

2022 IEEE International Parallel and Distributed Processing Symposium Workshops (IPDPSW 2022)

**Virtual Conference
30 May - 3 June 2022**

Pages 1-646



**IEEE Catalog Number: CFP2251J-POD
ISBN: 978-1-6654-9748-0**

**Copyright © 2022 by the Institute of Electrical and Electronics Engineers, Inc.
All Rights Reserved**

Copyright and Reprint Permissions: Abstracting is permitted with credit to the source. Libraries are permitted to photocopy beyond the limit of U.S. copyright law for private use of patrons those articles in this volume that carry a code at the bottom of the first page, provided the per-copy fee indicated in the code is paid through Copyright Clearance Center, 222 Rosewood Drive, Danvers, MA 01923.

For other copying, reprint or republication permission, write to IEEE Copyrights Manager, IEEE Service Center, 445 Hoes Lane, Piscataway, NJ 08854. All rights reserved.

****** This is a print representation of what appears in the IEEE Digital Library. Some format issues inherent in the e-media version may also appear in this print version.***

IEEE Catalog Number:	CFP2251J-POD
ISBN (Print-On-Demand):	978-1-6654-9748-0
ISBN (Online):	978-1-6654-9747-3

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

2022 IEEE International Parallel and Distributed Processing Symposium Workshops (IPDPSW) **IPDPSW 2022**

Table of Contents

Message from the 2022 General Co-Chairs	xxviii
Message from the 2022 Workshops Chair and Vice-chair	xxx

HCW: Heterogeneity in Computing Workshop

Introduction to HCW 2022	1
<i>Ryan D. Friesse (Pacific Northwest National Laboratory Laboratory, USA) and Jong-Kook Kim (Korea University, Korea)</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>Ryan D. Friesse (Pacific Northwest National Laboratory, USA)</i>	
Message from the HCW Technical Program Committee Chair	4
<i>Jong-Kook Kim (Korea University, South Korea)</i>	
HCW 2022 Keynote Speaker	5
<i>Laurent White (Advanced Micro Devices, Inc.)</i>	
Heterogeneous Architecture for Sparse Data Processing	6
<i>Shashank Adavally (University of North Texas, USA), Alex Weaver (University of North Texas, USA), Pranathi Vasireddy (University of North Texas, USA), Krishna Kavi (University of North Texas, USA), Gayatri Mehta (University of North Texas, USA), and Nagendra Gulur (University of North Texas, USA)</i>	
Combined Application of Approximate Computing Techniques in DNN Hardware Accelerators	16
<i>Enrico Russo (University of Catania, Italy), Maurizio Palesi (University of Catania, Italy), Davide Patti (University of Catania, Italy), Habiba Lahdhiri (ESIEE, Paris), Salvatore Monteleone (Niccolo Cusano University, Italy), Giuseppe Ascia (University of Catania, Italy), and Vincenzo Catania (University of Catania, Italy)</i>	

Highly Efficient Alltoall and Alltoallv Communication Algorithms for GPU Systems	24
<i>Chen-Chun Chen (The Ohio State University, USA), Kawthar Shafie Khorassani (The Ohio State University, USA), Quentin G. Anthony (The Ohio State University, USA), Aamir Shafi (The Ohio State University, USA), Hari Subramoni (The Ohio State University, USA), and Dhableswar K. Panda (The Ohio State University, USA)</i>	
On Energy Nonproportionality of CPUs and GPUs	34
<i>Ravi Reddy Manumachu (University College Dublin, Ireland) and Alexey Lastovetsky (University College Dublin, Ireland)</i>	
Implementing Spatio-Temporal Graph Convolutional Networks on Graphcore IPUs	45
<i>Johannes Moe (Simula Research Laboratory, University of Oslo, Norway), Konstantin Pogorelov (Simula Research Laboratory, Norway), Daniel Thilo Schroeder (Simula Research Laboratory, Norway), and Johannes Langguth (Simula Research Laboratory, University of Bergen, Norway)</i>	
The Best of Many Worlds: Scheduling Machine Learning Inference on CPU-GPU Integrated Architectures	55
<i>Rafail Tsirbas (Foundation for Research and Technology - Hellas), Giorgos Vasiliadis (Foundation for Research and Technology - Hellas; Hellenic Mediterranean University), and Sotiris Ioannidis (Foundation for Research and Technology - Hellas; Technical University of Crete)</i>	

RAW: Reconfigurable Architectures Workshop

Introduction to RAW 2022	65
<i>Jürgen Becker (Karlsruhe Institute of Technology, Germany), Lana Josipović (ETH Zürich, Switzerland), Viktor K. Prasanna (University of Southern California, USA), Marco Santambrogio (Politecnico di Milano, Italy), and Ramachandran Vaidyanathan (Louisiana State University, USA)</i>	
RAW 2022 Keynote Speaker 1	67
<i>Gustavo Alonso (ETH Zurich, Switzerland)</i>	
RAW 2022 Keynote Speaker 2	N/A
<i>Stephen Neuendorffer (Xilinx Research Labs)</i>	
Online Learning RTL Synthesis for Automated Design Space Exploration	69
<i>Daniele Paletti (Politecnico di Milano, Italy), Francesco Peverelli (Politecnico di Milano, Italy), and Davide Conficconi (Politecnico di Milano, Italy)</i>	
Machine Learning Aided Hardware Resource Estimation for FPGA DNN Implementations	77
<i>Dana Diaconu (Northeastern University, USA), Lucian Petrica (Xilinx Research, Ireland), Michaela Blott (Xilinx Research, Ireland), and Miriam Leeser (Northeastern University, USA)</i>	
DECISION: Distributing OpenVX Applications on CPUs, GPUs and FPGAs using OpenCL	84
<i>Lester Kalms (Technische Universität Dresden, Germany), Tim Haering (Technische Universität Dresden, Germany), and Diana Goehring (Technische Universität Dresden, Germany)</i>	

A Hybrid Approach Combining ANN-Based and Conventional Demapping in Communication for Efficient FPGA-Implementation	92
<i>Jonas Ney (TU Kaiserslautern, Germany), Bilal Hammoud (TU Kaiserslautern, Germany), and Norbert Wehn (TU Kaiserslautern, Germany)</i>	
Optimal Schedules for High-Level Programming Environments on FPGAs with Constraint Programming	96
<i>Pascal Jungblut (Ludwig-Maximilians-Universität, Germany) and Dieter Kranzlmüller (Ludwig-Maximilians-Universität, Germany)</i>	
Optimization of Compiler-Generated OpenCL CNN Kernels and Runtime for FPGAs	100
<i>Seung-Hun Chung (University of Toronto) and Tarek S. Abdelrahman (University of Toronto)</i>	
On How to Push Efficient Medical Semantic Segmentation to the Edge: the SENECA Approach	104
<i>Raffaele Berzoini (Politecnico di Milano, Italy), Eleonora D'Arnese (Politecnico di Milano, Italy), and Davide Conficconi (Politecnico di Milano, Italy)</i>	
Exploiting High-Bandwidth Memory for FPGA-Acceleration of Inference on Sum-Product Networks	112
<i>Lukas Weber (Embedded Systems and Applications Group, TU Darmstadt, Germany), Johannes Wirth (Embedded Systems and Applications Group, TU Darmstadt, Germany), Lukas Sommer (Embedded Systems and Applications Group, TU Darmstadt, Germany), and Andreas Koch (Embedded Systems and Applications Group, TU Darmstadt, Germany)</i>	
ReconOS64: A Hardware Operating System for Modern Platform FPGAs with 64-Bit Support	120
<i>Lennart Clausing (Paderborn University, Germany) and Marco Platzner (Paderborn University, Germany)</i>	
An FPGA-Based IP Core Subscription-Oriented Fog Computing Platform	128
<i>Tze Hon Tan (Universiti Teknologi Malaysia, Malaysia), Chia Yee Ooi (Universiti Teknologi Malaysia, Malaysia), and M. N. Marsono (Universiti Teknologi Malaysia, Malaysia)</i>	
A SHA-512 Hardware Implementation Based on Block RAM Storage Structure	132
<i>Mingyuan Yang (Tsinghua University, China), Yemeng Zhang (Tsinghua University, China), Bohan Yang (Tsinghua University, China), Hanning Wang (Tsinghua University, China), Shouyi Yin (Tsinghua University, China), Shaojun Wei (Tsinghua University, China), and Leibo Liu (Tsinghua University, China)</i>	
Fast Genome Analysis Leveraging Exact String Matching	136
<i>Beatrice Branchini (Politecnico di Milano, Italy), Sofia Breschi (Politecnico di Milano, Italy), Alberto Zeni (Politecnico di Milano, Italy), and Marco D. Santambrogio (Politecnico di Milano, Italy)</i>	

