

# **2020 IEEE International Parallel and Distributed Processing Symposium (IPDPS 2020)**

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

**Pages 1-577**



**IEEE Catalog Number: CFP20023-POD  
ISBN: 978-1-7281-6877-7**

**Copyright © 2020 by the Institute of Electrical and Electronics Engineers, Inc.  
All Rights Reserved**

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

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

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

IEEE Catalog Number:	CFP20023-POD
ISBN (Print-On-Demand):	978-1-7281-6877-7
ISBN (Online):	978-1-7281-6876-0
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

# 2020 IEEE International Parallel and Distributed Processing Symposium (IPDPS) **IPDPS 2020**

## Table of Contents

Message from the 2020 General Co-Chairs .xix.....	
Message from the 2020 Program Chair .xxii.....	
Message from the Steering Chair .xxiii.....	
IPDPS 2020 Technical Program .xxiv.....	
IPDPS 2020 Organization .xxvi.....	
IPDPS 2020 Reviewers .xxxv.....	

### Session 1: Communication & NoCs

DozzNoC: Reducing Static and Dynamic Energy in NoCs with Low-Latency Voltage Regulators using Machine Learning .1.....	
<i>Mark Clark (Ohio University), Yingping Chen (University of Texas at Dallas), Avinash Karanth (Ohio University), Brian Ma (University of Texas at Dallas), and Ahmed Louri (George Washington University)</i>	
Nexus: An Interconnect for Heterogeneous System-In-Package Architectures .12.....	
<i>Vidushi Goyal (University of Michigan), Xiaowei Wang (University of Michigan), Valeria Bertacco (University of Michigan), and Reetuparna Das (University of Michigan)</i>	
Accelerated Reply Injection for Removing NoC Bottleneck in GPGPUs .22.....	
<i>Yunfan Li (Oregon State University) and Lizhong Chen (Oregon State University)</i>	
Machine-Agnostic and Communication-Aware Designs for MPI on Emerging Architectures .32.....	
<i>Jahanzeb Maqbool Hashmi (The Ohio State University), Shulei Xu (The Ohio State University), Bharath Ramesh (The Ohio State University), Mohammadreza Bayatpour (The Ohio State University), Hari Subramoni (The Ohio State University), and Dhabaleswar K. (DK) Panda (The Ohio State University)</i>	

### Session 2: Storage & IO

ClusterSR: Cluster-Aware Scattered Repair in Erasure-Coded Storage .42.....	
<i>Zhirong Shen (Xiamen University), Jiwu Shu (Tsinghua University), Zhijie Huang (The University of Texas at Arlington), and Yingxun Fu (North China University of Technology)</i>	

Stitch It Up: Using Progressive Data Storage to Scale Science .52.....	
	<i>Jay Lofstead (Sandia National Laboratories), John Mitchell (Sandia National Laboratories), and Enze Chen (UC Berkeley)</i>
HFetch: Hierarchical Data Prefetching for Scientific Workflows in Multi-tiered Storage Environments .62.....	
	<i>Hariharan Devarajan (Illinois Institute of Technology), Anthony Kougkas (Illinois Institute of Technology), and Xian-He Sun (Illinois Institute of Technology)</i>
CanarIO: Sounding the Alarm on IO-Related Performance Degradation .73.....	
	<i>Michael Wyatt (University of Tennessee Knoxville), Stephen Herbein (Lawrence Livermore National Laboratory), Kathleen Shoga (Lawrence Livermore National Laboratory), Todd Gamblin (Lawrence Livermore National Laboratory), and Michael Taufer (University of Tennessee Knoxville)</i>

## Session 3: Applications

A Study of Graph Analytics for Massive Datasets on Distributed Multi-GPUs .84.....	
	<i>Vishwesh Jatala (The University of Texas at Austin, USA), Roshan Dathathri (The University of Texas at Austin, USA), Gurbinder Gill (The University of Texas at Austin, USA), Loc Hoang (The University of Texas at Austin, USA), V. Krishna Nandivada (IIT Madras, India), and Keshav Pingali (The University of Texas at Austin, USA)</i>
A Highly Efficient Dynamical Core of Atmospheric General Circulation Model Based on Leap-Format .95.....	
	<i>Hang Cao (Institute of Computing Technology, Chinese Academy of Sciences, China), Liang Yuan (Institute of Computing Technology, Chinese Academy of Sciences, China), He Zhang (Institute of Atmospheric Physics, Chinese Academy of Sciences, China), Baodong Wu (Institute of Computing Technology, Chinese Academy of Sciences, China), Shigang Li (ETH Zurich, Switzerland), Pengqi Lu (Institute of Computing Technology, Chinese Academy of Sciences, China), Yunquan Zhang (Institute of Computing Technology, Chinese Academy of Sciences, China), Yongjun Xu (Institute of Computing Technology, Chinese Academy of Sciences, China), and Minghua Zhang (Institute of Atmospheric Physics, Chinese Academy of Sciences, China)</i>
Understanding GPU-Based Lossy Compression for Extreme-Scale Cosmological Simulations .105..	
	<i>Sian Jin (The University of Alabama, USA), Pascal Grosset (Los Alamos National Laboratory, USA), Christopher M. Biwer (Los Alamos National Laboratory, USA), Jesus Pulido (Los Alamos National Laboratory, USA), Jiannan Tian (The University of Alabama, USA), Dingwen Tao (The University of Alabama, USA), and James Ahrens (Los Alamos National Laboratory)</i>
Optimizing High Performance Markov Clustering for Pre-Exascale Architectures .116.....	
	<i>Oguz Selvitopi (Lawrence Berkeley National Laboratory), MD Taufique Hussain (Indiana University, Bloomington), Ariful Azad (Indiana University, Bloomington), and Aydın Buluç (Lawrence Berkeley National Laboratory)</i>

