

2024 IEEE 24th International Symposium on Cluster, Cloud and Internet Computing (CCGrid 2024)

**Philadelphia, Pennsylvania, USA
6-9 May 2024**



**IEEE Catalog Number: CFP24276-POD
ISBN: 979-8-3503-9567-9**

**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:	CFP24276-POD
ISBN (Print-On-Demand):	979-8-3503-9567-9
ISBN (Online):	979-8-3503-9566-2
ISSN:	2376-4414

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 24th International Symposium on Cluster, Cloud and Internet Computing (CCGrid) **CCGrid 2024**

Table of Contents

Message from the General Chairs	xvi
Message from the Program Chairs	xvii
Organizing Committee	xix
Program Committee	xx
Steering Committee	xxvii

Hardware Systems and Networking

A Bandwidth-Optimal All-to-All Communication in Two-Dimensional Fully Connected Network	1
<i>Kien Trung Pham (SOKENDAI; National Institute of Informatics, Japan), Truong Thao Nguyen (National Institute of Advanced Industrial Science and Technology), and Michihiro Koibuchi (National Institute of Advanced Industrial Science and Technology, Japan)</i>	
A Hybrid Solution to Provide End-to-End Flow Control and Congestion Management in High-Performance Interconnection Networks	8
<i>Alberto Merino (Universidad de Castilla-La Mancha, Spain), Jesus Escudero-Sahuquillo (Universidad de Castilla-La Mancha, Spain), Pedro Javier Garcia (Universidad de Castilla-La Mancha, Spain), Francisco J. Quiles (Universidad de Castilla-La Mancha, Spain), Fei Chen (Huawei Technologies Co., Ltd., China), Yunping Lyu (Huawei Technologies Co., Ltd., China), Long Yan (Huawei Technologies Co., Ltd., China), and Jose Duato (Royal Spanish Academy of Sciences, Spain)</i>	
AppSteer: Framework for Improving Multicore Scalability of Network Functions via Application-Aware Packet Steering	18
<i>Ashwin Kumar (IIT Bombay, India), Rajneesh Katkam (IIT Bombay, India), Pranav Chaudhary (IIT Bombay, India), Priyanka Naik (IBM Research, India), and Mythili Vutukuru (IIT Bombay, India)</i>	

CMB: A Configurable Messaging Benchmark to Explore Fine-Grained Communication	28
<i>W. Pepper Marts (Center for Computing Research, Sandia Natinal Laboratories, USA; University of New Mexico, USA), Donald A. Kruse (Center for Computing Research, Sandia Natinal Laboratories, USA), Matthew G. F. Dosanjh (Center for Computing Research, Sandia Natinal Laboratories, USA), Whit Schonbein (Center for Computing Research, Sandia Natinal Laboratories, USA), Scott Levy (Center for Computing Research, Sandia Natinal Laboratories, USA), and Patrick G. Bridges (University of New Mexico, USA)</i>	
FTTN: Feature-Targeted Testing for Numerical Properties of NVIDIA & AMD Matrix Accelerators	39
<i>Xinyi Li (University of Utah), Ang Li (Pacific Northwest National Laboratory), Bo Fang (Pacific Northwest National Laboratory), Katarzyna Swirydowicz (Pacific Northwest National Laboratory), Ignacio Laguna (Lawrence Livermore National Laboratory), and Ganesh Gopalakrishnan (University of Utah)</i>	
PireSPM: Efficient and Recoverable Secure Persistent Memory for Multi-Cores	47
<i>Weijie Huang (Xiamen University), Bohong Zhu (Xiamen University), Jiwu Shu (Xiamen University; Minjiang University), Shu Li (Alibaba group), Zhengyong Wang (Alibaba group), and Yu Gao (Alibaba group)</i>	
Towards Better QoS and Lower Costs of P4 EIP Gateway at the Edge	57
<i>Ming Yang (Cloud & Network Product Division, China Telecom Cloud Technology Co., Ltd.), Xin Yang (Cloud & Network Product Division, China Telecom Cloud Technology Co., Ltd.), Liang Xu (Cloud & Network Product Division, China Telecom Cloud Technology Co., Ltd.), Yefei Hou (Cloud & Network Product Division, China Telecom Cloud Technology Co., Ltd.), Long Xie (Cloud & Network Product Division, China Telecom Cloud Technology Co., Ltd.), and Yong Wang (Cloud & Network Product Division, China Telecom Cloud Technology Co., Ltd.)</i>	