HiCOMB: High Performance Computational Biology

Introduction to HiCOMB 2022	140
<i>Alba Cristina M. A. de Melo (University of Brasilia, Brazil) and Ananth Kalyanaraman (Washington State University, USA)</i>	

HiCOMB 2022 Keynote Speaker	142
<i>Christina Boucher (University of Florida)</i>	
HiCOMB 2022 Invited Speaker	143
<i>Yatish Turakhia (University of California, San Diego, USA)</i>	
Optimizing the Accuracy of Randomized Embedding for Sequence Alignment	144
<i>Yiqing Yan (EURECOM, France), Nimisha Chaturvedi (ACCELOM, France), and Raja Appuswamy (EURECOM, France)</i>	
On using Consistency Consistently in Multiple Sequence Alignments	152
<i>Mario João Jr. (Fluminense Federal University, Brazil), Alexandre C. Sena (State University of Rio de Janeiro, Brazil), and Vinod E. F. Rebello (Fluminense Federal University, Brazil)</i>	
Algorithmic Improvement and GPU Acceleration of the GenASM Algorithm	162
<i>Joël Lindegger (ETH Zürich), Damla Senol Cali (Bionano Genomics), Mohammed Alser (ETH Zürich), Juan Gómez-Luna (ETH Zürich), and Onur Mutlu (ETH Zürich)</i>	
High-Throughput Pairwise Alignment with the Wavefront Algorithm using Processing-in-Memory....	163
<i>Safaa Diab (American University of Beirut, Lebanon), Amir Nassereldine (American University of Beirut, Lebanon), Mohammed Alser (ETH Zürich, Switzerland), Juan Gómez Luna (ETH Zürich, Switzerland), Onur Mutlu (ETH Zürich, Switzerland), and Izzat El Hajj (American University of Beirut, Lebanon)</i>	
Sequire: A high-Performance Framework for Rapid Development of Secure Bioinformatics Pipelines	164
<i>Haris Smajlović (University of Victoria, Canada), Ariya Shajii (Massachusetts Institute of Technology, USA), Bonnie Berger (Massachusetts Institute of Technology, USA), Hyunghoon Cho (Broad Institute of MIT and Harvard, USA), and Ibrahim Numanagić (University of Victoria, Canada)</i>	
Scalable and Extensible Robinson-Foulds for Comparative Phylogenetics	166
<i>Alvin Chon (Iowa State University, Bioinformatics and Computational Biology, USA), Paweł Górecki (University of Poland, Institute of Informatics, Poland), Oliver Eulenstein (Iowa State University, Bioinformatics and Computational Biology, USA), Xiaoqiu Huang (Iowa State University, Bioinformatics and Computational Biology, USA), and Ali Jannesari (Iowa State University, Bioinformatics and Computational Biology, USA)</i>	
Accelerating Deep Learning Based Identification of Chromatin Accessibility from Noisy ATAC-seq Data	176
<i>Narendra Chaudhary (Parallel Computing Lab, Intel, India), Sanchit Misra (Parallel Computing Lab, Intel, India), Dhiraj Kalamkar (Parallel Computing Lab, Intel, India), Alexander Heinecke (Parallel Computing Lab, Intel Corporation, USA), Evangelos Georganas (Parallel Computing Lab, Intel Corporation, USA), Barukh Ziv (Intel Corporation, Israel), Menachem Adelman (Intel Corporation, Israel), and Bharat Kaul (Parallel Computing Lab, Intel, India)</i>	

Graph Convolutional Neural Networks for Alzheimer's Classification with Transfer Learning and HPC Methods	186
<i>Anoop Kumar (Indian Institute of Science, India), Vibha Balaji (Indian Institute of Science, India), Chandrashekar M.A. (Indian Institute of Science, India), Ambedkar Dukkipati (Indian Institute of Science, India), and Sathish Vadhiyar (Indian Institute of Science, India)</i>	
Accelerated LD-Based Selective Sweep Detection using GPUs and FPGAs	196
<i>Reinout Corts (University of Twente, Netherlands), Niek Sterenborg (University of Twente, Netherlands), and Nikolaos Alachiotis (University of Twente, Netherlands)</i>	
Proteome-Scale Deployment of Protein Structure Prediction Workflows on the Summit Supercomputer	206
<i>Mu Gao (Georgia Institute of Technology, USA), Mark Coletti (Oak Ridge National Laboratory, USA), Russell B. Davidson (Oak Ridge National Laboratory, USA), Ryan Prout (Oak Ridge National Laboratory, USA), Subil Abraham (Oak Ridge National Laboratory, USA), Benjamin Hernández (Oak Ridge National Laboratory, USA), and Ada Sedova (Oak Ridge National Laboratory, USA)</i>	
Reproducibility of Bioinformatics Tools	216
<i>Pelin Icer Baykal (ETH Zurich SIB Swiss Institute of Bioinformatics Basel, Switzerland), Niko Beerenwinkel (ETH Zurich SIB Swiss Institute of Bioinformatics Basel, Switzerland), and Serghei Mangul (University of Southern California, USA)</i>	
TAMPA: Interpretable Analysis and Visualization of Metagenomics-Based Taxon Abundance Profiles	217
<i>Varuni Sarwal (University of California, USA), Serghei Mangul (University of Southern California, USA), and David Koslicki (Penn State University, USA)</i>	

GrAPL: Graphs, Architectures, Programming, and Learning

Introduction to GrAPL 2022	218
<i>Manoj Kumar (IBM) and Scott McMillan (CMU SEI)</i>	
GrAPL 2022 Keynote Speaker	220
<i>Tim Mattson (Intel Labs, USA)</i>	
High-Performance GraphBLAS Backend Prototype for NEC SX-Aurora TSUBASA	221
<i>Ilya Afanasyev (Lomonosov Moscow State University, Russia), Kazuhiko Komatsu (Tohoku University, Japan), Dmitry Lichmanov (Lomonosov Moscow State University, Russia), Vadim Voevodin (Lomonosov Moscow State University, Russia), and Hiroaki Kobayashi (Tohoku University, Japan)</i>	
Nonblocking Execution in GraphBLAS	230
<i>Aristeidis Mastoras (Huawei Technologies, Switzerland), Sotiris Anagnostidis (ETH Zurich, Switzerland), and Albert-Jan N. Yzelman (Huawei Technologies, Switzerland)</i>	