## Session 4: Distributed Algorithms

- Tightening Up the Incentive Ratio for Resource Sharing over the Rings .127.....  
*Yukun Cheng (Suzhou University of Science and Technology, China),  
Xiaotie Deng (Peking University, China), and Yuhao Li (Peking  
University, China)*
- Communication-Efficient String Sorting .137.....  
*Timo Bingmann (Karlsruhe Institute of Technology), Peter Sanders  
(Karlsruhe Institute of Technology), and Matthias Schimek (Karlsruhe  
Institute of Technology)*
- SCSL: Optimizing Matching Algorithms to Improve Real-Time for Content-Based Pub/Sub  
Systems .148.....  
*Tianchen Ding (Shanghai Jiao Tong University, China), Shiyong Qian  
(Shanghai Jiao Tong University, China), Jian Cao (Shanghai Jiao Tong  
University, China), Guangtao Xue (Shanghai Jiao Tong University,  
China), and Minglu Li (Zhejiang Normal University, China)*
- Distributed Graph Realizations .158.....  
*John Augustine (Indian Institute of Technology Madras), Keerti  
Choudhary (Weizmann Institute of Science), Avi Cohen (Weizmann  
Institute of Science), David Peleg (Weizmann Institute of Science),  
Sumathi Sivasubramaniam (Indian Institute of Technology Madras), and  
Suman Sourav (National University of Singapore)*

## Session 5: Reliability and QoS

- Transaction-Based Core Reliability .168.....  
*Sang Wook Stephen Do (Futurewei Technologie, US) and Michel Dubois  
(University of Southern California)*
- Understanding the Interplay between Hardware Errors and User Job Characteristics on the  
Titan Supercomputer .180.....  
*Seung-Hwan Lim (Oak Ridge National Laboratory), Ross G. Miller (Oak  
Ridge National Laboratory), and Sudharshan S. Vazhkudai (Oak Ridge  
National Laboratory)*
- EC-Fusion: An Efficient Hybrid Erasure Coding Framework to Improve Both Application and  
Recovery Performance in Cloud Storage Systems .191.....  
*Han Qiu (Shanghai Jiao Tong University, China), Chentao Wu (Shanghai  
Jiao Tong University, China), Jie Li (Shanghai Jiao Tong University,  
China), Minyi Guo (Shanghai Jiao Tong University, China), Tong Liu  
(Temple University, USA), Xubin He (Temple University, USA), Yuanyuan  
Dong (Alibaba Group, China), and Yafei Zhao (Alibaba Group, China)*

## Session 6: Learning Algorithms

- Learning an Effective Charging Scheme for Mobile Devices .202.....  
*Tang Liu (Sichuan Normal University, China), Baijun Wu (University of Louisiana at Lafayette, USA), Wenzheng Xu (Sichuan University, China), Xianbo Cao (Sichuan University, China), Jian Peng (Sichuan University, China), and Hongyi Wu (Dominion University, USA)*
- Optimize Scheduling of Federated Learning on Battery-Powered Mobile Devices .212.....  
*Cong Wang (Old Dominion University), Xin Wei (Old Dominion University), and Pengzhan Zhou (Stony Brook University)*
- Harnessing Deep Learning via a Single Building Block .222.....  
*Evangelos Georganas (Intel), Kunal Banerjee (Intel), Dhiraj Kalamkar (Intel), Sasikanth Avancha (Intel), Anand Venkat (Intel), Michael Anderson (Intel), Greg Henry (Intel), Hans Pabst (Intel), and Alexander Heinecke (Intel)*
- Experience-Driven Computational Resource Allocation of Federated Learning by Deep Reinforcement Learning .234.....  
*Yufeng Zhan (The Hong Kong Polytechnic University, China), Peng Li (The University of Aizu, Japan), and Song Guo (The Hong Kong Polytechnic University, China)*

## Session 7: Data Analysis and Management

- An Active Learning Method for Empirical Modeling in Performance Tuning .244.....  
*Jiepeng Zhang (University of Science and Technology of China, China), Jingwei Sun (University of Science and Technology of China, China), Wenju Zhou (University of Science and Technology of China, China), and Guangzhong Sun (University of Science and Technology of China, China)*
- DASSA: Parallel DAS Data Storage and Analysis for Subsurface Event Detection .254.....  
*Bin Dong (Lawrence Berkeley National Laboratory, Berkeley, CA, USA), Veronica Rodriguez Tribaldos (Lawrence Berkeley National Laboratory, Berkeley, CA, USA), Xin Xing (Georgia Institute of Technology, Atlanta, GA, USA), Suren Byna (Lawrence Berkeley National Laboratory, Berkeley, CA, USA), Jonathan Ajo-Franklin (Rice University, Houston, TX, USA), and Kesheng Wu (Lawrence Berkeley National Laboratory, Berkeley, CA, USA)*
- Scaling of Union of Intersections for Inference of Granger Causal Networks from Observational Data .264.....  
*Mahesh Balasubramanian (Arizona State University, USA), Trevor Ruiz (Oregon State University, USA), Brandon Cook (Lawrence Berkeley National Laboratory, USA), Mr Prabhat (Lawrence Berkeley National Laboratory, USA), Sharmodeep Bhattacharyya (Oregon State University, USA), Aviral Shrivastava (Arizona State University, USA), and Kristofer Bouchard (Lawrence Berkeley National Laboratory, USA)*

