

2020 IEEE International Parallel and Distributed Processing Symposium Workshops (IPDPSW 2020)

**New Orleans, Louisiana, USA
18-22 May 2020**

Pages 1-566



**IEEE Catalog Number: CFP2051J-POD
ISBN: 978-1-7281-7457-0**

**Copyright © 2020 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:	CFP2051J-POD
ISBN (Print-On-Demand):	978-1-7281-7457-0
ISBN (Online):	978-1-7281-7445-7

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

CURRAN ASSOCIATES INC.
proceedings
.com

2020 IEEE International Parallel and Distributed Processing Symposium Workshops (IPDPSW) **IPDPSW 2020**

Table of Contents

Message from the 2020 General Co-Chairs .xxiv.....
Message from the 2020 Workshops Chair and Vice-chair .xxvii.....

HCW: Heterogeneity in Computing Workshop

Introduction to HCW 2020 .1.....
John K. Antonio (University of Oklahoma)

Message from the HCW Steering Committee Chair .2.....
Behrooz Shirazi (Washington State University)

Message from the HCW General Chair .3.....
John K. Antonio (University of Oklahoma)

Message from the HCW Technical Program Committee Chair .4.....
Florina M. Ciorba (University of Basel)

HCW 2020 Keynote Address .5.....
Albert Y. Zomaya (University of Sydney)

MigHEFT: DAG-Based Scheduling of Migratable Tasks on Heterogeneous Compute Nodes .6.....
*Achim Lösch (Computer Engineering Group, Paderborn University) and
Marco Platzner (Computer Engineering Group, Paderborn University)*

Autonomous Task Dropping Mechanism to Achieve Robustness in Heterogeneous Computing
Systems .17.....
*Ali Mokhtari (University of Louisiana at Lafayette), Chavit Denninnart
(University of Louisiana at Lafayette), and Mohsen Amini Salehi
(University of Louisiana at Lafayette)*

I/O Performance of the SX-Aurora TSUBASA .27.....
*Mitsuo Yokokawa (Kobe University), Ayano Nakai (Kobe University),
Kazuhiko Komatsu (Tohoku University), Yuta Watanabe (NEC Corporation),
Yasuhisa Masaoka (NEC Corporation), Yoko Isobe (NEC Corporation), and
Hiroaki Kobayashi (Tohoku University)*

Enabling Domain-Specific Architectures with an Open-Source Soft-Core GPGPU .36.....	
	<i>Marcelo Brandalero (Brandenburg University of Technology (B-TU)), Hector Gerardo Muñoz Hernandez (Brandenburg University of Technology (B-TU)), Mitko Veleski (Brandenburg University of Technology (B-TU)), Muhammed Al Kadi (Ruhr-University Bochum (RUB)), Paolo Rech (Universidade Federal do Rio Grande do Sul (UFRGS)), and Michael Hübner (Brandenburg University of Technology (B-TU))</i>
User-Space Emulation Framework for Domain-Specific SoC Design .44.....	
	<i>Joshua Mack (University of Arizona), Nirmal Kumbhare (University of Arizona), Anish NK (Arizona State University), Umit Y. Ogras (Arizona State University), and Ali Akoglu (University of Arizona)</i>
Improving Inference Latency and Energy of Network-on-Chip Based Convolutional Neural Networks through Weights Compression .54.....	
	<i>Giuseppe Ascia (University of Catania), Vincenzo Catania (University of Catania), John Jose (Indian Institute of Technology Guwahati), Salvatore Monteleone (CY Cergy Paris Université), Maurizio Palesi (University of Catania), and Davide Patti (University of Catania)</i>

RAW: Reconfigurable Architectures Workshop

RAW 2020 Introduction and Committees .64.....	
	<i>Marco D. Santambrogio (Politecnico di Milano) and Yu Wang (Tsinghua University)</i>
CNN-Based Monocular Decentralized SLAM on Embedded FPGA .66.....	
	<i>Jincheng Yu (Tsinghua University, Beijing, China), Feng Gao (Tsinghua University, Beijing, China), Jianfei Cao (Hefei University of Technology, Hefei, China), Chao Yu (Tsinghua University, Beijing, China), Zhaoliang Zhang (Tsinghua University, Beijing China), Zhengfeng Huang (Hefei University of Technology, Hefei, China), Yu Wang (Tsinghua University, Beijing, China), and Huazhong Yang (Tsinghua University, Beijing, China)</i>
Improving HLS Generated Accelerators Through Relaxed Memory Access Scheduling .74.....	
	<i>Johanna Rohde (Technische Universität Darmstadt), Karsten Müller (Technische Universität Darmstadt), and Christian Hochberger (Technische Universität Darmstadt)</i>
Real-Time Automatic Modulation Classification using RFSoc .82.....	
	<i>Stephen Tridgell (The University of Sydney), David Boland (The University of Sydney), Philip H.W. Leong (The University of Sydney), Ryan Kastner (University of California, San Diego), Alireza Khodamoradi (University of California, San Diego), and Siddhartha Siddhartha (The University of Sydney)</i>
FPGA Based Emulation Environment for Neuromorphic Architectures .90.....	
	<i>Spencer Valancius (University of Arizona), Edward Richter (University of Arizona), Ruben Purdy (University of Arizona), Kris Rockowitz (University of Arizona), Michael Inouye (University of Arizona), Joshua Mack (University of Arizona), Nirmal Kumbhare (University of Arizona), Kaitlin Fair (Air Force Research Labs), John Mixer (Raytheon Missile Systems), and Ali Akoglu (University of Arizona)</i>