Towards a GraphBLAS Implementation for Go	234
<i>Pascal Costanza (Intel, Extreme Scale Computing Group, Belgium), Ibrahim Hur (Intel, Extreme Scale Computing Group, USA), and Timothy G. Mattson (Intel, Parallel Computing Lab, USA)</i>	
GraphBLAS: C++ Iterators for Sparse Matrices	238
<i>Benjamin Brock (University of California, Berkeley, USA), Scott McMillan (Carnegie Mellon University, USA), Aydın Buluç (Lawrence Berkeley National Laboratory, USA), Timothy G. Mattson (Parallel Computing Lab, Intel, USA), and José E. Moreira (IBM Thomas J. Watson Research Center, USA)</i>	
Temporal Correlation of Internet Observatories and Outposts	247
<i>Jeremy Kepner (MIT), Michael Jones (MIT), Daniel Andersen (CAIDA), Aydin Buluc (LBNL), Chansup Byun (MIT), Kimberly Claffy (CAIDA), Timothy Davis (Texas A&M), William Arcand (MIT), Jonathan Bernays (MIT), David Bestor (MIT), William Bergeron (MIT), Vijay Gadepally (MIT), Daniel Grant (GreyNoise), Micheal Houle (MIT), Matthew Hubbell (MIT), Hayden Jananthan (MIT), Anna Klein (MIT), Chad Meiners (MIT), Lauren Milechin (MIT), Andrew Morris (GreyNoise), Julie Mullen (MIT), Sandeep Pisharody (MIT), Andrew Prout (MIT), Albert Reuther (MIT), Antonio Rosa (MIT), Siddharth Samsi (MIT), Doug Stetson (MIT), Charles Yee (MIT), and Peter Michaleas (MIT)</i>	
Interactive Visualization of Protein RINs using NetworKit in the Cloud	255
<i>Eugenio Angriman (Humboldt-Universität zu Berlin, Germany), Fabian Brandt-Tumescheit (Humboldt-Universität zu Berlin, Germany), Leon Franke (University of Konstanz, Germany), Alexander van der Grinten (Humboldt-Universität zu Berlin, Germany), and Henning Meyerhenke (Humboldt-Universität zu Berlin, Germany)</i>	
An Efficient Parallel Implementation of a Perfect Hashing Method for Hypergraphs	265
<i>Somesh Singh (INRIA and LIP (CNRS - INRIA - ENS Lyon - UCBL), France) and Bora Ucar (CNRS and LIP (CNRS - INRIA - ENS Lyon - UCBL), France)</i>	
NWHy: A Framework for Hypergraph Analytics: Representations, Data Structures, and Algorithms	275
<i>Xu T. Liu (University of Washington, USA), Jesun Firoz (Pacific Northwest National Lab, USA), Assefaw H. Gebremedhin (Washington State University, USA), and Andrew Lumsdaine (University of Washington, USA)</i>	
Parallel Algorithms for Adding a Collection of Sparse Matrices	285
<i>Md Taufique Hussain (Indiana University, USA), Guttu Sai Abhishek (Indian Institute of Technology Bombay, India), Aydın Buluç (Lawrence Berkeley National Laboratory, USA), and Ariful Azad (Indiana University, USA)</i>	
Multi-view Learning for Parallelism Discovery of Sequential Programs	295
<i>Le Chen (Iowa State University, USA), Quazi Ishtiaque Mahmud (Iowa State University, USA), and Ali Jannesari (Iowa State University, USA)</i>	
Families of Butterfly Counting Algorithms for Bipartite Graphs	304
<i>Jay A. Acosta (The University of Texas at Austin, USA), Tze Meng Low (Carnegie Mellon University, USA), and Devangi N. Parikh (The University of Texas at Austin, USA)</i>	

Essentials of Parallel Graph Analytics	314
<i>Muhammad Osama (University of California, Davis), Serban D. Porumbescu (University of California, Davis), and John D. Owens (University of California, Davis)</i>	

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

Introduction to EduPar 2022	318
<i>Sushil Prasad (University of Texas at San Antonio), David Brown (Elmhurst University), Steven Bogaerts (DePauw University), David Bunde (Knox College), and Satish Puri (Marquette University)</i>	
EduPar 2022 Keynote	320
<i>Rajendra K. Raj (Rochester Institute of Technology, USA)</i>	
Introducing Parallel Computing in a Second CS Course	321
<i>Tia Newhall (Swarthmore College, USA), Kevin C. Webb (Swarthmore College, USA), Vasanta Chaganti (Swarthmore College, USA), and Andrew Danner (Swarthmore College, USA)</i>	
Feedback from a Data Center for Education at CentraleSupélec Engineering School	330
<i>Jeremy Fix (LORIA, CNRS, Université de Lorraine, CentraleSupélec Gif-sur-Yvette, France), Stephane Vialle (LISN, CNRS, Université Paris-Saclay CentraleSupélec Gif-sur-Yvette, France), Rémi Hellequin (Information systems department CentraleSupélec Gif-sur-Yvette, France), Claudine Mercier (Information systems department CentraleSupélec Gif-sur-Yvette, France), Patrick Mercier (Information systems department CentraleSupélec Gif-sur-Yvette, France), and Jean-Baptiste Tavernier (Information systems department CentraleSupélec Gif-sur-Yvette, France)</i>	
Teaching High-Performance Computing in Developing Countries: A Case Study in Mexican Universities	338
<i>Joel Antonio Trejo-Sánchez (CONACYT-Center for Research in Mathematics, México), Francisco Javier Hernández-López (CONACYT-Center for Research in Mathematics, México), Miguel Ángel Uh-Zapata (CONACYT-Center for Research in Mathematics, México), Jose Luis Lopez-Martinez (UADY, México), Daniel Fajardo-Delgado (Tecnologico Nacional de Mexico, México), and Julio César Ramirez-Pacheco (U Qroo, México)</i>	
A Research-Based Course Module to Study Non-Determinism in High Performance Applications .	346
<i>Patrick Bell (University of Tennessee Knoxville, USA), Kae Suarez (University of Tennessee Knoxville, USA), Barbara Fossum (University of Tennessee Knoxville, USA), Dylan Chapp (University of Tennessee Knoxville, USA), Sanjukta Bhowmick (University of North Texas, USA), and Michela Taufer (University of Tennessee Knoxville, USA)</i>	
Teaching Heterogeneous Computing using DPC++	354
<i>Joel Fuentes (Universidad del Bío-Bío, Chile), Daniel López (Universidad del Bío-Bío, Chile), and Sebastián González (Universidad del Bío-Bío, Chile)</i>	

Peachy Parallel Assignments (EduPar 2022)	361
<i>H. Martin Bückner (Friedrich Schiller University Jena, Germany), Henri Casanova (University of Hawaii, USA), Rafael Ferreira da Silva (Oak Ridge National Laboratories, USA), Alice Lasserre (University of Bordeaux, France), Derrick Luyen (University of Hawaii, USA), Raymond Namyst (University of Bordeaux, France), Johannes Schoder (Friedrich Schiller University Jena, Germany), Pierre-André Wacrenier (University of Bordeaux, France), and David P. Bunde (Knox College, USA)</i>	

AsHES: Accelerators and Hybrid Emerging Systems

Introduction to AsHES 2022	369
<i>Lena Oden (FernUni Hagen, Germany)</i>	
AsHES 2022 Keynote Speaker	371
<i>Estela Suarez (Jülich Supercomputing Centre (JSC) and University of Bonn)</i>	
Performance Analysis of Parallel FFT on Large Multi-GPU Systems	372
<i>Alan Ayala (University of Tennessee, USA), Stan Tomov (University of Tennessee, USA), Miroslav Stoyanov (Oak Ridge National Laboratory, USA), Azzam Haidar (Nvidia Corporation, USA), and Jack Dongarra (University of Tennessee, USA; Oak Ridge National Laboratory, USA; University of Manchester, UK)</i>	
Heterogeneous GPU and FPGA Computing: a VexCL Case-Study	382
<i>Tristan Laan (University of Amsterdam, The Netherlands) and Ana-Lucia Varbanescu (University of Amsterdam, The Netherlands)</i>	
COMPOFF: A Compiler Cost Model using Machine Learning to Predict the Cost of OpenMP Offloading	391
<i>Alok Mishra (Stony Brook University, USA), Smeet Chheda (Stony Brook University, USA), Carlos Soto (Brookhaven National Laboratory, USA), Abid M. Malik (Brookhaven National Laboratory, USA), Meifeng Lin (Brookhaven National Laboratory, USA), and Barbara Chapman (Stony Brook University, USA; Brookhaven National Laboratory, USA)</i>	
A Novel Set of Directives for Multi-Device Programming with OpenMP	401
<i>Raul Torres (Barcelona Supercomputing Center, Spain), Roger Ferrer (Barcelona Supercomputing Center, Spain), and Xavier Teruel (Barcelona Supercomputing Center, Spain)</i>	

APDCM: Advances in Parallel and Distributed Computational Models

Introduction to APDCM 2022	411
<i>Jacir L. Bordim (University of Brasilia, Brazil) and Koji Nakano (Hiroshima University, Japan)</i>	
APDCM 2022 Keynote Talk	413
<i>Thorsten Koch (ZIB), Daniel Rehfeldt (ZIB), and Yuji Shinano (ZIB)</i>	