GPU-Based Static Data-Flow Analysis for Fast and Scalable Android App Vetting .274.....	
	<i>Xiaodong Yu (Argonne National Laboratory, USA), Fengguo Wei (University of South Florida, USA), Xinming Ou (University of South Florida, USA), Michela Becchi (North Carolina State University, USA), Tekin Bicer (Argonne National Laboratory, USA), and Danfeng (Daphne) Yao (Virginia Tech, USA)</i>

## Session 8: Edge Computing

Robust Server Placement for Edge Computing .285.....	
	<i>Dongyu Lu (Shanghai Jiao Tong University, China), Yuben Qu (Shanghai Jiao Tong University, China), Fan Wu (Shanghai Jiao Tong University, China), Haipeng Dai (Nanjing University, China), Chao Dong (Nanjing University of Aeronautics and Astronautics, China), and Guihai Chen (Shanghai Jiao Tong University, China)</i>
EdgeIso: Effective Performance Isolation for Edge Devices .295.....	
	<i>Yoonsung Nam (Seoul National University, South Korea), Yongjun Choi (Seoul National University, South Korea), Byeonghun Yoo (Seoul National University, South Korea), Yongseok Son (Chung-Ang University, South Korea), and Hyeonsang Eom (Seoul National University, South Korea)</i>
Busy-Time Scheduling on Heterogeneous Machines .306.....	
	<i>Runtian Ren (Nanyang Technological University, Singapore) and Xueyan Tang (Nanyang Technological University, Singapore)</i>
Scheduling Malleable Jobs under Topological Constraints .316.....	
	<i>Evrpidis Bampis (Sorbonne Université, CNRS, LIP6, F-75005 Paris, France), Konstantinos Dogeas (Sorbonne Université, CNRS, LIP6, F-75005 Paris, France), Alexander Kononov (Sobolev Institute of Mathematics, Novosibirsk State University, Novosibirsk, Russia), Giorgio Lucarelli (University of Lorraine, LCOMS, Metz, France), and Fanny Pascual (Sorbonne Université, CNRS, LIP6, F-75005 Paris, France)</i>

## Best Papers

XSP: Across-Stack Profiling and Analysis of Machine Learning Models on GPUs .326.....	
	<i>Cheng Li (University of Illinois Urbana-Champaign), Abdul Dakkak (University of Illinois Urbana-Champaign), Jinjun Xiong (IBM T. J. Watson Research Center), Wei Wei (Alibaba Group), Lingjie Xu (Alibaba Group), and Wen-mei Hwu (University of Illinois Urbana-Champaign)</i>
Exploring the Binary Precision Capabilities of Tensor Cores for Epistasis Detection .338.....	
	<i>Ricardo Nobre (INESC-ID, Instituto Superior Técnico, Universidade de Lisboa, Portugal), Aleksandar Ilic (INESC-ID, Instituto Superior Técnico, Universidade de Lisboa, Portugal), Sergio Santander-Jiménez (INESC-ID, Instituto Superior Técnico, Universidade de Lisboa, Portugal), and Leonel Sousa (INESC-ID, Instituto Superior Técnico, Universidade de Lisboa, Portugal)</i>

- Understanding and Improving Persistent Transactions on Optane (TM) DC Memory .348.....  
*Pantea Zardoshti (Lehigh University), Michael Spear (Lehigh University), Aida Vosoughi (Oracle Corp), and Garret Swart (Oracle Corp)*
- CycLedger: A Scalable and Secure Parallel Protocol for Distributed Ledger via Sharding .358.....  
*Mengqian Zhang (Shanghai Jiao Tong University), Jichen Li (Peking University), Zhaohua Chen (Peking University), Hongyin Chen (Peking University), and Xiaotie Deng (Peking University)*

## Session 9: Cloud Technology

- Mitigating Large Response Time Fluctuations through Fast Concurrency Adapting in Clouds .368.  
*Jianshu Liu (Louisiana State University–Baton Rouge), Shungeng Zhang (Louisiana State University–Baton Rouge), Qingyang Wang (Louisiana State University–Baton Rouge), and Jinpeng Wei (University of North Carolina–Charlotte)*
- DAG-Aware Joint Task Scheduling and Cache Management in Spark Clusters .378.....  
*Yinggen Xu (Tongji University, China), Liu Liu (Tongji University, China), and Zhijun Ding (Tongji University, China)*
- Solving the Container Explosion Problem for Distributed High Throughput Computing .388.....  
*Tim Shaffer (University of Notre Dame), Nicholas Hazekamp (University of Notre Dame), Jakob Blomer (European Laboratory for Particle Physics (CERN)), and Douglas Thain (University of Notre Dame)*
- Amoeba: QoS-Awareness and Reduced Resource Usage of Microservices with Serverless Computing .399.....  
*Zijun Li (Shanghai Jiao Tong University, China), Quan Chen (Shanghai Institute for Advanced Communication and Data Science, Shanghai Jiao Tong University, China), Shuai Xue (Shanghai Jiao Tong University, China), Tao Ma (Alibaba Group, China), Yong Yang (Alibaba Group, China), Zhuo Song (Alibaba Group, China), and Minyi Guo (Shanghai Institute for Advanced Communication and Data Science, Shanghai Jiao Tong University, China)*