PHRYCTORIA: A Messaging System for Transprecision OpenCAPI-Attached FPGA Accelerators .98	
	<i>Dionysios Diamantopoulos (IBM Research Europe), Mitra Purandare (IBM Research Europe), Burkhard Ringlein (IBM Research Europe), and Christoph Hagleitner (IBM Research Europe)</i>
QTAccel: A Generic FPGA Based Design for Q-Table Based Reinforcement Learning Accelerators.107	
	<i>Yuan Meng (University of Southern California), Sanmukh Kuppannagari (University of Southern California), Rachit Rajat (University of Southern California), Ajitesh Srivastava (University of Southern California), Rajgopal Kannan (US Army Research Lab-West), and Viktor Prasanna (University of Southern California)</i>
SpiderWeb - High Performance FPGA NoC .115.....	
	<i>Martin Langhammer (Intel Corporation), Gregg Baeckler (Intel Corporation), and Sergey Gribok (Intel Corporation)</i>
Secure Acceleration on Cloud-Based FPGAs - FPGA Enclaves .119.....	
	<i>Niklas Lindskog (Ericsson) and Håkan Englund (Ericsson)</i>
A FPGA-Based Post-Processing and Validation Platform for Random Number Generators .123.....	
	<i>Laurent Gantel (University of Applied Sciences and Arts Western Switzerland (HES-SO/HEPIA)), Alexandre Duc (University of Applied Sciences and Arts Western Switzerland (HES-SO/HEIG-VD)), Lucie Steiner (University of Applied Sciences and Arts Western Switzerland (HES-SO/HEIG-VD)), Fabien Vannel (University of Applied Sciences and Arts Western Switzerland (HES-SO/HEPIA)), Andres Upegui (University of Applied Sciences and Arts Western Switzerland (HES-SO/HEPIA)), and Florent Gluck (University of Applied Sciences and Arts Western Switzerland (HES-SO/HEPIA))</i>
An Interval-Based Mapping Algorithm for Multi-shape Tasks on Dynamic Partial Reconfigurable FPGAs .127.....	
	<i>Tingyu Zhou (Waseda University, Japan), Tiejuan Pan (DENSO Corporation, Kariya, Japan), Michael Conrad Meyer (Waseda University, Japan), Yiping Dong (China Key System & Integrated Circuit Co., Ltd, Wuxi, China), and Takahiro Watanabe (Waseda University, Japan)</i>
EMPhASIS: An EMbedded Public Attention Stress Identification System .131.....	
	<i>Jessica Leoni (Politecnico di Milano), Asia Ciallella (Politecnico di Milano), Luca Stornaiuolo (Politecnico di Milano), Marco Santambrogio (Politecnico di Milano), and Donatella Sciuto (Politecnico di Milano)</i>
Hardware Resources Analysis of BNNs Splitting for FARD-Based Multi-FPGAs Distributed Systems .135.....	
	<i>Giorgia Fiscaletti (Politecnico di Milano), Marco Speziali (Politecnico di Milano), Luca Stornaiuolo (Politecnico di Milano), Marco Santambrogio (Politecnico di Milano), and Donatella Sciuto (Politecnico di Milano)</i>
A Microcode-Based Control Unit for Deep Learning Processors .139.....	
	<i>Qian Zhao (Kyushu Institute of Technology), Yasuhiro Nakahara (Kumamoto University), Motoki Amagasaki (Kumamoto University), Masahiro Iida (Kumamoto University), and Takaichi Yoshida (Kyushu Institute of Technology)</i>

Fast Monocular Depth Estimation on an FPGA .143.....	
	<i>Youki Sada (Tokyo Institute of Technology), Naoto Soga (Tokyo Institute of Technology), Masayuki Shimoda (Tokyo Institute of Technology), Akira Jinguji (Tokyo Institute of Technology), Shimpei Sato (Tokyo Institute of Technology), and Hiroki Nakahara (Tokyo Institute of Technology)</i>
SALSA: A Domain Specific Architecture for Sequence Alignment .147.....	
	<i>Lorenzo Di Tucci (Politecnico di Milano), Riyadh Baghdadi (Massachusetts Institute of Technology), Saman Amarasinghe (Massachusetts Institute of Technology), and Marco Domenico Santambrogio (Politecnico di Milano)</i>
Optimizing OpenCL Kernels and Runtime for DNN Inference on FPGAs .151.....	
	<i>Seung-Hun Chung (University of Toronto) and Tarek S. Abdelrahman (University of Toronto)</i>
Leveraging Succinct Data Structures for DNA Sequence Mapping on FPGA .155.....	
	<i>Guido Walter Di Donato (Politecnico di Milano), Alberto Zeni (Politecnico di Milano), Lorenzo Di Tucci (Politecnico di Milano), and Marco Domenico Santambrogio (Politecnico di Milano)</i>

HiCOMB: High Performance Computational Biology

HiCOMB 2020 Introduction and Committees .159.....	
	<i>Alba Cristina M. A. de Melo (University of Brasilia) and Ananth Kalyanaraman (Washington State University)</i>
A Tropical Semiring Multiple Matrix-Product Library on GPUs: (not Just) a Step Towards RNA-RNA Interaction Computations .160.....	
	<i>Brandon Gildemaster (Colorado State University, Fort Collins, CO, USA), Prerana Ghalsasi (Colorado State University), and Sanjay Rajopadhye (Colorado State University)</i>
Fast and High Quality Graph Alignment via Treelets .170.....	
	<i>Morgan Lee (Rensselaer Polytechnic Institute) and George M Slota (Rensselaer Polytechnic Institute)</i>
GPU Accelerated Partial Order Multiple Sequence Alignment for Long Reads Self-Correction .174..	
	<i>Francesco Peverelli (Politecnico di Milano), Lorenzo Di Tucci (Politecnico di Milano), Marco Domenico Santambrogio (Politecnico di Milano), Nan Ding (Lawrence Berkeley National Laboratory), Steven Hofmeyr (Lawrence Berkeley National Laboratory), Aydin Buluc (University of California at Berkeley), Leonid Olikier (Lawrence Berkeley National Laboratory), and Katherine Yelick (University of California at Berkeley)</i>
Optimizing High-Performance Computing Systems for Biomedical Workloads .183.....	
	<i>Patricia Kovatch (Icahn School of Medicine at Mount Sinai), Lili Gai (Icahn School of Medicine at Mount Sinai), Hyung Min Cho (Icahn School of Medicine at Mount Sinai), Eugene Fluder (Icahn School of Medicine at Mount Sinai), and Dansha Jiang (Icahn School of Medicine at Mount Sinai)</i>

Kcollections: A Fast and Efficient Library for K-Mers .193.....
M. Stanley Fujimoto (Brigham Young University), Cole A. Lyman (Brigham Young University), and Mark J. Clement (Brigham Young University)

GrAPL: Graphs, Architectures, Programming, and Learning

GrAPL 2020 Introduction and Committees .199.....
Scott McMillan (CMU SEI) and Manoj Kumar (IBM)

GrAPL Keynote 1 .201.....
George Karypis (AWS Deep Learning Science & University of Minnesota)

GrAPL Keynote 2 .202.....
Saman Amarasinghe (Massachusetts Institute of Technology)

An Incremental GraphBLAS Solution for the 2018 TTC Social Media Case Study .203.....
Márton Elekes (Budapest University of Technology and Economics) and Gábor Szárnyas (Budapest University of Technology and Economics)

75,000,000,000 Streaming Inserts/Second Using Hierarchical Hypersparse GraphBLAS Matrices .207
Jeremy Kepner (MIT Lincoln Laboratory Supercomputing Center, MIT Computer Science & AI Laboratory, MIT Mathematics Department), Tim Davis (Texas A&M), Chansup Byun (MIT Lincoln Laboratory Supercomputing Center), William Arcand (MIT Lincoln Laboratory Supercomputing Center), David Bestor (MIT Lincoln Laboratory Supercomputing Center), William Bergeron (MIT Lincoln Laboratory Supercomputing Center), Vijay Gadepally (MIT Lincoln Laboratory Supercomputing Center, MIT Computer Science & AI Laboratory), Matthew Hubbell (MIT Lincoln Laboratory Supercomputing Center), Michael Houle (MIT Lincoln Laboratory Supercomputing Center), Michael Jones (MIT Lincoln Laboratory Supercomputing Center), Anna Klein (MIT Lincoln Laboratory Supercomputing Center), Peter Michaleas (MIT Lincoln Laboratory Supercomputing Center), Lauren Milechin (MIT), Julie Mullen (MIT Lincoln Laboratory Supercomputing Center), Andrew Prout (MIT Lincoln Laboratory Supercomputing Center), Antonio Rosa (MIT Lincoln Laboratory Supercomputing Center), Siddharth Samsi (MIT Lincoln Laboratory Supercomputing Center), Charles Yee (MIT Lincoln Laboratory Supercomputing Center), and Albert Reuther (MIT Lincoln Laboratory Supercomputing Center)