APC-SCA: A Fully-Parallel Annealing Algorithm with Autonomous Pinning Effect Control	414
<i>Daiki Okonogi (Tokyo Institute of Technology, Japan), Satoru Jimbo (Tokyo Institute of Technology, Japan), Kota Ando (Tokyo Institute of Technology, Japan), Thiem Van Chu (Tokyo Institute of Technology, Japan), Jaehoon Yu (Tokyo Institute of Technology, Japan), Masato Motomura (Tokyo Institute of Technology, Japan), and Kazushi Kawamura (Tokyo Institute of Technology, Japan)</i>	
Practical Effectiveness of Quantum Annealing for Shift Scheduling Problem	421
<i>Natsuki Hamada (Keio University, Japan), Kazuhiro Saito (KDDI Research, Inc., Japan), and Hideyuki Kawashima (Keio University, Japan)</i>	
Graph-Theoretic Formulation of QUBO for Scalable Local Search on GPUs	425
<i>Ryota Yasudo (Hiroshima University, Japan; Kyoto University, Japan), Koji Nakano (Hiroshima University, Japan), Yasuaki Ito (Hiroshima University, Japan), Yuya Kawamata (NTT DATA Corporation, Japan), Ryota Katsuki (NTT DATA Corporation, Japan), Shiro Ozaki (NTT DATA Corporation, Japan), Takashi Yazane (NTT DATA Corporation, Japan), and Kenichiro Hamano (NTT DATA Corporation, Japan)</i>	
Performance Evaluations of Noisy Approximate Quantum Fourier Arithmetic	435
<i>Robert Basili (Iowa State University, USA), Wenyang Qian (Iowa State University, USA), Shuo Tang (Iowa State University, USA), Austin Castellino (Iowa State University, USA), Mary Eshaghian-Wilner (Iowa State University, USA), Ashfaq Khokhar (Iowa State University, USA), Glenn Luecke (Iowa State University, USA), and James P. Vary (Iowa State University, USA)</i>	
Performance Evaluation of Data Transfer API for Rank Level Approximate Computing on HPC Systems	445
<i>Yoshiyuki Morie (Teikyo University, Japan), Yasutaka Wada (Meisei University, Japan), Ryohei Kobayashi (University of Tsukuba, Japan), and Ryuichi Sakamoto (Tokyo Institute of Technology, Japan)</i>	
Arm Meets Cloud: A Case Study of MPI Library Performance on AWS Arm-Based HPC Cloud with Elastic Fabric Adapter	449
<i>Shulei Xu (The Ohio State University, USA), Aamir Shaf (The Ohio State University, USA), Hari Subramoni (The Ohio State University, USA), and Dhableswar K. Panda (The Ohio State University, USA)</i>	
Aspect-Oriented Programming Based Building Block Platform to Construct Domain-Specific Language for HPC Application	457
<i>Osamu Ishimura (The University of Tokyo, Japan) and Yoshihide Yoshimoto (The University of Tokyo, Japan)</i>	
Optimizing Non-Commutative Allreduce Over Virtualized, Migratable MPI Ranks	467
<i>Sam White (University of Illinois at Urbana-Champaign, USA) and Laxmikant V. Kale (University of Illinois at Urbana-Champaign, USA)</i>	
Modeling Memory Contention Between Communications and Computations in Distributed HPC Systems	476
<i>Alexandre Denis (Inria Bordeaux Sud-Ouest, France), Emmanuel Jeannot (Inria Bordeaux Sud-Ouest, France), and Philippe Swartvagher (Inria Bordeaux Sud-Ouest, France)</i>	

Fully Dynamic Line Maintenance by Hybrid Programmable Matter	486
<i>Nooshin Nokhanji (Carleton University, Canada), Paola Flocchini (University of Ottawa, Canada), and Nicola Santoro (Carleton University, Canada)</i>	
Integer Sum Reduction with OpenMP on an AMD MI100 GPU	496
<i>Zheming Jin (Oak Ridge National Laboratory) and Jeffrey S. Vetter (Oak Ridge National Laboratory)</i>	
Optimal Triangulation on the High Bandwidth Memory Model	500
<i>Koji Nakano (Hiroshima University, Japan) and Victor Poupet (University of Montpellier, France)</i>	

HIPS: High-level Parallel Programming Models and Supportive Environments

Introduction to HIPS 2022	508
<i>Jiajia Li (College of William and Mary) and Martin Ruefenacht (Leibniz Supercomputing Centre)</i>	
Towards Java-Based HPC using the MVAPICH2 Library: Early Experiences	510
<i>Kinan Al-Attar (The Ohio State University, USA), Aamir Shafi (The Ohio State University, USA), Hari Subramoni (The Ohio State University, USA), and Dhabaleswar K. Panda (The Ohio State University, USA)</i>	
mpisee: MPI Profiling for Communication and Communicator Structure	520
<i>Ioannis Vardas (TU Wien, Faculty of Informatics, Austria), Sascha Hunold (TU Wien, Faculty of Informatics, Austria), Jordy I. Ajanohoun (TU Wien, Faculty of Informatics, Austria), and Jesper Larsson Träff (TU Wien, Faculty of Informatics, Austria)</i>	
An On-the-Fly Method to Exchange Vector Clocks in Distributed-Memory Programs	530
<i>Simon Schwaitanski (IT Center, RWTH Aachen University, Aachen, Germany), Felix Tomski (IT Center, RWTH Aachen University, Aachen, Germany), Joachim Protze (IT Center, RWTH Aachen University, Aachen, Germany), Christian Terboven (IT Center, RWTH Aachen University, Aachen, Germany), and Matthias S. Müller (IT Center, RWTH Aachen University, Aachen, Germany)</i>	
Automatic Parallelization of Programs via Software Stream Rewriting	541
<i>Tao Tao (University of North Carolina at Chapel Hill, USA) and David Plaisted (University of North Carolina at Chapel Hill, USA)</i>	
Decentralized in-Order Execution of a Sequential Task-Based Code for Shared-Memory Architectures	552
<i>Charly Castes (Inria - LaBRI, EPFL, France), Emmanuel Agullo (Inria - LaBRI, France), Olivier Aumage (Inria - LaBRI, France), and Emmanuelle Saillard (Inria - LaBRI, France)</i>	
Evaluating Unified Memory Performance in HIP	562
<i>Zheming Jin (Oak Ridge National Laboratory, USA) and Jeffrey S. Vetter (Oak Ridge National Laboratory, USA)</i>	

Improving Scalability with GPU-Aware Asynchronous Tasks	569
<i>Jaemin Choi (University of Illinois at Urbana-Champaign, USA), David F. Richards (Lawrence Livermore National Laboratory, USA), and Laxmikant V. Kale (University of Illinois at Urbana-Champaign, USA)</i>	
A Customizable Lightweight STM for Irregular Algorithms on GPU	579
<i>Shayan Manoochehri (Concordia University, Canada), Patrick Cristofaro (Concordia University, Canada), and Dhrubajyoti Goswami (Concordia University, Canada)</i>	
Concurrent CPU-GPU Task Programming using Modern C++	588
<i>Tsung-Wei Huang (University of Utah, USA) and Yibo Lin (Peking University, China)</i>	

QCCC: Quantum Classical Cooperative Computing

Introduction to QCCC 2022	598
<i>Ang Li (Pacific Northwest National Laboratory, USA) and Qiang Guan (Kent State University)</i>	
QCCC 2022 Keynote Talk	599
<i>Nathan Wiebe (University of Toronto, Canada)</i>	
Methods and Results for Quantum Optimal Pulse Control on Superconducting Qubit Systems	600
<i>Elisha Siddiqui Matekole (Computational Science Initiative, Brookhaven National Laboratory, USA), Yao-Lung Leo Fang (Computational Science Initiative, Brookhaven National Laboratory, USA), and Meifeng Lin (Computational Science Initiative, Brookhaven National Laboratory, USA)</i>	
Locality-Aware Qubit Routing for the Grid Architecture	607
<i>Avah Banerjee (Missouri S&T, USA), Xin Liang (Missouri S&T, USA), and R. Tohid (Louisiana State University, USA)</i>	
SQCC: Smart Quantum Circuit Cutting	614
<i>Betis Baheri (Kent State University, USA), Qiang Guan (Kent State University, USA), Shuai Xu (Case Western Reserve University, USA), and Vipin Chaudhary (Case Western Reserve University, USA)</i>	
Improving Variational Quantum Algorithms Performance Through Weighted Quantum Ensembles ...	616
<i>Samuel Alexander Stein (Pacific Northwest National Laboratory, USA), Nathan Wiebe (Pacific Northwest National Laboratory, USA), James Ang (Pacific Northwest National Laboratory, USA), and Ang Li (Pacific Northwest National Laboratory, USA)</i>	
Benchmarking Quantum Processor Performance Through Quantum Distance Metrics Over an Algorithm Suite	618
<i>Samuel Stein (Pacific Northwest National Laboratory, USA), Nathan Wiebe (Pacific Northwest National Laboratory, USA), James Ang (Pacific Northwest National Laboratory, USA), and Ang Li (Pacific Northwest National Laboratory, USA)</i>	

