

2023 IEEE International Parallel and Distributed Processing Symposium (IPDPS 2023)

**St. Petersburg, Florida, USA
15-19 May 2023**

Pages 1-523



**IEEE Catalog Number: CFP23023-POD
ISBN: 979-8-3503-3767-9**

**Copyright © 2023 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:	CFP23023-POD
ISBN (Print-On-Demand):	979-8-3503-3767-9
ISBN (Online):	979-8-3503-3766-2
ISSN:	1530-2075

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

2023 IEEE International Parallel and Distributed Processing Symposium (IPDPS) **IPDPS 2023**

Table of Contents

Message from the IPDPS 2023 General Co-chairs	xix
Message from the IPDPS 2023 Program Chairs	xxi
IPDPS 2023 Technical Program	xxiii
IPDPS 2023 Organization	xxv
IPDPS 2023 Technical Program Committee	xxvii

Keynote Talk I

Keynote: Fifty Years of Parallel Programming: Ieri, Oggi, Domani or Yesterday, Today, Tomorrow	1
<i>Keshav Pingali (The University of Texas at Austin and Katana Graph Inc.)</i>	

SESSION 1A: Graphs Processing

GraphTensor: Comprehensive GNN-Acceleration Framework for Efficient Parallel Processing of Massive Datasets	2
<i>Junhyeok Jang (Korea Advanced Institute of Science and Technology (KAIST)), Miryeong Kwon (Korea Advanced Institute of Science and Technology (KAIST)), Donghyun Gouk (Korea Advanced Institute of Science and Technology (KAIST)), Hanyeoreum Bae (Korea Advanced Institute of Science and Technology (KAIST)), and Myoungsoo Jung (Korea Advanced Institute of Science and Technology (KAIST))</i>	
GraphMetaP: Efficient MetaPath Generation for Dynamic Heterogeneous Graph Models	13
<i>Haiheng He (Huazhong University of Science and Technology, China; Zhejiang Lab, China), Dan Chen (Huazhong University of Science and Technology, China; Zhejiang Lab, China), Long Zheng (Huazhong University of Science and Technology, China; Zhejiang Lab, China), Yu Huang (Huazhong University of Science and Technology, China; Zhejiang Lab, China), Haifeng Liu (Huazhong University of Science and Technology, China; Zhejiang Lab, China), Chaoqiang Liu (Huazhong University of Science and Technology, China; Zhejiang Lab, China), Xiaofei Liao (Huazhong University of Science and Technology, China), and Hai Jin (Huazhong University of Science and Technology, China)</i>	

Traversing Large Compressed Graphs on GPUs	25
<i>Prasun Gera (Cerebras Systems) and Hyesoon Kim (Georgia Tech)</i>	
Distributed Sparse Random Projection Trees for Constructing K-Nearest Neighbor Graphs	36
<i>Isuru Ranawaka (Indiana University, USA), Md Khaledur Rahman (Meta Inc.), and Ariful Azad (Indiana University, USA)</i>	
Fast Deterministic Gathering with Detection on Arbitrary Graphs: The Power of Many Robots	47
<i>Anisur Rahaman Molla (Indian Statistical Institute, India), Kaushik Mondal (Indian Institute of Technology Ropar, India), and William K. Moses Jr. (Durham University, UK)</i>	
Accurate and Efficient Distributed COVID-19 Spread Prediction Based on a Large-Scale Time-Varying People Mobility Graph	58
<i>Sudipta Saha Shubha (University of Virginia, USA), Shohaib Mahmud (University of Virginia, USA), Haiying Shen (University of Virginia, USA), Geoffrey C. Fox (University of Virginia, USA), and Madhav Marathe (University of Virginia, USA)</i>	

SESSION 1B: Architectural Advances

H-Cache: Traffic-Aware Hybrid Rule-Caching in Software-Defined Networks	69
<i>Zeyu Luan (Tsinghua University, China; Peng Cheng Laboratory (PCL), China), Qing Li (Peng Cheng Laboratory (PCL), China), Yi Wang (Tsinghua University, China), and Yong Jiang (Tsinghua University, China; Peng Cheng Laboratory (PCL), China)</i>	
Accelerating Packet Processing in Container Overlay Networks via Packet-Level Parallelism	79
<i>Jiaxin Lei (Binghamton University, USA), Manish Munikar (The University of Texas at Arlington, USA), Hui Lu (Binghamton University, USA), and Rao Jia (The University of Texas at Arlington, USA)</i>	
Software-Defined, Fast and Strongly-Consistent Data Replication for RDMA-Based PM Datastores	90
<i>Haodi Lu (Huazhong University of Science and Technology, China), Haikun Liu (Huazhong University of Science and Technology, China), Chencheng Ye (Huazhong University of Science and Technology, China), Xiaofei Liao (Huazhong University of Science and Technology, China), Fubing Mao (Huazhong University of Science and Technology, China), Yu Zhang (Huazhong University of Science and Technology, China), and Hai Jin (Huazhong University of Science and Technology, China)</i>	
Signal Detection for Large MIMO Systems using Sphere Decoding on FPGAs	102
<i>Mohamed W. Hassan (King Abdullah University of Science and Technology (KAUST), Saudi Arabia), Adel Dabah (King Abdullah University of Science and Technology (KAUST), Saudi Arabia), Hatem Ltaief (King Abdullah University of Science and Technology (KAUST), Saudi Arabia), and Suhaib A. Fahmy (King Abdullah University of Science and Technology (KAUST), Saudi Arabia)</i>	
Efficient Hardware Primitives for Immediate Memory Reclamation in Optimistic Data Structures	112
<i>Ajay Singh (University of Waterloo), Trevor Brown (University of Waterloo), and Michael Spear (Lehigh University)</i>	

