

2023 IEEE/ACM 23rd International Symposium on Cluster, Cloud and Internet Computing (CCGrid 2023)

**Bangalore, India
1-4 May 2023**



**IEEE Catalog Number: CFP23276-POD
ISBN: 979-8-3503-0120-5**

**Copyright © 2023 by the Institute of Electrical and Electronics Engineers, Inc.
All Rights Reserved**

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

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

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

IEEE Catalog Number:	CFP23276-POD
ISBN (Print-On-Demand):	979-8-3503-0120-5
ISBN (Online):	979-8-3503-0119-9

Additional Copies of This Publication Are Available From:

Curran Associates, Inc
57 Morehouse Lane
Red Hook, NY 12571 USA
Phone: (845) 758-0400
Fax: (845) 758-2633
E-mail: curran@proceedings.com
Web: www.proceedings.com

CURRAN ASSOCIATES INC.
proceedings
.com

2023 IEEE/ACM 23rd International Symposium on Cluster, Cloud and Internet Computing (CCGrid) **CCGrid 2023**

Table of Contents

Message from the General Chairs	xiv
Message from the Program Chairs	xviii
Organizing Committee Members	xxi
Program Committee Members	xxiv
Artifact Evaluation Committee Members	xxviii
Steering Committee Members	xxix

Hardware Systems and Networking Track

<p>An Optical Transceiver Reliability Study Based on SFP Monitoring and OS-Level Metric Data 1</p> <p><i>Paolo Notaro (Huawei Technologies Duesseldorf GmbH, Germany; Technical University of Munich, Germany), Qiao Yu (Huawei Technologies Duesseldorf GmbH, Germany; Technical University of Berlin, Germany), Soroush Haeri (Huawei Technologies Duesseldorf GmbH, Germany), Jorge Cardoso (Huawei Technologies Duesseldorf GmbH, Germany; University of Coimbra, Portugal), and Michael Gerndt (Technical University of Munich, Germany)</i></p> <p>HyQ: Hybrid I/O Queue Architecture for NVMe over Fabrics to Enable High-Performance</p> <p>Hardware Offloading 13</p> <p><i>Yiquan Chen (Zhejiang University & Alibaba Group, China), Jinlong Chen (Zhejiang University, China), Yijing Wang (Alibaba Group, China), Yi Chen (Zhejiang University, China), Zhen Jin (Zhejiang University, China), Jiexiong Xu (Zhejiang University, China), Guoju Fang (Alibaba Group, China), Wenhai Lin (Zhejiang University, China), Chengkun Wei (Zhejiang University, China), and Wenzhi Chen (Zhejiang University, China)</i></p>	<p>1</p> <p>13</p>
--	--------------------

