerated Distributed Algorithms	104
<i>Lorenzo Farinelli (Politecnico di Milano, Italy), Daniele Valentino De Vincenti (Politecnico di Milano, Italy), Andrea Damiani (Politecnico di Milano, Italy), Luca Stornaiuolo (Politecnico di Milano, Italy), Rolando Brondolin (Politecnico di Milano, Italy), Marco D. Santambrogio (Politecnico di Milano, Italy), and Donatella Sciuto (Politecnico di Milano, Italy)</i>	
BinaryCoP: Binary Neural Network-Based COVID-19 Face-Mask Wear and Positioning Predictor on Edge Devices	108
<i>Nael Fafous (Technical University of Munich, Germany), Manoj-Rohit Vemparala (BMW Group, Germany), Alexander Frickenstein (BMW Group, Germany), Lukas Frickenstein (BMW Group, Germany), Mohamed Badawy (Technical University of Munich, Germany), and Walter Stechele (Technical University of Munich, Germany)</i>	
Exploring a Layer-Based Pre-Implemented Flow for Mapping CNN on FPGA	116
<i>Danielle Tchuinkou Kwadjo (University of Florida, USA), Joel Mandebi Mbongue (University of Florida, USA), and Christophe Bobda (University of Florida, USA)</i>	
A Machine Learning Approach to Predict Timing Delays During FPGA Placement	124
<i>T. Martin (University of Guelph, Canada), G. Grewal (University of Guelph, Canada), and S. Areibi (University of Guelph, Canada)</i>	
Dovado: An Open-Source Design Space Exploration Framework	128
<i>Daniele Paletti (Politecnico di Milano, Italy), Davide Conficconi (Politecnico di Milano, Italy), and Marco D. Santambrogio (Politecnico di Milano, Italy)</i>	
A Framework for the Automatic Generation of FPGA-Based Near-Data Processing Accelerators in Smart Storage Systems	136
<i>Lukas Weber (Embedded Systems and Applications Group, TU Darmstadt, Germany), Lukas Sommer (Embedded Systems and Applications Group, TU Darmstadt, Germany), Leonardo Solis-Vasquez (Embedded Systems and Applications Group, TU Darmstadt, Germany), Tobias Vinçon (DBlab, Reutlingen University, Germany), Christian Knödler (DBlab, Reutlingen University, Germany), Arthur Bernhardt (DBlab, Reutlingen University, Germany), Ilia Petrov (DBlab, Reutlingen University, Germany), and Andreas Koch (Embedded Systems and Applications Group, TU Darmstadt, Germany)</i>	
On Data Parallelism Code Restructuring for HLS Targeting FPGAs	144
<i>Renato Campos (Faculty of Engineering of the University of Porto, and INESC TEC, Portugal) and João M.P. Cardoso (Faculty of Engineering of the University of Porto, and INESC TEC, Portugal)</i>	

Fast HBM Access with FPGAs: Analysis, Architectures, and Applications	152
<i>Philipp Holzinger (Friedrich-Alexander University Erlangen-Nürnberg, Germany), Daniel Reiser (Friedrich-Alexander University Erlangen-Nürnberg, Germany), Tobias Hahn (Friedrich-Alexander University Erlangen-Nürnberg, Germany), and Marc Reichenbach (Friedrich-Alexander University Erlangen-Nürnberg, Germany)</i>	
Graph Analytics on Hybrid System (GAHS) Case Study: PageRank	160
<i>Mohamed W. Hassan (Virginia Tech, USA) and Peter M. Athanas (Virginia Tech, USA)</i>	
Performance Study of Multi-tenant Cloud FPGAs	168
<i>Joel Mandebi Mbongue (University of Florida, USA), Sujan Kumar Saha (University of Florida, USA), and Christophe Bobda (University of Florida, USA)</i>	
RV-CAP: Enabling Dynamic Partial Reconfiguration for FPGA-Based RISC-V System-on-Chip	172
<i>Najdet Charaf (Chair of Adaptive Dynamic Systems, Technische Universität Dresden, Germany), Ahmed Kamaleldin (Chair of Adaptive Dynamic Systems, Technische Universität Dresden, Germany), Martin Thümmler (Chair of Adaptive Dynamic Systems, Technische Universität Dresden, Germany), and Diana Göhringer (Chair of Adaptive Dynamic Systems, Technische Universität Dresden, Germany)</i>	
An Area-Efficient SPHINCS+ Post-Quantum Signature Coprocessor	180
<i>Quentin Berthet (inIT, HEPIA, University of Applied Sciences of Western Switzerland, Switzerland), Andres Upegui (inIT, HEPIA, University of Applied Sciences of Western Switzerland, Switzerland), Laurent Gantel (inIT, HEPIA, University of Applied Sciences of Western Switzerland, Switzerland), Alexandre Duc (IICT, HEIG-VD, University of Applied Sciences of Western Switzerland, Switzerland), and Giulia Traverso (CYSEC SA, Switzerland)</i>	
FPGA Acceleration of Zstd Compression Algorithm	188
<i>Jianyu Chen (Accelerated Big Data Systems, Delft University of Technology, The Netherlands; Optiver, The Netherlands), Maurice Daverveldt (Optiver, The Netherlands), and Zaid Al-Ars (Accelerated Big Data Systems, Delft University of Technology, The Netherlands)</i>	

HiCOMB: High Performance Computational Biology

Introduction to HiCOMB 2021	192
<i>Alba Cristina M. A. de Melo (University of Brasilia, Brazil), Ananth Kalyanaraman (Washington State University, USA), and Mehmet Koyutürk (Case Western Reserve University, USA)</i>	