Software Systems and Platforms

Brug: An Adaptive Memory (Re-)Allocator	67
<i>Weikang Weng (LIACS, Leiden University), Alexandru Uta (DFINITY, Zürich), and Jan S. Rellermeyer (Leibniz University Hannover)</i>	
Causality Enhanced Graph Representation Learning for Alert-Based Root Cause Analysis	77
<i>Zhaoyang Yu (Tsinghua University, BNRist), Qianyu Ouyang (Tsinghua University, BNRist), Changhua Pei (Computer Network Information Center, Chinese Academy of Sciences), Xin Wang (Stony Brook University), Wenxiao Chen (Huawei), Liangfei Su (eBay Inc), Huai Jiang (eBay Inc), Xuanrun Wang (Lingjun Investment), Jianhui Li (Computer Network Information Center, Chinese Academy of Sciences), and Dan Pei (Tsinghua University, BNRist)</i>	
COTuner: Joint Optimization of Resource Configuration and Software Parameters for Recurring Streaming Jobs on the Cloud	87
<i>Hui Dou (Anhui University, China), Shanshan Zhu (Anhui University, China), Yuxuan Zhou (Anhui University, China), Yiwen Zhang (Anhui University, China), Jicheng Mei (Anhui University, China), Yang Wu (Anhui University, China), and Jiaqing Dai (Anhui University, China)</i>	

Elasticity in a Task-Based Dataflow Runtime Through Inter-Node GPU Work Stealing	97
<i>Joseph John (National Computational Infrastructure, Australia; Australian National University, Australia) and Josh Milthorpe (Oak Ridge National Laboratory, USA; Australian National University, Australia)</i>	
Fair, Efficient Multi-Resource Scheduling for Stateless Serverless Functions with Anubis	106
<i>Amit Samanta (University of Utah) and Ryan Stutsman (University of Utah)</i>	
HUILLY: A Non-Blocking Ingestion Buffer for Timestepped Simulation Analytics	113
<i>Xiaorui Du (Technical University of Munich, Germany), Andrea Piccione (Huawei Munich Research Center, Germany), Adriano Pimpini (Sapienza, University of Rome, Italy), Stefano Bortoli (Huawei Munich Research Center, Germany), Alois Knoll (Technical University of Munich, Germany), and Alessandro Pellegrini (University of Rome Tor Vergata, Italy)</i>	
Improved Parallel Application Performance and Makespan by Colocation and Topology-Aware Process Mapping	119
<i>Ioannis Vardas (TU Wien, Austria), Sascha Hunold (TU Wien, Austria), Philippe Swartvagher (Univ. Bordeaux Talence, France), and Jesper Larsson Träff (TU Wien, Austria)</i>	
Improving the Efficiency of Serverless Computing via Core-Level Power Management	125
<i>Du Liu (Shanghai Jiao Tong University, China), Jing Wang (Shanghai Jiao Tong University, China), Xinkai Wang (Shanghai Jiao Tong University, China), Chao Li (Shanghai Jiao Tong University, China), Lu Zhang (Shanghai Jiao Tong University, China), Xiaofeng Hou (Shanghai Jiao Tong University, China), Xiaoxiang Shi (Shanghai Jiao Tong University, China), and Minyi Guo (Shanghai Jiao Tong University, China)</i>	
SLO-Power: SLO and Power-Aware Elastic Scaling for Web Services	136
<i>Mehmet Savasci (University of Massachusetts Amherst, USA), Abel Souza (University of Massachusetts Amherst, USA), Li Wu (University of Massachusetts Amherst, USA), David Irwin (University of Massachusetts Amherst, USA), Ahmed Ali-Eldin (Chalmers University of Technology, Sweden), and Prashant Shenoy (University of Massachusetts Amherst, USA)</i>	
SweetSpotVM: Oversubscribing CPU Without Sacrificing VM Performance	148
<i>Pierre Jacquet (Inria, Univ. Lille, France), Thomas Ledoux (IMT Atlantique, Inria, France), and Romain Rouvoy (Univ. Lille, Inria, France)</i>	
Tackling Memory Footprint Expansion During Live Migration of Virtual Machines	158
<i>Roja Eswaran (Binghamton University, USA), Mingjie Yan (Binghamton University, USA), and Kartik Gopalan (Binghamton University, USA)</i>	

