

# **2022 IEEE International Parallel and Distributed Processing Symposium (IPDPS 2022)**

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

**Pages 1-682**



**IEEE Catalog Number: CFP22023-POD  
ISBN: 978-1-6654-8107-6**

**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:	CFP22023-POD
ISBN (Print-On-Demand):	978-1-6654-8107-6
ISBN (Online):	978-1-6654-8106-9
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](mailto:curran@proceedings.com)  
Web: [www.proceedings.com](http://www.proceedings.com)

CURRAN ASSOCIATES INC.  
**proceedings**  
.com

# 2022 IEEE International Parallel and Distributed Processing Symposium (IPDPS) **IPDPS 2022**

## Table of Contents

Message from the 2022 General Co-Chairs .....	xxiii
Message from the 2022 Program Chairs .....	xxv
IPDPS 2022 Technical Program .....	xxvii
IPDPS 2022 Organization .....	xxix

### Session 1: Welcome to IPDPS 2022

#### Session 2: Keynote 1

Challenges and Opportunities in Designing High-Performance and Scalable Middleware for HPC and AI: Past, Present, and Future .....	1
<i>DK Panda (The Ohio State University, USA)</i>	

#### Session 3: Algo-1 Linear Algebra

HTS: A Threaded Multilevel Sparse Hybrid Solver .....	2
<i>Joshua Dennis Booth (University of Alabama in Huntsville, USA)</i>	
A Scalable Adaptive-Matrix SpMV for Heterogeneous Architectures .....	13
<i>Han D. Tran (University of Utah, USA), Milinda Fernando (University of Texas at Austin, USA), Kumar Saurabh (Iowa State University, USA), Baskar Ganapathysubramanian (Iowa State University, USA), Robert M. Kirby (University of Utah, USA), and Hari Sundar (University of Utah, USA)</i>	
Direct Solution of Larger Coupled Sparse/dense Linear Systems using Low-Rank Compression on Single-Node Multi-core Machines in an Industrial Context .....	25
<i>Emmanuel Agullo (Inria Bordeaux Sud-Ouest, France), Marek Felsoci (Inria Bordeaux Sud-Ouest, France), and Guillaume Sylvand (Airbus Central R&amp;T, France)</i>	
I/O-Optimal Cache-Oblivious Sparse Matrix-Sparse Matrix Multiplication .....	36
<i>Niels Gleinig (ETH Zurich), Maciej Besta (ETH Zurich), and Torsten Hoefler (ETH Zurich)</i>	

Distributed-Memory Sparse Kernels for Machine Learning .....	47
<i>Vivek Bharadwaj (University of California, Berkeley, USA), Aydın Buluç (Computational Research Division, Lawrence Berkeley National Laboratory, USA; University of California, Berkeley, USA), and James Demmel (University of California, Berkeley, USA)</i>	

## Session 4: SysSoft-1 Memory Management

PokéMem: Taming Wild Memory Consumers in Apache Spark .....	59
<i>Minhyeok Kweun (Samsung Research, Korea), Goeun Kim (Samsung Research, Korea), Byungsoo Oh (Samsung Research, Korea), Seongho Jung (Samsung Research, Korea), Taegeon Um (Samsung Research, Korea), and Woo-Yeon Lee (Samsung Research, Korea)</i>	
SSB-Tree: Making Persistent Memory B+-Trees Crash-Consistent and Concurrent by Lazy-Box .....	70
<i>Tongliang Li (Tsinghua University, China), Haixia Wang (Tsinghua University, China), Airan Shao (Tsinghua University, China), and Dongsheng Wang (Tsinghua University, China)</i>	
FAM-Graph: Graph Analytics on Disaggregated Memory .....	81
<i>Daniel Zahka (Georgia Institute of Technology, USA) and Ada Gavrilovska (Georgia Institute of Technology, USA)</i>	
Scalable Multi-versioning Ordered Key-Value Stores with Persistent Memory Support .....	93
<i>Bogdan Nicolae (Argonne National Laboratory, USA)</i>	
In-Memory Indexed Caching for Distributed Data Processing .....	104
<i>Alexandru Uta (Leiden University), Bogdan Ghit (Databricks), Ankur Dave (UC Berkeley), Jan Rellermeyer (TU Delft), and Peter Boncz (CWI Amsterdam)</i>	

## Session 5: Multi-1 GPUs for Mutlidisciplinary Applications

Landau Collision Operator in the CUDA Programming Model Applied to Thermal Quench Plasmas ...	115
<i>Mark F. Adams (Lawrence Berkeley National Laboratory, USA), Dylan P. Brennan (Princeton University, USA), Matthew G. Knepley (University at Buffalo, USA), and Peng Wang (NVIDIA Corporation, USA)</i>	
Exploiting Reduced Precision for GPU-Based Time Series Mining .....	124
<i>Yi Ju (Max Plank Computing and Data Facility, Germany), Amir Raoofy (Technical University of Munich, Germany), Dai Yang (NVIDIA GmbH), Erwin Laure (Max Plank Computing and Data Facility, Germany), and Martin Schulz (Technical University of Munich, Germany)</i>	
MICCO: An Enhanced Multi-GPU Scheduling Framework for Many-Body Correlation Functions ..	135
<i>Qihan Wang (William &amp; Mary, USA), Bin Ren (William &amp; Mary, USA), Jie Chen (Jefferson Lab, USA), and Robert G. Edwards (Jefferson Lab, USA)</i>	
Unlocking Personalized Healthcare on Modern CPUs/GPUs: Three-way Gene Interaction Study ...	146
<i>Diogo Marques (INESC-ID, Portugal), Rafael Campos (INESC-ID, Portugal), Sergio Santander-Jiménez (University of Extremadura, Spain), Zakhar Matveev (Intel Corporation, Russia), Leonel Sousa (INESC-ID, Portugal), and Aleksandar Ilic (INESC-ID, Portugal)</i>	

Batched Sparse Iterative Solvers on GPU for the Collision Operator for Fusion Plasma Simulations .....	157
<i>Aditya Kashi (Karlsruhe Institute of Technology, Karlsruhe, Germany), Pratik Nayak (Karlsruhe Institute of Technology, Karlsruhe, Germany), Dhruva Kulkarni (Lawrence Berkeley National Laboratory, Berkeley, USA), Aaron Scheinberg (Jubilee Development, Cambridge, USA), Paul Lin (Lawrence Berkeley National Laboratory, Berkeley, USA), and Hartwig Anzt (Karlsruhe Institute of Technology, Karlsruhe, Germany; University of Tennessee, Knoxville, USA)</i>	

