

2024 IEEE International Parallel and Distributed Processing Symposium (IPDPS 2024)

**San Francisco, California, USA
27-31 May 2024**

Pages 1-543



IEEE Catalog Number: CFP24023-POD
ISBN: 979-8-3503-8712-4

**Copyright © 2024 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:	CFP24023-POD
ISBN (Print-On-Demand):	979-8-3503-8712-4
ISBN (Online):	979-8-3503-8711-7
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

2024 IEEE International Parallel and Distributed Processing Symposium (IPDPS) **IPDPS 2024**

Table of Contents

Message from the IPDPS 2024 General Co-chairs	xx
Message from the IPDPS 2024 Program Chairs	xxii
IPDPS 2024 Technical Program	xxiv
IPDPS 2024 Organization	xxvi

Keynote 1

AuroraGPT: Exploring AI Assistant for Science	1
<i>Franck Cappello (Argonne National Laboratory)</i>	

Session 1A: Numerical Linear Algebra

PckGNN: Optimizing Aggregation Operators with Packing Strategies in Graph Neural Networks	2
<i>Zhengding Hu (University of Science and Technology of China, China), Jingwei Sun (University of Science and Technology of China, China), Zhongyang Li (University of Science and Technology of China, China), and Guangzhong Sun (University of Science and Technology of China, China)</i>	
VNEC: A Vectorized Non-Empty Column Format for SpMV on CPUs	14
<i>Luhan Wang (Chinese Academy of Sciences, China; University of Chinese Academy of Sciences, China), Haipeng Jia (Chinese Academy of Sciences, China), Lei Xu (Chinese Academy of Sciences, China; University of Chinese Academy of Sciences, China), Cunyang Wei (Chinese Academy of Sciences, China; University of Chinese Academy of Sciences, China), Kun Li (Microsoft Research), Xianmeng Jiang (Chinese Academy of Sciences, China; University of Chinese Academy of Sciences, China), and Yunquan Zhang (Chinese Academy of Sciences, China)</i>	
Two-Stage Block Orthogonalization to Improve Performance of S-Step GMRES	26
<i>Ichitaro Yamazaki (Sandia National Laboratories, USA), Andrew J. Higgins (Temple University, USA), Erik G. Boman (Sandia National Laboratories, USA), and Daniel B. Szyld (Temple University, USA)</i>	

Alternative Basis Matrix Multiplication is Fast and Stable	38
<i>Oded Schwartz (The Hebrew University of Jerusalem, Israel), Sivan Toledo (Tel Aviv University, Israel), Noa Vaknin (The Hebrew University of Jerusalem, Israel), and Gal Wiernik (Tel Aviv University, Israel)</i>	
Fast Multiplication of Random Dense Matrices with Sparse Matrices	52
<i>Tianyu Liang (UC Berkeley; Lawrence Berkeley National Lab), Riley Murray (UC Berkeley; Lawrence Berkeley National Lab; International Computer Science Institute; Sandia National Laboratories), Aydin Buluç (UC Berkeley; Lawrence Berkeley National Lab), and James Demmel (UC Berkeley; Lawrence Berkeley National Lab)</i>	
A Cholesky QR Type Algorithm for Computing Tall-Skinny QR Factorization with Column Pivoting	63
<i>Takeshi Fukaya (Hokkaido University, Japan), Yuji Nakatsukasa (University of Oxford, UK), and Yusaku Yamamoto (The University of Electro-Communications, Japan)</i>	

Session 1B: Containers and Serverless Computing

CKSM: An Efficient Memory Deduplication Method for Container-Based Cloud Computing Systems... 76	
<i>Yunfei Gu (Shanghai Jiao Tong University, China), Yihui Lu (Shanghai Jiao Tong University, China), Chentao Wu (Shanghai Jiao Tong University, China), Jie Li (Shanghai Jiao Tong University, China), and Minyi Guo (Shanghai Jiao Tong University, China)</i>	
Tackling Cold Start in Serverless Computing with Multi-level Container Reuse	89
<i>Amelie Chi Zhou (Hong Kong Baptist University), Rongzheng Huang (Shenzhen University), Zhoubin Ke (Shenzhen University), Yusen Li (Nankai University), Yi Wang (Shenzhen University), and Rui Mao (Shenzhen University)</i>	
PALDIA: Enabling SLO-Compliant and Cost-Effective Serverless Computing on Heterogeneous Hardware	100
<i>Vivek M. Bhasi (The Pennsylvania State University), Aakash Sharma (The Pennsylvania State University), Shruti Mohanty (The Pennsylvania State University), Mahmut Taylan Kandemir (The Pennsylvania State University), and Chita R. Das (The Pennsylvania State University)</i>	
Application-Attuned Memory Management for Containerized HPC Workflows	114
<i>Moiz Arif (Rochester Institute of Technology), Avinash Maurya (Rochester Institute of Technology), M. Mustafa Rafique (Rochester Institute of Technology), Dimitrios S. Nikolopoulos (Virginia Tech), and Ali R. Butt (Virginia Tech)</i>	
FEDGE: An Interference-Aware QoS Prediction Framework for Black-Box Scenario in IaaS Clouds with Domain Generalization	128
<i>Yunlong Cheng (Shanghai Jiao Tong University, China), Xiuqi Huang (Shanghai Jiao Tong University, China), Zifeng Liu (Shanghai Jiao Tong University, China), Jiadong Chen (Shanghai Jiao Tong University, China), Xiaofeng Gao (Shanghai Jiao Tong University, China), Zhen Fang (Huawei Cloud Computing Technology Co., Ltd, China), and Yongqiang Yang (Huawei Cloud Computing Technology Co., Ltd, China)</i>	