GYAN: Accelerating Bioinformatics Tools in Galaxy with GPU-Aware Computation Mapping	194
<i>Gulsum Gudukbay (Pennsylvania State University, USA), Jashwant Raj Gunasekaran (Pennsylvania State University, USA), Yilin Feng (Pennsylvania State University, USA), Mahmut T. Kandemir (Pennsylvania State University, USA), Anton Nekrutenko (Pennsylvania State University, USA), Chita R. Das (Pennsylvania State University, USA), Paul Medvedev (Pennsylvania State University, USA), Björn Grüning (University of Freiburg, Germany), Nate Coraor (Pennsylvania State University, USA), Nathan Roach (Galaxyworks, Johns Hopkins University, USA), and Enis Afgan (Galaxyworks, Johns Hopkins University, USA)</i>	
Accelerating SARS-CoV-2 low Frequency Variant Calling on Ultra Deep Sequencing Datasets	204
<i>Bryce Kille (Rice University, USA), Yunxi Liu (Rice University, USA), Nicolae Sapoval (Rice University, USA), Michael Nute (Rice University, USA), Lawrence Rauchwerger (University of Illinois at Urbana-Champaign, USA), Nancy Amato (University of Illinois at Urbana-Champaign, USA), and Todd J. Treangen (Rice University, USA)</i>	
GateKeeper-GPU: Fast and Accurate Pre-Alignment Filtering in Short Read Mapping	209
<i>Zülal Bingöl (Bilkent University, Turkey), Mohammed Alser (ETH Zürich, Switzerland), Onur Mutlu (ETH Zürich, Switzerland), Ozcan Ozturk (Bilkent University, Turkey), and Can Alkan (Bilkent University, Turkey)</i>	
GPU Acceleration of 3D Agent-Based Biological Simulations	210
<i>Ahmad Hesam (Delft University of Technology, Netherlands), Lukas Breitwieser (CERN, Switzerland), Fons Rademakers (CERN, Switzerland), and Zaid Al-Ars (Delft University of Technology, Netherlands)</i>	
Efficient Memory Management in Likelihood-Based Phylogenetic Placement	218
<i>Pierre Barbera (Heidelberg Institute for Theoretical Studies, Germany) and Alexandros Stamatakis (Karlsruhe Institute of Technology, Germany; Heidelberg Institute for Theoretical Studies, Germany)</i>	
Accelerating the BPMax Algorithm for RNA-RNA Interaction	228
<i>Chiranjeb Mondal (Colorado State University, USA) and Sanjay Rajopadhye (Colorado State University, USA)</i>	

GrAPL: Graphs, Architectures, Programming, and Learning

Message from the GrAPL 2021 Workshop Chairs	238
<i>Scott McMillan (CMU SEI, USA), Manoj Kumar (IBM, USA), and Nesreen K. Ahmed (Intel Labs, USA)</i>	
GrAPL 2021 Keynote 1: Sparse Adjacency Matrices at the Core of Graph Databases: GraphBLAS the Engine Behind RedisGraph Property Graph Database	240
<i>Roi Lipman (Redis Labs, USA)</i>	
GrAPL 2021 Keynote 2: Label Propagation and Graph Neural Networks	242
<i>Austin Benson (Cornell University, USA)</i>	

LAGraph: Linear Algebra, Network Analysis Libraries, and the Study of Graph Algorithms	243
<i>Gábor Szárnyas (CWI Amsterdam, the Netherlands), David A. Bader (New Jersey Institute of Technology, USA), Timothy A. Davis (Texas A&M University, USA), James Kitchen (Anaconda, Inc., USA), Timothy G. Mattson (Intel Corp., USA), Scott McMillan (Software Engineering Institute, Carnegie Mellon University, USA), and Erik Welch (Anaconda, Inc., USA)</i>	
Introduction to GraphBLAS 2.0	253
<i>Benjamin Brock (University of California, Berkeley, USA), Aydin Buluc (Lawrence Berkeley National Laboratory, USA), Timothy G. Mattson (Parallel Computing Labs, Intel, USA), Scott McMillan (Software Engineering Institute, Carnegie Mellon University, USA), and Jose E. Moreira (IBM Thomas J. Watson Research Center, USA)</i>	
Mathematics of Digital Hyperspace	263
<i>Jeremy Kepner (MIT Lincoln Laboratory Supercomputing Center; MIT Computer Science & AI Laboratory, USA), Timothy Davis (Texas A&M University, USA), Vijay Gadepally (MIT Lincoln Laboratory Supercomputing Center; MIT Computer Science & AI Laboratory, USA), Hayden Jananathan (MIT Lincoln Laboratory Supercomputing Center USA; Vanderbilt, USA), and Lauren Milechin (MIT, USA)</i>	
SPbLA: The Library of GPGPU-Powered Sparse Boolean Linear Algebra Operations	272
<i>Egor Orachev (Saint Petersburg State University, JetBrains Research, Russia), Maria Karpenko (ITMO University, Russia), Artem Khoroshev (Computation Biology Department BIOCAD, Russia), and Semyon Grigorev (Saint Petersburg State University, JetBrains Research, Russia)</i>	
PIGO: A Parallel Graph Input/Output Library	276
<i>Kasimir Gabert (Georgia Institute of Technology, USA) and Ümit V. Çatalyürek (Georgia Institute of Technology, USA)</i>	
Hybrid Power-Law Models of Network Traffic	280
<i>Pat Deolin (Yale University, USA), Jeremy Kepner (MIT, USA), Ashley Luo (MIT, USA), and Erin Meger (Université du Québec à Montréal, Canada)</i>	
Characterizing Job-Task Dependency in Cloud Workloads Using Graph Learning	288
<i>Zhaochen Gu (University of North Texas, USA), Sihai Tang (University of North Texas, USA), Beilei Jiang (University of North Texas, USA), Song Huang (Allstate, USA), Qiang Guan (Kent State University, USA), and Song Fu (University of North Texas, USA)</i>	
Co-Design of Advanced Architectures for Graph Analytics using Machine Learning	298
<i>Kuldeep Kurte (Oak Ridge National Laboratory, USA), Neena Imam (Oak Ridge National Laboratory, USA), Ramakrishnan Kannan (Oak Ridge National Laboratory, USA), S.M.Shamimul Hasan (Oak Ridge National Laboratory, USA), and Srikanth Yoginath (Oak Ridge National Laboratory, USA)</i>	
Sparse Binary Matrix-Vector Multiplication on Neuromorphic Computers	308
<i>Catherine D. Schuman (Oak Ridge National Laboratory, USA), Bill Kay (Oak Ridge National Laboratory, USA), Prasanna Date (Oak Ridge National Laboratory, USA), Ramakrishnan Kannan (Oak Ridge National Laboratory, USA), Piyush Sao (Oak Ridge National Laboratory, USA), and Thomas E. Potok (Oak Ridge National Laboratory, USA)</i>	