A Novel Framework for Efficient Offloading of Communication Operations to Bluefield SmartNICs	123
<i>Kaushik Kandadi Suresh (The Ohio State University, USA), Benjamin Michalowicz (The Ohio State University, USA), Bharath Ramesh (The Ohio State University, USA), Nick Contini (The Ohio State University, USA), Jinghan Yao (The Ohio State University, USA), Shulei Xu (The Ohio State University, USA), Aamir Shafi (The Ohio State University, USA), Hari Subramoni (The Ohio State University, USA), and Dhabaleswar Panda (The Ohio State University, USA)</i>	

SESSION 2A: HPC Optimizations for ML

Accelerating Distributed Deep Learning Training with Compression Assisted Allgather and Reduce-Scatter Communication	134
<i>Qinghua Zhou (The Ohio State University, USA), Quentin Anthony (The Ohio State University, USA), Lang Xu (The Ohio State University, USA), Aamir Shafi (The Ohio State University, USA), Mustafa Abduljabbar (The Ohio State University, USA), Hari Subramoni (The Ohio State University, USA), and Dhabaleswar K. Panda (The Ohio State University, USA)</i>	
Accelerating CNN Inference on Long Vector Architectures via co-Design	145
<i>Sonia Rani Gupta (Chalmers University of Technology, Sweden), Nikela Papadopoulou (Chalmers University of Technology, Sweden), and Miquel Pericàs (Chalmers University of Technology, Sweden)</i>	
Exploiting Input Tensor Dynamics in Activation Checkpointing for Efficient Training on GPU.....	156
<i>Jianjin Liao (Beihang University, China), Mingzhen Li (Beihang University, China), Hailong Yang (Beihang University, China), Qingxiao Sun (Beihang University, China), Biao Sun (Beihang University, China), Jiwei Hao (Beihang University, China), Tianyu Feng (Beihang University, China), Fengwei Yu (SenseTime Research, China), Shengdong Chen (SenseTime Research, China), Ye Tao (SenseTime Research, China), Zicheng Zhang (SenseTime Research, China), Zhongzhi Luan (Beihang University, China), and Depei Qian (Beihang University, China)</i>	
MPipeMoE: Memory Efficient MoE for Pre-Trained Models with Adaptive Pipeline Parallelism	167
<i>Zheng Zhang (Wuhan University), Donglin Yang (Nvidia Corp.), Yaqi Xia (Wuhan University), Liang Ding (JD Explore Academy, JD.com Inc.), Dacheng Tao (JD Explore Academy, JD.com Inc.), Xiaobo Zhou (University of Macau), and Dazhao Cheng (Wuhan University)</i>	

SESSION 2B: I/O Optimizations

Mimir: Extending I/O Interfaces to Express User Intent for Complex Workloads in HPC	178
<i>Hariharaan Devarajan (Lawrence Livermore National Laboratory) and Kathryn Mohror (Lawrence Livermore National Laboratory)</i>	

Drill: Log-Based Anomaly Detection for Large-Scale Storage Systems using Source Code Analysis	189
<i>Di Zhang (University of North Carolina at Charlotte), Chris Egersdoerfer (University of North Carolina at Charlotte), Tabassum Mahmud (Iowa State University), Mai Zheng (Iowa State University), and Dong Dai (University of North Carolina at Charlotte)</i>	
FaultyRank: A Graph-Based Parallel File System Checker	200
<i>Saisha Kamat (University of North Carolina at Charlotte), Abdullah Al Raqibul Islam (University of North Carolina at Charlotte), Mai Zheng (Iowa State University), and Dong Dai (University of North Carolina at Charlotte)</i>	
Evaluating Asynchronous Parallel I/O on HPC Systems	211
<i>John Ravi (North Carolina State University), Suren Byna (Lawrence Berkeley National Laboratory), Quincey Koziol (Lawrence Berkeley National Laboratory), Houjun Tang (Lawrence Berkeley National Laboratory), and Michela Becchi (North Carolina State University)</i>	

SESSION 3A: Large Scale ML

An Efficient 2D Method for Training Super-Large Deep Learning Models	222
<i>Qifan Xu (University of California, USA) and Yang You (National University of Singapore, Singapore)</i>	
Dynaspase: Accelerating GNN Inference Through Dynamic Sparsity Exploitation	233
<i>Bingyi Zhang (University of Southern California, USA) and Viktor Prasanna (University of Southern California, USA)</i>	
Exploiting Sparsity in Pruned Neural Networks to Optimize Large Model Training	245
<i>Siddharth Singh (University of Maryland, USA) and Abhinav Bhatle (University of Maryland, USA)</i>	
Asynch-SGBDT: Train Stochastic Gradient Boosting Decision Trees in an Asynchronous Parallel Manner	256
<i>Daning Cheng (Tsinghua University, China), Shigang Li (Beijing University of Posts and Telecommunications, China), and Yunquan Zhang (Chinese Academy of Sciences, China)</i>	