Rethinking Design Paradigm of Graph Processing System with a CXL-Like Memory Semantic Fabric	25
<i>Xu Zhang (Institute of Computing Technology, CAS, China; University of Chinese Academy of Sciences, China), Yisong Chang (Institute of Computing Technology, CAS, China; University of Chinese Academy of Sciences, China; Zhongguancun Laboratory, China), Tianyue Lu (Institute of Computing Technology, CAS, China; University of Chinese Academy of Sciences, China), Ke Zhang (Institute of Computing Technology, CAS, China; University of Chinese Academy of Sciences, China), and Mingyu Chen (Institute of Computing Technology, CAS, China; University of Chinese Academy of Sciences, China; Zhongguancun Laboratory, China)</i>	
RoUD: Scalable RDMA over UD in Lossy Data Center Networks	36
<i>Zhiqiang He (University of Science and Technology of China, China), Yuxin Chen (University of Science and Technology of China, China), and Bei Hua (University of Science and Technology of China, China)</i>	
Taming Metadata-Intensive HPC Jobs Through Dynamic, Application-Agnostic QoS Control	47
<i>Ricardo Macedo (INESC TEC & University of Minho), Mariana Miranda (INESC TEC & University of Minho), Yusuke Tanimura (AIST), Jason Haga (AIST), Amit Ruhela (TACC & UTAustin), Stephen Lien Harrell (TACC & UTAustin), Richard Todd Evans (Intel), José Pereira (INESC TEC & University of Minho), and João Paulo (INESC TEC & University of Minho)</i>	
hsSpMV: A Heterogeneous and SPM-Aggregated SpMV for SW26010-Pro Many-Core Processor	62
<i>JingShan Pan (Shandong Computer Science Center (National Supercomputer Center in Jinan) Qilu University of Technology (Shandong Academy of Sciences)), Lei Xiao (Shandong Computer Science Center (National Supercomputer Center in Jinan) Qilu University of Technology (Shandong Academy of Sciences)), Min Tian (Shandong Computer Science Center (National Supercomputer Center in Jinan) Qilu University of Technology (Shandong Academy of Sciences)), Li Wang (Shandong Computer Science Center (National Supercomputer Center in Jinan) Qilu University of Technology (Shandong Academy of Sciences)), Chaochao Yang (Shandong Computer Science Center (National Supercomputer Center in Jinan) Qilu University of Technology (Shandong Academy of Sciences)), Renjiang Chen (Shandong Computer Science Center (National Supercomputer Center in Jinan) Qilu University of Technology (Shandong Academy of Sciences)), Zenghui Ren (Shandong Computer Science Center (National Supercomputer Center in Jinan) Qilu University of Technology (Shandong Academy of Sciences)), Anjun Liu (Shandong Computer Science Center (National Supercomputer Center in Jinan) Qilu University of Technology (Shandong Academy of Sciences)), and Guanghui Zhu (Shandong Computer Science Center (National Supercomputer Center in Jinan) Qilu University of Technology (Shandong Academy of Sciences))</i>	

Software Systems and Platforms Track

A Case Study of Data Management Challenges Presented in Large-Scale Machine Learning Workflows	71
<i>Claire Songhyun Lee (Northwestern University), V. Hewes (University of Cincinnati), Giuseppe Cerati (Fermi National Accelerator Laboratory), Jim Kowalkowski (Fermi National Accelerator Laboratory), Adam Aurisano (University of Cincinnati), Ankit Agrawal (Northwestern University), Alok Choudhary (Northwestern University), and Wei-keng Liao (Northwestern University)</i>	
An Asynchronous Dataflow-Driven Execution Model for Distributed Accelerator Computing	82
<i>Philip Salzmann (University of Innsbruck, Austria), Fabian Knorr (University of Innsbruck, Austria), Peter Thoman (University of Innsbruck, Austria), Philipp Gschwandtner (University of Innsbruck, Austria), Biagio Cosenza (University of Salerno, Italy), and Thomas Fahringer (University of Innsbruck, Austria)</i>	
An Empirical Study of Container Image Configurations and Their Impact on Start Times	94
<i>Martin Straesser (University of Würzburg, Germany), André Bauer (University of Chicago, USA), Robert Leppich (University of Würzburg, Germany), Nikolas Herbst (University of Würzburg, Germany), Kyle Chard (University of Chicago, USA), Ian Foster (University of Chicago, USA), and Samuel Kounev (University of Würzburg, Germany)</i>	
An Experimental Comparison of Software-Based Power Meters: Focus on CPU and GPU	106
<i>Mathilde Jay (Université Grenoble Alpes, CNRS, Inria, Grenoble INP, LIG, France), Vladimir Ostapenko (Univ. Lyon, EnsL, UCBL, CNRS, Inria, LIP, France), Laurent Lefevre (Univ. Lyon, EnsL, UCBL, CNRS, Inria, LIP, France), Denis Trytram (Université Grenoble Alpes, CNRS, Inria, Grenoble INP, LIG, France), Anne-Cécile Orgerie (University of Rennes, Inria, CNRS, IRISA, France), and Benjamin Fichel (OVHCloud, France)</i>	
CUDArap: Statically-Determined Execution Statistics as Alternative to Execution-Based Profiling	119
<i>Yannick Emonds (Heidelberg University, Germany), Lorenz Braun (Heidelberg University, Germany), and Holger Fröning (Heidelberg University, Germany)</i>	
Implementing and Optimizing a GPU-aware MPI Library for Intel GPUs: Early Experiences	131
<i>Chen-Chun Chen (The Ohio State University, USA), Kawthar Shafie Khorassani (The Ohio State University, USA), Rahul Vaidya (The Ohio State University, USA), Goutham Kalikrishna Reddy Kuncham (The Ohio State University, USA), Mustafa Abduljabbar (The Ohio State University, USA), Aamir Shafi (The Ohio State University, USA), Hari Subramoni (The Ohio State University, USA), and Dhabaleswar K. Panda (The Ohio State University, USA)</i>	
EMPI: Enhanced Message Passing Interface in Modern C++	141
<i>Majid Salimi Beni (University of Salerno, Italy), Luigi Crisci (University of Salerno, Italy), and Biagio Cosenza (University of Salerno, Italy)</i>	