Software Resource Disaggregation for HPC with Serverless Computing	139
<i>Marcin Copik (ETH Zürich, Switzerland), Marcin Chrapek (ETH Zürich, Switzerland), Larissa Schmid (Karlsruhe Institute of Technology, Germany), Alexandru Calotoiu (ETH Zürich, Switzerland), and Torsten Hoefler (ETH Zürich, Switzerland)</i>	

Session 2A: Algorithms on Trees

AMST: Accelerating Large-Scale Graph Minimum Spanning Tree Computation on FPGA	157
<i>Haishuang Fan (State Key Laboratory of Processors, Institute of Computing Technology, Chinese Academy of Sciences; University of Chinese Academy of Sciences), Rui Meng (State Key Laboratory of Processors, Institute of Computing Technology, Chinese Academy of Sciences; University of Chinese Academy of Sciences), Qichu Sun (State Key Laboratory of Processors, Institute of Computing Technology, Chinese Academy of Sciences; University of Chinese Academy of Sciences), Jingya Wu (State Key Laboratory of Processors, Institute of Computing Technology, Chinese Academy of Sciences), Wenyan Lu (State Key Laboratory of Processors, Institute of Computing Technology, Chinese Academy of Sciences; YUSUR Technology Co., Ltd.), Xiaowei Li (State Key Laboratory of Processors, Institute of Computing Technology, Chinese Academy of Sciences; University of Chinese Academy of Sciences), and Guihai Yan (State Key Laboratory of Processors, Institute of Computing Technology, Chinese Academy of Sciences; University of Chinese Academy of Sciences; YUSUR Technology Co., Ltd)</i>	
Wait-Free Trees with Asymptotically-Efficient Range Queries	169
<i>Ilya Kokorin (ITMO University, Russia), Victor Yudov (ITMO University, Russia), Vitaly Aksenov (University of London, UK), and Dan Alistarh (IST Austria, Austria)</i>	
Low-Depth Spatial Tree Algorithms	180
<i>Yves Baumann (ETH Zurich, Switzerland), Tal Ben-Nun (ETH Zurich, Switzerland), Maciej Besta (ETH Zurich, Switzerland), Lukas Gianinazzi (ETH Zurich, Switzerland), Torsten Hoefler (ETH Zurich, Switzerland), and Piotr Luczynski (ETH Zurich, Switzerland)</i>	

Session 2B: Federated and Distributed Learning

QSync: Quantization-Minimized Synchronous Distributed Training Across Hybrid Devices	193
<i>Juntao Zhao (The University of Hong Kong, Hong Kong), Borui Wan (The University of Hong Kong, Hong Kong), Yanghua Peng (ByteDance Inc., USA), Haibin Lin (ByteDance Inc., USA), Yibo Zhu (ByteDance Inc., USA), and Chuan Wu (The University of Hong Kong, Hong Kong)</i>	
Enhancing the Generalization of Personalized Federated Learning with Multi-head Model and Ensemble Voting	205
<i>Van An Le (National Institute of Advanced Industrial Science and Technology, Japan), Nam Duong Tran (Hanoi University of Science and Technology, Vietnam), Phuong Nam Nguyen (Hanoi University of Science and Technology, Vietnam), Thanh Hung Nguyen (Hanoi University of Science and Technology, Vietnam), Phi Le Nguyen (Hanoi University of Science and Technology, Vietnam), Thao Nguyen Truong (National Institute of Advanced Industrial Science and Technology, Japan), and Yusheng Ji (National Institute of Informatics, Japan)</i>	

UniFaaS: Programming Across Distributed Cyberinfrastructure with Federated Function Serving	217
<i>Yifei Li (Southern University of Science and Technology, China), Ryan Chard (Argonne National Laboratory, USA), Yadu Babuji (University of Chicago, USA), Kyle Chard (University of Chicago, USA; Argonne National Laboratory, USA), Ian Foster (Argonne National Laboratory, USA; University of Chicago, USA), and Zhuozhao Li (Southern University of Science and Technology, China)</i>	