## Session 10: Machine Learning

- Efficient I/O for Neural Network Training with Compressed Data .409.....  
*Zhao Zhang (Texas Advanced Computing Center), Lei Huang (Texas Advanced Computing Center), J. Gregory Pauloski (University of Texas at Austin), and Ian T. Foster (University of Chicago and Argonne National Laboratory)*
- Not All Explorations are Equal: Harnessing Heterogeneous Profiling Cost for Efficient MLaaS Training .419.....  
*Jun Yi (University of Nevada Reno), Chengliang Zhang (Hong Kong University of Science and Technology), Wei Wang (Hong Kong University of Science and Technology), Cheng Li (University of Science and Technology of China), and Feng Yan (University of Nevada Reno)*



ASYNC: A Cloud Engine with Asynchrony and History for Distributed Machine Learning .429.....  
*Saeed Soori (University of Toronto), Bugra Can (Rutgers University),  
Mert Gurbuzbalaban (Rutgers University), and Maryam Mehri Dehnavi  
(University of Toronto)*

Benanza: Automatic uBenchmark Generation to Compute "Lower-Bound" Latency and Inform  
Optimizations of Deep Learning Models on GPUs .440.....  
*Cheng Li (University of Illinois Urbana-Champaign), Abdul Dakkak  
(University of Illinois Urbana-Champaign), Jinjun Xiong (IBM T. J.  
Watson Research Center), and Wen-mei Hwu (University of Illinois  
Urbana-Champaign)*

## Session 11: GPUs

Adaptive Page Migration for Irregular Data-Intensive Applications under GPU Memory  
Oversubscription .451.....  
*Debashis Ganguly (University of Pittsburgh), Ziyu Zhang (University of  
Pittsburgh), Jun Yang (Electrical and Computer Engineering, University  
of Pittsburgh), and Rami Melhem (University of Pittsburgh)*

LOGAN: High-Performance GPU-Based X-Drop Long-Read Alignment .462.....  
*Alberto Zeni (Politecnico Di Milano, Italy), Giulia Guidi (University  
of California at Berkeley, USA), Marquita Ellis (University of  
California at Berkeley, USA), Nan Ding (Lawrence Berkeley National  
Laboratory, USA), Marco Santambrogio (Politecnico Di Milano, Italy),  
Steven Hofmeyr (Lawrence Berkeley National Laboratory, USA), Aydın  
Buluç (University of California at Berkeley, USA), Leonid Oliker  
(Lawrence Berkeley National Laboratory, USA), and Katherine Yelick  
(University of California at Berkeley, USA)*

Coordinated Page Prefetch and Eviction for Memory Oversubscription Management in GPUs .472  
*Qi Yu (National University of Defense Technology, China), Bruce  
Childers (University of Pittsburgh, United States), Libo Huang  
(National University of Defense Technology, China), Cheng Qian  
(National University of Defense Technology, China), Hui Guo (National  
University of Defense Technology, China), and Zhiying Wang (National  
University of Defense Technology, China)*

A Study of Single and Multi-device Synchronization Methods in Nvidia GPUs .483.....  
*Lingqi Zhang (Tokyo Institute of Technology), Mohamed Wahib (National  
Institute of Advanced Industrial Science and Technology), Haoyu Zhang  
(miHoYo Inc), and Satoshi Matsuoka (RIKEN Center for Computational  
Science)*

## Session 12: Applications

- DPF-ECC: Accelerating Elliptic Curve Cryptography with Floating-Point Computing Power of GPUs .494.....  
*Lili Gao (Institute of Information Engineering, Chinese Academy of Sciences), Fangyu Zheng (Institute of Information Engineering, Chinese Academy of Sciences), Niall Emmart (University of Massachusetts), Jiankuo Dong (Institute of Information Engineering, Chinese Academy of Sciences), Jingqiang Lin (Institute of Information Engineering, Chinese Academy of Sciences), and Charles Weems (University of Massachusetts)*
- Scalability Challenges of an Industrial Implicit Finite Element Code .505.....  
*François-Henry Rouet (Livermore Software Technology, An ANSYS Company), Cleve Ashcraft (Livermore Software Technology, An ANSYS Company), Jef Dawson (Cray, Inc.), Roger Grimes (Livermore Software Technology, An ANSYS Company), Erman Guleryuz (National Center for Supercomputing Applications, University of Illinois), Seid Koric (National Center for Supercomputing Applications, University of Illinois), Robert F. Lucas (Livermore Software Technology, An ANSYS Company), James S. Ong (Rolls-Royce), Todd A. Simons (Rolls-Royce), and Ting-Ting Zhu (Cray, Inc.)*
- ETH: An Architecture for Exploring the Design Space of In-Situ Scientific Visualization .515.....  
*Gregory Abram (Texas Advanced Computing Center (TACC)), Vignesh Adhinarayanan (Virginia Tech), Wu-chun Feng (Virginia Tech), David Rogers (Los Alamos National Laboratory (LANL)), and James Ahrens (Los Alamos National Laboratory (LANL))*
- Scaling Betweenness Approximation to Billions of Edges by MPI-Based Adaptive Sampling .527...  
*Alexander van der Grinten (Humboldt-Universität zu Berlin, Germany) and Henning Meyerhenke (Humboldt-Universität zu Berlin, Germany)*