HeROfake: Heterogeneous Resources Orchestration in a Serverless Cloud – An Application to Deepfake Detection	154
<i>Vincent Lannurien (b◁com Institute of Research and Technology; ENSTA Bretagne, Lab-STICC, CNRS, France), Laurent D’Orazio (Univ. Rennes, Inria, CNRS, IRISA; b◁com Institute of Research and Technology, France), Olivier Barais (Univ. Rennes, Inria, CNRS, IRISA; b◁com Institute of Research and Technology, France), Esther Bernard (b◁com Institute of Research and Technology, France), Olivier Weppe (b◁com Institute of Research and Technology, France), Laurent Beaulieu (b◁com Institute of Research and Technology, France), Amine Kacete (b◁com Institute of Research and Technology, France), Stéphane Paquelet (b◁com Institute of Research and Technology, France), and Jalil Boukhobza (b◁com Institute of Research and Technology; ENSTA Bretagne, Lab-STICC, CNRS, France)</i>	
How Workflow Engines Should Talk to Resource Managers: A Proposal for a Common Workflow Scheduling Interface	166
<i>Fabian Lehmann (Humboldt-Universität zu Berlin, Germany), Jonathan Bader (Technische Universität Berlin, Germany), Friedrich Tschirpke (Humboldt-Universität zu Berlin, Germany), Lauritz Thamsen (University of Glasgow, United Kingdom), and Ulf Leser (Humboldt-Universität zu Berlin, Germany)</i>	
KalpaVriksh: Efficient and Cost-Effective GUI Application Hosting using Singleton Snapshots	180
<i>Sumaiya Shaikh (IIT Kanpur, India), Saurabh Kumar (IIT Kanpur, India), and Debadatta Mishra (IIT Kanpur, India)</i>	
LayerCake: Efficient Inference Serving with Cloud and Mobile Resources	191
<i>Samuel S. Ogden (California State University Monterey Bay) and Tian Guo (Worcester Polytechnic Institute)</i>	
Optimal Sizing of a Globally Distributed low Carbon Cloud Federation	203
<i>Miguel Vasconcelos (Université Grenoble Alpes, France; University of São Paulo, Brazil), Daniel Cordeiro (University of São Paulo, Brazil), Georges Da Costa (IRIT, Université de Toulouse, France), Fanny Dufossé (University Grenoble Alpes, France), Jean-Marc Nicod (University Bourgogne Franche-Comté, France), and Veronika Rehn-Sonigo (University Bourgogne Franche-Comté, France)</i>	
Predicting the Performance-Cost Trade-off of Applications Across Multiple Systems	216
<i>Amir Nassereldine (American University of Beirut, Lebanon), Safaa Diab (American University of Beirut, Lebanon), Mohammed Baydoun (American University of Beirut, Lebanon), Kenneth Leach (Hewlett Packard Enterprise, USA), Maxim Alt (Hewlett Packard Enterprise, USA), Dejan Milojevic (Hewlett Packard Enterprise, USA), and Izzat El Hajj (American University of Beirut, Lebanon)</i>	
Runway: In-Transit Data Compression on Heterogeneous HPC Systems	229
<i>John Ravi (North Carolina State University Raleigh, USA), Suren Byna (The Ohio State University Columbus, USA), and Michela Becchi (North Carolina State University Raleigh, USA)</i>	