vASP: Full VM Life-Cycle Protection Based on Active Security Processor Architecture	168
<i>Jiayun Chen (Institute of Information Engineering, Chinese Academy of Sciences, China; University of Chinese Academy of Sciences, China), Qihang Zhou (Institute of Information Engineering, Chinese Academy of Sciences, China), Weijuan Zhang (Institute of Information Engineering, Chinese Academy of Sciences, China), Nan Jiang (Institute of Information Engineering, Chinese Academy of Sciences, China; University of Chinese Academy of Sciences, China), Yamin Xie (Institute of Information Engineering, Chinese Academy of Sciences, China), and Xiaoqi Jia (Institute of Information Engineering, Chinese Academy of Sciences, China; University of Chinese Academy of Sciences, China)</i>	
Workload-Aware Live Migratable Cloud Instance Detector	178
<i>Junho Lim (Kookmin Univ., South Korea), KyungHwan Kim (Kookmin Univ., South Korea), and Kyungyong Lee (Kookmin Univ, South Korea)</i>	

AI/ML for Systems and Systems for AI/ML

A Clustering-Oriented Method for Open-Domain Named Entity Recognition	189
<i>Jiahui Li (China university of Geoscience, China), Diange Zhou (China University of Geoscience, China), Yilin Duan (China University of Geoscience, China), Xinchuan Li (China University of Geoscience, China), and Hong Yao (China University of Geoscience, China)</i>	
Accelerating Large Language Model Training with Hybrid GPU-Based Compression	196
<i>Lang Xu (The Ohio State University, Ohio), Quentin Anthony (The Ohio State University, Ohio), Qinghua Zhou (The Ohio State University, Ohio), Nawras Alnaasan (The Ohio State University, Ohio), Radha Gulhane (The Ohio State University, Ohio), Aamir Shafi (The Ohio State University, Ohio), Hari Subramoni (The Ohio State University, Ohio), and Dhabaleswar K. Panda (The Ohio State University, Ohio)</i>	
Apodotiko: Enabling Efficient Serverless Federated Learning in Heterogeneous Environments	206
<i>Mohak Chadha (Technische Universität München, Germany), Alexander Jensen (Technische Universität München, Germany), Jianfeng Gu (Technische Universität München, Germany), Osama Abboud (Huawei Technologies, Germany), and Michael Gerndt (Technische Universität München, Germany)</i>	
CNN Training Latency Prediction Using Hardware Metrics on Cloud GPUs	216
<i>Yoonseo Hur (Kookmin Univ. South Korea) and Kyungyong Lee (Kookmin Univ. South Korea)</i>	
DeepVM: Integrating Spot and On-Demand VMs for Cost-Efficient Deep Learning Clusters in the Cloud	227
<i>Yoochan Kim (Sogang University, Republic of Korea), Kihyun Kim (Sogang University, Republic of Korea), Yonghyeon Cho (Sogang University, Republic of Korea; LG Electronics), Jinwoo Kim (Sogang University, Republic of Korea), Awais Khan (Oak Ridge National Laboratory, USA), Ki-Dong Kang (ETRI, Daejeon, Republic of Korea), Baik-Song An (ETRI, Daejeon, Republic of Korea), Myung-Hoon Cha (ETRI, Daejeon, Republic of Korea), Hong-Yeon Kim (ETRI, Daejeon, Republic of Korea), and Youngjae Kim (Sogang University, Republic of Korea)</i>	