## Session 6: Exp-1 GPU Applications

PARSEC: PARallel Subgraph Enumeration in CUDA .....	168
<i>Vibhor Dodeja (University of Illinois Urbana-Champaign, USA), Mohammad Almasri (University of Illinois Urbana-Champaign, USA), Rakesh Nagi (University of Illinois Urbana-Champaign, USA), Jinjun Xiong (University of Buffalo, USA; University at Buffalo, USA), and Wen-Mei Hwu (Nvidia Corporation, USA; Nvidia Corporation, USA)</i>	
Top-Down Performance Profiling on NVIDIA's GPUs .....	179
<i>Alvaro Saiz (Universidad de Cantabria, Spain), Pablo Prieto (Universidad de Cantabria, Spain), Pablo Abad (Universidad de Cantabria, Spain), Jose Angel Gregorio (Universidad de Cantabria, Spain), and Valentin Puente (Universidad de Cantabria, Spain)</i>	
Scaling and Selecting GPU Methods for All Pairs Shortest Paths (APSP) Computations .....	190
<i>Yang Xia (Ohio State University, USA), Peng Jiang (University of Iowa, USA), Gagan Agrawal (Augusta University, USA), and Rajiv Ramnath (Ohio State University, USA)</i>	
Parallel Vertex Cover Algorithms on GPUs .....	201
<i>Peter Yamout (American University of Beirut, Lebanon), Karim Barada (American University of Beirut, Lebanon), Adnan Jaljuli (American University of Beirut, Lebanon), Amer E. Mouawad (American University of Beirut, Lebanon), and Izzat El Hajj (American University of Beirut, Lebanon)</i>	

## Session 7: Arch-1 Memory Management

SecFortress: Securing Hypervisor using Cross-Layer Isolation .....	212
<i>Qihang Zhou (Institute of Information Engineering, Chinese Academy of Sciences; University of Chinese Academy of Sciences), Xiaoqi Jia (Institute of Information Engineering, Chinese Academy of Sciences; University of Chinese Academy of Sciences), Shengzhi Zhang (Metropolitan College, Boston University, USA), Nan Jiang (Institute of Information Engineering, Chinese Academy of Sciences; University of Chinese Academy of Sciences), Jiayun Chen (Institute of Information Engineering, Chinese Academy of Sciences; University of Chinese Academy of Sciences), and Weijuan Zhang (Institute of Information Engineering, Chinese Academy of Sciences)</i>	

Exploring Efficient Microservice Level Parallelism .....	223
<i>Xinkai Wang (Shanghai Jiao Tong University, China), Chao Li (Shanghai Jiao Tong University, China), Lu Zhang (Shanghai Jiao Tong University, China), Xiaofeng Hou (Hong Kong University of Science and Technology, China), Quan Chen (Shanghai Jiao Tong University, China), and Minyi Guo (Shanghai Jiao Tong University, China)</i>	
Scalable Low-Latency Inter-FPGA Networks .....	234
<i>Kien Trung Pham (Graduate University for Advanced Studies SOKENDAI, Japan), Thao Nguyen Truong (National Institute of Advanced Industrial Science and Technology (AIST), Japan), Hiroshi Yamaguchi (Photonics Electronics Technology Research Association, Japan), Yutaka Urino (Photonics Electronics Technology Research Association, Japan), and Michihiro Koibuchi (National Institute of Informatics, Japan)</i>	
A General Offloading Approach for Near-DRAM Processing-In-Memory Architectures .....	246
<i>Dan Chen (Huazhong University of Science and Technology, China), Hai Jin (Huazhong University of Science and Technology, China), Long Zheng (Huazhong University of Science and Technology, China), Yu Huang (Huazhong University of Science and Technology, China), Pengcheng Yao (Huazhong University of Science and Technology, China), Chuangyi Gui (Huazhong University of Science and Technology, China), Qinggang Wang (Huazhong University of Science and Technology, China), Haifeng Liu (Huazhong University of Science and Technology, China), Haiheng He (Huazhong University of Science and Technology, China), Xiaofei Liao (Huazhong University of Science and Technology, China), and Ran Zheng (Huazhong University of Science and Technology, China)</i>	
Minerva: Rethinking Secure Architectures for the Era of Fabric-Attached Memory Architectures .....	258
<i>Mazen Alwadi (Jordan University of Science and Technology), Rujia Wang (Illinois Institute of Technology, USA), David Mohaisen (University of Central Florida, USA), Clayton Hughes (Sandia National Laboratories, USA), Simon Hammond (Sandia National Laboratories, USA), and Amro Awad (North Carolina State University, USA)</i>	

## Session 8: Algo-2 Graphs

Parallel Global Edge Switching for the Uniform Sampling of Simple Graphs with Prescribed Degrees .....	269
<i>Daniel Allendorf (Goethe-Universität-Frankfurt, Germany), Ulrich Meyer (Goethe-Universität-Frankfurt, Germany), Manuel Penschuck (Goethe-Universität-Frankfurt, Germany), and Hung Tran (Goethe-Universität-Frankfurt, Germany)</i>	
Parallel, Portable Algorithms for Distance-2 Maximal Independent Set and Graph Coarsening .....	280
<i>Brian Kelley (Sandia National Laboratories, USA) and Sivasankaran Rajamanickam (Sandia National Laboratories, USA)</i>	
Asynchronous Distributed-Memory Triangle Counting and LCC with RMA Caching .....	291
<i>András Strausz (ETH Zürich, Switzerland), Flavio Vella (University of Trento, Italy), Salvatore Di Girolamo (ETH Zürich, Switzerland), Maciej Besta (ETH Zürich, Switzerland), and Torsten Hoefler (ETH Zürich, Switzerland)</i>	

Communication-Efficient Massively Distributed Connected Components .....	302
<i>Sebastian Lamm (Karlsruhe Institute of Technology, Germany) and Peter Sanders (Karlsruhe Institute of Technology, Germany)</i>	
Mnemonic: A Parallel Subgraph Matching System for Streaming Graphs .....	313
<i>Bibek Bhattacharai (George Washington University, USA) and Howie Huang (George Washington University, USA)</i>	