ML for Systems, and Systems for ML Track

A Deep Learning Pipeline Parallel Optimization Method	240
<i>Tiantian Lv (Qilu University of Technology(Shandong Academy of Sciences), China), Lu Wu (Qilu University of Technology(Shandong Academy of Sciences), China), Zhigang Zhao (Qilu University of Technology(Shandong Academy of Sciences), China), Chunxiao Wang (Qilu University of Technology(Shandong Academy of Sciences), China), and Chuantao Li (Qilu University of Technology(Shandong Academy of Sciences), China)</i>	
CADIS: Handling Cluster-Skewed Non-IID Data in Federated Learning with Clustered Aggregation and Knowledge DISTilled Regularization	249
<i>Nang Hung Nguyen (Hanoi University of Science and Technology, Vietnam), Duc Long Nguyen (Hanoi University of Science and Technology, Vietnam), Trong Bang Nguyen (Hanoi University of Science and Technology, Vietnam), Thanh-Hung Nguyen (Hanoi University of Science and Technology, Vietnam), Huy Hieu Pham (VinUniversity, Vietnam), Thao Nguyen Truong (National Institute of Advanced Industrial Science and Technology (AIST), Japan), and Phi Le Nguyen (Hanoi University of Science and Technology, Vietnam)</i>	
Chronica: A Data-Imbalance-Aware Scheduler for Distributed Deep Learning	262
<i>Sanha Maeng (Sogang University, Republic of Korea), Gordon Euhyun Moon (Sogang University, Republic of Korea), and Sungyong Park (Sogang University, Republic of Korea)</i>	
Control Channel Isolation in SDN Virtualization: A Machine Learning Approach	273
<i>Yeonho Yoo (Korea University, Republic of Korea), Gyeongsik Yang (Korea University, Republic of Korea), Changyong Shin (Korea University, Republic of Korea), Jeunghwan Lee (Korea University, Republic of Korea), and Chuck Yoo (Korea University, Republic of Korea)</i>	
COUNSEL: Cloud Resource Configuration Management using Deep Reinforcement Learning	286
<i>Adithya Hegde (The National Institute of Engineering, India), Sameer G. Kulkarni (Indian Institute of Technology Gandhinagar, India), and Abhinandan S.Prasad (The National Institute of Engineering, India)</i>	
FreeTrain: A Framework to Utilize Unused Supercomputer Nodes for Training Neural Networks .	299
<i>Zhengchun Liu (Argonne National Laboratory, USA), Rajkumar Kettimuthu (Argonne National Laboratory, USA), Michael Papka (Argonne National Laboratory, USA), and Ian Foster (Argonne National Laboratory, USA)</i>	
HDFL: A Heterogeneity and Client Dropout-Aware Federated Learning Framework	311
<i>Syed Zawad (IBM Research Almaden, USA), Ali Anwar (University of Minnesota, USA), Yi Zhou (IBM Research Almaden, USA), Nathalie Baracaldo (IBM Research Almaden, USA), and Feng Yan (University of Houston, USA)</i>	
Heterogeneous Federated Learning using Dynamic Model Pruning and Adaptive Gradient	322
<i>Sixing Yu (Iowa State University), Phuong Nguyen (Iowa State University), Ali Anwar (University of Minnesota), and Ali Jannesari (Iowa State University)</i>	