Dual Adaptive Compression for Efficient Communication in Heterogeneous Federated Learning ..	236
<i>Yingchi Mao (Hohai University, China), Zibo Wang (Hohai University, China), Chenxin Li (Hohai University, China), Jiakai Zhang (Hohai University, China), Shufang Xu (Hohai University, China), and Jie Wu (Temple University, USA)</i>	
Efficient Data-Parallel Continual Learning with Asynchronous Distributed Rehearsal Buffers.....	245
<i>Thomas Bouvier (University of Rennes, Inria, CNRS, IRISA), Bogdan Nicolae (Argonne National Laboratory), Hugo Chaugier (University of Rennes, Inria, CNRS, IRISA), Alexandru Costan (University of Rennes, Inria, CNRS, IRISA), Ian Foster (Argonne National Laboratory), and Gabriel Antoniu (University of Rennes, Inria, CNRS, IRISA)</i>	
Federated Semi-Supervised Learning with Local and Global Updating Frequency Optimization	255
<i>Xin Hang (University of Science and Technology of China, China), Yang Xu (University of Science and Technology of China, China), Hongli Xu (University of Science and Technology of China, China), Yunming Liao (University of Science and Technology of China, China), and Lun Wang (University of Science and Technology of China, China)</i>	
Intelligent Data Source Emission Rate Control for Optimising the Performance of Streaming Applications	266
<i>Ziren Xiao (The University of Melbourne, Australia), Christopher Leckie (The University of Melbourne, Australia), and Maria A. Rodriguez (The University of Melbourne, Australia)</i>	
Is the Powersave Governor Really Saving Power?	273
<i>Darong Huang (Embedded Systems Laboratory (ESL), EPFL, Switzerland), Luis Costero (UCM, Spain), and David Atienza (Embedded Systems Laboratory (ESL), EPFL, Switzerland)</i>	
MDSTGCN : Multi-Scale Dynamic Spatial-Temporal Graph Convolution Network With Edge Feature Embedding for Traffic Forecasting	284
<i>Sijia Liu (Heilongjiang University, China), Hui Xu (Heilongjiang University, China), Fanyu Meng (Heilongjiang University, China), and Qianqian Ren (Heilongjiang University, China)</i>	
Multi-view Negative-Free Contrastive Learning on Adaptive Graph Augmentation	291
<i>Xingyue Wang (Hainan University, China), Huazhong Liu (Hainan University, China), Jihong Ding (Hainan University, China), and Peng Tan (Hainan Credit Investigation Company Limited, China)</i>	
Opportunistic Energy-Aware Scheduling for Container Orchestration Platforms Using Graph Neural Networks	299
<i>Philipp Raith (TU Wien, Austria), Gourav Rattihalli (Hewlett Packard Labs, USA), Aditya Dhakal (Hewlett Packard Labs, USA), Sai Rahul Chalamalasetti (Hewlett Packard Labs, USA), Dejan Milojevic (Hewlett Packard Labs, USA), Eitan Frachtenberg (Hewlett Packard Labs, USA), Stefan Nastic (TU Wien, Austria), and Schahram Dustdar (TU Wien, Austria)</i>	
Preserving Near-Optimal Gradient Sparsification Cost for Scalable Distributed Deep Learning	307
<i>Daegun Yoon (ETRI Daejeon, Republic of Korea) and Sangyoon Oh (Ajou University, Republic of Korea)</i>	

Rapid Deployment of DNNs for Edge Computing via Structured Pruning at Initialization	317
<i>Bailey J. Eccles (University of St Andrews, UK), Leon Wong (Rakuten Mobile, Japan), and Blesson Varghese (University of St Andrews, UK)</i>	
Scheduling with Fully Compressible Tasks: Application to Deep Learning Inference with Neural Network Compression	327
<i>Tiago Da Silva Barros (Université Côte d'Azur, France), Frédéric Giroire (Université Côte d'Azur, France), Ramon Aparicio-Pardo (Université Côte d'Azur, France), Stéphanne Perennes (Université Côte d'Azur, France), and Emanuele Natale (Université Côte d'Azur, France)</i>	