CGRA4HPC: Coarse-Grained Reconfigurable Architectures for HPC

Introduction to CGRA4HPC 2022	625
<i>Artur Podobas (KTH Royal Institute of Technology, Sweden), Kentaro Sano (RIKEN Center for Computational Science, Japan), and Jason Anderson (University of Toronto, Canada)</i>	
CGRA4HPC 2022 Invited Speaker 1	627
<i>Raghu Prabhakar (SambaNova Systems)</i>	
CGRA4HPC 2022 Invited Speaker 2	628
<i>Elliott Delaye (AMD, USA)</i>	
CGRA4HPC 2022 Invited Speaker 3	629
<i>Martin Snelgrove (Untether AI)</i>	
CGRA4HPC 2022 Invited Speaker 4	630
<i>Ilan Tayari (NextSilicon)</i>	
An Architecture-Independent CGRA Compiler Enabling OpenMP Applications	631
<i>Takuya Kojima (The University of Tokyo, Japan), Boma Adhi (Center for Computational Science (R-CCS), Japan), Carlos Cortes (Center for Computational Science (R-CCS), Japan), Yiyu Tan (Center for Computational Science (R-CCS), Japan), and Kentaro Sano (Center for Computational Science (R-CCS), Japan)</i>	
Exploration Framework for Synthesizable CGRAs Targeting HPC: Initial Design and Evaluation... ..	639
<i>Boma Adhi (Center for Computational Science (R-CCS), Japan), Carlos Cortes (Center for Computational Science (R-CCS), Japan), Yiyu Tan (Center for Computational Science (R-CCS), Japan), Takuya Kojima (Center for Computational Science (R-CCS), Japan; The University of Tokyo, Japan), Artur Podobas (KTH Royal Institute of Technology, Sweden), and Kentaro Sano (Center for Computational Science (R-CCS), Japan)</i>	
An Analysis of Mapping Polybench Kernels to HPC CGRAs	647
<i>Markus Weinhardt (Osnabrueck University of Applied Sciences, Germany)</i>	
Elastic Multi-context CGRAs	655
<i>Omar Ragheb (University of Toronto, Canada), Tianyi Yu (University of Toronto, Canada), Rami Beidas (University of Toronto, Canada), and Jason Anderson (University of Toronto, Canada)</i>	
Accelerating SLIDE: Exploiting Sparsity on Accelerator Architectures	663
<i>Sho Ko (Stanford University, USA), Alexander Rucker (Stanford University, USA), Yaqi Zhang (SambaNova Systems, USA), Paul Mure (Stanford University, USA), and Kunle Olukotun (Stanford University, USA)</i>	
A Coarse Grained Reconfigurable Architecture for SHA-2 Acceleration	671
<i>Hoi Luan Pham (Nara Institute of Science and Technology (NAIST), Japan), Thi Hong Tran (Osaka City University (OCU), Japan; Nara Institute of Science and Technology (NAIST), Japan), Vu Trung Duong Le (Nara Institute of Science and Technology (NAIST), Japan), and Yasuhiko Nakashima (Nara Institute of Science and Technology (NAIST), Japan)</i>	
Twenty Years of Automated Methods for Mapping Applications on CGRA	679
<i>Kevin J. M. Martin (Université de Bretagne Sud, France)</i>	

ADOPT: AI for Datacenter Operations

Introduction to ADOPT 2022	687
<i>Stephanie Brink (LLNL)</i>	
When and How to Retrain Machine Learning-Based Cloud Management Systems	688
<i>Lidia Kidane (Umeå University, Sweden), Paul Townend (Umeå University, Sweden), Thijs Metsch (Intel Corporation, Germany), and Erik Elmroth (Umeå University, Sweden)</i>	
Scalable Data Parallel Distributed Training for Graph Neural Networks	699
<i>Sohei Koyama (University of Tsukuba, Japan) and Osamu Tatebe (University of Tsukuba, Japan)</i>	
The MIT Supercloud Workload Classification Challenge	708
<i>Benny J. Tang (MIT, USA), Qiqi Chen (MIT, USA), Matthew L. Weiss (MIT Lincoln Laboratory, USA), Nathan Frey (MIT Lincoln Laboratory, USA), Joseph McDonald (MIT Lincoln Laboratory, USA), David Bestor (MIT Lincoln Laboratory, USA), Charles Yee (MIT Lincoln Laboratory, USA), William Arcand (MIT Lincoln Laboratory, USA), William Bergeron (MIT Lincoln Laboratory, USA), Chansup Byun (MIT Lincoln Laboratory, USA), Daneil Edelman (MIT, USA), Michael Houle (MIT Lincoln Laboratory, USA), Matthew Hubbell (MIT Lincoln Laboratory, USA), Michael Jones (MIT Lincoln Laboratory, USA), Jeremy Kepner (MIT Lincoln Laboratory, USA), Anna Klein (MIT Lincoln Laboratory, USA), Adam Michaleas (MIT Lincoln Laboratory, USA), Peter Michaleas (MIT Lincoln Laboratory, USA), Lauren Milechin (MIT, USA), Julia Mullen (MIT Lincoln Laboratory, USA), Andrew Prout (MIT Lincoln Laboratory, USA), Albert Reuther (MIT Lincoln Laboratory, USA), Antonio Rosa (MIT Lincoln Laboratory, USA), Andrew Bowne (US Air Force, USA), Lindsey McEvoy (US Air Force, USA), Baolin Li (Northeastern University, USA), Devesh Tiwari (Northeastern University, USA), Vijay Gadepally (MIT Lincoln Laboratory, USA), and Siddharth Samsi (MIT Lincoln Laboratory, USA)</i>	
Loss Curve Approximations for Fast Neural Architecture Ranking & Training Elasticity Estimation	715
<i>Dan Zhao (Massachusetts Institute of Technology, USA), Nathan C. Frey (Massachusetts Institute of Technology, USA), Vijay Gadepally (Massachusetts Institute of Technology, USA), and Siddharth Samsi (Massachusetts Institute of Technology, USA)</i>	
Characterizing Multi-instance GPU for Machine Learning Workloads	724
<i>Baolin Li (Northeastern University, USA), Vijay Gadepally (MIT Lincoln Laboratory, USA), Siddharth Samsi (MIT Lincoln Laboratory, USA), and Devesh Tiwari (Northeastern University, USA)</i>	

Energy-Aware Neural Architecture Selection and Hyperparameter Optimization	732
<i>Nathan C. Frey (Massachusetts Institute of Technology, USA), Dan Zhao (Massachusetts Institute of Technology, USA), Simon Axelrod (Massachusetts Institute of Technology, USA; Harvard University, USA), Michael Jones (Massachusetts Institute of Technology, USA), David Bestor (Massachusetts Institute of Technology, USA), Vijay Gadepally (Massachusetts Institute of Technology, USA), Rafael Gómez-Bombarelli (Massachusetts Institute of Technology, USA), and Siddharth Samsi (Massachusetts Institute of Technology, USA)</i>	
A Green(er) World for A.I.	742
<i>Dan Zhao (MIT Lincoln Laboratory, USA), Nathan C. Frey (MIT Lincoln Laboratory, USA), Joseph McDonald (MIT Lincoln Laboratory, USA), Matthew Hubbell (MIT Lincoln Laboratory, USA), David Bestor (MIT Lincoln Laboratory, USA), Michael Jones (MIT Lincoln Laboratory, USA), Andrew Prout (MIT Lincoln Laboratory, USA), Vijay Gadepally (MIT Lincoln Laboratory, USA), and Siddharth Samsi (MIT Lincoln Laboratory, USA)</i>	