Implementing Reinforcement Learning Datacenter Congestion Control in NVIDIA NICs	331
<i>Benjamin Fuhrer (NVIDIA Networking), Yuval Shpigelman (NVIDIA Networking), Chen Tessler (NVIDIA Research), Shie Mannor (NVIDIA Research; Technion Institute of Technology), Gal Chechik (NVIDIA Research; Bar-Ilan University), Eitan Zahavi (NVIDIA Networking), and Gal Dalal (NVIDIA Research)</i>	
Measuring the Impact of Gradient Accumulation on Cloud-Based Distributed Training	344
<i>Zimeng Huang (Shanghai Jiao Tong University, China), Bo Jiang (Shanghai Jiao Tong University, China), Tian Guo (Worcester Polytechnic Institute, USA), and Yunzhuo Liu (Shanghai Jiao Tong University, China)</i>	
Optimizing Decentralized Learning with Local Heterogeneity using Topology Morphing and Clustering	355
<i>Waqwoya Abebe (Iowa State University, USA) and Ali Jannesari (Iowa State University, USA)</i>	
Overcoming Noisy Labels in Federated Learning Through Local Self-Guiding	367
<i>Daokuan Bai (University of Jinan, China), Shanshan Wang (University of Jinan, China), Wenyue Wang (University of Jinan, China), Hua Wang (Victoria University, Australia), Chuan Zhao (University of Jinan, China; Quan Cheng Laboratory, China), Peng Yuan (University of Jinan, China), and Zhenxiang Chen (University of Jinan, China)</i>	
PFSL: Personalized & Fair Split Learning with Data & Label Privacy for thin Clients	377
<i>Manas Wadhwa (Indian Institute of Technology, India), Gagan Raj Gupta (Indian Institute of Technology, India), Ashutosh Sahu (Indian Institute of Technology, India), Rahul Saini (Indian Institute of Technology, India), and Vidhi Mittal (Indian Institute of Technology, India)</i>	
ScaMP: Scalable Meta-Parallelism for Deep Learning Search	391
<i>Quentin Anthony (The Ohio State University, USA), Lang Xu (The Ohio State University, USA), Aamir Shafi (The Ohio State University, USA), Hari Subramoni (The Ohio State University, USA), and Panda Dhabaleswar (The Ohio State University, USA)</i>	
Scavenger: A Cloud Service For Optimizing Cost and Performance of ML Training	403
<i>Sahil Tyagi (Indiana University Bloomington) and Prateek Sharma (Indiana University Bloomington)</i>	

Future Compute Continuum Track

AggFirstJoin: Optimizing Geo-Distributed Joins using Aggregation-Based Transformations	414
<i>Dhruv Kumar (IIIT Delhi, India), Sohaib Ahmad (University of Massachusetts, Amherst, USA), Abhishek Chandra (University of Minnesota, Twin Cities, USA), and Ramesh K. Sitaraman (University of Massachusetts, Amherst, USA)</i>	
Bottleneck Identification and Failure Prevention with Procedural Learning in 5G RAN	426
<i>Tobias Sundqvist (Umeå University, Sweden), Monowar Bhuyan (Umeå University, Sweden), and Erik Elmroth (Umeå University, Sweden)</i>	