EduPar: NSF/TCPP Workshop on Parallel and Distributed Computing Education

Message from the EduPar-21 Workshop Chair	312
<i>Sushil K. Prasad (The University of Texas at San Antonio, USA)</i>	
EduPar-21 Keynote: Using Smartphones to Teach Parallelism to First Year Students	314
<i>David H. Albonesi (Cornell University, USA)</i>	
Let's Put the Memory Model Front and Center when Teaching Parallel Programming in C++	315
<i>Jiri Dokulil (University of Vienna, Austria)</i>	
Teaching Complex Scheduling Algorithms	321
<i>Sascha Hunold (Tu Wien, Austria) and Bartłomiej Przybylski (Adam Mickiewicz University, Poland)</i>	
ABET Accreditation: A Way Forward for PDC Education	328
<i>Sherif G. Aly (The American University in Cairo, Egypt), Haidar Harmanani (Lebanese American University, Lebanon), Rajendra K. Raj (Rochester Institute of Technology, USA), and Sanaa Sharafeddine (Lebanese American University, Lebanon)</i>	
EduPar Virtual Poster Session	336
<i>Jesús Cámara (University of Murcia, Spain), José-Carlos Cano (University of Murcia, Spain), Javier Cuenca (University of Murcia, Spain), Toshiyuki Maeda (Hannan University, Japan), Mariano Saura-Sánchez (Technical University of Cartagena, Spain), Lewis Tseng (Boston College, USA), Akiyoshi Wakatani (Konan University, Japan), and Martina Barnas (Indiana University Bloomington, USA)</i>	
Teaching PDC in the Time of COVID: Hands-on Materials for Remote Learning	342
<i>Joel C. Adams (Calvin University, USA), Richard Brown (St. Olaf College, USA), Suzanne J. Matthews (U.S. Military Academy, USA), and Elizabeth Shoop (Macalester College, USA)</i>	
Data-Intensive Computing Modules for Teaching Parallel and Distributed Computing	350
<i>Michael Gowanlock (Northern Arizona University, USA) and Benoît Gallet (Northern Arizona University, USA)</i>	

HIPS: High-level Parallel Programming Models and Supportive Environments

Message from the HIPS 2021 Workshop Co-Chairs	358
<i>Julian Kunkel (University of Reading, UK) and Bin Ren (William & Mary, USA)</i>	

Developing Medical Ultrasound Beamforming Application on GPU and FPGA using oneAPI	360
<i>Yong Wang (Intel China Research Center Ltd., China), Yongfa Zhou (Intel China Research Center Ltd., China), Qi (Scott) Wang (Intel China Research Center Ltd., China), Yang Wang (Intel China Research Center Ltd., China), Qing Xu (Intel China Research Center Ltd., China), Chen Wang (Intel China Research Center Ltd., China), Bo Peng (Intel China Research Center Ltd., China), Zhaojun Zhu (Intel China Research Center Ltd., China), Katayama Takuya (Intel China Research Center Ltd., China), and Dylan Wang (Intel China Research Center Ltd., China)</i>	
Evaluating CUDA Portability with HIPCL and DPCT	371
<i>Zheming Jin (Oak Ridge National Laboratory, USA) and Jeffrey Vetter (Oak Ridge National Laboratory, USA)</i>	
Beyond Fork-Join: Integration of Performance Portable Kokkos Kernels with HPX	377
<i>Gregor Daiß (University of Stuttgart, Institute for Parallel and Distributed Systems, Scientific Computing, Germany), Mikael Simberg (Swiss National Supercomputing Centre, Switzerland), Auriane Reverdell (Swiss National Supercomputing Centre, Switzerland), John Biddiscombe (Swiss National Supercomputing Centre, Switzerland), Theresa Pollinger (University of Stuttgart, Institute for Parallel and Distributed Systems, Scientific Computing, Germany), Hartmut Kaiser (Louisiana State University, CCT, USA), and Dirk Pflüger (University of Stuttgart, Institute for Parallel and Distributed Systems, Scientific Computing, Germany)</i>	
An Efficient Approach for Image Border Handling on GPUs via Iteration Space Partitioning	387
<i>Bo Qiao (Friedrich-Alexander University Erlangen-Nürnberg, Germany), Jürgen Teich (Friedrich-Alexander University Erlangen-Nürnberg, Germany), and Frank Hannig (Friedrich-Alexander University Erlangen-Nürnberg, Germany)</i>	
CUDAMicroBench: Microbenchmarks to Assist CUDA Performance Programming	397
<i>Xinyao Yi (University of North Carolina at Charlotte, USA), David Stokes (University of North Carolina at Charlotte, USA), Yonghong Yan (University of North Carolina at Charlotte, USA), and Chunhua Liao (Lawrence Livermore National Laboratory, USA)</i>	
Understanding Recursive Divide-and-Conquer Dynamic Programs in Fork-Join and Data-Flow Execution Models	407
<i>Poornima Nookala (Illinois Institute of Technology, USA), Zafar Ahmad (Stony Brook University, USA), Mohammad Mahdi Javanmard (Stony Brook University, USA), Martin Kong (University of Oklahoma, USA), Rezaul Chowdhury (Stony Brook University, USA), and Robert Harrison (Stony Brook University, USA)</i>	
Measuring Cache Complexity Using Data Movement Distance (DMD) (Position Paper)	417
<i>Donovan Snyder (University of Rochester, USA) and Chen Ding (University of Rochester, USA)</i>	
Combining Static and Dynamic Analysis to Query Characteristics of HPC Applications	420
<i>Aaron Welch (University of Houston, USA), Oscar Hernandez (Oak Ridge National Laboratory, USA), and Barbara Chapman (Stony Brook University, USA)</i>	