Parallelizing Maximal Clique Enumeration on Modern Manycore Processors .211.....
Jovan Blanuša (IBM Research Zurich; EPFL), Radu Stoica (IBM Research Zurich), Paolo Ienne (EPFL), and Kubilay Atasu (IBM Research Zurich)

Considerations for a Distributed GraphBLAS API .215.....	
	<i>Benjamin Brock (University of California, Berkeley, CA), Aydın Buluc (Computational Research Department, Lawrence Berkeley National Laboratory, Berkeley, CA), Timothy Mattson (Parallel Computing Labs, Intel, Hillsboro, OR), Scott McMillan (Software Engineering Institute, Carnegie Mellon University, Pittsburg, PA), Jose Moreira (IBM Thomas J. Watson Research Center, Yorktown Heights, NY), Roger Pearce (Center for Applied Scientific Computing, Lawrence Livermore National Laboratory, Livermore, CA), Oguz Selvitopi (Computational Research Department, Lawrence Berkeley National Laboratory, Berkeley, CA), and Trevor Steil (University of Minnesota, Minneapolis, MN)</i>
A Roadmap for the GraphBLAS C++ API .219.....	
	<i>Benjamin Brock (University of California, Berkeley, CA), Aydın Buluc (Lawrence Berkeley National Laboratory, Berkeley, CA), Timothy Mattson (Parallel Computing Labs, Intel, Hillsboro, OR), Scott McMillan (Software Engineering Institute, Carnegie Mellon University, Pittsburg, PA), and Jose Moreira (IBM Thomas J. Watson Research Center, Yorktown Heights, NY)</i>
Linear Algebraic Louvain Method in Python .223.....	
	<i>Tze Meng Low (Carnegie Mellon University), Daniele G. Spampinato (Carnegie Mellon University), Scott McMillan (Software Engineering Institute, Carnegie Mellon University), and Michel Pelletier (Graphegon, Inc.)</i>
A Scalable Graph Generation Algorithm to Sample over a Given Shell Distribution .227.....	
	<i>M. Yusuf Özkaya (Georgia Institute of Technology), M. Fatih Balm (Georgia Institute of Technology), Ali Pınar (Sandia National Laboratories), and Ümit V. Çatalyürek (Georgia Institute of Technology)</i>
Kronecker Graph Generation with Ground Truth for 4-Cycles and Dense Structure in Bipartite Graphs .237.....	
	<i>Trevor Steil (University of Minnesota), Scott McMillan (Carnegie Mellon University), Geoffrey Sanders (Lawrence Livermore National Laboratory), Roger Pearce (Lawrence Livermore National Laboratory), and Benjamin Priest (Lawrence Livermore National Laboratory)</i>

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

EduPar 2020 Introduction and Committees .247.....	
	<i>Sushil K. Prasad (The University of Texas at San Antonio)</i>
EduPar-20 Keynote Speaker .250.....	
	<i>Margaret Martonosi (National Science Foundation and Princeton University)</i>
EduPar-20 Invited Panel .251.....	
	<i>Henry A. Gabb (Intel), Andrew Lumsdaine (Northwest Institute for Advanced Computing (NIAC)), Margaret Martonosi (US National Science Foundation's (NSF)), and Arnold L. Rosenberg (Northeastern University)</i>

Retrospective: A Look Back at 20+ Years of Experience in Parallel Computing Education .252.....	
	<i>Joel Adams (Calvin University)</i>
A Framework for the Evaluation of Parallel and Distributed Computing Educational Resources.261	
	<i>David W. Brown (Elmhurst College), Vitaly Ford (Arcadia University), and Sheikh K. Ghafoor (Tennessee Tech University)</i>
NumbaSummarizer: A Python Library for Simplified Vectorization Reports .269.....	
	<i>Neftali Watkinson (University of California, Irvine), Preston Tai (University of California, Irvine), Alexandru Nicolau (University of California, Irvine), and Alexander Veidenbaum (University of California, Irvine)</i>
EasyPAP: A Framework for Learning Parallel Programming .276.....	
	<i>Alice Lasserre (University of Bordeaux), Raymond Namyst (University of Bordeaux, Inria), and Pierre-Andre Wacrenier (University of Bordeaux, Inria)</i>
PDUnplugged: A Free Repository of Unplugged Parallel & Distributed Computing Activities .284	
	<i>Suzanne Matthews (United States Military Academy)</i>
Teaching Modern Multithreading in CS2 with Actors .292.....	
	<i>Mark Lewis (Trinity University) and Lisa Lacher (University of Houston-Clear Lake)</i>
Teaching Cloud Computing: Motivations, Challenges and Tools .300.....	
	<i>Cosimo Anglano (University of Piemonte Orientale), Massimo Canonico (University of Piemonte Orientale), and Marco Guazzone (University of Piemonte Orientale)</i>
Using Embedded Xinu and the Raspberry Pi 3 to Teach Operating Systems .307.....	
	<i>Patrick McGee (Marquette University), Rade Latinovich (Marquette University), and Dennis Brylow (Marquette University)</i>

HIPS: High-level Parallel Programming Models and Supportive Environments

HIPS 2020 Introduction and Committees .316.....	
	<i>Dong Li (University of California, Merced) and Heike Jagode (University of Tennessee, Knoxville)</i>
Compile-Time Parallelization of Subscripted Subscript Patterns .317.....	
	<i>Akshay Bhosale (University of Delaware) and Rudolf Eigenmann (University of Delaware)</i>
Online Scheduling with Redirection for Parallel Jobs .326.....	
	<i>Adrien Faure (Atos, Univ. Grenoble Alpes, CNRS, INRIA, Grenoble INP, LIG), Giorgio Lucarelli (LCOMS, University of Lorraine), Olivier Richard (Univ. Grenoble Alpes, CNRS, INRIA, Grenoble INP, LIG), and Denis Trystram (Univ. Grenoble Alpes, CNRS, INRIA, Grenoble INP, LIG)</i>