SESSION 3B: New Systems for Storage

SRC: Mitigate I/O Throughput Degradation in Network Congestion Control of Disaggregated Storage Systems	268
<i>Danlin Jia (Samsung Semiconductor Inc., USA), Yiming Xie (Northeastern University, USA), Li Wang (Northeastern University, USA), Xiaoqian Zhang (University of Massachusetts Boston, USA), Allen Yang (University of Massachusetts Boston, USA), Xuebin Yao (Samsung Semiconductor Inc., USA), Mahsa Bayati (Samsung Semiconductor Inc., USA), Pradeep Subedi (Samsung Semiconductor Inc., USA), Bo Sheng (University of Massachusetts Boston, USA), and Ningfang Mi (Northeastern University, USA)</i>	

Boosting Multi-Block Repair in Cloud Storage Systems with Wide-Stripe Erasure Coding	279
<i>Qi Yu (Huazhong University of Science and Technology), Lin Wang (Huazhong University of Science and Technology), Yuchong Hu (Huazhong University of Science and Technology), Yumeng Xu (Huazhong University of Science and Technology), Dan Feng (Huazhong University of Science and Technology), Jie Fu (Huawei Technologies Co., Ltd., China), Xia Zhu (Huawei Technologies Co., Ltd., China), Zhen Yao (Huawei Technologies Co., Ltd., China), and Wenjia Wei (Huawei Technologies Co., Ltd., China)</i>	
UnifyFS: A User-Level Shared File System for Unified Access to Distributed Local Storage	290
<i>Michael J. Brim (Oak Ridge National Laboratory, USA), Adam T. Moody (Lawrence Livermore National Laboratory, USA), Seung-Hwan Lim (Oak Ridge National Laboratory, USA), Ross Miller (Oak Ridge National Laboratory, USA), Swen Boehm (Oak Ridge National Laboratory, USA), Cameron Stanavige (Lawrence Livermore National Laboratory, USA), Kathryn M. Mohror (Lawrence Livermore National Laboratory, USA), and Sarp Oral (Oak Ridge National Laboratory, USA)</i>	
ArkFS: A Distributed File System on Object Storage for Archiving Data in HPC Environment	301
<i>Kyu-Jin Cho (Seoul National University, South Korea), Injae Kang (Seoul National University, South Korea), and Jin-Soo Kim (Seoul National University, South Korea)</i>	

Best Paper Nominees - Plenary

On Doorway Egress by Autonomous Robots	312
<i>Rory Hector (Louisiana State University, USA), Ramachandran Vaidyanathan (Louisiana State University, USA), Gokarna Sharma (Kent State University, USA), and Jerry L. Trahan (Louisiana State University, USA)</i>	
PAQR: Pivoting Avoiding QR Factorization	322
<i>Wissam Sid-Lakhdar (University of Tennessee), Sebastien Cayrols (University of Tennessee), Daniel Bielich (University of Tennessee), Ahmad Abdelfattah (University of Tennessee), Piotr Luszczek (University of Tennessee), Mark Gates (University of Tennessee), Stanimire Tomov (University of Tennessee), Hans Johansen (Lawrence Berkeley National Laboratory), David Williams-Young (Lawrence Berkeley National Laboratory), Timothy Davis (Texas A&M University), Jack Dongarra (University of Tennessee; Oak Ridge National Laboratory), and Hartwig Anzt (University of Tennessee)</i>	
DeepThermo: Deep Learning Accelerated Parallel Monte Carlo Sampling for Thermodynamics Evaluation of High Entropy Alloys	333
<i>Junqi Yin (Oak Ridge National Laboratory), Feiyi Wang (Oak Ridge National Laboratory), and Mallikarjun Shankar (Oak Ridge National Laboratory)</i>	

ByteTransformer: A High-Performance Transformer Boosted for Variable-Length Inputs	344
<i>Yujia Zhai (University of California, Riverside), Chengquan Jiang (ByteDance Ltd.), Leyuan Wang (ByteDance Ltd.), Xiaoying Jia (ByteDance Ltd.), Shang Zhang (NVIDIA Corporation), Zizhong Chen (University of California, Riverside), Xin Liu (ByteDance Ltd.), and Yibo Zhu (ByteDance Ltd.)</i>	

Keynote Talk II

Keynote: The Adventurous Life of a System Software Researcher	356
<i>Dilma Da Silva (National Science Foundation & Texas A&M University)</i>	