AsHES: Accelerators and Hybrid Emerging Systems

Introduction to AsHES 2021	430
<i>Min Si (Argonne National Laboratory, USA)</i>	
AsHES 2021 Keynote	432
<i>Dong Li (University of California, Merced, USA)</i>	
Time-Division Multiplexing for FPGA Considering CNN Model Switch Time	433
<i>Tetsuro Nakamura (NTT Network Service Systems Laboratories, Japan), Shogo Saito (NTT Network Service Systems Laboratories, Japan), Kei Fujimoto (NTT Network Service Systems Laboratories, Japan), Masashi Kaneko (NTT Network Service Systems Laboratories, Japan), and Akinori Shiraga (NTT Network Service Systems Laboratories, Japan)</i>	
Design Space Exploration of Emerging Memory Technologies for Machine Learning Applications ..	439
<i>S.M.Shamimul Hasan (Oak Ridge National Laboratory, USA), Neena Imam (Oak Ridge National Laboratory, USA), Ramakrishnan Kannan (Oak Ridge National Laboratory, USA), Srikanth Yoginath (Oak Ridge National Laboratory, USA), and Kuldeep Kurte (Oak Ridge National Laboratory, USA)</i>	
Accelerating Radiation Therapy Dose Calculation with Nvidia GPUs	449
<i>Felix Liu (KTH Royal Institute of Technology, RaySearch Laboratories, Sweden), Niclas Jansson (KTH Royal Institute of Technology, Sweden), Artur Podobas (KTH Royal Institute of Technology, Sweden), Albin Fredriksson (RaySearch Laboratories, Sweden), and Stefano Markidis (KTH Royal Institute of Technology, Sweden)</i>	
Improving Cryptanalytic Applications with Stochastic Runtimes on GPUs	459
<i>Lena Oden (FernUniversität in Hagen, Germany) and Jörg Keller (FernUniversität in Hagen, Germany)</i>	
Experimental Evaluation of Multiprecision Strategies for GMRES on GPUs	469
<i>Jennifer A. Loe (Center for Computing Research, Sandia National Laboratories, USA), Christian A. Glusa (Center for Computing Research, Sandia National Laboratories, USA), Ichitaro Yamazaki (Center for Computing Research, Sandia National Laboratories, USA), Erik G. Boman (Center for Computing Research, Sandia National Laboratories, USA), and Sivasankaran Rajamanickam (Center for Computing Research, Sandia National Laboratories, USA)</i>	
GPU-Aware Communication with UCX in Parallel Programming Models: Charm++, MPI, and Python	479
<i>Jaemin Choi (University of Illinois at Urbana-Champaign, USA), Zane Fink (University of Illinois at Urbana-Champaign, USA), Sam White (University of Illinois at Urbana-Champaign, USA), Nitin Bhat (Charmworks, Inc., USA), David F. Richards (Center for Applied Scientific Computing, Lawrence Livermore National Laboratory, USA), and Laxmikant V. Kale (University of Illinois at Urbana-Champaign, USA; Charmworks, Inc., USA)</i>	

PDCO: Parallel / Distributed Combinatorics and Optimization

Introduction to PDCO 2021	489
<i>Grégoire Danoy (University of Luxembourg, Luxembourg) and Didier El Baz (team CDA, LAAS-CNRS, France)</i>	
CPRIC: Collaborative Parallelism for Randomized Incremental Constructions	490
<i>Florian Fey (University of Muenster, Germany) and Sergei Gortlatch (University of Muenster, Germany)</i>	
Characters Recognition Based on CNN-RNN Architecture and Metaheuristic	500
<i>F. Keddous (ATOS (Bull SAS), France, UPEC, LISSI, Vitry Sur Seine, France), H-N. Nguyen (ATOS (Bull) MVS Service, Les Clayes-Sous-Bois, France), and A. Nakib (UPEC, LISSI Vitry-sur seine, France)</i>	
Linearizing Computing the Power Set with OpenMP	508
<i>Roger L Goodwin (Independent Researcher)</i>	
TurboBFS: GPU Based Breadth-First Search (BFS) Algorithms in the Language of Linear Algebra	520
<i>Oswaldo Artiles (Florida International University, USA) and Fahad Saeed (Florida International University, USA)</i>	
A Parallel Meta-Solver for the Multi-objective Set Covering Problem	529
<i>Ryan J. Marshall (UTC SimCenter, University of Tennessee, USA), Lakmali Weerasena (University of Tennessee, USA), and Anthony Skjellum (UTC SimCenter, University of Tennessee, USA)</i>	
Leveraging High Dimensional Spatial Graph Embedding as a Heuristic for Graph Algorithms	539
<i>Peter Oostema (Carnegie Mellon University, USA) and Franz Franchetti (Carnegie Mellon University, USA)</i>	
RRNS Base Extension Error-Correcting Code for Performance Optimization of Scalable Reliable Distributed Cloud Data Storage	548
<i>Mikhail Babenko (North-Caucasus Federal University, Russia; Ivannikov Institute for System Programming, Russia), Andrei Tchernykh (CICESE Research Center, Mexico; South Ural State University, Russia; Ivannikov Institute for System Programming, Russia), Bernardo Pulido-Gaytan (CICESE Research Center, Mexico), Jorge M. Cortés-Mendoza (South Ural State University, Russia), Egor Shiryayev (North-Caucasus Federal University, Russia), Elena Golimblevskaia (North-Caucasus Federal University, Russia), Arutyun Avetisyan (Ivannikov Institute for System Programming, Russia), and Sergio Nesmachnow (Universidad de la Republica, Uruguay)</i>	

APDCM: Advances in Parallel and Distributed Computational Models

Introduction to APDCM 2021	554
<i>Jacir L. Bordim (University of Brasilia, Brazil) and Koji Nakano (Hiroshima University, Japan)</i>	