Performance Portability Evaluation of OpenCL Benchmarks Across Intel and NVIDIA Platforms .330	
<i>Colleen Bertoni (Argonne National Laboratory), JaeHyuk Kwack (Argonne National Laboratory), Thomas Applencourt (Argonne National Laboratory), Yasaman Ghadar (Argonne National Laboratory), Brian Homerding (Argonne National Laboratory), Christopher Knight (Argonne National Laboratory), Brice Videau (Argonne National Laboratory), Huihuo Zheng (Argonne National Laboratory), Vitali Morozov (Argonne National Laboratory), and Scott Parker (Argonne National Laboratory)</i>	
Scalable Crash Consistency for Staging-Based In-Situ Scientific Workflows .340.....	
<i>Shaohua Duan (Rutgers Discovery Informatics Institute, Rutgers University) and Manish Parashar (Rutgers Discovery Informatics Institute, Rutgers University)</i>	
Automatic Selection of Tuning Plugins in PTF Using Machine Learning .349.....	
<i>Robert Mijakovic (Leibniz Supercomputing Centre (LRZ) of the Bavarian Academy of Sciences and Humanities) and Michael Gerndt (Technische Universität München, Institut für Informatik, Garching, Germany)</i>	
Porting a Legacy CUDA Stencil Code to oneAPI .359.....	
<i>Steffen Christgau (Zuse Institute Berlin) and Thomas Steinke (Zuse Institute Berlin)</i>	
A Case Study on the HACCmk Routine in SYCL on Integrated Graphics .368.....	
<i>Zheming Jin (Argonne National Lab), Vitali Morozov (Argonne National Lab), and Hal Finkel (Argonne National Lab)</i>	
Enhancing Java Streams API with PowerList Computation .375.....	
<i>Virginia Niculescu (Babes-Bolyai University of Cluj-Napoca, Romania), Darius Bufnea (Babes-Bolyai University of Cluj-Napoca, Romania), and Adrian Sterca (Babes-Bolyai University of Cluj-Napoca, Romania)</i>	

HPBDC: High-Performance Big Data and Cloud Computing

HPBDC 2020 Introduction and Committees .385.....	
<i>Xiaoyi Lu (The Ohio State University) and Jianfeng Zhan (Institute of Computing Technology, Chinese Academy of Sciences, China)</i>	
Two-Pass Softmax Algorithm .386.....	
<i>Marat Dukhan (Google Research) and Artsiom Ablavatski (Google Research)</i>	
Smart Streaming: A High-Throughput Fault-Tolerant Online Processing System .396.....	
<i>Jia Guo (the Ohio State University) and Gagan Agrawal (Augusta University)</i>	
Parallel Query Service for Object-Centric Data Management Systems .406.....	
<i>Houjun Tang (Lawrence Berkeley National Laboratory), Suren Byna (Lawrence Berkeley National Laboratory), Bin Dong (Lawrence Berkeley National Laboratory), and Quincey Koziol (Lawrence Berkeley National Laboratory)</i>	

Pinocchio: A Blockchain-Based Algorithm for Sensor Fault Tolerance in Low Trust Environment .416.....
Chen Zeng (SKL of Computer Architecture, Institute of Computing Technology, CAS; University of Chinese Academy of Sciences), Yifan Wang (SKL of Computer Architecture, Institute of Computing Technology, CAS; University of Chinese Academy of Sciences), Fan Liang (SKL of Computer Architecture, Institute of Computing Technology, CAS, Beijing, China), and Xiaohui Peng (SKL of Computer Architecture, Institute of Computing Technology, CAS; University of Chinese Academy of Sciences)

Scaling Optimizations for Large-Scale Distributed Data with Lightweight Coresets .426.....
Daniel Pinheiro (Universidade Federal do Rio Grande do Norte), Samuel Xavier-de-Souza (Universidade Federal do Rio Grande do Norte), and Daniel Aloise (Polytechnique Montreal)

AsHES: Accelerators and Hybrid Exascale Systems

AsHES 2020 Introduction and Committees .430.....
Min Si (Argonne National Laboratory)

AsHES 2020 Keynote Speaker .431.....
Taisuke Boku (University of Tsukuba)

Population Count on Intel® CPU, GPU, and FPGA .432.....
Zheming Jin (Argonne National Lab) and Hal Finkel (Argonne National Lab)

SPHYNX: Spectral Partitioning for HYbrid aNd aXelerator-Enabled Systems .440.....
Seher Acer (Sandia National Laboratories), Erik G. Boman (Sandia National Laboratories), and Sivasanakaran Rajamanickam (Sandia National Laboratories)

Performance Evaluation of Pipelined Communication Combined with Computation in OpenCL Programming on FPGA .450.....
Norihisa Fujita (Center for Computational Sciences University of Tsukuba, Tsukuba, Japan), Ryohei Kobayashi (Center for Computational Sciences University of Tsukuba, Tsukuba, Japan), Yoshiki Yamaguchi (University of Tsukuba, Tsukuba, Japan), Tomohiro Ueno (RIKEN Center for Computational Sciences, Kobe, Japan), Kentaro Sano (RIKEN Center for Computational Sciences, Kobe, Japan), and Taisuke Boku (Center for Computational Sciences University of Tsukuba, Tsukuba, Japan)

In-depth Optimization with the OpenACC-to-FPGA Framework on an Arria 10 FPGA .460.....
Jacob Lambert (University of Oregon), Seyong Lee (Oak Ridge National Laboratory), Jeffrey S. Vetter (Oak Ridge National Laboratory), and Allen Malony (University of Oregon)

Unified Data Movement for Offloading Charm++ Applications .471.....
Matthias Diener (University of Illinois at Urbana-Champaign) and Laxmikant Kale (University of Illinois at Urbana-Champaign)

Towards Automated Kernel Selection in Machine Learning Systems: A SYCL Case Study .475.....
John Lawson (Codeplay Software)

Understanding the Performance of Elementary NLA Kernels in FPGAs .479.....	
	<i>Federico Favaro (Universidad de la República, Uruguay), Juan Pablo Oliver (Universidad de la República, Uruguay), Ernesto Dufrechou (Universidad de la República, Uruguay), and Pablo Ezzatti (Universidad de la República, Uruguay)</i>
Scalability of Sparse Matrix Dense Vector Multiply (SpMV) on a Migrating Thread Architecture .483.....	
	<i>Brian A. Page (University of Notre Dame) and Peter Kogge (University of Notre Dame)</i>