## Session 9: SysSoft-2 System Scheduling

QoS-Awareness of Microservices with Excessive Loads via Inter-Datacenter Scheduling .....	324
<i>Jiuchen Shi (Shanghai Jiao Tong University, China), Jiawen Wang (Shanghai Jiao Tong University, China), Kaihua Fu (Shanghai Jiao Tong University, China), Quan Chen (Shanghai Jiao Tong University, China), Deze Zeng (China University of Geosciences, China), and Minyi Guo (Shanghai Jiao Tong University, China)</i>	
Resource Utilization Aware Job Scheduling to Mitigate Performance Variability .....	335
<i>Daniel Nichols (University of Maryland, USA), Aniruddha Marathe (Lawrence Livermore National Laboratory, USA), Kathleen Shoga (Lawrence Livermore National Laboratory, USA), Todd Gamblin (Lawrence Livermore National Laboratory, USA), and Abhinav Bhatele (University of Maryland, USA)</i>	
Dynamic Task Shaping for High Throughput Data Analysis Applications in High Energy Physics..	346
<i>Benjamin Tovar (University of Notre Dame, USA), Benjamin Lyons (University of Notre Dame, USA), Kelci Mohrman (University of Notre Dame, USA), Barry Sly-Delgado (University of Notre Dame, USA), Kevin Lannon (University of Notre Dame, USA), and Douglas Thain (University of Notre Dame, USA)</i>	
Multi-phase Task-Based HPC Applications: Quickly Learning how to Run Fast .....	357
<i>Lucas Leandro Nesi (Institute of Informatics, PPGC/UFRGS, Brazil), Lucas Mello Schnorr (Institute of Informatics, PPGC/UFRGS, Brazil), and Arnaud Legrand (Université Grenoble Alpes, CNRS, Inria, France)</i>	
DFMan: A Graph-based Optimization of Dataflow Scheduling on High-Performance Computing Systems .....	368
<i>Fahim Chowdhury (Florida State University, USA), Francesco Di Natale (Lawrence Livermore National Laboratory, USA), Adam Moody (Lawrence Livermore National Laboratory, USA), Kathryn Mohror (Lawrence Livermore National Laboratory, USA), and Weikuan Yu (Florida State University, USA)</i>	

## Session 10: Multi-2 Scheduling Applications

Parallel Approximations of the Tukey g-and-h Likelihoods and Predictions for Non-Gaussian Geostatistics .....	379
<i>Sagnik Mondal (King Abdullah University of Science and Technology, Saudi Arabia), Sameh Abdulah (King Abdullah University of Science and Technology, Saudi Arabia), 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>	
Parallelizing and Balancing Coupled DSMC/PIC for Large-Scale Particle Simulations .....	390
<i>Haozhong Qiu (National University of Defense Technology, China), Chuanfu Xu (National University of Defense Technology, China), Dali Li (National University of Defense Technology, China), Haoyu Wang (National University of Defense Technology, China), Jie Li (National University of Defense Technology, China), and Zheng Wang (University of Leeds, United Kingdom)</i>	
Next-Generation Local Time Stepping for the ADER-DG Finite Element Method .....	402
<i>Alexander Breuer (Friedrich Schiller University Jena, Germany) and Alexander Heinecke (Intel Corporation, USA)</i>	
A Framework to Exploit Data Sparsity in Tile Low-Rank Cholesky Factorization .....	414
<i>Qinglei Cao (University of Tennessee, USA), Rabab Alomairy (King Abdullah University of Science and Technology (KAUST), Saudi Arabia), Yu Pei (University of Tennessee, USA), George Bosilca (University of Tennessee, USA), Hatem Ltaief (King Abdullah University of Science and Technology (KAUST), Saudi Arabia), David Keyes (King Abdullah University of Science and Technology (KAUST), Saudi Arabia), and Jack Dongarra (University of Tennessee, USA)</i>	
On the Parallel Reconstruction from Pooled Data .....	425
<i>Oliver Gebhard (TU Dortmund University, Germany), Max Hahn-Klimroth (TU Dortmund University, Germany), Dominik Kaaser (Universität Hamburg, Germany), and Philipp Loick (Goethe University Frankfurt, Germany)</i>	

## Session 11: Exp-2 Scheduling & Optimization

P-ckpt: Coordinated Prioritized Checkpointing .....	436
<i>Subhendu Behera (North Carolina State University, USA), Lipeng Wan (Oak Ridge National Laboratory, USA), Frank Mueller (North Carolina State University, USA), Matthew Wolf (Oak Ridge National Laboratory, USA), and Scott Klasky (Oak Ridge National Laboratory, USA)</i>	
TEE-Based Decentralized Recommender Systems: The Raw Data Sharing Redemption .....	447
<i>Akash Dhasade (EPFL - Swiss Federal Institute of Technology, Switzerland), Nevena Dresevic (EPFL - Swiss Federal Institute of Technology, Switzerland), Anne-Marie Kermarrec (EPFL - Swiss Federal Institute of Technology, Switzerland), and Rafael Pires (EPFL - Swiss Federal Institute of Technology, Switzerland)</i>	



Accuracy vs. Cost in Parallel Fixed-Precision Low-Rank Approximations of Sparse Matrices .....	459
<i>Robert Ernstbrunner (University of Vienna, Austria), Viktoria Mayer (University of Vienna, Austria), and Wilfried Gansterer (University of Vienna, Austria)</i>	
Hybrid Workload Scheduling on HPC Systems .....	470
<i>Yuping Fan (Illinois Institute of Technology, USA), Zhiling Lan (Illinois Institute of Technology, USA), Paul Rich (Argonne National Laboratory, USA), William Allcock (Argonne National Laboratory, USA), and Michael E. Papka (Argonne National Laboratory, USA)</i>	

## Session 12: Models-1 GPU Programming

GSpecPal: Speculation-Centric Finite State Machine Parallelization on GPUs .....	481
<i>Yuguang Wang (Michigan Technological University, USA), Robbie Watling (Michigan Technological University, USA), Junqiao Qiu (Michigan Technological University, USA), and Zhenlin Wang (Michigan Technological University, USA)</i>	
Lightning: Scaling the GPU Programming Model Beyond a Single GPU .....	492
<i>Stijn Heldens (Netherlands eScience Center, the Netherlands; University of Amsterdam, the Netherlands), Pieter Hijma (University of Amsterdam; Vrije Universiteit Amsterdam, the Netherlands), Ben van Werkhoven (Netherlands eScience Center, the Netherlands), Jason Maassen (Netherlands eScience Center, the Netherlands), and Rob V. van Nieuwpoort (Netherlands eScience Center, the Netherlands; University of Amsterdam, the Netherlands)</i>	
Co-designing an OpenMP GPU Runtime and Optimizations for Near-Zero Overhead Execution ...	504
<i>Johannes Doerfert (Argonne National Laboratory, USA), Atemn Patel (University of Waterloo, Canada), Joseph Huber (Oak Ridge National Laboratory, USA), Shilei Tian (Stony Brook University, USA), Jose M Monsalve Diaz (Argonne National Laboratory, USA), Barbara Chapman (Stony Brook University, USA), and Giorgis Georgakoudis (Lawrence Livermore National Laboratory, USA)</i>	
Bit-GraphBLAS: Bit-Level Optimizations of Matrix-Centric Graph Processing on GPU .....	515
<i>Jou-An Chen (North Carolina State University, USA), Hsin-Hsuan Sung (North Carolina State University, USA), Xipeng Shen (North Carolina State University, USA), Nathan Tallent (Pacific Northwest National Laboratory, USA), Kevin Barker (Pacific Northwest National Laboratory, USA), and Ang Li (Pacific Northwest National Laboratory, USA)</i>	
CSMV: A Highly Scalable Multi-versioned Software Transactional Memory for GPUs .....	526
<i>Diogo Nunes (University of Lisbon - Instituto Superior Técnico &amp; INESC-ID, Portugal), Daniel Castro (University of Lisbon - Instituto Superior Técnico &amp; INESC-ID, Portugal), and Paolo Romano (University of Lisbon - Instituto Superior Técnico &amp; INESC-ID, Portugal)</i>	