Checkpointing vs. Supervision Resilience Approaches for Dynamic Independent Tasks	556
<i>Jonas Posner (Research Group Programming Languages / Methodologies, University of Kassel, Germany), Lukas Reitz (Research Group Programming Languages / Methodologies, University of Kassel, Germany), and Claudia Fohry (Research Group Programming Languages / Methodologies, University of Kassel, Germany)</i>	
Gathering of Seven Autonomous Mobile Robots on Triangular Grids	566
<i>Masahiro Shibata (Kyushu Institute of Technology, Japan), Masaki Ohyabu (Nagoya Institute of Technology, Japan), Yuichi Sudo (Osaka University, Japan), Junya Nakamura (Toyoashi University of Technology, Japan), Yonghwan Kim (Nagoya Institute of Technology, Japan), and Yoshiaki Katayama (Nagoya Institute of Technology, Japan)</i>	
Autonomous Mobile Robots: Refining the Computational Landscape	576
<i>Kevin Buchin (TU Eindhoven, The Netherlands), Paola Flocchini (University of Ottawa, Canada), Irina Kostitsyna (TU Eindhoven, The Netherlands), Tom Peters (TU Eindhoven, The Netherlands), Nicola Santoro (Carleton University, Canada), and Koichi Wada (Hosei University, Japan)</i>	
Terminating grid Exploration with Myopic Luminous Robots	586
<i>Shota Nagahama (Nara Institute of Science and Technology, Japan), Fukuhito Ooshita (Nara Institute of Science and Technology, Japan), and Michiko Inoue (Nara Institute of Science and Technology, Japan)</i>	
A self-Stabilizing Token Circulation with Graceful Handover on Bidirectional Ring Networks.....	596
<i>Hirotsugu Kakugawa (Ryukoku University, Japan) and Sayaka Kamei (Hiroshima University, Japan)</i>	
Scalable and Highly Available Multi-objective Neural Architecture Search in Bare Metal Kubernetes Cluster	605
<i>Andreas Klos (FernUniversität in Hagen, Germany), Marius Rosenbaum (FernUniversität in Hagen, Germany), and Wolfram Schiffmann (FernUniversität in Hagen, Germany)</i>	
Revisiting Credit Distribution Algorithms for Distributed Termination Detection	611
<i>George Bosilca (University Tennessee Knoxville, USA), Aurélien Bouteiller (University Tennessee Knoxville, USA), Thomas Herault (University Tennessee Knoxville, USA), Valentin Le Fèvre (Barcelona Supercomputing Center, Spain), Yves Robert (Laboratoire LIP, ENS Lyon, France), and Jack Dongarra (University Tennessee Knoxville, USA)</i>	
Efficient and Eventually Consistent Collective Operations	621
<i>Roman Iakymchuk (Fraunhofer ITWM / Sorbonne University, Germany; Sorbonne Université, France), Amandio Faustino (INESC-ID & IST (ULisboa), Portugal), Andrew Emerson (CINECA, Italy), Joao Barreto (INESC-ID & IST (ULisboa), Portugal), Valeria Bartsch (Fraunhofer ITWM, Germany), Rodrigo Rodrigues (INESC-ID & IST (ULisboa), Portugal), and Jose C. Monteiro (INESC-ID & IST (ULisboa), Portugal)</i>	
Autonomous Load Balancing in Distributed Hash Tables Using Churn and the Sybil Attack	631
<i>Andrew Rosen (Temple University, USA), Benjamin Levin (Temple University, USA), and Anu G. Bourgeois (Georgia State University, USA)</i>	
Performance Models for Hybrid Programs Accelerated by GPUs	641
<i>Aparna Sasidharan (Ansys Inc, USA)</i>	

Evaluating the Performance of Integer Sum Reduction on an Intel GPU	652
<i>Zheming Jin (Oak Ridge National Laboratory, USA) and Jeffrey Vetter (Oak Ridge National Laboratory, USA)</i>	
On the Computational Power of Convolution Pooling: A Theoretical Approach for Deep Learning	656
<i>Koji Nakano (Hiroshima University, Japan), Shotaro Aoki (Hiroshima University, Japan), Yasuaki Ito (Hiroshima University, Japan), and Akihiko Kasagi (Fujitsu Laboratories Ltd., Japan)</i>	

JSSPP: Job Scheduling Strategies for Parallel Processing

Introduction to JSSPP 2021	666
<i>Dalibor Klusáček (CESNET a.l.e.), Walfredo Cirne (Google), and Gonzalo P. Rodrigo (Apple)</i>	

PDSEC: Parallel and Distributed Scientific and Engineering Computing

Message from the PDSEC-21 Workshop Chairs	667
<i>Sabine Roller (German Aerospace Center, Germany), Peter Strazdins (The Australian National University, Australia), Srishti Srivastava (University of Southern Indiana, USA), Raphael Couturier (University Bourgogne Franche-Comte, France), Neda Ebrahimi Pour (German Aerospace Center, Germany), Suzanne Michelle Shontz (University of Kansas, USA), Thomas Rauber (University of Bayreuth, Germany), Gudula Runger (Chemnitz University of Technology, Germany), and Laurence T. Yang (St. Francis Xavier University, Canada)</i>	
Load Balancing for Distributed Nonlocal Models within Asynchronous Many-Task Systems	669
<i>Pranav Gadikar (Indian Institute of Technology Madras, India), Patrick Diehl (Center for Computation and Technology, Louisiana State University, USA), and Prashant K. Jha (Oden Institute for Computational Engineering and Sciences, The University of Texas at Austin, USA)</i>	
Scalable Hybrid Loop-and Task-Parallel Matrix Inversion for Multicore Processors	679
<i>Sandra Catalán (Universidad Complutense de Madrid, Spain), Francisco D. Igual (Universidad Complutense de Madrid, Spain), Rafael Rodríguez-Sánchez (Universidad Complutense de Madrid, Spain), and Enrique S. Quintana-Ortí (Universitat Politècnica de València, Spain)</i>	
cuFINUFFT: A Load-Balanced GPU Library for General-Purpose Nonuniform FFTs	688
<i>Yu-hsuan Shih (Courant Institute of Mathematical Sciences, New York University, USA), Garrett Wright (PACM, Princeton University, USA), Joakim Andén (KTH Royal Institute of Technology, Sweden), Johannes Blaschke (National Energy Research Scientific Computing Center, Lawrence Berkeley National Laboratory, USA), and Alex H. Barnett (Center for Computational Mathematics, Flatiron Institute, USA)</i>	