PDCO: Parallel / Distributed Combinatorics and Optimization

PDCO 2020 Introduction and Committees .489.....	
	<i>Grégoire Danoy (University of Luxembourg) and Didier El Baz (LAAS-CNRS)</i>
Load Balancing Run-Times and Space Usage for Computing the Power Set .490.....	
	<i>Roger L Goodwin (Retired)</i>
Implementing Central Force Optimization on the Intel Xeon Phi .502.....	
	<i>Thomas Charest (Bowling Green State University) and Robert Green (Bowling Green State University)</i>
Parallel/Distributed Implementation of Cellular Training for Generative Adversarial Neural Networks .512.....	
	<i>Emiliano Perez (Universidad de la Republica), Sergio Nesmachnow (Universidad de la Republica), Jamal Toutouh (Massachusetts Institute of Technology), Erik Hemberg (Massachusetts Institute of Technology), and Una-May O'Reilly (Massachusetts Institute of Technology)</i>
Predicting Near-Optimal Skin Distance in Verlet Buffer Approach for Discrete Element Method .519.....	
	<i>Abdoul Wahid Mainassara Chekaraou (University of Luxembourg.), Xavier Besseron (University of Luxembourg), Alban Rousset (University of Luxembourg), Emmanuel Kieffer (University of Luxembourg), and Bernhard Peters (University of Luxembourg)</i>
Competitive Evolution of a UAV Swarm for Improving Intruder Detection Rates .528.....	
	<i>Daniel H. Stolfi (SnT, University of Luxembourg), Matthias R. Brust (SnT, University of Luxembourg), Grégoire Danoy (SnT and the FSTM/DCS, University of Luxembourg), and Pascal Bouvry (SnT and the FSTM/DCS, University of Luxembourg)</i>

APDCM: Advances in Parallel and Distributed Computational Models

APDCM 2020 Introduction and Committees .536.....	
	<i>Jacir L. Bordim (University of Brasilia) and Koji Nakano (Hiroshima University)</i>
Debugging Strongly-Compartmentalized Distributed Systems .538.....	
	<i>Henry Zhu (University of Pennsylvania), Nik Sultana (University of Pennsylvania), and Boon Thau Loo (University of Pennsylvania)</i>

An Efficient Multicore CPU Implementation for Convolution-Pooling Computation in CNNs .548.	
	<i>Hiroki Kataoka (Hiroshima University), Kohei Yamashita (Hiroshima University), Yasuaki Ito (Hiroshima University), Koji Nakano (Hiroshima University), Akihiko Kasagi (Fujitsu Laboratories Ltd.), and Tsuguchika Tabaru (Fujitsu Laboratory. Ltd.)</i>
A Work-Time Optimal Parallel Exhaustive Search Algorithm for the QUBO and the Ising Model, with GPU Implementation .557.....	
	<i>Masaki Tao (Hiroshima University), Koji Nakano (Hiroshima University), Yasuaki Ito (Hiroshima University), Ryota Yasudo (Hiroshima University), Masaru Tatekawa (NTT DATA Corporation), Ryota Katsuki (NTT DATA Corporation), Takashi Yazane (NTT DATA Corporation), and Yoko Inaba (NTT DATA Corporation)</i>
Design and Comparison of Resilient Scheduling Heuristics for Parallel Jobs .567.....	
	<i>Anne Benoit (ENS Lyon, France), Valentin Le Fèvre (ENS Lyon, France), Padma Raghavan (Vanderbilt University, USA), Yves Robert (ENS Lyon, France), and Hongyang Sun (Vanderbilt University, USA)</i>
Optimizing Memory Access in TCF Processors with Compute-Update Operations .577.....	
	<i>Martti Forsell (VTT), Jussi Roivainen (VTT), and Jesper Larsson Träff (Vienna University of Technology)</i>
TOSS: A Topology-Based Scheduler for Storm Clusters .587.....	
	<i>Yi Zhou (Columbus State University), Yangyang Liu (Auburn University), Chaowei Zhang (Auburn University), Xiaopu Peng (Auburn University), and Xiao Qin (Auburn University)</i>
Revisiting Dynamic DAG Scheduling Under Memory Constraints for Shared-Memory Platforms .597	
	<i>Gabriel Bathie (ENS Lyon), Loris Marchal (ENS Lyon), Yves Robert (ENS Lyon), and Samuel Thibault (Inria)</i>
Optimal Randomized Complete Visibility on a Grid for Asynchronous Robots with Lights .607.....	
	<i>Gokarna Sharma (Kent State University), Ramachandran Vaidyanathan (Louisiana State University), and Jerry L. Trahan (Louisiana State University)</i>
An Initial Assessment of NVSHMEM for High Performance Computing .617.....	
	<i>Chung-Hsing Hsu (Oak Ridge National Laboratory), Neena Imam (Oak Ridge National Laboratory), Akhil Langer (NVIDIA Corporation), Sreeram Potluri (NVIDIA Corporation), and Chris J. Newburn (NVIDIA Corporation)</i>
A Model Checking Method for Secure Routing Protocols by SPIN with State Space Reduction .627..	
	<i>Hideharu Kojima (Osaka University) and Naoto Yanai (Osaka University)</i>
Methods and Experiences for Developing Abstractions for Data-Intensive, Scientific Applications .636.....	
	<i>Andre Luckow (Ludwig-Maximilian University, Munich, Germany) and Shantenu Jha (Rutgers University)</i>

JSSPP 2020 – 23rd Workshop on Job Scheduling Strategies for Parallel Processing

JSSPP 2020 Introduction and Committees .646.....
Dalibor Klusáček (CESNET), Walfredo Cirne (Google), and Narayan Desai (Google)

CHIUIW: Chapel Implementers and Users Workshop

CHIUIW 2020 Introduction and Committees .648.....
Benjamin Robbins (Cray, a Hewlett Packard Enterprise Company)

CHIUIW 2020 Keynote .650.....
William Reus (U. S. Department of Defense)

Development of Parallel CFD Applications on Distributed Memory with Chapel .651.....
Mathieu Parenteau (Polytechnique Montreal), Simon Bourgault-Cote (Polytechnique Montreal), Frederic Plante (Polytechnique Montreal), and Eric Laurendeau (Polytechnique Montreal)

Paving the way for Distributed Non-blocking Algorithms and Data Structures in the Partitioned Global Address Space Model .659.....
Garvit Dewan (Indian Institute of Technology Roorkee (IITR)) and Louis Jenkins (University of Rochester)

Computing Hypergraph Homology in Chapel .667.....
Jesun Firoz (Pacific Northwest National Laboratory), Louis Jenkins (University of Rochester), Cliff Joslyn (Pacific Northwest National Laboratory), Brenda Praggastis (Pacific Northwest National Laboratory), Emilie Purvine (Pacific Northwest National Laboratory), and Mark Raugas (Pacific Northwest National Laboratory)

An Automated Machine Learning Approach for Data Locality Optimizations in Chapel .671.....
Engin Kayraklioglu (Hewlett Packard Enterprise) and Tarek El-Ghazawi (The George Washington University)

Exploring Chapel Productivity Using Some Graph Algorithms .672.....
Richard F. Barrett (Sandia National Laboratories), Jeanine Cook (Sandia National Laboratories), Stephen L. Oliver (Sandia National Laboratories), Omar Aaziz (Sandia National Laboratories), Christipher D. Jenkins (Sandia National Laboratories), and Courtenay T. Vaughan (Sandia National Laboratories)