Session 3A: Applications I

Scalable and Differentiable Simulator for Quantum Computational Chemistry	230
<i>Zhiqian Xu (Institute of Computing Technology, Chinese Academy of Sciences, China), Honghui Shang (University of Science and Technology of China, China), Yi Fan (University of Science and Technology of China, China), Xiongzhi Zeng (University of Science and Technology of China, China), Yunquan Zhang (Institute of Computing Technology, Chinese Academy of Sciences, China), and Chu Guo (Hunan normal University, China)</i>	
Picasso: Memory-Efficient Graph Coloring Using Palettes With Applications in Quantum Computing	241
<i>S M Ferdous (Pacific Northwest National Laboratory, USA), Reece Neff (North Carolina State University, USA; Pacific Northwest National Laboratory, USA), Bo Peng (Pacific Northwest National Laboratory, USA), Salman Shuvo (Pacific Northwest National Laboratory, USA), Marco Minutoli (Pacific Northwest National Laboratory, USA), Sayak Mukherjee (Pacific Northwest National Laboratory, USA), Karol Kowalski (Pacific Northwest National Laboratory, USA), Michela Becchi (North Carolina State University, USA; Pacific Northwest National Laboratory, USA), and Mahantesh Halappanavar (Pacific Northwest National Laboratory, USA)</i>	
Optimizing and Scaling the 3D Reconstruction of Single-Particle Imaging	253
<i>Niteya Shah (Virginia Tech, USA), Christine Sweeney (Los Alamos National Lab, USA), Vinay Ramakrishnaiah (Los Alamos National Lab, USA), Jeffrey Donatelli (Lawrence Berkeley National Lab, USA), and Wu-chun Feng (Virginia Tech, USA)</i>	
Parallel Approximations for High-Dimensional Multivariate Normal Probability Computation in Confidence Region Detection Applications	265
<i>Xiran Zhang (King Abdullah University of Science and Technology, Saudi Arabia), Sameh Abdulah (King Abdullah University of Science and Technology, Saudi Arabia), Jian Cao (The University of Houston, USA), Hatem Ltaief (King Abdullah University of Science and Technology, Saudi Arabia), Ying Sun (King Abdullah University of Science and Technology, Saudi Arabia), Marc G. Genton (King Abdullah University of Science and Technology, Saudi Arabia), and David E. Keyes (King Abdullah University of Science and Technology, Saudi Arabia)</i>	

Enabling High-Performance Physical Based Rendering on New Sunway Supercomputer	277
<i>Zeyu Song (Tsinghua University, China), Lin Gan (Tsinghua University, China; Beijing National Research Center for Information Science and Technology, China; National Supercomputing Center in Wuxi, China), Shengye Xiang (Tsinghua University, China), Yinuo Wang (Tsinghua University, China), Xiaohui Duan (Shandong University, China), and Guangwen Yang (Tsinghua University, China; Beijing National Research Center for Information Science and Technology, China; National Supercomputing Center in Wuxi, China; Shandong University, China; Zhejiang Lab, China)</i>	

Session 3B: Scheduling I

CoCG: Fine-Grained Cloud Game Co-Location on Heterogeneous Platform	289
<i>Taolei Wang (Shanghai Jiao Tong University, China), Chao Li (Shanghai Jiao Tong University, China), Jing Wang (Shanghai Jiao Tong University, China), Cheng Xu (Shanghai Jiao Tong University, China), Xiaofeng Hou (Shanghai Jiao Tong University, China), and Minyi Guo (Shanghai Jiao Tong University, China)</i>	
Adaptive Task-Oriented Resource Allocation for Large Dynamic Workflows on Opportunistic Resources	300
<i>Thanh Son Phung (University of Notre Dame) and Douglas Thain (University of Notre Dame)</i>	
nOS-V: Co-Executing HPC Applications Using System-Wide Task Scheduling	312
<i>David Álvarez (Barcelona Supercomputing Center, Spain), Kevin Sala (Barcelona Supercomputing Center, Spain), and Vicenç Beltran (Barcelona Supercomputing Center, Spain)</i>	
SWEEP: Adaptive Task Scheduling for Exploring Energy Performance Trade-Offs	325
<i>Jing Chen (Chalmers University of Technology), Madhavan Manivannan (Chalmers University of Technology), Bhavishya Goel (Chalmers University of Technology), and Miquel Pericàs (Chalmers University of Technology)</i>	
Interpretable Analysis of Production GPU Clusters Monitoring Data via Association Rule Mining	337
<i>Baolin Li (Northeastern University), Siddharth Samsi (MIT Lincoln Laboratory), Vijay Gadepally (MIT Lincoln Laboratory), and Devesh Tiwari (Northeastern University)</i>	

Best Paper Nominees

CloverLeaf on Intel Multi-core CPUs: A Case Study in Write-Allocate Evasion	350
<i>Jan Laukemann (Friedrich-Alexander-Universität Erlangen-Nürnberg, Germany), Thomas Gruber (Friedrich-Alexander-Universität Erlangen-Nürnberg, Germany), Georg Hager (Friedrich-Alexander-Universität Erlangen-Nürnberg, Germany), Dossay Oryspayev (Brookhaven National Laboratory, USA), and Gerhard Wellein (Friedrich-Alexander-Universität Erlangen-Nürnberg, Germany)</i>	