Future Compute Continuum and Seamless Ecosystems

A Cross-Architecture Evaluation of WebAssembly in the Cloud-Edge Continuum	337
<i>Sangeeta Kakati (University of Luxembourg, Luxembourg) and Mats Brorsson (University of Luxembourg, Luxembourg)</i>	
A Formal Modeling and Verification Approach for IoT-Cloud Resource-Oriented Applications	347
<i>Yasmine Gara Hellal (University Of Monastir, Tunisia), Lazhar Hamel (University Of Monastir, Tunisia), Mohamed Graiet (Inria, LS2N Laboratory, IMT Atlantique, France), and Daniel Balouek (Inria, LS2N Laboratory, IMT Atlantique, France)</i>	
Benchmarking Performance of Various MQTT Broker Implementations in a Compute Continuum	357
<i>Jasenska Dizdarević (Technische Universität Braunschweig, Germany), Marc Michalke (Technische Universität Braunschweig, Germany), Admela Jukan (Technische Universität Braunschweig, Germany), Xavi Masip-Bruin (Universitat Politècnica de Catalunya, Spain), and Francesco D'Andria (Eviden, Spain)</i>	
Demystifying Swarm Learning: An Emerging Decentralized Federated Learning System	367
<i>Jialiang Han (Peking University, China), Yudong Han (Peking University, China), Ying Zhang (Peking University, China), Xiang Jing (Peking University, China), and Yun Ma (Peking University, China)</i>	
Efficient and Budget-Balanced Decentralized Management of Federated Cloud and Edge Providers	374
<i>George Darzanos (Athens University of Economics and Business (AUEB), Greece), Thanasis G. Papaioannou (Athens University of Economics and Business (AUEB), Greece; National and Kapodistrian University of Athens (NKUA), Greece), and George D. Stamoulis (Athens University of Economics and Business (AUEB), Greece)</i>	
High-Throughput Real-Time Edge Stream Processing with Topology-Aware Resource Matching ..	385
<i>Peng Kang (The University of Texas at San Antonio), Samee U. Khan (Mississippi State University), Xiaobo Zhou (The University of Colorado Colorado Springs), and Palden Lama (University of Texas at San Antonio)</i>	
Jingle: IoT-Informed Autoscaling for Efficient Resource Management in Edge Computing	395
<i>Yixuan Wang (University of Minnesota, USA), Abhishek Chandra (University of Minnesota, USA), and Jon Weissman (University of Minnesota, USA)</i>	

Leveraging Multi-Modal Data for Efficient Edge Inference Serving	408
<i>Joel Wolfrath (University of Minnesota, USA), Anirudh Achanta (University of Minnesota, USA), and Abhishek Chandra (University of Minnesota, USA)</i>	
MQTT2EdgePeer: a Robust and Scalable Brokerless Peer-to-Peer Edge Middleware for Topic-Based Publish/Subscribe	419
<i>Saeed Rahmani (ÉTS Montréal / Université du Québec, Canada), Amir Ali-Pour (ÉTS Montréal / Université du Québec, Canada), Camille Coti (ÉTS Montréal / Université du Québec, Canada), and Julien Gascon-Samson (ÉTS Montréal / Université du Québec, Canada)</i>	

Applications and Workflows

EnC-IoT: An Efficient Encryption and Access Control Framework Based on IPFS for Decentralized IoT	425
<i>Mansub Song (Seoul National University, South Korea), Minji Lee (Seoul National University, South Korea), Sunggon Kim (Seoul National University of Science and Technology, South Korea), Hyeonsang Eom (Seoul National University, South Korea), and Yongseok Son (Chung-Ang University, South Korea)</i>	
Multi-Objective Optimization for Joint Task Scheduling and Data Placement in Edge-Based AIoT Systems: A Learning-Based Approach	435
<i>Mingyan Fang (Anhui University, China), Xiao Liu (Deakin University, Australia), Jia Xu (Anhui University, China), Aiting Yao (Anhui University, China), Fengjie Tang (Anhui University, China), and Xuejun Li (Anhui University, China)</i>	
SMART: Serverless Module Analysis and Recognition Technique for Managed Applications	442
<i>Adi Ashkenazi (Ben-Gurion University of the Negev, Israel), Edita Grolman (Ben-Gurion University of the Negev, Israel), Aviad Elyashar (Ben-Gurion University of the Negev, Israel; Shamoon College of Engineering, Israel), Dudu Mimran (Ben-Gurion University of the Negev, Israel), Oleg Brodt (Ben-Gurion University of the Negev, Israel), Yuval Elovici (Ben-Gurion University of the Negev, Israel), and Asaf Shabtai (Ben-Gurion University of the Negev, Israel)</i>	
STRonG: System Topology Risk Analysis on Graphs	453
<i>Lars Schneidenbach (IBM T.J. Watson Research Center, USA), Sandhya Koteswara (IBM T.J. Watson Research Center, USA), Martin Ohmacht (IBM T.J. Watson Research Center, USA), Apoorve Mohan (IBM T.J. Watson Research Center, USA), and Eun Kyung Lee (IBM T.J. Watson Research Center, USA)</i>	
Workflow Mini-Apps: Portable, Scalable, Tunable & Faithful Representations of Scientific Workflows	465
<i>Ozgur O. Kilic (Brookhaven National Laboratory, USA), Tianle Wang (Brookhaven National Laboratory, USA), Matteo Turilli (Brookhaven National Laboratory, USA; Rutgers-New Brunswick, USA), Mikhail Titov (Brookhaven National Laboratory, USA), Andre Merzky (RADICAL-Computing Inc, USA), Line Pouchard (Brookhaven National Laboratory, USA), and Shantenu Jha (Rutgers-New Brunswick, USA; Brookhaven National Laboratory, USA)</i>	