Visibility Control: Use and Import Statement Improvements .673.....
Lydia Duncan (Hewlett Packard Enterprise)

Towards Stability in the Chapel Language .674.....
Michael P. Ferguson (Hewlett Packard Enterprise)

Exploring a Multi-resolution GPU Programming Model for Chapel .675.....
Akihiro Hayashi (Georgia Institute of Technology), Sri Raj Paul (Georgia Institute of Technology), and Vivek Sarkar (Georgia Institute of Technology)

Random Forests in Chapel .676.....
Benjamin Albrecht (Hewlett Packard Enterprise)

Squeezing Performance out of Arkouda .677.....	<i>Elliot Ronaghan (Hewlett Packard Enterprise)</i>
Simulating Ultralight Dark Matter in Chapel .678.....	<i>Nikhil Padmanabhan (Yale University), Elliot Ronaghan (Cray/HPE), J. Luna Zagorac (Yale University), and Richard Easther (Univ. Auckland)</i>
Chapel on Accelerators .679.....	<i>Rahul Ghangas (The Australian National University) and Josh Milthorpe (The Australian National University)</i>

PDSEC: Parallel and Distributed Scientific and Engineering Computing

PDSEC 2020 Introduction and Committees .680.....	<i>Raphaël Couturier (University of Franche-Comté) and Peter Strazdins (The Australian National University)</i>
Comparison of MPI and Spark for Data Science Applications .682.....	<i>Manvi Saxena (University of Houston), Shweta Jha (University of Houston), Saba Khan (University of Houston), John Rodgers (University of Houston), Peggy Lindner (University of Houston), and Edgar Gabriel (University of Houston)</i>
Improving MPI Application Communication Time with an Introspection Monitoring Library .691..	<i>Emmanuel Jeannot (Inria) and Richard Sartori (Inria)</i>
A GPU-Accelerated Barycentric Lagrange Treecode .701.....	<i>Nathan Vaughn (University of Michigan), Leighton Wilson (University of Michigan), and Robert Krasny (University of Michigan)</i>
Vectorization and Minimization of Memory Footprint for Linear High-Order Discontinuous Galerkin Schemes .711.....	<i>Jean-Matthieu Gallard (Technical University of Munich), Leonhard Rannabauer (Technical University of Munich), Anne Reinartz (Technical University of Munich), and Michael Bader (Technical University of Munich)</i>
Communication Avoiding 2D Stencil Implementations over PARSEC Task-Based Runtime .721.....	<i>Yu Pei (University of Tennessee), Qinglei Cao (University of Tennessee), George Bosilca (University of Tennessee), Piotr Luszczek (University of Tennessee), Victor Eijkhout (Texas Advanced Computing Center University of Texas), and Jack Dongarra (University of Tennessee)</i>
Implementing an Attack Graph Generator in CUDA .730.....	<i>Ming Li (The University of Tulsa), Peter Hawrylak (The University of Tulsa), and John Hale (The University of Tulsa)</i>
Tri-Objective Workflow Scheduling and Optimization in Heterogeneous Cloud Environments .739.	<i>Huda Alrammah (Middle TN State University), Yi Gu (Middle TN State University), and Zhifeng Liu (Jiangsu University)</i>
Identifying Optimization Opportunities Using Memory Access Tracing in OpenSHMEM Runtimes with the TAU Performance System .749.....	<i>Nicholas Chaimov (ParaTools, Inc.), Sameer Shende (ParaTools, Inc.), Allen Malony (ParaTools, Inc.), and Neena Imam (Oak Ridge National Laboratory)</i>

Tiled Algorithms for Efficient Task-Parallel H-Matrix Solvers .757.....
Rocío Carratalá-Sáez (Universitat Jaume I), Mathieu Faverge (Bordeaux INP, CNRS, Inria, Univ. Bordeaux), Grégoire Pichon (Univ Lyon, EnsL, UCBL, CNRS, Inria, LIP), Guillaume Sylvoand (Airbus Central R&T, Inria), and Enrique S. Quintana-Ortí (Universitat Politècnica de València)

iWAPT: Automatic Performance Tuning

iWAPT 2020 Introduction and Committees .767.....
I-Hsin Chung (IBM T. J. Watson Research Center)

Machine Learning-Based Prefetching for SCM Main Memory System .769.....
Mayuko Koezuka (Toshiba Corporation), Yusuke Shiota (Toshiba Corporation), Satoshi Shirai (Toshiba Corporation), and Tatsunori Kanai (Toshiba Corporation)

Acceleration of Structural Analysis Simulations using CNN- Based Auto-Tuning of Solver Tolerance .777.....
Amir Haderbache (Fujitsu Laboratories Ltd.), Koichi Shirahata (Fujitsu Laboratories Ltd.), Takuji Yamamoto (Fujitsu Laboratories Ltd.), Yasumoto Tomita (Fujitsu Laboratories Ltd.), and Hiroshi Okuda (The University of Tokyo)

Using Small-Scale History Data to Predict Large-Scale Performance of HPC Application .787.....
Wenju Zhou (University of Science and Technology of China (USTC)), Jingwei Sun (University of Science and Technology of China (USTC)), Jiepeng Zhang (University of Science and Technology of China (USTC)), and Guangzhong Sun (University of Science and Technology of China (USTC))

Node-Aware Stencil Communication for Heterogeneous Supercomputers .796.....
Carl Pearson (University of Illinois Urbana-Champaign), Mert Hidayetoglu (University of Illinois Urbana-Champaign), Mohammad Almasri (University of Illinois Urbana-Champaign), Omer Anjum (University of Illinois Urbana-Champaign), I-Hsin Chung (IBM T. J. Watson Research), Jinjun Xiong (IBM T. J. Watson Research), and Wen-Mei Hwu (University of Illinois Urbana-Champaign)

Task Priority Control for the HPX Runtime System .806.....
Suhang Jiang (Tohoku University), Mulya Agung (Tohoku University), Ryusuke Egawa (Tohoku University), and Hiroyuki Takizawa (Tohoku University)

Improving Collective I/O Performance with Machine Learning Supported Auto-Tuning .814.....
Ayşe Bağbaba (HLRS)

Automatically Avoiding Memory Access Conflicts on SX-Aurora TSUBASA .822.....
Naoki Ebata (Tohoku University), Ryusuke Egawa (Tohoku University), Yoko Isobe (Tohoku University, NEC Corporation), Ryoji Takaki (Japan Aerospace Exploration Agency), and Hiroyuki Takizawa (Tohoku University)

Importance of Selecting Data Layouts in the Tsunami Simulation Code .830.....
Takumi Kishitani (Tohoku University), Kazuhiko Komatsu (Tohoku University), Masayuki Sato (Tohoku University), Akihiro Musa (Tohoku University, NEC Corporation), and Hiroaki Kobayashi (Tohoku University)

MPP: Parallel Programming Models - Special Edition Machine Learning Performance and Security