ARGO: An Auto-Tuning Runtime System for Scalable GNN Training on Multi-Core Processor	361
<i>Yi-Chien Lin (University of Southern California), Yuyang Chen (Tsinghua University), Sameh Gobriel (Intel Labs), Nilesh Jain (Intel Labs), Gopi Krishna Jha (Intel Labs), and Viktor Prasanna (University of Southern California)</i>	
Accelerating Lossy and Lossless Compression on Emerging BlueField DPU Architectures	373
<i>Li Yuke (University of California, Merced, USA), Kashyap Arjun (University of California, Merced, USA), Chen Weicong (University of California, Merced, USA), Guo Yanfei (Argonne National Laboratory, USA), and Lu Xiaoyi (University of California, Merced, USA)</i>	
Performance-Portable Multiphase Flow Solutions with Discontinuous Galerkin Methods	386
<i>Tobias S. Flynn (University of Warwick, UK), Robert Manson-Sawoko (IBM Research Europe), and Gihan R. Mudalige (University of Warwick, UK)</i>	

Session 4A: Applications II

Optimized GPU Implementation of Grid Refinement in Lattice Boltzmann Method	398
<i>Ahmed H. Mahmoud (Autodesk Research, Canada; University of California, USA), Hesam Salehipour (Autodesk Research, Canada), and Massimiliano Meneghin (Autodesk Research, Canada)</i>	
Alya Towards Exascale: Optimal OpenACC Performance of the Navier-Stokes Finite Element Assembly on GPUs	408
<i>Herbert Owen (Barcelona Supercomputing Center, Spain), Dominik Ernst (Friedrich-Alexander-Universität Erlangen-Nürnberg, Germany), Thomas Gruber (Friedrich-Alexander-Universität Erlangen-Nürnberg, Germany), Oriol Lemkuhl (Barcelona Supercomputing Center, Spain), Guillaume Houzeaux (Barcelona Supercomputing Center, Spain), Lucas Gasparino (Barcelona Supercomputing Center, Spain), and Gerhard Wellein (Friedrich-Alexander-Universität Erlangen-Nürnberg, Germany)</i>	
CliZ: Optimizing Lossy Compression for Climate Datasets with Adaptive Fine-Tuned Data Prediction	417
<i>Zizhe Jian (University of California, USA), Sheng Di (Argonne National Laboratory, USA), Jinyang Liu (University of California, USA), Kai Zhao (Florida State University, USA), Xin Liang (University of Kentucky, USA), Haiying Xu (National Center for Atmospheric Research, USA), Robert Underwood (Argonne National Laboratory, USA), Shixun Wu (University of California, USA), Jiajun Huang (University of California, USA), Zizhong Chen (University of California, USA), and Franck Cappello (Argonne National Laboratory, USA)</i>	
Automating GPU Scalability for Complex Scientific Models: Phonon Boltzmann Transport Equation	430
<i>Eric Heisler (University of Utah, USA), Siddharth Saurav (The Ohio State University, USA), Aadesh Deshmukh (University of Utah, USA), Sandip Mazumder (The Ohio State University, USA), and Hari Sundar (University of Utah, USA)</i>	

An O(N) Distributed-Memory Parallel Direct Solver for Planar Integral Equations	440
<i>Tianyu Liang (University of California, USA), Chao Chen (North Carolina State University, USA), Per-Gunnar Martinsson (The University of Texas at Austin, USA), and George Biros (The University of Texas at Austin, USA)</i>	
Exploiting Long Vectors with a CFD Code: A Co-design Show Case	453
<i>Marc Blancafort (Barcelona Supercomputing Center, Spain), Roger Ferrer (Barcelona Supercomputing Center, Spain), Guillaume Houzeaux (Barcelona Supercomputing Center, Spain), Marta Garcia-Gasulla (Barcelona Supercomputing Center, Spain), and Filippo Mantovani (Barcelona Supercomputing Center, Spain)</i>	

Session 4B: I/O and Storage Systems

Capturing Periodic I/O Using Frequency Techniques	465
<i>Ahmad Tarraf (Technical University of Darmstadt, Germany), Alexis Bandet (Univ. Bordeaux, France), Francieli Boito (Univ. Bordeaux, France), Guillaume Pallez (Inria, France), and Felix Wolf (Technical University of Darmstadt, Germany)</i>	
To Store or Not to Store: a Graph Theoretical Approach for Dataset Versioning	479
<i>Anxin Guo (Northwestern University, USA; Columbia University, USA), Jingwei Li (Columbia University, USA; Northwestern University, USA; Databricks, USA), Pattara Sukprasert (Columbia University, USA; Databricks, USA), Samir Khuller (Northwestern University, USA), Amol Deshpande (University of Maryland, USA), and Koyel Mukherjee (Adobe Research, India)</i>	
TunIO: An AI-Powered Framework for Optimizing HPC I/O	494
<i>Neeraj Rajesh (Illinois Institute of Technology), Keith Bateman (Illinois Institute of Technology), Jean Luca Bez (Lawrence Berkeley National Laboratory), Suren Byna (Ohio State University and Lawrence Berkeley National Laboratory), Anthony Kougkas (Illinois Institute of Technology), and Xian-He Sun (Illinois Institute of Technology)</i>	
A2FL: Autonomous and Adaptive File Layout in HPC through Real-Time Access Pattern Analysis	506
<i>Dong Kyu Sung (Seoul National University, South Korea), Yongseok Son (Chung-Ang University, South Korea), Alex Sim (Lawrence Berkeley National Laboratory, USA), Kesheng Wu (Lawrence Berkeley National Laboratory, USA), Suren Byna (Lawrence Berkeley National Laboratory, USA; The Ohio State University, USA), Houjun Tang (Lawrence Berkeley National Laboratory, USA), Hyeonsang Eom (Seoul National University, South Korea), Changjong Kim (Seoul National University of Science and Technology, South Korea), and Sunggon Kim (Seoul National University of Science and Technology, South Korea)</i>	
NVMe-oPF: Designing Efficient Priority Schemes for NVMe-Over-Fabrics with Multi-Tenancy Support	519
<i>Darren Ng (University of California, Merced, USA), Andrew Lin (University of California, Merced, USA), Arjun Kashyap (University of California, Merced, USA), Guanpeng Li (University of Iowa, USA), and Xiaoyi Lu (University of California, Merced, USA)</i>	