Parallel Machine Learning of Partial Differential Equations	698
<i>Amin Totounferoush (Institute for Parallel and Distributed Systems, University of Stuttgart, Germany), Neda Ebrahimi Pour (Chair of Simulation Techniques and Scientific Computing, University of Siegen, Germany), Sabine Roller (Chair of Simulation Techniques and Scientific Computing, University of Siegen, Germany), and Miriam Mehl (Institute for Parallel and Distributed Systems, University of Stuttgart, Germany)</i>	
Improving Workload Balance of a Marine CSEM Inversion Application	704
<i>Jessica Imlau Dagostini (Institute of Informatics, Federal University of Rio Grande do Sul, Brazil), Henrique Corrêa Pereira da Silva (Institute of Informatics, Federal University of Rio Grande do Sul, Brazil), Vinicius Garcia Pinto (Institute of Informatics, Federal University of Rio Grande do Sul, Brazil), Roberto Machado Velho (Institute of Informatics, Federal University of Rio Grande do Sul, Brazil), Eduardo S. L. Gastal (Institute of Informatics, Federal University of Rio Grande do Sul, Brazil), and Lucas Mello Schnorr (Institute of Informatics, Federal University of Rio Grande do Sul, Brazil)</i>	
Performance Modeling and Tuning for DFT Calculations on Heterogeneous Architectures	714
<i>Hadia Ahmed (Lawrence Berkeley National Laboratory, USA), David B. Williams-Young (Lawrence Berkeley National Laboratory, USA), Khaled Z. Ibrahim (Lawrence Berkeley National Laboratory, USA), and Chao Yang (Lawrence Berkeley National Laboratory, USA)</i>	
Parallelization of GKV Benchmark using OpenACC	723
<i>Makoto Morishita (Nagoya University, Japan), Satoshi Ohshima (Information Technology Center, Nagoya University, Japan), Takahiro Katagiri (Information Technology Center, Nagoya University, Japan), and Toru Nagai (Information Technology Center, Nagoya University, Japan)</i>	
A Flexible Research-Oriented Framework for Distributed Training of Deep Neural Networks	730
<i>Sergio Barrachina (Universitat Jaume I, Spain), Adrián Castelló (Universitat Jaume I, Spain), Mar Catalán (Universitat Jaume I, Spain), Manuel F. Dolz (Universitat Jaume I, Spain), and Jose I. Mestre (Universitat Jaume I, Spain)</i>	
Accelerated Polynomial Evaluation and Differentiation at Power Series in Multiple Double Precision	740
<i>Jan Verschelde (University of Illinois at Chicago, USA)</i>	
 iWAPT: Automatic Performance Tuning	
Introduction to iWAPT 2021	750
<i>Kazuhiko Komatsu (Tohoku University, Japan)</i>	

Evaluating I/O Acceleration Mechanisms of SX-Aurora TSUBASA	752
<i>Yuta Sasaki (Graduate School of Information Sciences, Tohoku University, Japan), Ayumu Ishizuka (Graduate School of Information Sciences, Tohoku University, Japan), Mulya Agung (Cyberscience Center, Tohoku University, Japan), and Hiroyuki Takizawa (Cyberscience Center, Tohoku University, Japan; Graduate School of Information Sciences, Tohoku University, Japan)</i>	
Efficient Parallel Multigrid Methods on Manycore Clusters with Double/Single Precision Computing	760
<i>Kengo Nakajima (Information Technology Center, The University of Tokyo/RIKEN, Japan), Takeshi Ogita (Div. of Mathematical Sciences, Tokyo Woman's Christian University, Japan), and Masatoshi Kawai (Information Technology Center, The University of Tokyo, Japan)</i>	
Automatic Selection of Tensor Decomposition for Compressing Convolutional Neural Networks: A Case Study on VGG-type Networks	770
<i>Chia-Chun Liang (National Tsing Hua University, Taiwan) and Che-Rung Lee (National Tsing Hua University, Taiwan)</i>	
A Processor Selection Method Based on Execution Time Estimation for Machine Learning Programs	779
<i>Kou Murakami (Graduate school of Information Sciences, Tohoku University, Japan), Kazuhiko Komatsu (Cyberscience Center, Tohoku University, Japan), Masayuki Sato (Graduate school of Information Sciences, Tohoku University, Japan), and Hiroaki Kobayashi (Graduate school of Information Sciences, Tohoku University, Japan)</i>	
An Auto-Tuning with Adaptation of A64 Scalable Vector Extension for SPIRAL	789
<i>Naruya Kitai (Graduate School of Informatics, Nagoya University, Japan), Daisuke Takahashi (Center for Computational Sciences, University of Tsukuba, Japan), Franz Franchetti (Carnegie Mellon University, USA), Takahiro Katagiri (Information Technology Center, Nagoya University, Japan), Satoshi Ohshima (Information Technology Center, Nagoya University, Japan), and Toru Nagai (Information Technology Center, Nagoya University, Japan)</i>	
Improving the MPI-IO Performance of Applications with Genetic Algorithm Based Auto-Tuning ..	798
<i>Ayşe Bağbaba (The High-Performance Computing Center Stuttgart (HLRS), University of Stuttgart, Germany) and Xuan Wang (The High-Performance Computing Center Stuttgart (HLRS), University of Stuttgart, Germany)</i>	
Autotuning Benchmarking Techniques: A Roofline Model Case Study	806
<i>Jacob O. Tørring (Norwegian University of Science and Technology (NTNU), Norway), Jan Christian Meyer (Norwegian University of Science and Technology (NTNU), Norway), and Anne C. Elster (Norwegian University of Science and Technology (NTNU), Norway)</i>	
Scalable Performance Prediction of Irregular Workloads in Multi-phase Particle-in-Cell Applications	816
<i>Sai P. Chenna (University of Florida, USA), Herman Lam (University of Florida, USA), Greg Stitt (University of Florida, USA), and S. Balachandar (University of Florida, USA)</i>	