CacheIn: A Secure Distributed Multi-layer Mobility-Assisted Edge Intelligence Based Caching for Internet of Vehicles	437
<i>Ankur Nahar (Indian Institute of Technology Jodhpur, India), Himani Sikarwar (Indian Institute of Technology Jodhpur, India), Sanyam Jain (Indian Institute of Technology Jodhpur, India), and Debasis Das (Indian Institute of Technology Jodhpur, India)</i>	
PrivFlow: Secure and Privacy Preserving Serverless Workflows on Cloud	447
<i>Surabhi Garg (Tata Consultancy Services, India), Meena Singh Dilip Thakur (Tata Consultancy Services, India), Rajan M A (Tata Consultancy Services, India), Lakshmi Padmaja Maddali (Tata Consultancy Services, India), and Vigneswaran Ramachandran (Tata Consultancy Services, India)</i>	
Soft Reliability Aware Scheduling of Real-Time Applications on Cloud with MTTF Constraints.....	459
<i>Manojit Ghose (Indian Institute of Information Technology Guwahati, India), Krishna Prabin Pandey (IIT Guwahati, India), Niyati Chaudhari (IIT Guwahati, India), and Aryabartta Sahu (Indian Institute of Technology Guwahati, India)</i>	
The SPEC-RG Reference Architecture for the Compute Continuum	469
<i>Matthijs Jansen (Vrije Universiteit Amsterdam, The Netherlands), Auday Al-Dulaimy (Mälardalen University, Sweden), Alessandro V. Papadopoulos (Mälardalen University, Sweden), Animesh Trivedi (Vrije Universiteit Amsterdam, The Netherlands), and Alexandru Iosup (Vrije Universiteit Amsterdam, The Netherlands)</i>	
Towards a Multi-objective Scheduling Policy for Serverless-Based Edge-Cloud Continuum	485
<i>Luc Angelelli (Univ. Grenoble Alpes, France), Anderson Andrei Da Silva (Ryax Technologies, France; Univ. Grenoble Alpes, France), Yiannis Georgiou (Ryax Technologies, France), Michael Mercier (Ryax Technologies, France), Gregory Mounié (Univ. Grenoble Alpes, France), and Denis Trystram (Univ. Grenoble Alpes, France)</i>	
XFaaS: Cross-platform Orchestration of FaaS Workflows on Hybrid Clouds	498
<i>Aakash Khochare (Indian Institute of Science), Tuhin Khare (Indian Institute of Science), Varad Kulkarni (Indian Institute of Science), and Yogesh Simmhan (Indian Institute of Science)</i>	

Applications and Translation Track

A Cloud-Fog Architecture for Video Analytics on Large Scale Camera Networks using Semantic Scene Analysis	513
<i>Kunal Jain (International Institute of Information Technology, India), Kishan Sairam Adapa (International Institute of Information Technology, India), Singh Grover (International Institute of Information Technology, India), Ravi Kiran Sarvadevabhatla (International Institute of Information Technology, India), and Suresh Purini (International Institute of Information Technology, India)</i>	
Accelerating Hybrid DFT Simulations using Performance Modeling on Supercomputers	524
<i>Yosuke Oyama (Fujitsu Limited, Japan), Takumi Honda (Fujitsu Limited, Japan), Atsushi Ishikawa (National Institute for Materials Science, Japan), and Koichi Shirahata (Fujitsu Limited, Japan)</i>	

Balancing Computation and Communication in Distributed Sparse Matrix-Vector Multiplication...	535
<i>Hongli Mi (China University of Petroleum-Beijing, China), Xiangrui Yu (China University of Petroleum-Beijing, China), Xiaosong Yu (China University of Petroleum-Beijing, China), Shuangyuan Wu (China University of Petroleum-Beijing, China), and Weifeng Liu (China University of Petroleum-Beijing, China)</i>	
Blockchain Proportional Governance Reconfiguration: Mitigating a Governance Oligarchy	545
<i>Deepal Tennakoon (University of Sydney, Australia) and Vincent Gramoli (University of Sydney, Australia)</i>	
Congestion Minimization using Fog-Deployed DRL-Agent Feedback Enabled Traffic Light Cooperative Framework	557
<i>Anuj Sachan (Indian Institute of Technology Roorkee, India), Nisha Singh Chauhan (Indian Institute of Technology Roorkee, India), and Neetesh Kumar (Indian Institute of Technology Roorkee, India)</i>	
CrossLedger: A Pioneer Cross-Chain Asset Transfer Protocol	568
<i>Lokendra Vishwakarma (Indian Institute of Technology Jodhpur, India), Amritesh Kumar (Indian Institute of Technology Jodhpur, India), and Debasis Das (Indian Institute of Technology Jodhpur, India)</i>	
Efficient PRAM and Practical GPU Algorithms for Large Polygon Clipping with Degenerate Cases	579
<i>Buddhi Ashan M. K. (University of Texas at San Antonio, USA), Satish Puri (Marquette University, USA), and Sushil K. Prasad (University of Texas at San Antonio, USA)</i>	
Enabling Fast, Effective Visualization of Voluminous Gridded Spatial Datasets	592
<i>Paahuni Khandelwal (Colorado State University, USA), Menuka Warushavithana (Colorado State University, USA), Sangmi Lee Pallickara (Colorado State University, USA), and Shrideep Pallickara (Colorado State University, USA)</i>	
Mixed Precision Based Parallel Optimization of Tensor Mathematical Operations on a New-Generation Sunway Processor	605
<i>Shuwei Fan (East China Normal University, China), Yao Liu (East China Normal University, China), Juliang Su (East China Normal University, China), Xianyou Wu (East China Normal University, China), and Qiong Jiang (East China Normal University, China)</i>	
Scheduling DNN Inferencing on Edge and Cloud for Personalized UAV Fleets	615
<i>Suman Raj (Indian Institute of Science, Bangalore, India), Harshil Gupta (Indian Institute of Science, Bangalore, India), and Yogesh Simran (Indian Institute of Science, Bangalore, India)</i>	