## Session 13: Data Management

- Improved Intermediate Data Management for MapReduce Frameworks .536.....  
*Haoyu Wang (University of Virginia), Haiying Shen (University of Virginia), Charles Reiss (University of Virginia), Arnim Jain (University of Virginia), and Yunqiao Zhang (Facebook)*
- Bandwidth-Aware Page Placement in NUMA .546.....  
*David Gureya (INESC-ID, Instituto Superior Técnico, University of Lisbon, Portugal), João Neto (INESC-ID, Instituto Superior Técnico, University of Lisbon, Portugal), Reza Karimi (Emory University, USA), João Barreto (INESC-ID, Instituto Superior Técnico, University of Lisbon, Portugal), Pramod Bhatotia (University of Edinburgh, United Kingdom), Vivien Quema (Grenoble INP/ENSIMAG, France), Rodrigo Rodrigues (INESC-ID, Instituto Superior Técnico, University of Lisbon, Portugal), Paolo Romano (INESC-ID, Instituto Superior Técnico, University of Lisbon, Portugal), and Vladimir Vlassov (KTH Royal Institute of Technology, Sweden)*

HCompress: Hierarchical Data Compression for Multi-tiered Storage Environments .557.....	
<i>Hariharan Devarajan (Illinois Institute of Technology), Anthony Kougkas (Illinois Institute of Technology), Luke Logan (Illinois Institute of Technology), and Xian-He Sun (Illinois Institute of Technology)</i>	
FRaZ: A Generic High-Fidelity Fixed-Ratio Lossy Compression Framework for Scientific Floating-Point Data .567.....	
<i>Rober Underwood (Clemson University and Argonne National Laboratory), Sheng Di (Argonne National Laboratory), Jon Calhoun (Clemson University), and Franck Cappello (Argonne National Laboratory)</i>	

## Session 14: Storage & Caching

DELTA: Distributed Locality-Aware Cache Partitioning for Tile-Based Chip Multiprocessors .578..	
<i>Nadja Holtryd (Chalmers University of Technology, Sweden), Madhavan Manivannan (Chalmers University of Technology, Sweden), Per Stenström (Chalmers University of Technology, Sweden), and Miquel Pericàs (Chalmers University of Technology, Sweden)</i>	
Coordinated Management of Processor Configuration and Cache Partitioning to Optimize Energy under QoS Constraints .590.....	
<i>Mehrzad Nejat (Chalmers University of Technology), Madhavan Manivannan (Chalmers University of Technology), Miquel Pericàs (Chalmers University of Technology), and Per Stenstrom (Chalmers University of Technology)</i>	
StragglerHelper: Alleviating Straggling in Computing Clusters via Sharing Memory Access Patterns .602.....	
<i>Wenjie Liu (Temple University), Ping Huang (Temple University), and Xubin He (Temple University)</i>	

## Session 15: Numerics

Evaluating the Numerical Stability of Posit Arithmetic .612.....	
<i>Nicholas Buoncristiani (Lawrence Berkeley National Laboratory and University of California), Sanjana Shah (Lawrence Berkeley National Laboratory and University of California), David Donofrio (Lawrence Berkeley National Laboratory Berkeley), and John Shalf (Lawrence Berkeley National Laboratory)</i>	
Varity: Quantifying Floating-Point Variations in HPC Systems Through Randomized Testing .622..	
<i>Ignacio Laguna (Lawrence Livermore National Laboratory)</i>	
Demystifying Tensor Cores to Optimize Half-Precision Matrix Multiply .634.....	
<i>Da Yan (Hong Kong University of Science and Technology), Wei Wang (Hong Kong University of Science and Technology), and Xiaowen Chu (Hong Kong Baptist University)</i>	

## Session 16: IoT and Consensus

- Data Collection of IoT Devices Using an Energy-Constrained UAV .644.....  
*Yuchen Li (Australian National University, Australia), Weifa Liang (Australian National University, Australia), Wenzheng Xu (Sichuan University, China), and Xiaohua Jia (City University of Hong Kong, China)*
- Argus: Multi-level Service Visibility Scoping for Internet-of-Things in Enterprise Environments .654.....  
*Qian Zhou (Stony Brook University, USA), Omkant Pandey (Stony Brook University, USA), and Fan Ye (Stony Brook University, USA)*
- G-PBFT: A Location-Based and Scalable Consensus Protocol for IoT-Blockchain Applications .664.  
*Laphou Lao (The Hong Kong Polytechnic University, Hong Kong), Xiaohai Dai (The Hong Kong Polytechnic University, Hong Kong, Huazhong University of Science and Technology, China), Bin Xiao (The Hong Kong Polytechnic University, Hong Kong), and Songtao Guo (Chongqing University, China)*
- Byzantine Generalized Lattice Agreement .674.....  
*Giuseppe Antonio Di Luna (Sapienza, University of Rome, Italy), Emmanuelle Anceaume (CNRS, Univ Rennes, Inria, IRISA, France), and Leonardo Querzoni (Sapienza University of Rome, Italy)*