## Session 13: Keynote 2

Resilience at Extreme Scale and Connections with Other Domains .....	537
<i>Leonardo Bautista Gomez (Barcelona Supercomputing Center, Spain)</i>	

## Session 14: Best Paper Candidates

Colza: Enabling Elastic In Situ Visualization for High-Performance Computing Simulations .....	538
<i>Matthieu Dorier (Argonne National Laboratory, USA), Zhe Wang (Rutgers University, USA), Utkarsh Ayachit (Kitware, Inc., USA), Shane Snyder (Argonne National Laboratory, USA), Rob Ross (Argonne National Laboratory, USA), and Manish Parashar (University of Utah, USA)</i>	
Towards Distributed 2-Approximation Steiner Minimal Trees in Billion-Edge Graphs .....	549
<i>Tahsin Reza (Lawrence Livermore National Laboratory, USA), Geoffrey Sanders (Lawrence Livermore National Laboratory, USA), and Roger Pearce (Lawrence Livermore National Laboratory, USA)</i>	
As Easy as ABC: Optimal (A) Countable (B) Yzantine (C) Onsensus is Easy! .....	560
<i>Pierre Civiit (Sorbonne University, CNRS, LIP6, France), Seth Gilbert (NUS Singapore, Singapore), Vincent Gramoli (University of Sydney, Australia; Ecole Polytechnique Federale de Lausanne (EPFL) Switzerland), Rachid Guerraoui (Ecole Polytechnique Federale de Lausanne (EPFL), Switzerland), and Jovan Komatovic (Ecole Polytechnique Federale de Lausanne (EPFL), Switzerland)</i>	
Understanding the Design-Space of Sparse/Dense Multiphase GNN Dataflows on Spatial Accelerators .....	571
<i>Raveesh Garg (Georgia Institute of Technology, USA), Eric Qin (Georgia Institute of Technology, USA), Francisco Muñoz-Martínez (Universidad de Murcia, Spain), Robert Guirado (Universitat Politècnica de Catalunya, Spain), Akshay Jain (Neutron), Sergi Abadal (Universitat Politècnica de Catalunya, Spain), José L. Abellán (Universidad Católica de Murcia, Spain), Manuel E. Acacio (Universidad de Murcia, Spain), Eduard Alarcón (Universitat Politècnica de Catalunya, Spain), Sivasankaran Rajamanickam (Sandia National Laboratories, USA), and Tushar Krishna (Georgia Institute of Technology, USA)</i>	
“Smarter” NICs for Faster Molecular Dynamics: A Case Study .....	583
<i>Sara Karamati (Georgia Institute of Technology, USA), Clayton Hughes (Sandia National Laboratories, USA), K. Scott Hemmert (Sandia National Laboratories, USA), Ryan E. Grant (Queen’s University, Canada), W. Whit Schonbein (Sandia National Laboratories, USA), Scott Levy (Sandia National Laboratories, USA), Thomas M. Conte (Georgia Institute of Technology, USA), Jeffrey Young (Georgia Institute of Technology, USA), and Richard W. Vuduc (Georgia Institute of Technology, USA)</i>	

## Session 15: Algo-3 Machine Learning

RLRP: High-Efficient Data Placement with Reinforcement Learning for Modern Distributed Storage Systems .....	595
<i>Kai Lu (Huazhong University of Science and Technology, China), Nannan Zhao (Northwestern Polytechnical University, China), Jiguang Wan (Huazhong University of Science and Technology, China), Changhong Fei (Huazhong University of Science and Technology, China), Wei Zhao (SenseTime Research, China), and Tongliang Deng (SenseTime Research, China)</i>	

AxoNN: An Asynchronous, Message-Driven Parallel Framework for Extreme-Scale Deep Learning ....	606
<i>Siddharth Singh (University of Maryland, USA) and Abhinav Bhatle (University of Maryland, USA)</i>	
Fast Parallel Bayesian Network Structure Learning .....	617
<i>Jiantong Jiang (The University of Western Australia, Australia), Zeyi Wen (The University of Western Australia, Australia), and Ajmal Mian (The University of Western Australia, Australia)</i>	
Adaptive Verifiable Coded Computing: Towards Fast, Secure and Private Distributed Machine Learning .....	628
<i>Tingting Tang (University of Southern California, USA), Ramy E. Ali (University of Southern California, USA), Hanieh Hashemi (University of Southern California, USA), Tynan Gangwani (University of Southern California, USA), Salman Avestimehr (University of Southern California, USA), and Murali Annavaram (University of Southern California, USA)</i>	
pFedGF: Enabling Personalized Federated Learning via Gradient Fusion .....	639
<i>Xinghao Wu (Beihang University, China), Jianwei Niu (Beihang University, China), Xuefeng Liu (Beihang University, China), Tao Ren (Beihang University, China), Zhangmin Huang (Beihang University, China), and Zhetao Li (Xiangtan University, China)</i>	

## Session 16: Algo-4 Scheduling

An Efficient Vectorization Scheme for Stencil Computation .....	650
<i>Kun Li (Institute of Computing Technology of Chinese Academy of Sciences; University of Chinese Academy of Sciences), Liang Yuan (Institute of Computing Technology of Chinese Academy of Sciences), Yunquan Zhang (Institute of Computing Technology of Chinese Academy of Sciences), Yue Yue (Institute of Computing Technology of Chinese Academy of Sciences; University of Chinese Academy of Sciences), and Hang Cao (Institute of Computing Technology of Chinese Academy of Sciences; University of Chinese Academy of Sciences)</i>	
Scheduling on Uniform and Unrelated Machines with Bipartite Incompatibility Graphs .....	661
<i>Tytus Pikies (Gdańsk University of Technology, Poland) and Hanna Furmańczyk (University of Gdańsk, Poland)</i>	
SPIDER: An Effective, Efficient and Robust Load Scheduler for Real-Time Split Frame Rendering .....	672
<i>Bingzheng Ma (Nankai University, China), Ziqiang Zhang (Nankai University, China), Yusen Li (Nankai University, China), Wentong Cai (Nanyang Technological University, Singapore), Gang Wang (Nankai University, China), and Xiaoguang Liu (Nankai University, China)</i>	
Bounding the Flow Time in Online Scheduling with Structured Processing Sets .....	683
<i>Louis-Claude Canon (FEMTO-ST Institute, Univ. Franche-Comté, France), Anthony Dugois (LIP, ENS Lyon, Inria, France), and Loris Marchal (LIP, ENS Lyon, CNRS, Inria, France)</i>	