SESSION 4A: Linear Algebra Algorithms

On the Arithmetic Intensity of Distributed-Memory Dense Matrix Multiplication Involving a Symmetric Input Matrix (SYMM)	357
<i>Emmanuel Agullo (Inria), Alfredo Buttari (CNRS), Olivier Coulaud (Inria), Lionel Eyraud-Dubois (Inria), Mathieu Faverge (Bordeaux INP), Alain Franc (INRAE), Abdou Guermouche (Université de Bordeaux), Antoine Jego (INPT), Romain Peressoni (Inria), and Florent Pruvost (Inria)</i>	
A Novel Triangular Space-Filling Curve for Cache-Oblivious In-Place Transposition of Square Matrices	368
<i>João N.F. Alves (University of Vienna; Universidade de Lisboa), Luís M.S. Russo (Universidade de Lisboa), Alexandre P. Francisco (Universidade de Lisboa), and Siegfried Benkner (University of Vienna)</i>	
Memory-Aware Optimization for Sequences of Sparse Matrix-Vector Multiplications	379
<i>Yichen Zhang (National University of Defense Technology, China), Shengguo Li (National University of Defense Technology, China), Fan Yuan (Xiangtan University, China), Dezun Dong (National University of Defense Technology, China), Xiaojian Yang (National University of Defense Technology, China), Tiejun Li (National University of Defense Technology, China), and Zheng Wang (University of Leeds, United Kingdom)</i>	
Data Distribution Schemes for Dense Linear Algebra Factorizations on Any Number of Nodes	390
<i>Olivier Beaumont (Inria Center of the University of Bordeaux, France), Jean-Alexandre Collin (Inria Center of the University of Bordeaux, France), Lionel Eyraud-Dubois (Inria Center of the University of Bordeaux, France), and Mathieu Vérité (Inria Center of the University of Bordeaux, France)</i>	
Dynamic Tensor Linearization and Time Slicing for Efficient Factorization of Infinite Data Streams	402
<i>Yongseok Soh (University of Oregon), Ahmed E. Helal (Intel Labs), Fabio Checconi (Intel Labs), Jan Laukemann (University of Erlangen-Nürnberg), Jesmin Jahan Tithi (Intel Labs), Teresa Ranadive (Laboratory for Physical Sciences), Fabrizio Petrini (Intel Labs), and Jee W. Choi (University of Oregon)</i>	

SESSION 4B: Resource Management

Scheduling with Many Shared Resources	413
<i>Max A. Deppert (Kiel University, Germany), Klaus Jansen (Kiel University, Germany), Marten Maack (Paderborn University, Germany), Simon Pukrop (Paderborn University, Germany), and Malin Rau (University of Hamburg, Germany)</i>	
Chic-Sched: a HPC Placement-Group Scheduler on Hierarchical Topologies with Constraints	424
<i>Laurent Schares (IBM Thomas J. Watson Research Center, USA), Asser Tantawi (IBM Thomas J. Watson Research Center, USA), Pavlos Maniotis (IBM Thomas J. Watson Research Center, USA), Ming-Hung Chen (IBM Thomas J. Watson Research Center, USA), Claudia Misale (IBM Thomas J. Watson Research Center, USA), Seetharami Seelam (IBM Thomas J. Watson Research Center, USA), and Hao Yu (IBM Thomas J. Watson Research Center, USA)</i>	
Generalizable Reinforcement Learning-Based Coarsening Model for Resource Allocation over Large and Diverse Stream Processing Graphs	435
<i>Lanshun Nie (Harbin Institute of Technology), Yuqi Qiu (Harbin Institute of Technology), Fei Meng (Harbin Institute of Technology), Mo Yu (IBM Research), and Jing Li (New Jersey Institute of Technology)</i>	
RLP: Power Management Based on a Latency-Aware Roofline Model	446
<i>Bo Wang (RWTH Aachen University, Germany), Anara Kozhokanova (RWTH Aachen University, Germany), Christian Terboven (RWTH Aachen University, Germany), and Matthias Mueller (RWTH Aachen University, Germany)</i>	
SLAP: An Adaptive, Learned Admission Policy for Content Delivery Network Caching	457
<i>Ke Liu (Huazhong University of Science and Technology, China), Kan Wu (University of Wisconsin-Madison), Hua Wang (Huazhong University of Science and Technology, China), Ke Zhou (Huazhong University of Science and Technology, China), Ji Zhang (Huazhong University of Science and Technology, China), and Cong Li (Tencent Technology (Shenzhen) Co., Ltd., China)</i>	
Proactive SLA-Aware Application Placement in the Computing Continuum	468
<i>Zahra Najafabadi Samani (Alpen-Adria-Universität Klagenfurt, Austria), Narges Mehran (Alpen-Adria-Universität Klagenfurt, Austria), Dragi Kimovski (Alpen-Adria-Universität Klagenfurt, Austria), and Radu Prodan (Alpen-Adria-Universität Klagenfurt, Austria)</i>	

SESSION 5A: Federated and Graph Learning

PFedSA: Personalized Federated Multi-Task Learning via Similarity Awareness	480
<i>Chuyao Ye (Central South University, China), Hao Zheng (Central South University, China), Zhigang Hu (Central South University, China), and Meiguang Zheng (Central South University, China)</i>	