Performance Monitoring, Modeling, Analysis, and Benchmarking

A Multi-Level, Multi-Scale Visual Analytics Approach to Assessment of Multifidelity HPC Systems	478
<i>Shilpika Shilpika (University of California), Bethany Lusch (Argonne Leadership Computing Facility, Argonne National Laboratory), Murali Emani (Argonne Leadership Computing Facility, Argonne National Laboratory), Filippo Simini (Argonne Leadership Computing Facility, Argonne National Laboratory), Venkatram Vishwanath (Argonne Leadership Computing Facility, Argonne National Laboratory), Michael E. Papka (University of Illinois Chicago), and Kwan-Liu Ma (University of California, Davis)</i>	
BlueJay: A Platform to Quantifying the Impact of Memory Latency on Datacenter Application Performance	489
<i>Jingchang Qin (Zhejiang University), Yiquan Chen (Zhejiang University; Alibaba Group), Shishun Cai (Alibaba Group), Wenhai Lin (Zhejiang University), Jiexiong Xu (Zhejiang University), Zhen Jin (Zhejiang University), Lifa Cao (Alibaba Group), Zijie Zheng (Alibaba Group), Yuzhong Zhang (Alibaba Group), Yi Chen (Zhejiang University), and Wenzhi Chen (Zhejiang University)</i>	
TacVar: Tackling Variability in Short-Interval Timing Measurements on X86 Processors	496
<i>Qiucheng Liao (Shanghai Jiao Tong University, China) and James Lin (Shanghai Jiao Tong University, China)</i>	
CINDA: Don't Ignore Instructions When Cloning Memory Access Behavior	507
<i>Wenhai Lin (Zhejiang University), Yiquan Chen (Zhejiang University; Alibaba Group), Jiexiong Xu (Zhejiang University), Zhen Jin (Zhejiang University), Peiyu Liu (Zhejiang University), Shishun Cai (Alibaba Group), Yuzhong Zhang (Alibaba Group), Jingchang Qin (Zhejiang University), Yiquan Lin (Zhejiang University), and Wenzhi Chen (Zhejiang University)</i>	
HAPPIES: a History-Aware Efficient Cloud Resource Overcommitment System	514
<i>Ziwei Huang (State Key Lab of Processors, Institute of Computing Technology, China; University of Chinese Academy of Sciences, China), Shibo Tang (State Key Lab of Processors, Institute of Computing Technology, China; University of Chinese Academy of Sciences, China), Zihao Chang (State Key Lab of Processors, Institute of Computing Technology, China; University of Chinese Academy of Sciences, China), Lin Tan (Meituan Inc., China), Qichao Lu (Meituan Inc., China), Jian Ouyang (Meituan Inc., China), Wenbin Lv (State Key Lab of Processors, Institute of Computing Technology, China; University of Chinese Academy of Sciences, China), Zhicheng Yao (State Key Lab of Processors, Institute of Computing Technology, China; University of Chinese Academy of Sciences, China), Yungang Bao (State Key Lab of Processors, Institute of Computing Technology, China; University of Chinese Academy of Sciences, China), and Sa Wang (State Key Lab of Processors, Institute of Computing Technology, China; University of Chinese Academy of Sciences, China)</i>	