Memory-Aware Scheduling of Tasks Sharing Data on Multiple GPUs with Dynamic Runtime Systems .....	694
<i>Maxime Gonthier (ENS-Lyon, France), Loris Marchal (ENS-Lyon, France), and Samuel Thibault (University of Bordeaux, France)</i>	

## Session 17: SysSoft-3 GPU Systems

Accelerating Encrypted Computing on Intel GPUs .....	705
<i>Yujia Zhai (University of California, USA), Mohammad Ibrahim (North Carolina State University, USA), Yiqin Qiu (Intel Corporation, USA), Fabian Boemer (Intel Corporation, USA), Zizhong Chen (University of California, USA), Alexey Titov (Intel Corporation, USA), and Alexander Lyashevsky (Intel Corporation, USA)</i>	
Optimizing Huffman Decoding for Error-Bounded Lossy Compression on GPUs .....	717
<i>Cody Rivera (University of Alabama, USA), Sheng Di (Argonne National Laboratory, USA), Jiannan Tian (Washington State University, USA), Xiaodong Yu (Argonne National Laboratory, USA), Dingwen Tao (Washington State University, USA), and Franck Cappello (Argonne National Laboratory, USA)</i>	
SALoBa: Maximizing Data Locality and Workload Balance for Fast Sequence Alignment on GPUs	728
<i>Seongyeon Park (Yonsei University, South Korea), Hajin Kim (Yonsei University, South Korea), Tanveer Ahmad (TU Delft, Netherlands), Nauman Ahmed (TU Delft, Netherlands), Zaid Al-Ars (TU Delft, Netherlands), H. Peter Hofstee (TU Delft, Netherlands, IBM, United States of America), Youngsok Kim (Yonsei University, South Korea), and Jinho Lee (Yonsei University, South Korea)</i>	
DGSF: Disaggregated GPUs for Serverless Functions .....	739
<i>Henrique Fingler (The University of Texas at Austin, USA), Zhiting Zhu (The University of Texas at Austin, USA), Esther Yoon (The University of Texas at Austin, USA), Zhipeng Jia (The University of Texas at Austin, USA), Emmett Witchel (The University of Texas at Austin, USA), and Christopher Rossbach (The University of Texas at Austin, USA)</i>	
Compiler-Directed Incremental Checkpointing for Low Latency GPU Preemption .....	751
<i>Zhuoran Ji (The University of Hong Kong, China) and Cho-Li Wang (The University of Hong Kong, China)</i>	

## Session 18: Multi-3 Graphs & Massive Data

ParaTreeT: A Fast, General Framework for Spatial Tree Traversal .....	762
<i>Joseph Hutter (University of Illinois at Chicago, USA), Justin Szaday (University of Illinois at Chicago, USA), Jaemin Choi (University of Illinois at Chicago, USA), Simeng Liu (University of Illinois at Chicago, USA), Laxmikant Kale (University of Illinois at Chicago, USA), Spencer Wallace (University of Washington, USA), and Thomas Quinn (University of Washington, USA)</i>	

An Integral-Equation-Oriented Vectorized SpMV Algorithm and Its Application on CT Imaging Reconstruction .....	773
<i>Weicai Ye (Sun Yat-sen University, China), Chenghuan Huang (Sun Yat-sen University, China), Jiasheng Huang (Sun Yat-sen University, China), Jiajun Li (Sun Yat-sen University, China), Yao Lu (Sun Yat-sen University, China), and Ying Jiang (Sun Yat-sen University, China)</i>	
High-Order Line Graphs of Non-Uniform Hypergraphs: Algorithms, Applications, and Experimental Analysis .....	784
<i>Xu T. Liu (University of Washington, Washington State University, USA), Jesun Firoz (Pacific Northwest National Laboratory, USA), Sinan Akoy (Pacific Northwest National Laboratory, USA), Ilya Amburg (Pacific Northwest National Laboratory, USA), Andrew Lumsdaine (University of Washington, Pacific Northwest National Laboratory, USA), Cliff Joslyn (Pacific Northwest National Laboratory, USA), Brenda Praggastis (Pacific Northwest National Laboratory, USA), and Assefaw H. Gebremedhin (Washington State University, USA)</i>	
Topological Modeling and Parallelization of Multidimensional Data on Microelectrode Arrays.....	795
<i>Olamide Tawose (University of Nevada, Reno, USA), Bin Li (University of Nevada, Reno, USA), Lei Yang (University of Nevada, Reno, USA), Feng Yan (University of Nevada, Reno, USA), and Dongfang Zhao (University of Nevada, Reno, USA)</i>	
Coupling Streaming AI and HPC Ensembles to Achieve 100-1000× Faster Biomolecular Simulations .....	806
<i>Alexander Brace (Argonne National Laboratory, USA; University of Chicago, USA), Igor Yakushin (Argonne National Laboratory, USA), Heng Ma (Argonne National Laboratory, USA), Anda Trifan (Argonne National Laboratory, USA; University of Illinois Urbana-Champaign, USA), Todd Munson (Argonne National Laboratory, USA), Ian Foster (Argonne National Laboratory, USA; University of Chicago, USA), Arvind Ramanathan (Argonne National Laboratory, USA), Hyungro Lee (Rutgers University, USA), Matteo Turilli (Rutgers University, USA; Brookhaven National Laboratory, USA), and Shantenu Jha (Rutgers University, USA; Brookhaven National Laboratory, USA)</i>	

## Session 19: Models-2 Programming Models

Neon: A Multi-GPU Programming Model for Grid-Based Computations .....	817
<i>Massimiliano Meneghin (Autodesk Research, Canada), Ahmed H. Mahmoud (Autodesk Research, Canada; University of California, Davis, USA), Pradeep Kumar Jayaraman (Autodesk Research, Canada), and Nigel J. W. Morris (Autodesk Research, Canada)</i>	