SNACS: Scalable Networks for Advanced Computing Systems Workshop

Introduction to SNACS 2021	826
<i>Taylor Groves (Lawrence Berkeley National Laboratory, USA), Matthew Dosanjh (Sandia National Laboratories, USA), and Ryan Grant (Sandia National Laboratories, USA)</i>	
SNACS 2021 Keynote	827
<i>John Shalf (Lawrence Berkeley National Laboratory, USA)</i>	
User Allocation for Real-Time Applications with State Sharing in Fog Computing Networks	828
<i>Ryohei Sato (NTT Network Service Systems Laboratories, Japan), Hidetoshi Kawaguchi (NTT Network Service Systems Laboratories, Japan), and Yuichi Nakatani (NTT Network Service Systems Laboratories, Japan)</i>	
Multi-path Routing in the Jellyfish Network	832
<i>Zaid Alzaid (Florida State University, USA), Saptarshi Bhowmik (Florida State University, USA), and Xin Yuan (Florida State University, USA)</i>	

PAISE: Parallel AI and Systems for the Edge

Introduction to PAISE 2021	842
<i>Istemi Ekin Akkus (Nokia Bell Labs, Germany) and Nirmal V. Desai (IBM Research, USA)</i>	
Addressing the Constraints of Active Learning on the Edge	845
<i>Enrique Nueve (Northwestern University, USA), Sean Shahkarami (University of Chicago, USA), Seongha Park (Mathematics and Computer Science Division Argonne National Laboratory, USA), and Nicola Ferrier (Mathematics and Computer Science Division Argonne National Laboratory, USA)</i>	
Informed Prefetching in I/O Bounded Distributed Deep Learning	850
<i>Xiaojun Ruan (California State University, East Bay, USA) and Haiquan Chen (California State University, Sacramento, USA)</i>	
Performance Evaluation of Deep Learning Compilers for Edge Inference	858
<i>Gaurav Verma (Stony Brook University, USA), Yashi Gupta (Stony Brook University, USA), Abid M. Malik (Brookhaven National Laboratory, USA), and Barbara Chapman (Stony Brook University, USA; Brookhaven National Laboratory, USA)</i>	
DataVinci: Proactive Data Placement for Ad-Hoc Computing	866
<i>Martin Breitbach (University of Mannheim, Germany), Janick Edinger (University of Hamburg, Germany), Dominik Schäfer (University of Mannheim, Germany), and Christian Becker (University of Mannheim, Germany)</i>	
Pilot-Edge: Distributed Resource Management Along the Edge-to-Cloud Continuum	874
<i>Andre Luckow (RADICAL, ECE, Rutgers University, USA; Ludwig-Maximilian University, Munich, Germany; Clemson University, USA), Kartik Rattan (RADICAL, ECE, Rutgers University, USA), and Shantenu Jha (Brookhaven National Laboratory, USA; RADICAL, ECE, Rutgers University, USA)</i>	

INT Based Network-Aware Task Scheduling for Edge Computing	879
<i>Bibek Shrestha (University of Nevada, Reno, USA), Richard Cziva (Lawrence Berkeley National Laboratory, USA), and Engin Arslan (University of Nevada, Reno, USA)</i>	
Performance Comparison for Scientific Computations on the Edge via Relative Performance	887
<i>Aravind Sankaran (RWTH Aachen University, Germany) and Paolo Bientinesi (Umea Universitet, Sweden)</i>	

RADR: Resource Arbitration for Dynamic Runtimes

Introduction to RADR 2021	896
<i>Pete Beckman (Argonne National Laboratory, USA; Northwestern University, USA), Emmanuel Jeannot (TADaaM Team, Inria, France), and Swann Perarnau (Argonne National Laboratory, USA)</i>	
Dynamic Computing Resources Allocation for Multiple Deep Learning Tasks	899
<i>Liang Wei (Tokyo Institute of Technology, Japan) and Kazuyuki Shudo (Tokyo Institute of Technology, Japan)</i>	

ScaDL: Scalable Deep Learning over Parallel And Distributed Infrastructures

Message from the ScaDL 2021 Workshop Chairs	906
<i>Stacy Patterson (Rensselaer Polytechnic Institute, USA), Parijat Dube (IBM Research, USA), Danilo Ardagna (Politecnico di Milano, Italy), and Yogish Sabharwal (IBM Research, India)</i>	
ScaDL 2021 Invited Speaker-1	907
<i>Rosa Badia (Barcelona Supercomputing Center, Spain)</i>	
ScaDL 2021 Invited Speaker-2	908
<i>Torsten Hoefler (ETH Zurich, Switzerland)</i>	
ScaDL 2021 Invited Speaker-3	909
<i>Milind Tambe (Harvard University, USA; Google Research, India)</i>	
ScaDL 2021 Invited Speaker-4	910
<i>Chris Carothers (Rensselaer Polytechnic Institute, USA)</i>	
ScaDL 2021 Invited Speaker-5	911
<i>Rania Khalaf (IBM Research, USA)</i>	
ScaDL 2021 Invited Speaker-6	912
<i>Zhe Zhang (Anyscale Inc., USA)</i>	
A Distributed Multi-GPU System for Large-Scale Node Embedding at Tencent	913
<i>Wanjing Wei (Tencent Inc., China), Yangzihao Wang (Tencent Inc., China), Pin Gao (Tencent Inc., China), Shijie Sun (Tencent Inc., China), and Donghai Yu (Tencent Inc., China)</i>	
Scaling Single-Image Super-Resolution Training on Modern HPC Clusters: Early Experiences	923
<i>Quentin Anthony (The Ohio State University, USA), Lang Xu (The Ohio State University, USA), Hari Subramoni (The Ohio State University, USA), and Dhabaleswar K. Panda (The Ohio State University, USA)</i>	