Drilling Down I/O Bottlenecks with Cross-Layer I/O Profile Exploration	532
<i>Hammad Ather (University of Oregon; Lawrence Berkeley National Laboratory), Jean Luca Bez (Lawrence Berkeley National Laboratory), Yankun Xia (The Ohio State University), and Suren Byna (The Ohio State University; Lawrence Berkeley National Laboratory)</i>	

Keynote 2

PyTorch 2 and Its Compiler Technologies	544
<i>Peng Wu (Meta)</i>	

Session 5A: Performance

CachedArrays: Optimizing Data Movement for Heterogeneous Memory Systems	545
<i>Mark Hildebrand (University of California, Davis), Jason Lowe-Power (University of California, Davis), and Venkatesh Akella (University of California, Davis)</i>	
Comparative Study of Large Language Model Architectures on Frontier	556
<i>Junqi Yin (Oak Ridge National Laboratory, TN), Avishek Bose (Oak Ridge National Laboratory, TN), Guojing Cong (Oak Ridge National Laboratory, TN), Isaac Lyngaas (Oak Ridge National Laboratory, TN), and Quentin Anthony (Ohio State University, OH)</i>	
Predicting Cross-Architecture Performance of Parallel Programs	570
<i>Daniel Nichols (University of Maryland), Alexander Movsesyan (University of Maryland), Jae-Seung Yeom (Lawrence Livermore National Laboratory), Abhik Sarkar (Lawrence Livermore National Laboratory), Daniel Milroy (Lawrence Livermore National Laboratory), Tapasya Patki (Lawrence Livermore National Laboratory), and Abhinav Bhatele (University of Maryland)</i>	

Session 5B: Resilience

DRUTO: Upper-Bounding Silent Data Corruption Vulnerability in GPU Applications	582
<i>Hasanur Rahman (University of Iowa, USA), Sheng Di (Argonne National Laboratory, USA), Shengjian Guo (Amazon Web Services, USA), Xiaoyi Lu (University of California, USA), Guanpeng Li (University of Iowa, USA), and Franck Cappello (Argonne National Laboratory, USA)</i>	
MPI Errors Detection Using GNN Embedding and Vector Embedding over LLVM IR	595
<i>Jad El Karchi (Inria, France), Hanze Chen (Iowa State University, USA), Ali TehraniJamsaz (Iowa State University, USA), Ali Jannesari (Iowa State University, USA), Mihail Popov (Inria, France), and Emmanuelle Saillard (Inria, France)</i>	

A Parallel Partial Merge Repair Algorithm for Multi-block Failures for Erasure Storage Systems	608
<i>Shuaipeng Zhang (Harbin Institute of Technology, Shenzhen), Shiyi Li (Harbin Institute of Technology, Shenzhen; Guangdong Provincial Key Laboratory of Novel Security Intelligence Technologies; Wuhan National Laboratory for Optoelectronics), Chentao Wu (Shanghai Jiao Tong University), Ruobin Wu (Harbin Institute of Technology, Shenzhen), Saiqin Long (Jinan University), and Wen Xia (Harbin Institute of Technology, Shenzhen; Guangdong Provincial Key Laboratory of Novel Security Intelligence Technologies)</i>	

Session 6A: Accelerators

Harmonica: Hybrid Accelerator to Overcome Imperfections of Mixed-Signal DNN Accelerators	619
<i>Payman Behnam (Georgia Institute of Technology, USA), Uday Kamal (Georgia Institute of Technology, USA), Ali Shafiee (Independent Researcher), Alexey Tumanov (Georgia Institute of Technology, USA), and Saibal Mukhopadhyay (Georgia Institute of Technology, USA)</i>	
IPU-EpiDet: Identifying Gene Interactions on Massively Parallel Graph-Based AI Accelerators	631
<i>Ricardo Nobre (Universidade de Lisboa, Portugal), Aleksandar Ilic (Universidade de Lisboa, Portugal), Sergio Santander-Jiménez (University of Extremadura (UNEX), Spain), and Leonel Sousa (INESC-ID, Universidade de Lisboa, Portugal)</i>	
DEFCON: Deformable Convolutions Leveraging Interval Search and GPU Texture Hardware	644
<i>Malith Jayaweera (Northeastern University, USA), Yanyu Li (Northeastern University, USA), Yanzhi Wang (Northeastern University, USA), Bin Ren (William & Mary, USA), and David Kaeli (Northeastern University, USA)</i>	
Benchmarking and Dissecting the Nvidia Hopper GPU Architecture	656
<i>Weile Luo (The Hong Kong University of Science and Technology (Guangzhou)), Ruibo Fan (The Hong Kong University of Science and Technology (Guangzhou)), Zeyu Li (The Hong Kong University of Science and Technology (Guangzhou)), Dayou Du (The Hong Kong University of Science and Technology (Guangzhou)), Qiang Wang (Harbin Institute of Technology (Shenzhen)), and Xiaowen Chu (The Hong Kong University of Science and Technology (Guangzhou))</i>	
Exploration of Trade-Offs Between General-Purpose and Specialized Processing Elements in HPC-Oriented CGRA	668
<i>Emanuele Del Sozzo (Center for Computational Science (R-CCS), Japan), Xinyuan Wang (University of Toronto, Canada; Center for Computational Science (R-CCS), Japan), Boma Adhi (Center for Computational Science (R-CCS), Japan), Carlos Cortes (Center for Computational Science (R-CCS), Japan), Jason Anderson (University of Toronto, Canada), and Kentaro Sano (Center for Computational Science (R-CCS), Japan)</i>	