OmpSs@cloudFPGA: An FPGA Task-Based Programming Model with Message Passing .....	828
<i>Juan Miguel de Haro (Barcelona Supercomputing Center, Spain; Universitat Politècnica de Catalunya, Spain), Rubén Cano (Barcelona Supercomputing Center, Spain), Carlos Álvarez (Barcelona Supercomputing Center, Spain; Universitat Politècnica de Catalunya, Spain), Daniel Jiménez-González (Barcelona Supercomputing Center, Spain; Universitat Politècnica de Catalunya, Spain), Xavier Martorell (Barcelona Supercomputing Center, Spain; Universitat Politècnica de Catalunya, Spain), Eduard Ayguadé (Barcelona Supercomputing Center, Spain; Universitat Politècnica de Catalunya, Spain), Jesús Labarta (Barcelona Supercomputing Center, Spain; Universitat Politècnica de Catalunya, Spain), Francois Abel (IBM Research Europe, Switzerland), Burkhard Ringlein (IBM Research Europe, Switzerland), and Beat Weiss (IBM Research Europe, Switzerland)</i>	
Generalized Flow-Graph Programming using Template Task-Graphs: Initial Implementation and Assessment .....	839
<i>Joseph Schuchart (The University of Tennessee, USA), Poornima Nookala (Stony Brook University, USA), Mohammad Mahdi Javanmard (Meta Platforms, Inc, USA), Thomas Haurault (The University of Tennessee, USA), Edward F. Valeev (Virginia Polytechnic Institute and State University, USA), George Bosilca (The University of Tennessee, USA), and Robert J. Harrison (Stony Brook University, USA)</i>	
PINT: Parallel INTERVAL-Based Race Detector .....	850
<i>Yifan Xu (Washington University in St. Louis, USA), Anchengcheng Zhou (Washington University in St. Louis, USA), Kunal Agrawal (Washington University in St. Louis, USA), and I-Ting Angelina Lee (Washington University in St. Louis, USA)</i>	

## Session 20: Keynote 3

Frugal Decentralized Learning .....	862
<i>Anne-Marie Kermarrec (Ecole Polytechnique Fédérale de Lausanne, Switzerland)</i>	

## Session 21: Algo-5 GPU Algorithms

A Fine-Grained Prefetching Scheme for DGEMM Kernels on GPU with Auto-Tuning Compatibility ....	863
<i>Jialin Li (Computer Network Information Center, Chinese Academy of Sciences, China; University of Chinese Academy of Sciences, China), Huang Ye (Computer Network Information Center, Chinese Academy of Sciences, China), Shaobo Tian (Computer Network Information Center, Chinese Academy of Sciences, China; University of Chinese Academy of Sciences, China), Xinyuan Li (Computer Network Information Center, Chinese Academy of Sciences; Alibaba Group, China), and Jian Zhang (Computer Network Information Center, Chinese Academy of Sciences, China)</i>	

StencilMART: Predicting Optimization Selection for Stencil Computations Across GPUs .....	875
<i>Qingxiao Sun (Beihang University, China), Yi Liu (Beihang University, China), Hailong Yang (Beihang University, China), Zhonghui Jiang (Beihang University, China), Zhongzhi Luan (Beihang University, China), and Depei Qian (Beihang University, China)</i>	
Mixed Precision s-Step Conjugate Gradient with Residual Replacement on GPUs .....	886
<i>Ichitaro Yamazaki (Sandia National Laboratories, USA), Erin Carson (Charles University, Prague, Czech Republic), and Brian Kelley (Sandia National Laboratories, USA)</i>	
Degree-Aware Kernels for Computing Jaccard Weights on GPUs .....	897
<i>Amro Alabsi Aljundi (Sabanc University, Turkey), Taha Atahan Akyildiz (Sabanc University, Turkey), and Kamer Kaya (Sabanc University, Turkey)</i>	
Fast and High-Quality Influence Maximization on Multiple GPUs .....	908
<i>Gökhan Gökürk (Sabanci University, Turkey) and Kamer Kaya (Sabanci University, Turkey)</i>	

## Session 22: Algo-6 Communications & Tensors

Traffic-Optimal Virtual Network Function Placement and Migration in Dynamic Cloud Data Centers .....	919
<i>Vincent Tran (California State University Dominguez Hills, USA), Jingsong Sun (California State University Dominguez Hills, USA), Bin Tang (California State University Dominguez Hills, USA), and Deng Pan (Florida International University, USA)</i>	
Parallel Tensor Train Rounding using Gram SVD .....	930
<i>Hussam Al Daas (Rutherford Appleton Laboratory, United Kingdom), Grey Ballard (Wake Forest University, USA), and Lawton Manning (Wake Forest University, USA)</i>	
Task-Based Acceleration of Bidirectional Recurrent Neural Networks on Multi-core Architectures .....	941
<i>Robin Kumar Sharma (Barcelona Supercomputing Center, Spain) and Marc Casas (Barcelona Supercomputing Center, Spain)</i>	
Sparsity-Aware Tensor Decomposition .....	952
<i>Sureyya Emre Kurt (University of Utah), Saurabh Raje (University of Utah), Aravind Sukumaran-Rajam (Washington State University), and P. Sadayappan (University of Utah)</i>	
Coloring the Vertices of 9-pt and 27-pt Stencils with Intervals .....	963
<i>Dante Durrman (UNC Charlotte, USA) and Erik Saule (UNC Charlotte, USA)</i>	

## Session 23: SysSoft-4 Caching & Learning

Falcon: A Timestamp-Based Protocol to Maximize the Cache Efficiency in the Distributed Shared Memory .....	974
<i>Jin Zhang (Shanghai Jiao Tong University, China), Xiangyao Yu (University of Wisconsin–Madison, United States), Zhengwei Qi (Shanghai Jiao Tong University, China), and Haibing Guan (Shanghai Jiao Tong University, China)</i>	

HACCS: Heterogeneity-Aware Clustered Client Selection for Accelerated Federated Learning .....	985
<i>Joel Wolfrath (University of Minnesota, USA), Nikhil Sreekumar (University of Minnesota, USA), Dhruv Kumar (University of Minnesota, USA), Yuanli Wang (University of Minnesota, USA), and Abhishek Chandra (University of Minnesota, USA)</i>	
A Swap Dominated Tensor Re-Generation Strategy for Training Deep Learning Models .....	996
<i>Lijie Wen (Tsinghua University, China), Zan Zong (Tsinghua University, China), Li Lin (Tsinghua University, China), and Leilei Lin (Capital Normal University, China)</i>	
Fast Convergence to Fairness for Reduced Long Flow Tail Latency in Datacenter Networks .....	1007
<i>John Snyder (Duke University, USA) and Alvin R. Lebeck (Duke University, USA)</i>	
Dynamic Computation Offloading for Green Things-Edge-Cloud Computing with Local Caching .....	1018
<i>Xianzhong Tian (Zhejiang University of Technology, China), Huixiao Meng (Zhejiang University of Technology, China), Yanjun Li (Zhejiang University of Technology, China), Pingting Miao (Zhejiang University of Technology, China), and Pengcheng Xu (Zhejiang University of Technology, China)</i>	