MPP 2020 Introduction and Committees .838.....
Leandro A. J. Marzulo (Google) and Tiago A. O. Alves (Universidade do Estado do Rio de Janeiro (UERJ))

Enhancing the Utilization of Dot-Product Engines in Deep Learning Accelerators .840.....
Taha Soliman (Bosch GmbH), Armin Runge (Bosch GmbH), and Leonardo Ecco (Bosch GmbH)

Weightless Neural Networks Applied to Nonintrusive Load Monitoring .844.....
Guilherme C. De Lello (Universidade Federal do Rio de Janeiro), Juliano F. Caldeira (Universidade Federal do Rio de Janeiro), Mauricio Aredes (Universidade Federal do Rio de Janeiro), Felipe M. G. França (Universidade Federal do Rio de Janeiro), and Priscila M. V. Lima (Universidade Federal do Rio de Janeiro)

Tangle Ledger for Decentralized Learning .852.....
Robert Schmid (Hasso Plattner Institute, University of Potsdam, Germany), Bjarne Pfitzner (Digital Health Center, Hasso Plattner Institute, University of Potsdam, Germany), Jossekin Beilharz (Hasso Plattner Institute, University of Potsdam, Germany), Bert Arnrich (Digital Health Center, Hasso Plattner Institute, University of Potsdam, Germany), and Andreas Polze (Hasso Plattner Institute, University of Potsdam, Germany)

Regression WiSARD Application of Controller on DC STATCOM Converter Under Fault Conditions...
860
Raphael N. C. B. Rocha (Programa de Engenharia Eletrica – COPPE), Leopoldo L. Filho (Programa de Engenharia de Sistemas e Computação – COPPE), Mauricio Aredes (Programa de Engenharia Eletrica – COPPE), Felipe M. G. França (Programa de Engenharia de Sistemas e Computação – COPPE), and Priscila M. V. Lima (Programa de Engenharia de Sistemas e Computação – COPPE)

SNACS: Scalable Networks for Advanced Computing Systems

SNACS 2020 Introduction and Committees .868.....
Matthew G. F. Dosanjh (Sandia National Laboratories), Ryan E. Grant (Sandia National Laboratories), and Taylor Groves (Lawrence Berkeley National Laboratory)

Analyzing and Understanding the Impact of Interconnect Performance on HPC, Big Data, and Deep Learning Applications: A Case Study with InfiniBand EDR and HDR .869.....
*Amit Ruhela (TACC Austin), Shulei Xu (The Ohio State University),
Karthik Vadambacheri Manian (The Ohio State University), Hari
Subramoni (The Ohio State University), and Dhabaleswar K Panda (The
Ohio State University)*

The Case for Explicit Reuse Semantics for RDMA Communication .879.....
*Scott Levy (Sandia National Laboratories), Patrick Widener (Sandia
National Laboratories), Craig Ulmer (Sandia National Laboratories),
and Todd Kordenbrock (Perspecta, Inc.)*

Performance of MPI Sends of Non-contiguous Data .889.....
Victor Eijkhout (The University of Texas at Austin)

Performance Characterization of Network Mechanisms for Non-contiguous Data Transfers in MPI .896.....
*Kaushik Kandadi Suresh (The Ohio State University), Bharath Ramesh
(The Ohio State University), Seyedeh Mahdieh Ghazimirsaeed (The Ohio
State University), Mohammadreza Bayatpour (The Ohio State University),
Jahanzeb Hashmi (The Ohio State University), Hari Subramoni (The Ohio
State University), and Dhabaleswar K. (DK) Panda (The Ohio State
University)*

PAISE: Parallel AI and Systems for the Edge

PAISE 2020 Introduction and Committees .906.....
*Pete Beckman (Argonne National Laboratory) and Rajesh Sankaran
(Argonne National Laboratory)*

Analyzing Deep Learning Model Inferences for Image Classification using OpenVINO .908.....
*Zheming Jin (Argonne National Lab) and Hal Finkel (Argonne National
Lab)*

Energy-Efficient Machine Learning on the Edges .912.....
*Mohit Kumar (Wayne State University), Xingzhou Zhang (Wayne State
University), Liangkai Liu (Wayne State University), Yifan Wang (Wayne
State University), and Weisong Shi (Wayne State University)*

Indirect Deconvolution Algorithm .922.....
Marat Dukhan (Google Research)

Multiperspective Automotive Labeling .927.....
*Luke Jacobs (University of Illinois at Urbana-Champaign), Akhil
Kodumuri (University of Illinois at Urbana-Champaign), Jim James
(Metea Valley High School), Seongha Park (Argonne National
Laboratory), and Yongho Kim (Argonne National Laboratory)*

Integrating DOTS with Blockchain Can Secure Massive IoT Sensors .937.....
*Syed Badruddoja (University of North Texas), Ram Dantu (University of
North Texas), Logan Widick (University of North Texas), Zachary
Zaccagni (University of North Texas), and Kritagya Upadhyay
(University of North Texas)*

RADR: Resource Arbitration for Dynamic Runtimes

RADR 2020 Introduction and Committees .947.....	
<i>Pete Beckman (Argonne National Laboratory), Emmanuel Jeannot (TADaaM Team, Inria), and Swann Perarnau (Argonne National Laboratory)</i>	
NUMA-Aware CPU Core Allocation in Cooperating Dynamic Applications .950.....	
<i>Jiri Dokulil (University of Vienna, Austria) and Siegfried Benkner (University of Vienna, Ausitra)</i>	
Overlapping MPI Communications with Intel TBB Computation .958.....	
<i>Cassandra Rocha Barbosa (Atos, CReSTIC Laboratory EA3804), Pierre Lemarinier (Atos), Marc Sergent (Atos), Guillaume Papauré (Atos), and Marc Pérache (CEA DAM DIF, CReSTIC Laboratory EA3804)</i>	
System Software for Resource Arbitration on Future Many-* Architectures .967.....	
<i>Florian Schmaus (Friedrich-Alexander-Universität Erlangen-Nürnberg (FAU)), Sebastian Maier (Friedrich-Alexander-Universität Erlangen-Nürnberg (FAU)), Tobias Langer (Friedrich-Alexander-Universität Erlangen-Nürnberg (FAU)), Jonas Rabenstein (Friedrich-Alexander-Universität Erlangen-Nürnberg (FAU)), Timo Hönig (Friedrich-Alexander-Universität Erlangen-Nürnberg (FAU)), Lars Bauer (Karlsruhe Institute of Technology (KIT)), Jörg Henkel (Karlsruhe Institute of Technology (KIT)), and Wolfgang Schröder-Preikschat (Friedrich-Alexander-Universität Erlangen-Nürnberg (FAU))</i>	
An Implementation of User-Level Processes using Address Space Sharing .976.....	
<i>Atsushi Hori (Center for Computational Science, RIKEN, Tokyo, JAPAN), Balazs Gerofi (Riken), and Yutaka Ishikawa (Riken)</i>	