Session 6B: Scheduling II

Hadar: Heterogeneity-Aware Optimization-Based Online Scheduling for Deep Learning Cluster ...	681
<i>Abeda Sultana (University of Louisiana at Lafayette, USA), Fei Xu (East China Normal University, China), Xu Yuan (University of Delaware, USA), Li Chen (University of Louisiana at Lafayette, USA), and Nian-Feng Tzeng (University of Louisiana at Lafayette, USA)</i>	
Fast Abort-Freedom for Deterministic Transactions	692
<i>Chen Chen (University of Illinois at Chicago), Xingbo Wu (Microsoft Research), Wenshao Zhong (University of Illinois at Chicago), and Jakob Eriksson (University of Illinois at Chicago)</i>	
SYNPA: SMT Performance Analysis and Allocation of Threads to Cores in ARM Processors	705
<i>Marta Navarro (Universitat Politècnica de València, Spain), Josué Feliu (Universitat Politècnica de València, Spain), Salvador Petit (Universitat Politècnica de València, Spain), María E. Gómez (Universitat Politècnica de València, Spain), and Julio Sahuquillo (Universitat Politècnica de València, Spain)</i>	
Cross-System Analysis of Job Characterization and Scheduling in Large-Scale Computing Clusters	716
<i>Di Zhang (University of North Carolina at Charlotte, USA), Monish Soundar Raj (University of North Carolina at Charlotte, USA), Bing Xie (Microsoft, USA), Sheng Di (Argonne National Laboratory, USA), and Dong Dai (University of North Carolina at Charlotte, USA)</i>	
Automatic Task Parallelization of Dataflow Graphs in ML/DL Models	728
<i>Srinjoy Das (University of Illinois Urbana-Champaign, USA) and Lawrence Rauchwerger (University of Illinois Urbana-Champaign, USA)</i>	

Session 7A: Message Passing and Communication

Adaptive Prefetching for Fine-Grain Communication in PGAS Programs	740
<i>Thomas B. Rolinger (NVIDIA, USA) and Alan Sussman (University of Maryland, USA)</i>	
An Optimized Error-Controlled MPI Collective Framework Integrated with Lossy Compression ...	752
<i>Jiajun Huang (University of California, Riverside), Sheng Di (Argonne National Laboratory), Xiaodong Yu (Stevens Institute of Technology), Yujia Zhai (University of California, Riverside), Zhaorui Zhang (The Hong Kong Polytechnic University), Jinyang Liu (University of California, Riverside), Xiaoyi Lu (University of California, Merced), Ken Raffanetti (Argonne National Laboratory), Hui Zhou (Argonne National Laboratory), Kai Zhao (Florida State University), Zizhong Chen (University of California, Riverside), Franck Cappello (Argonne National Laboratory), Yanfei Guo (Argonne National Laboratory), and Rajeev Thakur (Argonne National Laboratory)</i>	

MUSE: A Runtime Incrementally Reconfigurable Network Adapting to HPC Real-Time Traffic	765
<i>Zijian Li (Fudan University), Zixuan Chen (Fudan University), Yiyang Tang (Fudan University), Xin Ai (Fudan University), Yuanyi Zhu (Fudan University), Zhigao Zhao (Fudan University), Jiang Shao (Fudan University), Guowei Liu (Tsinghua University), Sen Liu (Fudan University), Bin Liu (Tsinghua University), and Yang Xu (Fudan University)</i>	
Fast Policy Convergence for Traffic Engineering with Proactive Distributed Message-Passing	780
<i>Zicheng Wang (Beijing University of Posts and Telecommunications, China), Zirui Zhuang (Beijing University of Posts and Telecommunications, China), Jingyu Wang (Beijing University of Posts and Telecommunications, China), Qi Qi (Beijing University of Posts and Telecommunications, China), Haifeng Sun (Beijing University of Posts and Telecommunications, China), and Jianxin Liao (Beijing University of Posts and Telecommunications, China)</i>	
The Self-Adaptive and Topology-Aware MPI Bcast Leveraging Collective Offload on Tianhe Express Interconnect	791
<i>Chongshan Liang (National University of Defense Technology, China), Yi Dai (National University of Defense Technology, China), Jun Xia (Nanhu Lab, China), Jinbo Xu (National University of Defense Technology, China), Jintao Peng (National University of Defense Technology, China), Weixia Xu (National University of Defense Technology, China), Ming Xie (National University of Defense Technology, China), Jie Liu (National University of Defense Technology, China), Zhiquan Lai (National University of Defense Technology, China), Sheng Ma (National University of Defense Technology, China), and Qi Zhu (National University of Defense Technology, China)</i>	
HINT: Designing Cache-Efficient MPI_Alltoall Using Hybrid Memory Copy Ordering and Non-Temporal Instructions	802
<i>Bharath Ramesh (The Ohio State University, USA), Nick Contini (The Ohio State University, USA), Nawras Alnaasan (The Ohio State University, USA), Kaushik Kandadi Suresh (The Ohio State University, USA), Mustafa Abduljabbar (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>	