## Session 24: Arch-2 Networks

Excavating the Potential of Graph Workload on RDMA-Based Far Memory Architecture .....	1029
<i>Jing Wang (Shanghai Jiao Tong University, China), Chao Li (Shanghai Jiao Tong University, China), Taolei Wang (Shanghai Jiao Tong University, China), Lu Zhang (Shanghai Jiao Tong University, China), Pengyu Wang (Shanghai Jiao Tong University, China), Junyi Mei (Shanghai Jiao Tong University, China), and Minyi Guo (Shanghai Jiao Tong University, China)</i>	
SpectralFly: Ramanujan Graphs as Flexible and Efficient Interconnection Networks .....	1040
<i>Stephen Young (Pacific Northwest National Laboratory, USA), Sinan Aksoy (Pacific Northwest National Laboratory, USA), Jesun Firoz (Pacific Northwest National Laboratory, USA), Roberto Gioiosa (Pacific Northwest National Laboratory, USA), Tobias Hagge (Pacific Northwest National Laboratory, USA), Mark Kempton (Brigham Young University, USA), Juan Escobedo (Pacific Northwest National Laboratory, USA), and Mark Raugas (Pacific Northwest National Laboratory, USA)</i>	
Booster: An Accelerator for Gradient Boosting Decision Trees Training and Inference .....	1051
<i>Mingxuan He (Purdue University, USA), Mithuna Thottethodi (Purdue University, USA), and T. N. Vijaykumar (Purdue University, USA)</i>	
FlashWalker: An In-Storage Accelerator for Graph Random Walks .....	1063
<i>Fuping Niu (Huazhong University of Science and Technology, China), Jianhui Yue (Michigan Technological University, USA), Jiangqiu Shen (Michigan Technological University, USA), Xiaofei Liao (Huazhong University of Science and Technology, China), Haikun Liu (Huazhong University of Science and Technology, China), and Hai Jin (Huazhong University of Science and Technology, China)</i>	



Memory Access Granularity Aware Lossless Compression for GPUs .....	1074
<i>Sohan Lal (Technische Universität Hamburg, Germany), Manuel Renz (Technische Universität Berlin, Germany), Julian Hartmer (Technische Universität Berlin, Germany), and Ben Juurlink (Technische Universität Berlin, Germany)</i>	

## Session 25: Multi-4 Deep Learning

Why Globally Re-Shuffle? Revisiting Data Shuffling in Large Scale Deep Learning .....	1085
<i>Thao Nguyen Truong (National Institute of Advanced Science and Technology (AIST), Japan), François Trahay (Télécom SudParis, Institut Polytechnique de Paris, France), Jens Domke (RIKEN Center for Computational Science, Japan; Tokyo Institute of Technology, Japan), Aleksandr Drozd (RIKEN Center for Computational Science, Japan; Tokyo Institute of Technology, Japan; Amigawa GK, Japan), Emil Vatai (RIKEN Center for Computational Science, Japan; Tokyo Institute of Technology, Japan), Jianwei Liao (Southwest University of China, China), Mohamed Wahib (National Institute of Advanced Industrial Science &amp; Technology, Japan; RIKEN Center for Computational Science, Japan; Tokyo Institute of Technology, Japan), and Balazs Gerofi (RIKEN Center for Computational Science, Japan; Tokyo Institute of Technology, Japan)</i>	
DistrEdge: Speeding up Convolutional Neural Network Inference on Distributed Edge Devices ..	1097
<i>Xueyu Hou (New Jersey Institute of Technology, USA), Yongjie Guan (New Jersey Institute of Technology, USA), Tao Han (New Jersey Institute of Technology, USA), and Ning Zhang (Windsor University, Canada)</i>	
Model-Architecture Co-design for High Performance Temporal GNN Inference on FPGA .....	1108
<i>Hongkuan Zhou (University of Southern California, USA), Bingyi Zhang (University of Southern California, USA), Rajgopal Kannan (US Army Research Laboratory, USA), Viktor Prasanna (University of Southern California, USA), and Carl Busart (US Army Research Laboratory, USA)</i>	
Preprocessing Pipeline Optimization for Scientific Deep Learning Workloads .....	1118
<i>Khaled Z. Ibrahim (Lawrence Berkeley National Laboratory, USA) and Leonid Oliker (Lawrence Berkeley National Laboratory, USA)</i>	

## Session 26: Algo-7 Distributed Algorithms

Fault-Tolerant Snapshot Objects in Message Passing Systems .....	1129
<i>Vijay K. Garg (University of Texas at Austin, USA), Saptarni Kumar (Boston College, USA), Lewis Tseng (Boston College, USA), and Xiong Zheng (Google Inc, USA)</i>	
A Self-Stabilizing 2-Minimal Dominating set Algorithm Based on Loop Composition in Networks of Girth at Least 7 .....	1140
<i>Syohei Maruyama (Hiroshima University, Japan), Yuichi Sudo (Hosei University, Japan), Sayaka Kamei (Hiroshima University, Japan), and Hirotsugu Kakugawa (Ryukoku University, Japan)</i>	

Optimal Arbitrary Pattern Formation on a Grid by Asynchronous Autonomous Robots ..... 1151  
*Rory Hector (Louisiana State University, USA), Gokarna Sharma (Kent State University, USA), Ramachandran Vaidyanathan (Louisiana State University, USA), and Jerry L. Trahan (Louisiana State University, USA)*

The Universal Gossip Fighter ..... 1162  
*Anastasiia Gorbunova (École polytechnique fédérale de Lausanne, Switzerland), Rachid Guerraoui (École polytechnique fédérale de Lausanne, Switzerland), Anne-Marie Kermarrec (École polytechnique fédérale de Lausanne, Switzerland), Anastasiia Kucherenko (École polytechnique fédérale de Lausanne, Switzerland), and Rafaël Pinot (École polytechnique fédérale de Lausanne, Switzerland)*

## Session 27: Exp-3 Optimizing Applications

Modeling Matrix Engines for Portability and Performance ..... 1173  
*Nicholai Tukanov (Carnegie Mellon University, USA), Rajalakshmi Srinivasaraghavan (IBM Systems, USA), José E. Moreira (IBM Research, USA), and Tze Meng Low (Carnegie Mellon University, USA)*

MLCNN: Cross-Layer Cooperative Optimization and Accelerator Architecture for Speeding Up Deep Learning Applications ..... 1184  
*Beilei Jiang (University of North Texas, USA), Xianwei Cheng (University of North Texas, USA), Sihai Tang (University of North Texas, USA), Xu Ma (University of North Texas, USA), Zhaochen Gu (University of North Texas, USA), Song Fu (University of North Texas, USA), Qing Yang (University of North Texas, USA), and Mingxiong Liu (Los Alamos National Laboratory, USA)*