Quantifying and Modeling Irregular MPI Communication	525
<i>Carson Woods (Emory University, USA), Derek Schafer (University of New Mexico, USA), Patrick G. Bridges (University of New Mexico, USA), and Anthony Skjellum (Tennessee Technological University, USA)</i>	
TianheStar: Orchestrating SSSP Applications on Tianhe Supercomputer	534
<i>Xinbiao Gan (National University of Defense Technology), Qian Tang (National University of Defense Technology), Feng Xiong (National University of Defense Technology), Shijie Li (National University of Defense Technology), Bo Yang (National University of Defense Technology), and Tijun Li (National University of Defense Technology)</i>	
XFbench: A Cross-Cloud Benchmark Suite for Evaluating FaaS Workflow Platforms	543
<i>Varad Kulkarni (Indian Institute of Science, India), Nikhil Reddy (Indian Institute of Science, India), Tuhin Khare (Indian Institute of Science, India), Harini Mohan (Sri Sivasubramaniya Nadar College of Engineering, India), Jahnvi Murali (Sri Sivasubramaniya Nadar College of Engineering, India), Mohith A (Sri Sivasubramaniya Nadar College of Engineering, India), Ragul B (Sri Sivasubramaniya Nadar College of Engineering, India), Sanjai Balajee (Sri Sivasubramaniya Nadar College of Engineering, India), Sanjjit S (Sri Sivasubramaniya Nadar College of Engineering, India), Swathika D (Sri Sivasubramaniya Nadar College of Engineering, India), Vaishnavi S (Sri Sivasubramaniya Nadar College of Engineering, India), Yashasvee V (Sri Sivasubramaniya Nadar College of Engineering, India), Chitra Babu (Sri Sivasubramaniya Nadar College of Engineering, India), Abhinandan S. Prasad (Indian Institute of Technology, India), and Yogesh Simmhan (Indian Institute of Science, India)</i>	

Distributed and Parallel Storage Systems

Access-Based Carving of Data for Efficient Reproducibility of Containers	557
<i>Rohan Tikmany (DePaul University, Chicago, USA), Aniket Modi (DePaul University, Chicago, USA), Raffay Atiq (Lahore University of Management Sciences, Pakistan), Moaz Reyad (DePaul University, USA), Ashish Gehani (SRI International, USA), and Tanu Malik (DePaul University, Chicago, USA)</i>	
Dataplug: Unlocking Extreme Data Analytics with on-the-fly Dynamic Partitioning of Unstructured Data	567
<i>Aitor Arjona (Universitat Rovira i Virgili, Spain), Pedro García-López (Universitat Rovira i Virgili, Spain), and Daniel Barcelona-Pons (Universitat Rovira i Virgili, Spain)</i>	
Hades: A Context-Aware Active Storage Framework for Accelerating Large-Scale Data Analysis...	577
<i>Jaime Cernuda (Illinois Institute Of Techonogy), Luke Logan (Illinois Institute Of Techonogy), Ana Gainaru (Oak Ridge National Laboratory), Scott Klasky (Oak Ridge National Laboratory), Jay Lofstead (Sandia National Laboratory), Anthony Kougkas (Illinois Institute Of Techonogy), and Xian-He Sun (Illinois Institute Of Techonogy)</i>	

HeROcache: Storage-Aware Scheduling in Heterogeneous Serverless Edge – The Case of IDS	587
<i>Vincent Lannurien (bcom Institute of Research and Technology; ENSTA Bretagne, Lab-STICC, France), Camélia Slimani (ENSTA Bretagne, Lab-STICC, France), Laurent D’Orazio (bcom Institute of Research and Technology; Univ. Rennes, Inria, IRISA), Olivier Barais (bcom Institute of Research and Technology; Univ. Rennes, Inria, IRISA), Stéphane Paquelet (bcom Institute of Research and Technology), and Jalil Boukhobza (bcom Institute of Research and Technology; ENSTA Bretagne, Lab-STICC, France)</i>	
IDIOMS: Index-Powered Distributed Object-Centric Metadata Search for Scientific Data Management	598
<i>Wei Zhang (Lawrence Berkeley National Laboratory), Houjun Tang (Lawrence Berkeley National Laboratory), and Suren Byna (The Ohio State University; Lawrence Berkeley National Laboratory)</i>	
ZUFS: Enhancing Stability and Endurance in Mobile Devices with Integrated Zoned Namespaces in Universal Flash Storage	609
<i>Pengbo Yan (Xiamen University), Bohong Zhu (Xiamen University), Zhirong Shen (Xiamen University), Jiwu Shu (Xiamen University; Minjiang University), and Jiadong Yang (The 52nd Research Institute of China Electronics Technology Group Corporation)</i>	