PDCO: Parallel / Distributed Combinatorics and Optimization

Introduction to PDCO 2022	751
<i>Grégoire Danoy (University of Luxembourg, Luxembourg) and Didier El Baz (LAAS-CNRS, France)</i>	
PDCO 2022 Keynote Talk	753
<i>Georges Da Costa (Toulouse Institute for Research in Computer Science)</i>	
Exact k-way Sparse Matrix Partitioning	754
<i>Engelina L. Jenneskens (Utrecht University, The Netherlands) and Rob H. Bisseling (Utrecht University, The Netherlands)</i>	
A Family of Fast Parallel Greedy Algorithms for the Steiner Forest Problem	764
<i>Laleh Ghalami (Wayne State University, USA) and Daniel Grosu (Wayne State University, USA)</i>	
Parallel Minimum Spanning Tree Algorithms via Lattice Linear Predicate Detection	774
<i>David R. Alves (The University of Texas at Austin, USA) and Vijay K. Garg (The University of Texas at Austin, USA)</i>	
A Local Search for Automatic Parameterization of Distributed Tree Search Algorithms	783
<i>Tiago Carneiro (University of Luxembourg, Luxembourg), Loizos Koutsantonis (University of Luxembourg, Luxembourg), Nouredine Melab (Université de Lille, France), Emmanuel Kieffer (University of Luxembourg, Luxembourg), and Pascal Bouvry (University of Luxembourg, Luxembourg)</i>	
Parallel Bayesian Optimization for Optimal Scheduling of Underground Pumped Hydro-Energy Storage Systems	790
<i>Gobert Maxime (University of Mons, Belgium), Gmys Jan (Université de Lille, France), Toubreau Jean-François (University of Mons, Belgium), Melab Nouredine (Université de Lille, France), Tuyttens Daniel (University of Mons, Belgium), and Vallée François (University of Mons, Belgium)</i>	

A Parallel Novelty Search Metaheuristic Applied to a Wildfire Prediction System	798
<i>Jan Strappa (Consejo Nacional de Investigaciones Científicas y Técnicas (CONICET), Argentina; Universidad Tecnológica Nacional, Argentina), Paola Caymes-Scutari (Consejo Nacional de Investigaciones Científicas y Técnicas (CONICET), Argentina; Universidad Tecnológica Nacional, Argentina), and Germán Bianchini (Universidad Tecnológica Nacional, Argentina)</i>	
On Parallel or Distributed Asynchronous Iterations with Unbounded Delays and Possible Out of Order Messages or Flexible Communication for Convex Optimization Problems and Machine Learning	807
<i>Didier El-Baz (Université de Toulouse, France)</i>	

JSSPP: Job Scheduling Strategies for Parallel Processing

Introduction to JSSPP 2022	814
<i>Dalibor Klusáček (CESNET), Julita Corbalán (Barcelona Supercomputing Center, Spain), and Gonzalo P. Rodrigo (Apple)</i>	

PDSEC: Parallel and Distributed Scientific and Engineering Computing

Introduction to PDSEC 2022	816
<i>Sabine Roller (German Aerospace Center, Germany) and Peter Strazdins (The Australian National University, Australia)</i>	
PLSSVM: A (multi-)GPU-Accelerated Least Squares Support Vector Machine	818
<i>Alexander Van Craen (University of Stuttgart, Germany), Marcel Breyer (University of Stuttgart, Germany), and Dirk Pflüger (University of Stuttgart, Germany)</i>	
Least Squares on GPUs in Multiple Double Precision	828
<i>Jan Verschelde (University of Illinois at Chicago, USA)</i>	
A Simple, Fast, and GPU-Friendly Steiner-Tree Heuristic	838
<i>Alex Fallin (Texas State University, USA), Aarti Kothari (Texas State University, USA), Jiayuan He (University of Texas at Austin, USA), Christopher Yanez (Texas State University, USA), Keshav Pingali (University of Texas at Austin, USA), Rajit Manohar (Yale University, USA), and Martin Burtscher (Texas State University, USA)</i>	
Performance Evaluation of a Supercomputer Based on AMD Rome and Intel Cascade Lake Processors	848
<i>Subhash Saini (NASA Ames Research Center, USA), John Baron (Hewlett Packard Enterprise, USA), Johnny Chang (NASA Ames Research Center, USA), Robert Hood (NASA Ames Research Center, USA), and Haoqiang Jin (NASA Ames Research Center, USA)</i>	
A Scalable Parallel Partition Tridiagonal Solver for Many-Core and Low B/F Processors	860
<i>Tatsuya Mitsuda (Kyushu University, Japan) and Kenji Ono (Kyushu University, Japan)</i>	

OMB-Py: Python Micro-Benchmarks for Evaluating Performance of MPI Libraries on HPC Systems.....	870
<i>Nawras Alnaasan (The Ohio State University, USA), Arpan Jain (The Ohio State University, USA), Aamir Shafi (The Ohio State University, USA), Hari Subramoni (The Ohio State University, USA), and Dhabaleswar K. Panda (The Ohio State University, USA)</i>	
Machine Learning for CUDA+MPI Design Rules	880
<i>Carl Pearson (Sandia National Labs, USA), Aurya Javeed (Sandia National Labs, USA), and Karen Devine (Sandia National Labs, USA)</i>	
Using Performance Attributes for Managing Heterogeneous Memory in HPC Applications	890
<i>Brice Goglin (Inria, LaBRI, Univ. Bordeaux, France) and Andrès Rubio Proaño (Inria, LaBRI, Univ. Bordeaux, France)</i>	
Synchronous Parallel Multisplitting Method with Convergence Acceleration using a Local Krylov-Based Minimization for Solving Linear Systems	900
<i>Medane A. Tchakorom (Univ. Bourgogne Franche-Comte (UBFC), France), Raphaël Couturier (Univ. Bourgogne Franche-Comte (UBFC), France), and Jean-Claude Charr (Univ. Bourgogne Franche-Comte (UBFC), France)</i>	
MultiGrid on FPGA using Data Parallel C++	907
<i>Christopher Siefert (Sandia National Laboratories, USA), Stephen L. Olivier (Sandia National Laboratories, USA), Gwendolyn Voskuilen (Sandia National Laboratories, USA), and Jeffrey Young (Georgia Institute of Technology, USA)</i>	

iWAPT: Automatic Performance Tuning

Introduction to iWAPT 2022	911
<i>Che-Rung Lee (National Tsing Hua University, Taiwan) and Satoshi Ohshima (Nagoya University, Japan)</i>	
Qual ^L 2M: Learning Quantitative Performance of Latency-Sensitive Code	913
<i>Arun Sathanur (Pacific Northwest National Laboratory, USA), Nathan R. Tallent (Pacific Northwest National Laboratory, USA), Patrick Konsor (Intel Corporation, USA), Ken Koyanagi (Intel Corporation, USA), Ryan McLaughlin (Intel Corporation, USA), Joseph Olivas (Intel Corporation, USA), and Michael Chynoweth (Intel Corporation, USA)</i>	
Benchmarking the Linear Algebra Awareness of TensorFlow and PyTorch	924
<i>Aravind Sankaran (RWTH Aachen University, Germany), Navid Akbari Alashti (RWTH Aachen University, Germany), Christos Psarras (RWTH Aachen University, Germany), and Paolo Bientinesi (Umeå Universitet, Sweden)</i>	
Automated Selection of Build Configuration Based on Machine Learning	934
<i>Reo Furuhashi (Tohoku University, Japan), Minglu Zhao (Tohoku University, Japan), Keichi Takahashi (Tohoku University, Japan), Yoichi Shimomura (Tohoku University, Japan), and Hiroyuki Takizawa (Tohoku University, Japan)</i>	

A Cost Model for Compilers Based on Transfer Learning	942
<i>Yuta Sasaki (Tohoku University, Japan), Keichi Takahashi (Cyberscience Center, Tohoku University, Japan; Tohoku University, Japan), Yoichi Shimomura (Cyberscience Center, Tohoku University, Japan), and Hiroyuki Takizawa (Cyberscience Center, Tohoku University, Japan; Tohoku University, Japan)</i>	
Modeling pre-Exascale AMR Parallel I/O Workloads via Proxy Applications	952
<i>William F. Godoy (Oak Ridge National Laboratory, USA), Jenna Delozier (Georgia Institute of Technology, USA), and Gregory R. Watson (Oak Ridge National Laboratory, USA)</i>	
Smoothing on Dynamic Concurrency Throttling	962
<i>Janaina Schwarzrock (Federal University of Rio Grande do Sul, Brazil), Hiago Mayk G. de A. Rocha (Federal University of Rio Grande do Sul, Brazil), Arthur F. Lorenzon (Federal University of Pampa, Brazil), and Antonio Carlos S. Beck (Federal University of Rio Grande do Sul, Brazil)</i>	
Analyzing Search Techniques for Autotuning Image-Based GPU Kernels: The Impact of Sample Sizes	972
<i>Jacob O Tørring (Norwegian University of Science and Technology (NTNU), Norway) and Anne C. Elster (Norwegian University of Science and Technology (NTNU), Norway)</i>	