Shared-Memory Parallel Algorithms for Fully Dynamic Maintenance of 2-Connected Components ..... 1195  
*Chirayu Anant Haryan (IIT Tirupati, India), Ramakrishna G (IIT Tirupati, India), Kishore Kothapalli (IIIT Hyderabad, India), and Dip Sankar Banerjee (IIT Jodhpur, India)*

Learning Intermediate Representations using Graph Neural Networks for NUMA and Prefetchers Optimization ..... 1206  
*Ali TehraniJamsaz (Iowa State University, USA), Mihail Popov (Inria, France), Akash Dutta (Iowa State University, USA), Emmanuelle Saillard (Inria, France), and Ali Jannesari (Iowa State University, USA)*

HDagg: Hybrid Aggregation of Loop-Carried Dependence Iterations in Sparse Matrix Computations ..... 1217  
*Behrooz Zarebavani (University of Toronto, Canada), Kazem Cheshmi (University of Toronto, Canada), Bangtian Liu (University of Toronto, Canada), Michelle Mills Strout (University of Arizona, USA), and Maryam Mehri Dehnavi (University of Toronto, Canada)*

## Session 28: SysSoft-5 Blockchains

Alias-Chain: Improving Blockchain Scalability via Exploring Content Locality Among Transactions .....	1228
<i>Jintong Liu (Huazhong University of Science and Technology, China), Shenggang Wan (Huazhong University of Science and Technology, China), and Xubin He (Temple University, America)</i>	
SFP: Service Function Chain Provision on Programmable Switches for Cloud Tenants .....	1239
<i>Hongyi Huang (Tsinghua University, China), Wenfei Wu (Peking University, China), Yongchao He (Tsinghua University, China), and Zehua Guo (Beijing Institute of Technology, China)</i>	
An Efficient Block Validation Mechanism for UTXO-Based Blockchains .....	1250
<i>Xiaohai Dai (Huazhong University of Science and Technology, China), Bin Xiao (The Hong Kong Polytechnic University, Hong Kong), Jiang Xiao (Huazhong University of Science and Technology, China), and Hai Jin (Huazhong University of Science and Technology, China)</i>	
DEAN: A Lightweight and Resource-Efficient Blockchain Protocol for Reliable Edge Computing..	1261
<i>Abdullah Al-Mamun (University of Nevada, Reno, USA), Haoting Shen (University of Nevada, Reno, USA), and Dongfang Zhao (University of Nevada, Reno, USA)</i>	

## Session 29: SysSoft-6 System Optimization

PowerSpector: Towards Energy Efficiency with Calling-Context-Aware Profiling .....	1272
<i>Xin You (Beihang University, China), Hailong Yang (Beihang University, China), Zhibo Xuan (Beihang University, China), Zhongzhi Luan (Beihang University, China), and Depei Qian (Beihang University, China)</i>	
TagTree: Global Tagging Index with Efficient Querying for Time Series Databases .....	1283
<i>Jin Xue (The Chinese University of Hong Kong), Zhiqi Wang (The Chinese University of Hong Kong), Tianyu Wang (The Chinese University of Hong Kong), and Zili Shao (The Chinese University of Hong Kong)</i>	
An End-to-End and Adaptive I/O Optimization Tool for Modern HPC Storage Systems .....	1294
<i>Bin Yang (Shandong University, China; National Supercomputing Center in Wuxi, China), Yanliang Zou (ShanghaiTech University, China; National Supercomputing Center in Wuxi, China), Weiguo Liu (Shandong University, China; National Supercomputing Center in Wuxi, China), and Wei Xue (Tsinghua University, China; National Supercomputing Center in Wuxi, China)</i>	
The Fast and Scalable MPI Application Launch of the Tianhe HPC System .....	1305
<i>Yiqin Dai (National University of Defense Technology, China), Yong Dong (National University of Defense Technology, China), Min Xie (National University of Defense Technology, China), Kai Lu (National University of Defense Technology, China), Ruibo Wang (National University of Defense Technology, China), Mingtian Shao (National University of Defense Technology, China), and Juan Chen (National University of Defense Technology, China)</i>	

HRaft: Adaptive Erasure Coded Data Maintenance for Consensus in Distributed Networks .....	1316
<i>Yulei Jia (Tianjin University of Technology, China; Tianjin Key Laboratory of Intelligence Computing and Novel Software Technology, China), Guangping Xu (Tianjin University of Technology, China; Tianjin Key Laboratory of Intelligence Computing and Novel Software Technology, China), Chi Wan Sung (City University of Hong Kong, China), Salwa Mostafa (City University of Hong Kong, China), and Yulei Wu (University of Exeter, United Kingdom)</i>	

## Session 30: Arch-3 I/O Optimization

CSC: Collaborative System Configuration for I/O-Intensive Applications in Multi-Tenant Clouds .....	1327
<i>Haowei Huang (Shanghai Jiao Tong University, China), Pu Pang (Shanghai Jiao Tong University, China), Quan Chen (Shanghai Jiao Tong University, China), Jieru Zhao (Shanghai Jiao Tong University, China), Wenli Zheng (Shanghai Jiao Tong University, China), and Minyi Guo (Shanghai Jiao Tong University, China)</i>	
Archpipe: Fast and Flexible Pipelined Erasure-Coded Archival Scheme for Heterogeneous Networks .....	1338
<i>Bin Xu (Huazhong University of Sci.&amp; Tech, China), Jianzhong Huang (Huazhong University of Sci.&amp; Tech, China), Xiao Qin (Auburn University, USA), Qiang Cao (Huazhong University of Sci.&amp; Tech, China), Yuanyuan Dong (Alibaba Group, China), and Weikang Kong (Alibaba Group, China)</i>	
A Quantitative Study of the Spatiotemporal I/O Burstiness of HPC Application .....	1349
<i>Wenxiang Yang (National University of Defense Technology, China), Xiangke Liao (National University of Defense Technology, China), Dezun Dong (National University of Defense Technology, China), and Jie Yu (China Aerodynamics Research and Development Center, China)</i>	
DeNOVA: Deduplication Extended NOVA File System .....	1360
<i>Hyungjoon Kwon (Sogang University, Seoul, South Korea), Yonghyeon Cho (Sogang University, Seoul, South Korea), Awais Khan (Oak Ridge National Laboratory, USA), Yeohyeon Park (Sogang University, Seoul, South Korea), and Youngjae Kim (Sogang University, Seoul, South Korea)</i>	

## Author Index