## Session 17: Graph Processing & Coding

- A Heterogeneous PIM Hardware-Software Co-Design for Energy-Efficient Graph Processing .684.  
*Yu Huang (Huazhong University of Science and Technology), Long Zheng (Huazhong University of Science and Technology), Pengcheng Yao (Huazhong University of Science and Technology), Jieshan Zhao (Huazhong University of Science and Technology), Xiaofei Liao (Huazhong University of Science and Technology), Hai Jin (Huazhong University of Science and Technology), and Jingling Xue (The University of New South Wales)*
- Spara: An Energy-Efficient ReRAM-Based Accelerator for Sparse Graph Analytics Applications.696  
*Long Zheng (Huazhong University of Science and Technology), Jieshan Zhao (Huazhong University of Science and Technology), Yu Huang (Huazhong University of Science and Technology), Qinggang Wang (Huazhong University of Science and Technology), Zhen Zeng (Huazhong University of Science and Technology), Jingling Xue (The University of New South Wales), Xiaofei Liao (Huazhong University of Science and Technology), and Hai Jin (Huazhong University of Science and Technology)*
- Optimal Encoding and Decoding Algorithms for the RAID-6 Liberation Codes .708.....  
*Zhijie Huang (The University of Texas at Arlington, USA), Hong Jiang (The University of Texas at Arlington, USA), Zhirong Shen (Xiamen University, China), Hao Che (The University of Texas at Arlington, USA), Nong Xiao (Sun Yat-Sen University, China), and Ning Li (The University of Texas at Arlington, USA)*

Sturgeon: Preference-Aware Co-Location for Improving Utilization of Power Constrained Computers .718.....  
*Pu Pang (Shanghai Jiao Tong University, China), Quan Chen (Shanghai Jiao Tong University, China), Deze Zeng (China University of Geoscience, China), Chao Li (Shanghai Jiao Tong University), Jingwen Leng (Shanghai Jiao Tong University, China), Wenli Zheng (Shanghai Jiao Tong University, China), and Minyi Guo (Shanghai Jiao Tong University, China)*

## Session 18: Parallel Algorithms

A High-Throughput Solver for Marginalized Graph Kernels on GPU .728.....  
*Yu-Hang Tang (Lawrence Berkeley National Laboratory), Oguz Selvitopi (Lawrence Berkeley National Laboratory), Doru Popovici (Lawrence Berkeley National Laboratory), and Aydin Buluc (Lawrence Berkeley National Laboratory)*

Dynamic Graphs on the GPU .739.....  
*Muhammad A. Awad (UC Davis), Saman Ashkiani (UC Davis), Serban D. Porumbescu (UC Davis), and John D. Owens (UC Davis)*

Accelerating Parallel Hierarchical Matrix-Vector Products via Data-Driven Sampling .749.....  
*Lucas Erlandson (Georgia Institute of Technology, United States of America), Difeng Cai (Emory University, United States of America), Yuanzhe Xi (Emory University, United States of America), and Edmond Chow (Georgia Institute of Technology, United States of America)*

NC Algorithms for Popular Matchings in One-Sided Preference Systems and Related Problems .759  
*Changyong Hu (University of Texas at Austin) and Vijay Garg (University of Texas at Austin)*

## Session 19: Performance, Power, and Energy

Smartly Handling Renewable Energy Instability in Supporting A Cloud Datacenter .769.....  
*Jiechao Gao (University of Virginia), Haoyu Wang (University of Virginia), and Haiying Shen (University of Virginia)*

A Self-Optimized Generic Workload Prediction Framework for Cloud Computing .779.....  
*Vinodh Kumaran Jayakumar (The University of Texas at San Antonio), Jaewoo Lee (The University of Georgia), In Kee Kim (The University of Georgia), and Wei Wang (The University of Texas at San Antonio)*

SeeSAw: Optimizing Performance of In-Situ Analytics Applications under Power Constraints .789.  
*Ioana Marincic (University of Chicago), Venkatram Vishwanath (Argonne National Laboratory), and Henry Hoffmann (University of Chicago)*

## Session 20: Resource Management

- What Does Power Consumption Behavior of HPC Jobs Reveal? .799.....  
*Tirthak Patel (Northeastern University), Adam Wagenhäuser (Friedrich-Alexander University Erlangen-Nürnberg (FAU)), Christopher Eibel (Friedrich-Alexander University Erlangen-Nürnberg (FAU)), Timo Hönig (Friedrich-Alexander University Erlangen-Nürnberg (FAU)), Thomas Zeiser (Friedrich-Alexander University Erlangen-Nürnberg (FAU)), and Devesh Tiwari (Northeastern University)*
- Efficient Parallel and Adaptive Partitioning for Load-Balancing in Spatial Join .810.....  
*Jie Yang (Marquette University) and Satish Puri (Marquette University)*
- Union: An Automatic Workload Manager for Accelerating Network Simulation .821.....  
*Xin Wang (Illinois Institute of Technology), Misbah Mubarak (Argonne National Laboratory), Yao Kang (Illinois Institute of Technology), Robert B. Ross (Argonne National Laboratory), and Zhiling Lan (Illinois Institute of Technology)*
- Auto-Tuning Parameter Choices in HPC Applications using Bayesian Optimization .831.....  
*Harshitha Menon (Lawrence Livermore National Laboratory), Abhinav Bhatele (University of Maryland), and Todd Gamblin (Lawrence Livermore National Laboratory)*