Serverless Approach to Sensitivity Analysis of Computational Models	627
<i>Piotr Kica (Sano Centre for Computational Medicine, Poland; AGH University of Science and Technology, Poland), Magdalena Otta (Sano Centre for Computational Medicine, Poland; Insigneo Institute for in silico medicine, UK; University of Sheffield, UK), Krzysztof Czechowicz (Insigneo Institute for in silico medicine, UK; University of Sheffield, UK), Karol Zajac (Sano Centre for Computational Medicine, Poland; AGH University of Science and Technology, Poland), Piotr Nowakowski (Sano Centre for Computational Medicine, Poland), Andrew Narracott (Insigneo Institute for in silico medicine, UK; University of Sheffield, UK), Ian Halliday (Insigneo Institute for in silico medicine, UK; University of Sheffield, UK), and Maciej Malawski (Sano Centre for Computational Medicine, Poland; AGH University of Science and Technology, Poland)</i>	
Speaker Recognition System of Flexible Throat Microphone using Contrastive Learning	640
<i>Weiliang Zheng (University of Jinan, China), Zhenxiang Chen (University of Jinan, China), Yang Li (University of Jinan, China), Xiaoqing Jiang (University of Jinan, China), and Xueyang Cao (University of Jinan, China)</i>	
Towards Improving Reverse Time Migration Performance by High-Speed Lossy Compression	651
<i>Yafan Huang (University of Iowa, USA), Kai Zhao (University of Alabama at Birmingham, USA), Sheng Di (Argonne National Laboratory, USA), Guanpeng Li (University of Iowa, USA), Maxim Dmitriev (EXPEC Advanced Research Center, Saudi Arabia), Thierry-Laurent D. Tonellot (EXPEC Advanced Research Center, Saudi Arabia), and Franck Cappello (Argonne National Laboratory, USA)</i>	
Use of Cost Surface Analysis and Stream Order Analysis for Computing Shortest Paths	662
<i>Yogesh Dasgaonkar (BITS-Pilani, India)</i>	
WiDual: User Identified Gesture Recognition using Commercial WiFi	673
<i>Miaoling Dai (Shanghai University, China; Anhui Provincial Key Laboratory of Network and Information Security, Anhui Normal University, China), Chenhong Cao (Shanghai University, China; Anhui Provincial Key Laboratory of Network and Information Security, Anhui Normal University, China), Tong Liu (Shanghai University, China), Meijia Su (Shanghai University, China), Yufeng Li (Shanghai University, China), and Jiangtao Li (Shanghai University, China)</i>	

Diversity & Inclusion Track

CCGRID 2023: A Holistic Approach to Inclusion and Belonging	684
<i>Beth Plale (Indiana University, USA), Preeti Malakar (Indian Institute of Technology Kanpur, India), Meenakshi D' Souza (International Institute of Information Technology Bangalore, India), Hemangee K. Kapoor (Indian Institute of Technology Guwahati, India), Yogesh Simmhan (Indian Institute of Science Bangalore, India), Ilkay Altintas (San Diego Supercomputer Center, University of California San Diego, USA), and Manohar Swaminathan (Microsoft Research Lab – India)</i>	
Author Index	687