FedBIAD: Communication-Efficient and Accuracy-Guaranteed Federated Learning with Bayesian Inference-Based Adaptive Dropout	489
<i>Jingjing Xue (Institute of Computing Technology, China; University of Chinese Academy of Sciences, China), Min Liu (Institute of Computing Technology, China; University of Chinese Academy of Sciences, China; Zhongguancun Laboratory, China), Sheng Sun (Institute of Computing Technology, China), Yuwei Wang (Institute of Computing Technology, China), Hui Jiang (Institute of Computing Technology, China; University of Chinese Academy of Sciences, China), and Xuefeng Jiang (Institute of Computing Technology, China; University of Chinese Academy of Sciences, China)</i>	
Fast Sparse GPU Kernels for Accelerated Training of Graph Neural Networks	501
<i>Ruibo Fan (The Hong Kong University of Science and Technology (Guangzhou)), Wei Wang (The Hong Kong University of Science and Technology), and Xiaowen Chu (The Hong Kong University of Science and Technology (Guangzhou), China; The Hong Kong University of Science and Technology)</i>	
Communication Optimization for Distributed Execution of Graph Neural Networks	512
<i>Surreya Emre Kurt (University of Utah, USA), Jinghua Yan (University of Utah, USA), Aravind Sukumaran-Rajam (University of Utah, USA; Meta Platforms, USA), Prashant Pandey (University of Utah, USA), and P. Sadayappan (University of Utah, USA)</i>	

SESSION 5B: Systems and ML

A Machine Learning Approach Towards Runtime Optimisation of Matrix Multiplication	524
<i>Yufan Xia (Australian National University, Australia), Marco De La Pierre (Pawsey Supercomputing Centre, Australia), Amanda S. Barnard (Australian National University, Australia), and Giuseppe Maria Junior Barca (Australian National University, Australia)</i>	
Power Constrained Autotuning using Graph Neural Networks	535
<i>Akash Dutta (Iowa State University, USA), Jee Choi (University of Oregon, USA), and Ali Jannesari (Iowa State University, USA)</i>	
SCONNA: A Stochastic Computing Based Optical Accelerator for Ultra-Fast, Energy-Efficient Inference of Integer-Quantized CNNs	546
<i>Sairam Sri Vatsavai (University of Kentucky, USA), Venkata Sai Praneeth Karempudi (University of Kentucky, USA), Ishan Thakkar (University of Kentucky, USA), Ahmad Salehi (University of Kentucky, USA), and Todd Hastings (University of Kentucky, USA)</i>	
HyScale-GNN: A Scalable Hybrid GNN Training System on Single-Node Heterogeneous Architecture	557
<i>Yi-Chien Lin (University of Southern California, Los Angeles) and Viktor Prasanna (University of Southern California, Los Angeles)</i>	

SESSION 6A: Scientific Applications

Optimizing Cloud Computing Resource Usage for Hemodynamic Simulation	568
<i>William Ladd (Duke University, USA), Christopher Jensen (Duke University, USA), Madhurima Vardhan (Duke University, USA), Jeff Ames (Duke University, USA), Jeff R. Hammond (NVIDIA Helsinki Oy, Finland), Erik W. Draeger (Lawrence Livermore National Laboratory, USA), and Amanda Randles (Duke University, USA)</i>	
Predictive Analysis of Code Optimisations on Large-Scale Coupled CFD-Combustion Simulations using the CPX Mini-App	579
<i>A. Powell (University of Warwick, UK) and G.R. Mudalige (University of Warwick, UK)</i>	
Scalable Adaptive Algorithms for Next-Generation Multiphase flow Simulations	590
<i>Kumar Saurabh (Iowa State University, USA), Masado Ishii (University of Utah, USA), Makrand A. Khanwale (Stanford University, USA), Hari Sundar (University of Utah, USA), and Baskar Ganapathysubramanian (Iowa State University, USA)</i>	
Porting a Computational Fluid Dynamics Code with AMR to Large-Scale GPU Platforms	602
<i>Joshua H. Davis (University of Maryland, USA), Justin Shafner (University of Maryland, USA), Daniel Nichols (University of Maryland, USA), Nathan Grube (University of Maryland, USA), Pino Martin (University of Maryland, USA), and Abhinav Bhatele (University of Maryland, USA)</i>	
Neural Network Compiler for Parallel High-Throughput Simulation of Digital Circuits	613
<i>Ignacio Gavier (University of Massachusetts, USA), Joshua Russell (University of Massachusetts, USA), Devdhar Patel (University of Massachusetts, USA), Edward Rietman (University of Massachusetts, USA), and Hava Siegelmann (University of Massachusetts, USA)</i>	

SESSION 6B: Performance Engineering

Opportunities and Limitations of Hardware Timestamps in Concurrent Data Structures	624
<i>Olivia Grimes (Lehigh University, USA), Jacob Nelson-Slivon (Lehigh University, USA), Ahmed Hassan (Lehigh University, USA), and Roberto Palmieri (Lehigh University, USA)</i>	
Harnessing the Crowd for Autotuning High-Performance Computing Applications	635
<i>Younghyun Cho (University of California, Berkeley), James W. Demmel (University of California, Berkeley), Jacob King (Tech-X), Xiaoye S. Li (Lawrence Berkeley National Laboratory), Yang Liu (Lawrence Berkeley National Laboratory), and Hengrui Luo (Lawrence Berkeley National Laboratory)</i>	
Designing and Optimizing GPU-Aware Nonblocking MPI Neighborhood Collective Communication for PETSc	646
<i>Kawthar Shafie Khorassani (The Ohio State University, USA), Chen-Chun Chen (The Ohio State University, USA), Hari Subramoni (The Ohio State University, USA), and Dhabaleswar K. Panda (The Ohio State University, USA)</i>	