## Session 21: Runtime Systems

- Inter-Job Scheduling of High-Throughput Material Screening Applications .841.....  
*Zhihui Du (Tsinghua University), Xining Hui (Tsinghua University), Yurui Wang (Tsinghua Univ), Jun Jiang (Beijing Computing Center), Jason Liu (Florida International University), Baokun Lu (Tsinghua University), and Chongyu Wang (Tsinghua University)*
- Reservation and Checkpointing Strategies for Stochastic Jobs .853.....  
*Ana Gainaru (Vanderbilt University, Nashville, TN, USA), Brice Goglin (Inria, LaBRI, Univ. Bordeaux, Talence, France), Valentin Honoré (Inria, LaBRI, Univ. Bordeaux, Talence, France), Guillaume Pallez (Inria, LaBRI, Univ. Bordeaux, Talence, France), Padma Raghavan (Vanderbilt University, Nashville, TN, USA), Yves Robert (Laboratoire LIP, ENS Lyon, France & University of Tennessee Knoxville, USA), and Hongyang Sun (Vanderbilt University, Nashville, TN, USA)*
- A Scheduling Approach to Incremental Maintenance of Datalog Programs .864.....  
*Shikha Singh (Williams College), Sergey Madaminov (Stony Brook University), Michael A. Bender (Stony Brook University), Michael Ferdman (Stony Brook University), Ryan Johnson (Amazon, Inc.), Benjamin Moseley (Carnegie Mellon University), Hung Ngo (Relational AI), Dung Nguyen (Infor, Inc.), Soeren Olesen (Infor, Inc.), Kurt Stirewalt (Relational AI), and Geoffrey Washburn (Infor, Inc.)*
- Dynamic Scheduling in Distributed Transactional Memory .874.....  
*Costas Busch (Louisiana State University, USA), Maurice Herlihy (Brown University, USA), Miroslav Popovic (University of Novi Sad, Serbia), and Gokarna Sharma (Kent State University, USA)*

## Session 22: Performance Analysis

- Learning Cost-Effective Sampling Strategies for Empirical Performance Modeling .884.....  
*Marcus Ritter (Technical University of Darmstadt), Alexandru Calotoiu (Technical University of Darmstadt), Sebastian Rinke (Technical University of Darmstadt), Thorsten Reimann (Technical University of Darmstadt), Torsten Hoefler (ETH Zürich), and Felix Wolf (Technical University of Darmstadt)*
- The Case of Performance Variability on Dragonfly-Based Systems .896.....  
*Abhinav Bhatele (University of Maryland, College Park), Jayaraman J. Thiagarajan (Lawrence Livermore National Laboratory), Taylor Groves (Lawrence Berkeley National Laboratory), Rushil Anirudh (Lawrence Livermore National Laboratory), Staci A. Smith (The University of Arizona), Brandon Cook (Lawrence Berkeley National Laboratory), and David K. Lowenthal (The University of Arizona)*
- Predicting and Comparing the Performance of Array Management Libraries .906.....  
*Donghe Kang (The Ohio State University, US), Oliver Rübel (Lawrence Berkeley National Laboratory, US), Suren Byna (Lawrence Berkeley National Laboratory, US), and Spyros Blanas (The Ohio State University, US)*
- Demystifying the Performance of HPC Scientific Applications on NVM-Based Memory Systems .916  
*Ioy Peng (Lawrence Livermore National Laboratory), Kai Wu (University of California, Merced), Jie Ren (University of California, Merced), Dong Li (University of California, Merced), and Maya Gokhale (Lawrence Livermore National Laboratory)*

## Session 23: Communication

- Packet-in Request Redirection for Minimizing Control Plane Response Time .926.....  
*Rui Xia (Nanjing University, China), Haipeng Dai (Nanjing University), Jiaqi Zheng (Nanjing University, China), Hong Xu (City University of Hong Kong, China), Meng Li (Nanjing University, China), and Guihai Chen (Nanjing University, China)*
- PCGCN: Partition-Centric Processing for Accelerating Graph Convolutional Network .936.....  
*Chao Tian (Peking University), Lingxiao Ma (Peking University), Zhi Yang (Peking University), and Yafei Dai (Peking University)*
- ConMidbox: Consolidated Middleboxes Selection and Routing in SDN/NFV-Enabled Networks .946  
*Guiyan Liu (Southwest University), Songtao Guo (Chongqing University), Pan Li (Southwest University), and Liang Liu (Southwest University)*
- Scalable and Memory-Efficient Kernel Ridge Regression .956.....  
*Gustavo Chavez (Lawrence Berkeley National Laboratory), Yang Liu (Lawrence Berkeley National Laboratory), Pieter Ghysels (Lawrence Berkeley National Laboratory), Xiaoye Sherry Li (Lawrence Berkeley National Laboratory), and Elizaveta Rebrova (University of California, Los Angeles)*

## Session 24: Storage

- SSDKeeper: Self-Adapting Channel Allocation to Improve the Performance of SSD Devices .966....  
*Renping Liu (Chongqing University), Xianzhang Chen (Chongqing University), Yujuan Tan (Chongqing University), Runyu Zhang (Chongqing University), Liang Liang (Chongqing University), and Duo Liu (Chongqing University)*
- FlashKey: A High-Performance Flash Friendly Key-Value Store .976.....  
*Madhurima Ray (Temple University, USA), Krishna Kant (Temple University, USA), Peng Li (Intel Corporation, USA), and Sanjeev Trika (Intel Corporation, USA)*
- Pacon: Improving Scalability and Efficiency of Metadata Service through Partial Consistency .986.....  
*Yubo Liu (Sun Yat-sen University, China), Yutong Lu (Sun Yat-sen University, China), Zhiguang Chen (Sun Yat-sen University, China), and Ming Zhao (Arizona State University, United States of America)*