Session 7B: Communication Subsystems

Flexible NVMe Request Routing for Virtual Machines	814
<i>Tu Dinh Ngoc (Université de Toulouse, France), Boris Teabe (Université de Toulouse, France), Georges Da Costa (Université de Toulouse, France), and Daniel Hagimont (Université de Toulouse, France)</i>	

HA-CSD: Host and SSD Coordinated Compression for Capacity and Performance	825
<i>Xiang Chen (Huazhong University of Science and Technology, China; DapuStor Corporation, China), Tao Lu (DapuStor Corporation, China), Jiapin Wang (DapuStor Corporation, China), Yu Zhong (Huazhong University of Science and Technology, China), Guangchun Xie (DapuStor Corporation, China), Xueming Cao (DapuStor Corporation, China), Yuanpeng Ma (DapuStor Corporation, China), Bing Si (DapuStor Corporation, China), Feng Ding (DapuStor Corporation, China), Ying Yang (DapuStor Corporation, China), Yunxing Huang (DapuStor Corporation, China), Yafei Yang (DapuStor Corporation, China), You Zhou (Huazhong University of Science and Technology, China), and Fei Wu (Huazhong University of Science and Technology, China)</i>	
Graph Analytics on Jellyfish Topology	839
<i>Md Nahid Newaz (Oakland University), Sayan Ghosh (Pacific Northwest National Laboratory), Joshua Suetterlein (Pacific Northwest National Laboratory), Nathan R. Tallent (Pacific Northwest National Laboratory), Md Atiqul Mollah (Cornelis Networks), and Hua Ming (Oakland University)</i>	
TEEMO: Temperature Aware Energy Efficient Multi-Retention STT-RAM Cache Architecture	852
<i>Sukarn Agarwal (Indian Institute of Technology Mandi, India), Shounak Chakraborty (Norwegian University of Science and Technology, Norway), and Magnus Sjölander (Norwegian University of Science and Technology, Norway)</i>	
LockillerTM: Enhancing Performance Lower Bounds in Best-Effort Hardware Transactional Memory	865
<i>Li Wan (Fudan University, China), Chao Fu (Fudan University, China), Qiang Li (Fudan University, China), and Jun Han (Fudan University, China)</i>	
Attention, Distillation, and Tabularization: Towards Practical Neural Network-Based Prefetching	876
<i>Pengmiao Zhang (University of Southern California, USA), Neelesh Gupta (University of Southern California, USA), Rajgopal Kannan (DEVCOM Army Research Lab, USA), and Viktor K. Prasanna (University of Southern California, USA)</i>	

Keynote 3

Computing Systems in the Foundation Model Era	889
<i>Kunle Olukotun (Stanford University)</i>	

Session 8A: Graph and MoE Learning

Aurora: A Versatile and Flexible Accelerator for Graph Neural Networks	890
<i>Jiaqi Yang (George Washington University, USA), Hao Zheng (University of Central Florida, USA), and Ahmed Louri (George Washington University, USA)</i>	

cuKE: An Efficient Code Generator for Score Function Computation in Knowledge Graph Embedding	903
<i>Lihan Hu (University of Iowa), Jing Li (Nvidia), and Peng Jiang (University of Iowa)</i>	
Exploiting Inter-Layer Expert Affinity for Accelerating Mixture-of-Experts Model Inference.....	915
<i>Jinghan Yao (The Ohio State University, USA), Quentin Anthony (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>	
TASER: Temporal Adaptive Sampling for Fast and Accurate Dynamic Graph Representation Learning	926
<i>Gangda Deng (University of Southern California, USA), Hongkuan Zhou (University of Southern California, USA), Hanqing Zeng (Meta AI, USA), Yinglong Xia (Meta AI, USA), Christopher Leung (Meta AI, USA), Jianbo Li (Meta AI, USA), Rajgopal Kannan (US Army Research Lab, USA), and Viktor Prasanna (University of Southern California, USA)</i>	