ScaDL: Scalable Deep Learning over Parallel and Distributed Infrastructures

ScaDL 2020 Introduction and Committees .985.....	
<i>Ashish Verma (IBM Research), Christopher Carothers (RPI), K. R. Jayaram (IBM Research), and Parijat Dube (IBM Research)</i>	
ScaDL 2020 Invited Speaker 1 .987.....	
<i>Manish Gupta (Google Research, India)</i>	
ScaDL 2020 Invited Speaker 2 .988.....	
<i>Geoffrey Fox (Indiana University)</i>	
ScaDL 2020 Invited Speaker 3 .989.....	
<i>Wen-mei Hwu (UIUC)</i>	
ScaDL 2020 Invited Speaker 4 .990.....	
<i>Minsik Cho (IBM)</i>	

Neural Network Molecular Dynamics at Scale .991.....	
	<i>Pankaj Rajak (Argonne National Laboratory), Kuang Liu (University of Southern California), Aravind Krishnamoorthy (University of Southern California), Rajiv Kalia (University of Southern California), Aiichiro Nakano (University of Southern California), Ken-ichi Nomura (University of Southern California), Subodh Tiwari (University of Southern California), and Priya Vashishta (University of Southern California)</i>
Asynchronous SGD for DNN Training on Shared-Memory Parallel Architectures .995.....	
	<i>Florent Lopez (Innovative Computing Laboratory (ICL), University of Tennessee), Edmond Chow (Georgia Institute of Technology), Stanimire Tomov (Innovative Computing Laboratory (ICL), University of Tennessee), and Jack Dongarra (Innovative Computing Laboratory (ICL), University of Tennessee)</i>
Accelerating Towards Larger Deep Learning Models and Datasets – A System Platform View Point .999.....	
	<i>Saritha Vinod (IBM India Systems Development Lab, Bangalore, India), Naveen M (IBM India Systems Development Lab, Bangalore, India), Asis K Patra (IBM India Systems Development Lab, Bangalore, India), and Anto Ajay Raj John (IBM India Systems Development Lab, Bangalore, India)</i>
Data Parallel Large Sparse Deep Neural Network on GPU .1006.....	
	<i>Naw Safrin Sattar (University of New Orleans) and Shaikh Arifuzzaman (University of New Orleans)</i>
Efficient Training of Semantic Image Segmentation on Summit using Horovod and MVAPICH2-GDR. 1015	
	<i>Quentin Anthony (The Ohio State University), Ammar Awan (The Ohio State University), Arpan Jain (The Ohio State University), Hari Subramoni (The Ohio State University), and Dhabaleswar Panda (The Ohio State University)</i>

HPS: High-Performance Storage

HPS 2020 Introduction and Committees .1024.....	
	<i>Kathryn Mohror (Lawrence Livermore National Laboratory) and Marc Snir (University of Illinois at Urbana-Champaign)</i>
Dynamic Provisioning of Storage Resources: A Case Study with Burst Buffers .1027.....	
	<i>François Tessier (Swiss National Supercomputing Center, ETH Zurich), Maxime Martinasso (Swiss National Supercomputing Center, ETH Zurich), Matteo Chesi (Swiss National Supercomputing Center, ETH Zurich), Mark Klein (Swiss National Supercomputing Center, ETH Zurich), and Miguel Gila (Swiss National Supercomputing Center, ETH Zurich)</i>

Optimizing Asynchronous Multi-level Checkpoint/Restart Configurations with Machine Learning	1036
<i>Tonmoy Dey (Florida State University), Kento Sato (RIKEN Center for Computational Science), Bogdan Nicolae (Argonne National Laboratory), Jian Guo (RIKEN Center for Computational Science), Jens Domke (RIKEN Center for Computational Science), Weikuan Yu (Florida State University), Franck Cappello (Argonne National Laboratory), and Kathryn Mohror (Lawrence Livermore National Laboratory)</i>	
On Overlapping Communication and File I/O in Collective Write Operation	1044
<i>Raafat Feki (University of Houston) and Edgar Gabriel (University of Houston)</i>	
Recorder 2.0: Efficient Parallel I/O Tracing and Analysis	1052
<i>Chen Wang (University of Illinois at Urbana-Champaign), Jinghan Sun (University of Illinois at Urbana-Champaign), Marc Snir (University of Illinois at Urbana-Champaign), Kathryn Mohror (Lawrence Livermore National Laboratory), and Elsa Gonsiorowski (Lawrence Livermore National Laboratory)</i>	
Silent Data Access Protocol for NVRAM+RDMA Distributed Storage	1060
<i>Qingyue Liu (Rice University) and Peter Varman (Rice University)</i>	
Enhancing Endurance of SSD Based High-Performance Storage Systems using Emerging NVM Technologies	1070
<i>Tanaya Roy (Temple University) and Krishna Kant (Temple University)</i>	
Design of Locality-Aware MPI-IO for Scalable Shared File Write Performance	1080
<i>Kohei Sugihara (University of Tsukuba) and Osamu Tatebe (University of Tsukuba)</i>	

ParSocial: Parallel and Distributed Processing for Computational Social Systems

ParSocial 2020 Introduction and Committees	1090
<i>John Korah (California Polytechnic University Pomona) and Eunice E. Santos (University of Illinois at Urbana-Champaign)</i>	
Parallel Generation of Simple Null Graph Models	1091
<i>Jack Garbus (Rensselaer Polytechnic Institute), Christopher Brissette (Rensselaer Polytechnic Institute), and George M. Slota (Rensselaer Polytechnic Institute)</i>	
A System for High Performance Mining on GDELT Data	1101
<i>Konstantin Pogorelov (Simula Research Laboratory), Daniel Thilo Schroeder (Simula Metropolitan Center for Digital Engineering), Petra Filkukova (Simula Research Laboratory), and Johannes Langguth (Simula Research Laboratory)</i>	
A Parallel LFR-Like Benchmark for Evaluating Community Detection Algorithms	1112
<i>George M Slota (Rensselaer Polytechnic Institute) and Jack Garbus (Rensselaer Polytechnic Institute)</i>	
The Role of Artificial Intelligence and Cyber Security for Social Media	1116
<i>Bhavani Thuraisingham (The University of Texas at Dallas)</i>	
YouTube Data Collection Using Parallel Processing	1119
<i>Joseph Kready (University of Arkansas at Little Rock), Shishila Awung (University of Arkansas at Little Rock), Muhammad Nihal Hussain (University of Arkansas at Little Rock), and Nitin Agarwal (University of Arkansas at Little Rock)</i>	
New Approaches for Performance Optimization and Analysis of Large-Scale Dynamic Social Network Analysis using Anytime Anywhere Algorithms	1123
<i>Vairavan Murugappan (University of Illinois at Urbana-Champaign), Eunice Santos (University of Illinois at Urbana-Champaign), and John Korah (California State Polytechnic University, Pomona)</i>	