PAISE: Parallel AI and Systems for the Edge

Introduction to PAISE 2022	982
<i>Istemi Ekin Akkus (Nokia Bell Labs, Germany) and Nirmitt V. Desai (IBM Research, USA)</i>	
Don't Miss the Train: A Case for Systems Research into Training on the Edge	985
<i>Prashanthi S.K (Indian Institute of Science, India), Aakash Khochare (Indian Institute of Science, India), Sai Anuroop Kesanapalli (Indian Institute of Science, India), Rahul Bhope (Indian Institute of Science, India), and Yogesh Simmhan (Indian Institute of Science, India)</i>	
Litener: An Accelerator-Enabled Lightweight Container for Edge Computing	987
<i>Ryan Dyson (Queen's University Belfast, United Kingdom) and Carlos Reaño (Universitat de València, Spain)</i>	
Efficient Volume Estimation for Dynamic Environments using Deep Learning on the Edge	995
<i>Chandan Kumar (Iowa State University, USA), Yamini Mathur (Iowa State University, USA), and Ali Jannesari (Iowa State University, USA)</i>	
TinyMLOps: Operational Challenges for Widespread Edge AI Adoption	1003
<i>Sam Leroux (hotg.ai; Ghent University), Pieter Simoens (Ghent University), Meelis Lootus (hotg.ai), Kartik Thakore (hotg.ai), and Akshay Sharma (hotg.ai)</i>	

RADR: Resource Arbitration for Dynamic Runtimes

Introduction to RADR 2022	1011
<i>Pete Beckman (Argonne National Laboratory, USA), Emmanuel Jeannot (TADaaM Team, Inria, France), and Swann Perarnau (Argonne National Laboratory, USA)</i>	
Performance Analysis of Multi-containerized MD Simulations for Low-Level Resource Allocation	1014
<i>Shingo Okuno (FUJITSU LIMITED, Japan), Akira Hirai (FUJITSU LIMITED, Japan), and Naoto Fukumoto (FUJITSU LIMITED, Japan)</i>	
Operating System Convergence: An Example via the Maru OS Project	1018
<i>William White (Auburn University, USA) and Xiao Qin (Auburn University, USA)</i>	
Combining Uncore Frequency and Dynamic Power Capping to Improve Power Savings	1028
<i>Amina Guermouche (Universite de Bordeaux, Inria, France)</i>	

ScaDL: Scalable Deep Learning over Parallel And Distributed Infrastructures

Introduction to ScaDL 2022	1038
<i>Danilo Ardagna (Politecnico di Milano, Italy) and Stacy Patterson (Rensselaer Polytechnic Institute (RPI), USA)</i>	
ScaDL 2022 Invited Talk 1	1039
<i>William Fornaciari (Politecnico di Milano – DEIB, Italy)</i>	
ScaDL 2022 Invited Talk 2	1040
<i>Mudhakar Srivatsa (IBM Research, USA)</i>	
ScaDL 2022 Invited Talk 3	1041
<i>Anima Anandkumar (Caltech and NVIDIA)</i>	
ScaDL 2022 Invited Talk 4	1042
<i>Michael Gschwind (Meta AI and MLPerf)</i>	
ScaDL 2022 Invited Talk 5	1043
<i>David Kanter (MLCommons)</i>	
ScaDL 2022 Invited Talk 6	1044
<i>Barret Zoph (Google Brain)</i>	
Distributing Deep Learning Hyperparameter Tuning for 3D Medical Image Segmentation	1045
<i>Josep Ll. Berral Garcia (Barcelona Supercomputing Center, Spain), Oriol Aranda (Barcelona Supercomputing Center, Spain), Juan Luis Dominguez (Barcelona Supercomputing Center, Spain), and Jordi Torres (Barcelona Supercomputing Center, Universitat Politècnica de Catalunya, Spain)</i>	
A Methodology to Build Decision Analysis Tools Applied to Distributed Reinforcement Learning	1053
<i>Cédric Prigent (University of Rennes, Inria, CNRS, IRISA, Rennes, France), Loic Cudennec (DGA Maîtrise de l'Information, Rennes, France), Alexandru Costan (University of Rennes, Inria, CNRS, IRISA, Rennes, France), and Gabriel Antoniu (University of Rennes, Inria, CNRS, IRISA, Rennes, France)</i>	

MadPipe: Memory Aware Dynamic Programming Algorithm for Pipelined Model Parallelism	1063
<i>Olivier Beaumont (University of Bordeaux, France), Lionel Eyraud-Dubois (University of Bordeaux, France), and Alena Shilova (University of Lille, France)</i>	
APPFL: Open-Source Software Framework for Privacy-Preserving Federated Learning	1074
<i>Minseok Ryu (Argonne National Laboratory, USA), Youngdae Kim (Argonne National Laboratory, USA), Kibaek Kim (Argonne National Laboratory, USA), and Ravi K. Madduri (Argonne National Laboratory, USA)</i>	
Throughput-Oriented and Accuracy-Aware DNN Training with BFloat16 on GPU	1084
<i>Zhen Xie (Argonne National Laboratory), Siddhisanket Raskar (Argonne National Laboratory), and Murali Emani (Argonne National Laboratory)</i>	
Adaptive Optimization for Sparse Data on Heterogeneous GPUs	1088
<i>Yujing Ma (University of California Merced, USA), Florin Rusu (University of California Merced, USA), Kesheng Wu (Lawrence Berkeley National Laboratory, USA), and Alexander Sim (Lawrence Berkeley National Laboratory, USA)</i>	

ESSA: Extreme-Scale Storage and Analysis (formerly HPS)

Introduction to ESSA 2022	1098
<i>Osamu Tatebe (University of Tsukuba, Japan) and Gabriel Antoniu (Inria, France)</i>	
ESSA 2022 Keynote Speaker	1100
<i>Rob Ross (Argonne National Laboratory, USA)</i>	
ESSA 2022 Invited Speaker 1	1101
<i>Johann Lombardi (SCG Storage, Intel)</i>	
ESSA 2022 Invited Speaker 2	1102
<i>Lavanya Ramakrishnan (Lawrence Berkeley National Laboratory, USA)</i>	
Caching Support for CHFS Node-Local Persistent Memory File System	1103
<i>Osamu Tatebe (University of Tsukuba, Japan) and Hiroki Ohtsuji (Fujitsu Limited, Japan)</i>	
A Locality-Aware Cooperative Distributed Memory Caching for Parallel Data Analytic Applications	1111
<i>Chia-Ting Hung (National Tsing Hua University, Taiwan), Jerry Chou (National Tsing Hua University, Taiwan), Ming-Hung Chen (IBM T. J. Watson, USA), and I-Hsin Chung (IBM T. J. Watson, USA)</i>	
Modeling Power Consumption of Lossy Compressed I/O for Exascale HPC Systems	1118
<i>Grant Wilkins (Clemson University, USA) and Jon C. Calhoun (Clemson University, USA)</i>	

ParSocial: Parallel and Distributed Processing for Computational Social Systems

Introduction to ParSocial 2022	1127
<i>John Korah (California State Polytechnic University Pomona, USA) and Eunice E. Santos (University of Illinois at Urbana-Champaign, USA)</i>	

Dynamic Batch Parallel Algorithms for Updating PageRank	1129
<i>Subhajit Sahu (International Institute of Information Technology Hyderabad, India), Kishore Kothapalli (International Institute of Information Technology Hyderabad, India), and Dip Sankar Banerjee (Indian Institute of Technology Jodhpur, India)</i>	
Distributed Algorithms for the Graph Biconnectivity and Least Common Ancestor Problems	1139
<i>Ian Bogle (Rensselaer Polytechnic Institute) and George M. Slota (Rensselaer Polytechnic Institute)</i>	
Efficient Parallel PageRank Algorithm for Network Analysis	1143
<i>Maxence Vandromme (Univ. Lille, UMR 9189 - CRIStAL, CNRS, France) and Serge G. Petiton (Univ. Lille, UMR 9189 - CRIStAL, CNRS, France)</i>	
A Streaming System for Large-Scale Temporal Graph Mining of Reddit Data	1153
<i>Andreas Huber (Simula Research Laboratory, University of Oslo, Norway), Daniel Thilo Schroeder (Simula Research Laboratory, Norway), Konstantin Pogorelov (Simula Research Laboratory, Norway), Carsten Griwodz (University of Oslo, Norway), and Johannes Langguth (Simula Research Laboratory, Norway)</i>	
Unsupervised User Stance Detection on Tweets Against Web Articles using Sentence Transformers	1163
<i>Bhashithe Abeyasinghe (Georgia State University, USA), Gyandeep Reddy Vulupala (Georgia State University, USA), Anu G. Bourgeois (Georgia State University, USA), and Rajshekhar Sunderraman (Georgia State University, USA)</i>	
Effect of Community-Based Opinion Leaders on Guideline Dissemination in Large-Scale Physician Networks	1170
<i>Vairavan Murugappan (University of Illinois at Urbana-Champaign, USA), Suresh Subramanian (University of Illinois at Urbana-Champaign, USA), John Korah (California State Polytechnic University, Pomona, USA), Pranav Pamidighantam (University of Illinois at Urbana-Champaign, USA), and Eunice E. Santos (University of Illinois at Urbana-Champaign, USA)</i>	