SW-LCM: A Scalable and Weakly-Supervised Land Cover Mapping Method on a New Sunway Supercomputer	657
<i>Yi Zhao (Tsinghua University, China), Juepeng Zheng (Tsinghua University, China), Haohuan Fu (Tsinghua University, China), Wenzhao Wu (National Supercomputing Center in Wuxi, China), Jie Gao (National Research Center of Parallel Computer Engineering and Technology, China), Mengxuan Chen (Tsinghua University, China), Jinxiao Zhang (Tsinghua University, China), Lixian Zhang (Tsinghua University, China), Runmin Dong (Tsinghua University, China), Zhenrong Du (Tsinghua University, China), Sha Liu (National Research Center of Parallel Computer Engineering and Technology, China), Xin Liu (National Research Center of Parallel Computer Engineering and Technology, China), Shaoqing Zhang (Qingdao Pilot National Laboratory for Marine Science and Technology, China), and Le Yu (Tsinghua University, China)</i>	
Feature-Based SpMV Performance Analysis on Contemporary Devices	668
<i>Panagiotis Mpakos (National Technical University of Athens, Greece), Dimitrios Galanopoulos (National Technical University of Athens, Greece), Petros Anastasiadis (National Technical University of Athens, Greece), Nikela Papadopoulou (Chalmers University of Technology, Sweden), Nectarios Koziris (National Technical University of Athens, Greece), and Georgios Goumas (National Technical University of Athens, Greece)</i>	
An Experimental Study of Two-Level Schwarz Domain-Decomposition Preconditioners on GPUs .	680
<i>Ichitaro Yamazaki (Sandia National Laboratories, USA), Alexander Heinlein (Delft University of Technology, Netherlands), and Sivasankaran Rajamanickam (Sandia National Laboratories, USA)</i>	

Keynote Talk III

Keynote: Future Workloads Drive the Need for High Performant and Adaptive Computing Hardware	690
<i>Ivo Bolsens (Adaptive and Embedded Computing Group, AMD)</i>	

SESSION 7A: Combinatorial Algorithms

Engineering Massively Parallel MST Algorithms	691
<i>Peter Sanders (Karlsruhe Institute of Technology, Germany) and Matthias Schimek (Karlsruhe Institute of Technology, Germany)</i>	
Engineering a Distributed-Memory Triangle Counting Algorithm	702
<i>Peter Sanders (Karlsruhe Institute of Technology, Germany) and Tim Niklas Uhl (Karlsruhe Institute of Technology, Germany)</i>	
PRF: A Fast Parallel Relaxed Flooding Algorithm for Voronoi Diagram Generation on GPU	713
<i>Jue Wang (Osaka University, Japan), Fumihiko Ino (Osaka University, Japan), and Jing Ke (Shanghai Jiao Tong University, China)</i>	

Satellite Collision Detection using Spatial Data Structures	724
<i>Christian Hellwig (Technical University of Darmstadt, Germany), Fabian Czappa (Technical University of Darmstadt, Germany), Martin Michel (Technical University of Darmstadt, Germany), Reinhold Bertrand (Technical University of Darmstadt, Germany; European Space Agency ESA/ESOC), and Felix Wolf (Technical University of Darmstadt, Germany)</i>	
AnyQ: An Evaluation Framework for Massively-Parallel Queue Algorithms	736
<i>Michael Kenzel (Deutsches Forschungszentrum fur Kunstliche Intelligenz (DFKI), Saarland Informatics Campus; Graz University of Technology), Stefan Lemme (Deutsches Forschungszentrum fur Kunstliche Intelligenz (DFKI), Saarland Informatics Campus), Richard Membarth (Deutsches Forschungszentrum fur Kunstliche Intelligenz (DFKI), Saarland Informatics Campus; Technische Hochschule Ingolstadt (THI), Research Institute Almotion Bavaria), Matthias Kurtenacker (Deutsches Forschungszentrum fur Kunstliche Intelligenz (DFKI), Saarland Informatics Campus; Graz University of Technology), Hugo Devillers (Saarland University), Markus Steinberger (Graz University of Technology), and Philipp Slusallek (Deutsches Forschungszentrum fur Kunstliche Intelligenz (DFKI), Saarland Informatics Campus; Graz University of Technology; Saarland University)</i>	

SESSION 7B: Emerging Technology

qTask: Task-Parallel Quantum Circuit Simulation with Incrementality	746
<i>Tsung-Wei Huang (University of Utah)</i>	
GPU-Accelerated Error-Bounded Compression Framework for Quantum Circuit Simulations	757
<i>Milan Shah (North Carolina State University, USA), Xiaodong Yu (Argonne National Laboratory, USA), Sheng Di (Argonne National Laboratory, USA), Danylo Lykov (University of Chicago, USA), Yuri Alexeev (Argonne National Laboratory, USA), Michela Becchi (North Carolina State University, USA), and Franck Cappello (Argonne National Laboratory, USA)</i>	
An Adaptive Hybrid Quantum Algorithm for the Metric Traveling Salesman Problem	768
<i>Fei Li (George Mason University, USA) and Arul Mazumder (Massachusetts Academy of Math & Science at WPI, USA)</i>	
Stochastic Neuromorphic Circuits for Solving MAXCUT	779
<i>Bradley H. Theilman (Neural Exploration and Research Lab), Yipu Wang (Discrete Math and Optimization), Ojas Parekh (Discrete Math and Optimization), William Severa (Neural Exploration and Research Lab), J. Darby Smith (Neural Exploration and Research Lab), and James B. Aimone (Neural Exploration and Research Lab)</i>	
TurboHE: Accelerating Fully Homomorphic Encryption using FPGA Clusters	788
<i>Haohao Liao (Huawei Technologies, Canada), Mahmoud A. Elmohr (Huawei Technologies, Canada), Xuan Dong (Huawei Technologies, Canada), Yanjun Qian (Huawei Technologies, Canada), Wenzhe Yang (Huawei Technologies, Canada), Zhiwei Shang (Huawei Technologies, Canada), and Yin Tan (Huawei Technologies, Canada)</i>	