Education in Cloud, Cluster, and Internet Computing

Reviving Storage Systems Education in the 21st Century — An Experience Report	616
<i>Animesh Trivedi (Vrije Universiteit Amsterdam, The Netherlands), Matthijs Jansen (Vrije Universiteit Amsterdam, The Netherlands), Krijn Doekemeijer (Vrije Universiteit Amsterdam, The Netherlands), Sacheendra Talluri (Vrije Universiteit Amsterdam, The Netherlands), and Nick Tehrany (BlueOne Business Software LLC, Beverly Hills, USA)</i>	
Training Computer Scientists for the Challenges of Hybrid Quantum-Classical Computing	626
<i>Vincenzo De Maio (HPC Laboratory, Institute of Information Systems Engineering, TU Wien, Austria), Meerzhan Kanatbekova (HPC Laboratory, Institute of Information Systems Engineering, TU Wien, Austria), Felix Zilk (HPC Laboratory, Institute of Information Systems Engineering, TU Wien, Austria), Nicolai Friis (Atominstitut, TU Wien, Austria), Tobias Guggemos (University of Vienna, Austria), and Ivona Brandic (HPC Laboratory, Institute of Information Systems Engineering, TU Wien, Austria)</i>	

Posters

A Preparing Approach to Manipulating Nested Data Structures	636
<i>Jeffrey Myers (Western Kentucky University, Kentucky) and Yaser Mowafi (Western Kentucky University, Kentucky)</i>	
FBTuner: A Feedback-Directed Approach for Safe Mixed-Precision Tuning	638
<i>Xinyi Li (the University of Utah, Salt Lake City) and Ganesh Gopalakrishnan (the University of Utah, Salt Lake City)</i>	

Power of Insensitivity: Fixing Threshold Truncation of Switch Buffer Management Policies	640
<i>Yuan Lu (National University of Defense Technology, China), Dinghuang Hu (National University of Defense Technology, China), Guannan Zhang (National University of Defense Technology, China), Jie Shen (National University of Defense Technology, China), and Dezun Dong (National University of Defense Technology, China)</i>	
Reinforcement Learning Based Matching for Parallel Computation Offloading in Dynamic Fog Computing Networks	N/A
<i>Hoa Tran-Dang (n/a) and Dong-Seong Kim (n/a)</i>	

SCALE Challenge

A Distributed, Asynchronous Algorithm for Large-Scale Internet Network Topology Analysis	644
<i>Youssef Elmougy (Georgia Institute of Technology, USA), Akihiro Hayashi (Georgia Institute of Technology, USA), and Vivek Sarkar (Georgia Institute of Technology, USA)</i>	
Sub-Model Parallelism: A Scale-out Deployment Method for Large Multi-Modal DNNs	651
<i>Tianhao Huang (Shanghai Jiao Tong University, China), Lingyu Sun (Shanghai Jiao Tong University, China), Xiaofeng Hou (Shanghai Jiao Tong University, China), Xiaozhi Zhu (Shanghai Jiao Tong University, China), Xinfeng Xia (Shanghai Jiao Tong University, China), Yutong Wang (Shanghai Jiao Tong University, China), Mingxi Chen (Shanghai Jiao Tong University, China), and Chao Li (Shanghai Jiao Tong University, China)</i>	

Doctoral Symposium

Data Sharing-Aware Algorithms for Task Allocation in Edge Computing Systems	658
<i>Sanaz Rabinia (Wayne State University, USA) and Daniel Grosu (Wayne State University, USA)</i>	
Discovery of Floating-Point Differences Between NVIDIA and AMD GPUs	663
<i>Xinyi Li (University of Utah, USA), Ang Li (Pacific Northwest National Laboratory, USA), Bo Fang (Pacific Northwest National Laboratory, USA), Katarzyna Swirydowicz (Pacific Northwest National Laboratory, USA), Ignacio Laguna (Lawrence Livermore National Laboratory, USA), and Ganesh Gopalakrishnan (University of Utah, USA)</i>	
Incorporating Memory Sharing-Awareness in Multi-VM Live Migration	667
<i>Roja Eswaran (Binghamton University), Mingjie Yan (Binghamton University), and Kartik Gopalan (Binghamton University)</i>	
Towards an Intelligent Framework for Scientific Computational Steering in Big Data Systems.....	671
<i>Yijie Zhang (New Jersey Institute of Technology, USA) and Chase Q. Wu (New Jersey Institute of Technology, USA)</i>	

Author Index	677
--------------------	-----