## Session 25: Program Analysis and Runtime Library

- XPlacer: Automatic Analysis of Data Access Patterns on Heterogeneous CPU/GPU Systems .997....  
*Peter Pirkelbauer (Lawrence Livermore National Laboratory), Pei-Hung Lin (Lawrence Livermore National Laboratory), Tristan Vanderbruggen (Lawrence Livermore National Laboratory), and Chunhua Liao (Lawrence Livermore National Laboratory)*
- Improving Transactional Code Generation via Variable Annotation and Barrier Elision .1008.....  
*João Paulo Labegalini de Carvalho (University of Campinas (UNICAMP), Institute of Computing, Campinas), Bruno Chinelato Honorio (University of Campinas (UNICAMP), Institute of Computing, Campinas), Alexandro Baldassin (São Paulo State University (Unesp), Institute of Geosciences and Exact Sciences, Rio Claro), and Guido Araujo (University of Campinas (UNICAMP), Institute of Computing, Campinas)*
- Evaluating Thread Coarsening and Low-Cost Synchronization on Intel Xeon Phi .1018.....  
*Hancheng Wu (North Carolina State University) and Michela Becchi (North Carolina State University)*
- AnySeq: A High Performance Sequence Alignment Library Based on Partial Evaluation .1030.....  
*André Müller (Johannes Gutenberg University Mainz), Bertil Schmidt (Johannes Gutenberg University Mainz), Andreas Hildebrandt (Johannes Gutenberg University Mainz), Richard Membarth (DFKI), Matthias Kruse (Saarland University), Roland Leißa (Saarland University), and Sebastian Hack (Saarland University)*

## Session 26: Scheduling

- Analysis of a List Scheduling Algorithm for Task Graphs on Two Types of Resources .1041.....  
*Lionel Eyraud-Dubois (Inria Bordeaux – Sud-Ouest, Université de Bordeaux, France) and Suraj Kumar (Inria Paris, France)*



Optimal Convex Hull Formation on a Grid by Asynchronous Robots with Lights .....	1051
<i>Rory Hector (Louisiana State University), Ramachandran Vaidyanathan (Louisiana State University), Gokarna Sharma (Kent State University), and Jerry Trahan (Louisiana State University)</i>	
On the Complexity of Conditional DAG Scheduling in Multiprocessor Systems .....	1061
<i>Alberto Marchetti-Spaccamela (Sapienza University of Rome, Italy), Nicole Megow (University of Bremen, Germany), Jens Schlöter (University of Bremen, Germany), Martin Skutella (Technical University of Berlin, Germany), and Leen Stougie (CWI and Vrije Universiteit Amsterdam, Netherlands)</i>	
Weaver: Efficient Coflow Scheduling in Heterogeneous Parallel Networks .....	1071
<i>Xin Sunny Huang (Rice University), Yiting Xia (Facebook, Inc.), and T. S. Eugene Ng (Rice University)</i>	

## Session 27: Fault Tolerance

Fault-Tolerant Containers Using NiLiCon .....	1082
<i>Diyu Zhou (Computer Science Department, UCLA) and Yuval Tamir (Computer Science Department, UCLA)</i>	
Aarohi: Making Real-Time Node Failure Prediction Feasible .....	1092
<i>Anwesha Das (North Carolina State University), Frank Mueller (North Carolina State University), and Barry Rountree (Lawrence Livermore National Laboratory)</i>	
FP4S: Fragment-Based Parallel State Recovery for Stateful Stream Applications .....	1102
<i>Pinchao Liu (Florida International University), Hailu Xu (Florida International University), Dilma Da Silva (Texas A&amp;M University), Qingyang Wang (Louisiana State University), Sarker Tanzir Ahmed (Texas A&amp;M University), and Liting Hu (Florida International University)</i>	

## Session 28: Multidisciplinary

Implementation and Evaluation of a Hardware Decentralized Synchronization Lock for MPSoCs ..	1112
<i>Maxime France-Pillois (CEA, LETI, MINATEC Campus - Univ. Grenoble Alpes), Jérôme Martin (CEA, LETI, MINATEC Campus - Univ. Grenoble Alpes), and Frédéric Rousseau (Univ. Grenoble Alpes, CNRS, Grenoble INP, TIMA)</i>	
Communication-Efficient Jaccard Similarity for High-Performance Distributed Genome Comparisons .....	1122
<i>Maciej Besta (ETH Zurich, Switzerland), Raghavendra Kanakagiri (IIT Tirupati, India), Harun Mustafa (ETH Zurich, Switzerland), Mikhail Karasikov (ETH Zurich, Switzerland), Gunnar Ratsch (ETH Zurich, Switzerland), Torsten Hoefler (ETH Zurich, Switzerland), and Edgar Solomonik (University of Illinois at Urbana-Champaign, United States of America)</i>	
Engineering Worst-Case Inputs for Pairwise Merge Sort on GPUs .....	1133
<i>Kyle Berney (University of Hawaii at Manoa) and Nodari Sitchinava (University of Hawaii at Manoa)</i>	
The Impossibility of Fast Transactions .....	1143
<i>Karolos Antoniadis (EPFL), Diego Didona (IBM Zurich), Rachid Guerraoui (EPFL), and Willy Zwaenepoel (University of Sydney)</i>	

## Author Index