Towards Faster Fully Homomorphic Encryption Implementation with Integer and Floating-Point Computing Power of GPUs	798
<i>Guang Fan (Institute of Information Engineering, China; University of Chinese Academy of Sciences, China; Chinese Academy of Sciences, China), Fangyu Zheng (Institute of Information Engineering, China; University of Chinese Academy of Sciences, China; Chinese Academy of Sciences, China), Lipeng Wan (Institute of Information Engineering, China; University of Chinese Academy of Sciences, China; Chinese Academy of Sciences, China), Lily Gao (Nanjing University of Information Science and Technology, China), Yuan Zhao (Ant Group, China), Jiankuo Dong (Nanjing University of Posts and Telecommunications, China), Yixuan Song (Ant Group, China), Yuewu Wang (University of Chinese Academy of Sciences, China), and Jingqiang Lin (University of Science and Technology of China, China)</i>	

SESSION 8A: Data-Intensive Algorithms

FedTrip: A Resource-Efficient Federated Learning Method with Triplet Regularization	809
<i>Xujing Li (Institute of Computing Technology, China; University of Chinese Academy of Sciences, China), Min Liu (Institute of Computing Technology, China; University of Chinese Academy of Sciences, China; Zhongguancun Laboratory, China), Sheng Sun (Institute of Computing Technology, China), Yuwei Wang (Institute of Computing Technology, China), Hui Jiang (Institute of Computing Technology, China; University of Chinese Academy of Sciences, China), and Xuefeng Jiang (Institute of Computing Technology, China; University of Chinese Academy of Sciences, China)</i>	
A Guaranteed Approximation Algorithm for Scheduling Fork-Joins with Communication Delay ...	820
<i>Pierre-François Dutot (Univ. Grenoble Alpes, CNRS, Inria, France), Yeu-Shin Fu (University of Auckland, New Zealand), Nikhil Prasad (University of Auckland, New Zealand), and Oliver Sinnen (University of Auckland, New Zealand)</i>	
SelB-k-NN: A Mini-Batch K-Nearest Neighbors Algorithm on AI Processors	831
<i>Yifeng Tang (The University of Hong Kong, China) and Cho-li Wang (The University of Hong Kong, China)</i>	
Exact Fault-Tolerant Consensus with Voting Validity	842
<i>Zhangchen Xu (University of Washington, United States), Yuetai Li (University of Glasgow, United Kingdom), Chenglin Feng (University of Glasgow, United Kingdom), and Lei Zhang (University of Glasgow, United Kingdom)</i>	
k-Center Clustering with Outliers in the MPC and Streaming Model	853
<i>Mark de Berg (Eindhoven University of Technology, the Netherlands), Leyla Biabani (Eindhoven University of Technology, the Netherlands), and Morteza Monemizadeh (Eindhoven University of Technology, the Netherlands)</i>	

SESSION 8B: Serverless/Cloud Computing Systems

FIRST: Exploiting the Multi-Dimensional Attributes of Functions for Power-Aware Serverless Computing	864
<i>Lu Zhang (Shanghai Jiao Tong University, China), Chao Li (Shanghai Jiao Tong University, China), Xinkai Wang (Shanghai Jiao Tong University, China), Weiqi Feng (Shanghai Jiao Tong University, China), Zheng Yu (Shanghai Jiao Tong University, China), Quan Chen (Shanghai Jiao Tong University, China), Jingwen Leng (Shanghai Jiao Tong University, China), Minyi Guo (Shanghai Jiao Tong University, China), Pu Yang (Tencent, China), and Shang Yue (Tencent, China)</i>	
Duo: Improving Data Sharing of Stateful Serverless Applications by Efficiently Caching Multi-Read Data	875
<i>Zhuo Huang (Huazhong University of Science and Technology, China), Hao Fan (Huazhong University of Science and Technology, China), Chaoyi Cheng (Huazhong University of Science and Technology, China), Song Wu (Huazhong University of Science and Technology, China), and Hai Jin (Huazhong University of Science and Technology, China)</i>	
QoS-Aware and Cost-Efficient Dynamic Resource Allocation for Serverless ML Workflows	886
<i>Hao Wu (Huazhong University of Science and Technology, China), Junxiao Deng (Huazhong University of Science and Technology, China), Hao Fan (Huazhong University of Science and Technology, China), Shadi Ibrahim (Inria, Univ. Rennes, CNRS, IRISA, France), Song Wu (Huazhong University of Science and Technology, China), and Hai Jin (Huazhong University of Science and Technology, China)</i>	
rFaaS: Enabling High Performance Serverless with RDMA and Leases	897
<i>Marcin Copik (ETH Zurich, Switzerland), Konstantin Taranov (Microsoft), Alexandru Calotoiu (ETH Zurich, Switzerland), and Torsten Hoefler (ETH Zurich, Switzerland)</i>	
Alioth: A Machine Learning Based Interference-Aware Performance Monitor for Multi-Tenancy Applications in Public Cloud	908
<i>Tianyao Shi (Shanghai Jiao Tong University, China), Yingxuan Yang (Shanghai Jiao Tong University, China), Yunlong Cheng (Shanghai Jiao Tong University, China), Xiaofeng Gao (Shanghai Jiao Tong University, China), Zhen Fang (Huawei Technologies, China), and Yongqiang Yang (Huawei Technologies, China)</i>	
GPU-Enabled Function-as-a-Service for Machine Learning Inference	918
<i>Ming Zhao (Arizona State University), Kritshekhar Jha (Arizona State University), and Sungho Hong (Arizona State University)</i>	