Session 8B: Performance Optimization

OpenFFT-SME: An Efficient Outer Product Pattern FFT Library on ARM SME CPUs	938
<i>Ruge Zhang (Chinese Academy of Sciences, China; University of Chinese Academy of Sciences, China), Haipeng Jia (Chinese Academy of Sciences, China), Yunquan Zhang (Chinese Academy of Sciences, China), Baicheng Yan (Huawei Technologies Co. Ltd, China), Penghao Ma (Huawei Technologies Co. Ltd, China), Long Wang (Huawei Technologies Co. Ltd, China), and Wenxuan Zhao (Chinese Academy of Sciences, China; University of Chinese Academy of Sciences, China)</i>	
Harnessing Deep Learning and HPC Kernels via High-Level Loop and Tensor Abstractions on CPU Architectures	950
<i>Evangelos Georganas (Intel Corporation), Dhiraj Kalamkar (Intel Corporation), Kirill Voronin (Intel Corporation), Abhisek Kundu (Intel Corporation), Antonio Noack (Friedrich Schiller Universität, Jena), Hans Pabst (Intel Corporation), Alexander Breuer (Friedrich Schiller Universität, Jena), and Alexander Heinecke (Intel Corporation)</i>	
Optimizing General Matrix Multiplications on Modern Multi-core DSPs	964
<i>Kainan Yu (National University of Defense Technology, China), Xinxin Qi (National University of Defense Technology, China), Peng Zhang (National University of Defense Technology, China), Jianbin Fang (National University of Defense Technology, China), Dezun Dong (National University of Defense Technology, China), Ruiibo Wang (National University of Defense Technology, China), Tao Tang (National University of Defense Technology, China), Chun Huang (National University of Defense Technology, China), Yonggang Che (National University of Defense Technology, China), and Zheng Wang (Northwest University, China)</i>	
Machine-Learning-Driven Runtime Optimization of BLAS Level 3 on Modern Multi-Core Systems	976
<i>Yufan Xia (The Chinese University of Hong Kong, China) and Giuseppe Maria Junior Barca (The University Of Melbourne, Australia)</i>	

Session 9A: Distributed Algorithms

Time-Color Tradeoff on Uniform Circle Formation by Asynchronous Robots	987
<i>Debasish Pattanayak (Carleton University, Canada) and Gokarna Sharma (Kent State University, USA)</i>	
LightDAG: A Low-Latency DAG-Based BFT Consensus through Lightweight Broadcast	998
<i>Xiaohai Dai (National Engineering Research Center for Big Data Technology and System; Services Computing Technology and System Lab, Cluster and Grid Computing Lab; Huazhong University of Science and Technology, China), Guanxiong Wang (National Engineering Research Center for Big Data Technology and System; Services Computing Technology and System Lab, Cluster and Grid Computing Lab; Huazhong University of Science and Technology, China), Jiang Xiao (National Engineering Research Center for Big Data Technology and System; Services Computing Technology and System Lab, Cluster and Grid Computing Lab; Huazhong University of Science and Technology, China), Zhengxuan Guo (National Engineering Research Center for Big Data Technology and System; Services Computing Technology and System Lab, Cluster and Grid Computing Lab; Huazhong University of Science and Technology, China), Rui Hao (Wuhan University of Technology, China), Xia Xie (Hainan University, China), and Hai Jin (National Engineering Research Center for Big Data Technology and System; Services Computing Technology and System Lab, Cluster and Grid Computing Lab; Huazhong University of Science and Technology, China)</i>	
MAAD: A Distributed Anomaly Detection Architecture for Microservices Systems	1009
<i>Rongyuan Tan (Southern University of Science and Technology, China) and Zhuozhao Li (Southern University of Science and Technology, China)</i>	
OneShot: View-Adapting Streamlined BFT Protocols with Trusted Execution Environments	1022
<i>Jérémie Decouchant (TU Delft, The Netherlands), David Kozhaya (ABB Research, Switzerland), Vincent Rahli (University of Birmingham, UK), and Jiangshan Yu (University of Sydney, Australia)</i>	

Session 9B: Graph Algorithms

Practically Tackling Memory Bottlenecks of Graph-Processing Workloads	1034
<i>Alexandre Valentin Jamet (Universitat Politècnica de Catalunya (UPC)), Georgios Vavouliotis (Huawei Zurich Research Center), Daniel A. Jiménez (Texas A&M University), Lluc Alvarez (Universitat Politècnica de Catalunya (UPC)), and Marc Casas (Universitat Politècnica de Catalunya (UPC))</i>	
GCSM: GPU-Accelerated Continuous Subgraph Matching for Large Graphs	1046
<i>Yihua Wei (University of Iowa) and Peng Jiang (University of Iowa)</i>	
Parallel Derandomization for Coloring	1058
<i>Sam Coy (University of Warwick, United Kingdom), Artur Czumaj (University of Warwick, United Kingdom), Peter Davies-Peck (Durham University, United Kingdom), and Gopinath Mishra (National University of Singapore, Singapore)</i>	

A Comparative Study of Intersection-Based Triangle Counting Algorithms on GPUs 1070
Jiangbo Li (Nanchang University, China), Zichen Xu (Nanchang University, China), Minh Pham (University of South Florida, USA), Yicheng Tu (University of South Florida, USA), and Qihe Zhou (City University of Macau, USA)

Author Index