Distributed Deep Learning Using Volunteer Computing-Like Paradigm	933
<i>Medha Atre (Eydle Inc), Birendra Jha (Eydle Inc), and Ashwini Rao (Eydle Inc)</i>	
Ex-NNQMD: Extreme-Scale Neural Network Quantum Molecular Dynamics	943
<i>Pankaj Rajak (Argonne Leadership Computing Facility, Argonne National Laboratory, USA), Anikeya Aditya (Collaboratory for Advanced Computing and Simulations, University of Southern California, USA), Shogo Fukushima (Kumamoto University, Japan), Rajiv K. Kalia (Collaboratory for Advanced Computing and Simulations, University of Southern California, USA), Thomas Linker (Collaboratory for Advanced Computing and Simulations, University of Southern California, USA), Kuang Liu (Collaboratory for Advanced Computing and Simulations, University of Southern California, USA), Ye Luo (Argonne Leadership Computing Facility, Argonne National Laboratory, USA), Aiichiro Nakano (Collaboratory for Advanced Computing and Simulations, University of Southern California, USA), Ken-ichi Nomura (Collaboratory for Advanced Computing and Simulations, University of Southern California, USA), Kohei Shimamura (Kumamoto University, Japan), Fuyuki Shimojo (Kumamoto University, Japan), and Priya Vashishta (Collaboratory for Advanced Computing and Simulations, University of Southern California, USA)</i>	
Training EfficientNets at Supercomputer Scale: 83% ImageNet Top-1 Accuracy in One Hour	947
<i>Arissa Wongpanich (University of California, Berkeley, USA; Google Research, USA), Hieu Pham (Google Research, USA), James Demmel (University of California, Berkeley, USA), Mingxing Tan (Google Research, USA), Quoc Le (Google Research, USA), Yang You (National University of Singapore, Singapore), and Sameer Kumar (Google Research, USA)</i>	
Performance Analysis of Deep Learning Workloads on a Composable System	951
<i>Kaoutar El Maghraoui (IBM Research, USA), Lorraine M. Herger (IBM Research, USA), Chekuri Choudary (IBM Systems, USA), Kim Tran (IBM Systems, USA), Todd Deshane (IBM Research, USA), and David Hanson (IBM Research, USA)</i>	

HPS: High-Performance Storage

Message from the HPS 2021 Workshop Chairs	955
<i>Gabriel Antoniu (Inria, France) and Marc Snir (University of Illinois at Urbana Champaign, USA)</i>	
HPS 2021 Keynote Speaker	957
<i>Glenn K. Lockwood (National Energy Research Scientific Computing Center (NERSC) at Lawrence Berkeley National Laboratory, USA)</i>	
HPS 2021 Invited Speaker-1	958
<i>Takuya Okamoto (Fujitsu Ltd., Japan)</i>	
HPS 2021 Invited Speaker-2	959
<i>Jay Lofstead (Sandia National Laboratories, USA)</i>	

Facilitating Staging-Based Unstructured Mesh Processing to Support Hybrid In-Situ Workflows	960
<i>Zhe Wang (Rutgers University, USA), Pradeep Subedi (Rutgers University, USA), Matthieu Dorier (Argonne National Laboratory, USA), Philip E. Davis (Rutgers University, USA), and Manish Parashar (University of Utah, USA)</i>	

Exploring MPI Collective I/O and File-per-Process I/O for Checkpointing a Logical Inference Task	965
<i>Ke Fan (University of Alabama at Birmingham, USA), Kristopher Micinski (Syracuse University, USA), Thomas Gilray (University of Alabama at Birmingham, USA), and Sidharth Kumar (University of Alabama at Birmingham, USA)</i>	

ParSocial: Parallel and Distributed Processing for Computational Social Systems

Message from the ParSocial 2021 Workshop Co-Chairs	973
<i>John Korah (California State Polytechnic University Pomona, USA) and Eunice E. Santos (University of Illinois at Urbana-Champaign, USA)</i>	

Memory Efficient Edge Addition Designs for Large and Dynamic Social Networks	975
<i>Eunice E. Santos (University of Illinois at Urbana-Champaign, USA), Vairavan Murugappan (University of Illinois at Urbana-Champaign, USA), and John Korah (California State Polytechnic University, Pomona, USA)</i>	

Load Balancing Schemes for Large Synthetic Population-Based Complex Simulators	985
<i>Bogdan Mucenic (The University of Chicago, USA), Chaitanya Kaligotla (Simon Fraser University, Canada), Abby Stevens (The University of Chicago, USA), Jonathan Ozik (Argonne National Laboratory, USA), Nicholson Collier (Argonne National Laboratory, USA), and Charles Macal (Argonne National Laboratory, USA)</i>	

Application of Distributed Agent-Based Modeling to Investigate Opioid Use Outcomes in Justice Involved Populations	989
<i>Eric Tatara (University of Chicago, USA), John Schneider (University of Chicago, USA), Madeline Quasebarth (University of Chicago, USA), Nicholson Collier (University of Chicago, USA), Harold Pollack (University of Chicago, USA), Basmattee Boodram (University of Illinois at Chicago, USA), Sam Friedman (New York University, USA), Elizabeth Salisbury-Afshar (University of Wisconsin-Madison, USA), Mary Ellen Mackesy-Amiti (University of Illinois at Chicago, USA), and Jonathan Ozik (University of Chicago, USA)</i>	

Shared-Memory Scalable k-Core Maintenance on Dynamic Graphs and Hypergraphs	998
<i>Kasimir Gabert (Georgia Institute of Technology, USA), Ali Pinar (Sandia National Laboratories, USA), and Ümit V. Çatalyürek (Georgia Institute of Technology, USA)</i>	

P-Flee: An Efficient Parallel Algorithm for Simulating Human Migration	1008
<i>Petros Anastasiadis (Computing Systems Laboratory, National Technical University of Athens, Greece), Sergiy Gogolenko (High Performance Computing Center Stuttgart, Germany), Nikela Papadopoulou (Computing Systems Laboratory, National Technical University of Athens, Greece), Marcin Lawenda (Poznan Supercomputing and Networking Center, Poland), Hamid Arabnejad (Brunel University London, UK), Alireza Jahani (Brunel University London, UK), Imran Mahmood (Brunel University London, UK), and Derek Groen (Brunel University London, UK)</i>	

IPDPS 2021 PhD Forum

IPDPS 2021 PhD Forum Welcome and Abstracts	1012
<i>Sanjukta Bhowmick (University of North Texas, USA) and Akshaye Dhawan (Bloomberg L.P., USA)</i>	

Author Index