EDAML: Electronic Design Automation and Machine Learning

Introduction to EDAML 2022	1180
<i>Sudeep Pasricha (Colorado State University, USA) and Muhammad Shafique (New York University, Abu Dhabi)</i>	
EDAML 2022 Keynote Speaker	1181
<i>David Z. Pan (The University of Texas at Austin, USA)</i>	
EDAML 2022 Invited Speaker 1	1182
<i>Ankush Sood (Cadence Design Systems)</i>	
EDAML 2022 Invited Speaker 2	1183
<i>Deming Chen (University of Illinois at Urbana-Champaign)</i>	
EDAML 2022 Invited Speaker 3	1184
<i>R. Iris Bahar (Colorado School of Mines, USA)</i>	
EDAML 2022 Invited Speaker 4	1185
<i>Krishnendu Chakrabarty (Duke University)</i>	

EDAML 2022 Invited Speaker 5	1186
<i>Laleh Behjat (University of Calgary, Alberta, Canada)</i>	
EDAML 2022 Invited Speaker 6	1187
<i>Partha Pratim Pande (Washington State University, USA)</i>	
EDAML 2022 Invited Speaker 7	1188
<i>Sachin S. Sapatnekar (University of Minnesota, USA)</i>	
EDAML 2022 Invited Speaker 8	1189
<i>Muhammad Shafique (New York University (NYU) Abu Dhabi, UAE)</i>	
EDAML 2022 Invited Speaker 9	1190
<i>Sheldon Tan (University of California at Riverside)</i>	
EDAML 2022 Invited Speaker 10	1191
<i>Sudeep Pasricha (Colorado State University, USA)</i>	

COMPSYS: Composable Systems

Introduction to COMPSYS 2022	1192
<i>Christian Pinto (IBM Research Europe, Ireland), Felix Eberhardt (Hasso Plattner Institute, University of Potsdam, Germany), and Sergey Blagodurov (AMD, USA)</i>	
COMPSYS 2022 Keynote Talk	1194
<i>Paul Carpenter (Barcelona Supercomputing Center, Spain)</i>	
Quantifying Composable Data Center Utilization	1195
<i>Marc Taubenblatt (IBM Thomas J. Watson Research Center, USA) and Asser Tantawi (IBM Thomas J. Watson Research Center, USA)</i>	
Separated Allocator Metadata in Disaggregated In-Memory Databases: Friend or Foe?	1202
<i>Marcel Weisgut (Hasso Plattner Institute, University of Potsdam, Germany), Daniel Ritter (Hasso Plattner Institute, University of Potsdam, Germany), Martin Boissier (Hasso Plattner Institute, University of Potsdam, Germany), and Michael Perscheid (Hasso Plattner Institute, University of Potsdam, Germany)</i>	
Composable Infrastructures for an Academic Research Environment: Lessons Learned	1209
<i>Lance Long (University of Illinois Chicago, USA), Timothy Bargo (University of Illinois Chicago, USA), Luc Renambot (University of Illinois Chicago, USA), Maxine Brown (University of Illinois Chicago, USA), and Andrew Johnson (University of Illinois Chicago, USA)</i>	
Moving from Composable to Programmable	1215
<i>Zhongyi Chen (University of Illinois Chicago, USA), Luc Renambot (University of Illinois Chicago, USA), Lance Long (University of Illinois Chicago, USA), Maxine Brown (University of Illinois Chicago, USA), and Andrew Johnson (University of Illinois Chicago, USA)</i>	

Evaluating Hardware Memory Disaggregation Under Delay and Contention	1221
<i>Archit Patke (University of Illinois at Urbana-Champaign, USA), Haoran Qiu (University of Illinois at Urbana-Champaign, USA), Saurabh Jha (IBM Research, USA), Srikumar Venugopal (IBM Research, Ireland), Michele Gazzetti (IBM Research, Ireland), Christian Pinto (IBM Research, Ireland), Zbigniew Kalbarczyk (University of Illinois at Urbana-Champaign, USA), and Ravishankar Iyer (University of Illinois at Urbana-Champaign, USA)</i>	
Memory-Disaggregated In-Memory Object Store Framework for Big Data Applications	1228
<i>Robin Abrahamse (Delft University of Technology, The Netherlands), Ákos Hadnagy (Delft University of Technology, The Netherlands), and Zaid Al-Ars (Delft University of Technology, The Netherlands)</i>	

CORtEX: Computing using EmeRging Exotic AI-Inspired Systems

Introduction to CORtEX 2022	1235
<i>Artur Podobas (KTH Royal Institute of Technology, Sweden), Aleksandr Drozd (RIKEN Center for Computational Science, Japan), Catherine Schuman (University of Tennessee, USA), and Barry Devereux (Queen's University Belfast, United Kingdom)</i>	
CORtEX 2022 Invited Speaker 1	1237
<i>Thomas Nowotny (University of Sussex, United Kingdom) and James Knight (University of Sussex, United Kingdom)</i>	
CORtEX 2022 Invited Speaker 2	1238
<i>Anders Lansner (KTH Royal Institute of Technology, Sweden)</i>	
CORtEX 2022 Invited Speaker 3	1239
<i>Oliver Rhodes (University of Manchester, United Kingdom)</i>	
CORtEX 2022 Invited Speaker 4	1240
<i>Jun Igarashi (RIKEN Center for Computational Science, Japan)</i>	
CORtEX 2022 Invited Speaker 5	1241
<i>Lawrence Spracklen (Numenta, USA)</i>	

ExSAIS: Extreme Scaling of AI for Science

Introduction to ExSAIS 2022	1242
<i>Svitlana Volkova (PNNL) and Robert Rallo (PNNL)</i>	
ExSAIS 2022 Keynote Talk 1	1244
<i>Gennady Pekhimenko (University of Toronto, Canada)</i>	
ExSAIS 2022 Keynote Talk 2	1245
<i>Mostofa Patwary (NVIDIA)</i>	

Learning to Scale the Summit: AI for Science on a Leadership Supercomputer	1246
<i>Wayne Joubert (Oak Ridge National Laboratory, USA), Bronson Messer (Oak Ridge National Laboratory, USA), Philip C. Roth (Oak Ridge National Laboratory, USA), Antigoni Georgiadou (Oak Ridge National Laboratory, USA), Justin Lietz (Oak Ridge National Laboratory, USA), Markus Eisenbach (Oak Ridge National Laboratory, USA), and Junqi Yin (Oak Ridge National Laboratory, USA)</i>	
Strategies for Integrating Deep Learning Surrogate Models with HPC Simulation Applications....	1256
<i>Junqi Yin (Oak Ridge National Laboratory, USA), Feiyi Wang (Oak Ridge National Laboratory, USA), and Mallikarjun (Arjun) Shankar (Oak Ridge National Laboratory, USA)</i>	
A Scalable Pipeline for Gigapixel Whole Slide Imaging Analysis on Leadership Class HPC Systems	1266
<i>Sajal Dash (Oak Ridge National Laboratory, USA), Benjamín Hernández (Oak Ridge National Laboratory, USA), Aristeidis Tsaris (Oak Ridge National Laboratory, USA), Folami T. Alamudun (Oak Ridge National Laboratory, USA), Hong-Jun Yoon (Oak Ridge National Laboratory, USA), and Feiyi Wang (Oak Ridge National Laboratory, USA)</i>	

IPDPS 2022 PhD Forum

IPDPS 2022 PhD Forum Welcome and Abstracts	1275
<i>Sanjukta Bhowmick (University of North Texas, USA) and Anne-Cecile Orgerie (IRISA, France)</i>	

Author Index