SESSION 9A: Optimizations for New Applications

Lyra: Fast and Scalable Resilience to Reordering Attacks in Blockchains	929
<i>Pouriya Zarbafian (University of Sydney, Australia) and Vincent Gramoli (University of Sydney, Australia)</i>	
Smart Redbelly Blockchain: Reducing Congestion for Web3	940
<i>Deepal Tennakoon (University of Sydney, Australia), Yiding Hua (ETH Zurich, Switzerland), and Vincent Gramoli (University of Sydney, Australia)</i>	

SBGT: Scaling Bayesian-Based Group Testing for Disease Surveillance	951
<i>Weicong Chen (Case Western Reserve University, USA), Hao Qi (University of California Merced, USA), Xiaoyi Lu (University of California Merced, USA), and Curtis Tatsuoka (University of Pittsburgh, USA)</i>	
RT-DBSCAN: Accelerating DBSCAN using Ray Tracing Hardware	963
<i>Vani Nagarajan (Purdue University, USA) and Milind Kulkarni (Purdue University, USA)</i>	
Distributing Simplex-Shaped Nested for-Loops to Identify Carcinogenic Gene Combinations	974
<i>Sajal Dash (Oak Ridge National Laboratory, USA), Mohammad Alaul Haque Monil (Oak Ridge National Laboratory, USA), Junqi Yin (Oak Ridge National Laboratory, USA), Ramu Anandkrishnan (The Edward Via College of Osteopathic Medicine, USA), and Feiyi Wang (Oak Ridge National Laboratory, USA)</i>	

SESSION 9B: Big Data Management

LowFive: In Situ Data Transport for High-Performance Workflows	985
<i>Tom Peterka (Argonne National Laboratory, USA), Dmitriy Morozov (Lawrence Berkeley National Laboratory, USA), Arnur Nigmatov (Lawrence Berkeley National Laboratory, USA), Orcun Yildiz (Argonne National Laboratory, USA), Bogdan Nicolae (Argonne National Laboratory, USA), and Philip E. Davis (University of Utah, USA)</i>	
MCR-DL: Mix-and-Match Communication Runtime for Deep Learning	996
<i>Quentin Anthony (The Ohio State University, USA), Ammar Ahmad Awan (Microsoft Corporation, USA), Jeff Rasley (Microsoft Corporation, USA), Yuxiong He (Microsoft Corporation, USA), Aamir Shafi (The Ohio State University, USA), Mustafa Abduljabbar (The Ohio State University, USA), Hari Subramoni (The Ohio State University, USA), and Dhableswar Panda (The Ohio State University, USA)</i>	
Lossy Scientific Data Compression With SPERR	1007
<i>Shaomeng Li (Nat'l Center for Atmospheric Research), Peter Lindstrom (Lawrence Livermore Nat'l Laboratory), and John Clyne (Nat'l Center for Atmospheric Research)</i>	
Fast and Automatic Floating Point Error Analysis With CHEF-FP	1018
<i>Garima Singh (European Council for Nuclear Research, Switzerland; Princeton University, USA), Baidyanath Kundu (European Council for Nuclear Research, Switzerland; Princeton University, USA), Harshitha Menon (Center for Applied Scientific Computing, Lawrence Livermore National Laboratory, USA), Alexander Penev (University of Plovdiv, Bulgaria), David J. Lange (Princeton University, USA), and Vassil Vassilev (Princeton University, USA; European Council for Nuclear Research, Switzerland)</i>	

DAOS as HPC Storage: a View From Numerical Weather Prediction	1029
<i>Nicolau Manubens (European Centre for Medium-Range Weather Forecasts, Germany), Tiago Quintino (European Centre for Medium-Range Weather Forecasts, United Kingdom), Simon D. Smart (European Centre for Medium-Range Weather Forecasts, United Kingdom), Emanuele Danovaro (European Centre for Medium-Range Weather Forecasts, United Kingdom), and Adrian Jackson (The University of Edinburgh, United Kingdom)</i>	
ZFP-X: Efficient Embedded Coding for Accelerating Lossy Floating Point Compression	1041
<i>Bing Lu (Hunan University, China), Yida Li (Hunan University, China), Junqi Wang (Beijing Institute for General Artificial Intelligence (BIGAI), China), Huizhang Luo (Hunan University, China), and Kenli Li (Hunan University, China)</i>	

Author Index