

2024 IEEE International Conference on Cluster Computing (CLUSTER 2024)

**Kobe, Japan
24-27 September 2024**



**IEEE Catalog Number: CFP24235-POD
ISBN: 979-8-3503-5872-8**

**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:	CFP24235-POD
ISBN (Print-On-Demand):	979-8-3503-5872-8
ISBN (Online):	979-8-3503-5871-1
ISSN:	1552-5244

Additional Copies of This Publication Are Available From:

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

CURRAN ASSOCIATES INC.
proceedings
.com

2024 IEEE International Conference on Cluster Computing (CLUSTER) **CLUSTER 2024**

Table of Contents

Welcome Message from the IEEE CLUSTER 2024 General Co-Chairs	xii
Welcome Message from the IEEE CLUSTER 2024 Program Chairs	xiii
CLUSTER 2024 Committees	xiv

Best Paper Finalist

GPU Reliability Assessment: Insights Across the Abstraction Layers	1
<i>Lishan Yang (George Mason University), George Papadimitriou (University of Athens), Dimitris Sartzetakis (University of Athens), Adwait Jog (University of Virginia), Evgenia Smirni (William & Mary), and Dimitris Gizopoulos (University of Athens)</i>	
Siesta: Synthesizing Proxy Applications for MPI Programs	14
<i>Jiyu Luo (University of Science and Technology of China), Tao Yan (University of Science and Technology of China), Qingguo Xu (University of Science and Technology of China), Jingwei Sun (University of Science and Technology of China), and Guangzhong Sun (University of Science and Technology of China)</i>	
Distributed Order Recording Techniques for Efficient Record-and-Replay of Multi-threaded Programs	27
<i>Xiang Fu (Nanchang Hangkong University, China), Shiman Meng (Nanchang Hangkong University, China), Weiping Zhang (Nanchang Hangkong University, China), Luanzheng Guo (Pacific Northwest National Laboratory, USA), Kento Sato (RIKEN R-CCS, Japan), Dong H. Ahn (NVIDIA, USA), Ignacio Laguna (Lawrence Livermore National Laboratory, USA), Gregory L. Lee (Lawrence Livermore National Laboratory, USA), and Martin Schulz (Technical University of Munich, Germany)</i>	

Graph Algorithms & GNNs

FTGraph: A Flexible Tree-based Graph Store on Persistent Memory for Large-Scale Dynamic Graphs	39
<i>Gan Sun (Institute of Information Engineering, Chinese Academy of Sciences), Jiang Zhou (Institute of Information Engineering, Chinese Academy of Sciences), Shuibing He (Zhejiang University), Bo Li (Institute of Information Engineering, Chinese Academy of Sciences), Xiaoyan Gu (Institute of Information Engineering, Chinese Academy of Sciences), and Weiping Wang (Institute of Information Engineering, Chinese Academy of Sciences)</i>	
PGSampler: Accelerating GPU-based Graph Sampling in GNN Systems via Workload Fusion	51
<i>Xiaohui Wei (Jilin University, China), Weikai Tang (Jilin University, China), Hao Qi (Huazhong University of Science and Technology, China), and Hengshan Yue (Jilin University, China)</i>	
MassiveGNN: Efficient Training via Prefetching for Massively Connected Distributed Graphs	62
<i>Aishwarya Sarkar (Iowa State University, USA), Sayan Ghosh (Pacific Northwest National Laboratory, USA), Nathan R. Tallent (Pacific Northwest National Laboratory, USA), and Ali Jannesari (Iowa State University, USA)</i>	

Performance Modeling

A Protocol to Assess the Accuracy of Process-Level Power Models	74
<i>Emile Cadorel (Davidson consulting) and Dimitri Saingre (Davidson consulting)</i>	
Holistic Performance Analysis for Asynchronous Many-Task Runtimes	85
<i>Omri Mor (University of Illinois Urbana–Champaign), George Bosilca (Nvidia Corporation), and Marc Snir (University of Illinois Urbana–Champaign)</i>	
Automated approach for accurate CPU power modelling	97
<i>Tomé Maseda (Universidade da Coruña), Jonatan Enes (Universidade da Coruña), Roberto R. Expósito (Universidade da Coruña), and Juan Touriño (Universidade da Coruña)</i>	

Networks & Communication

MPI Collective Algorithm Selection in the Presence of Process Arrival Patterns	108
<i>Majid Salimi Beni (University of Salerno), Biagio Cosenza (University of Salerno), and Sascha Hunold (TU Wien)</i>	
Optimizing Neighbor Collectives with Topology Objects	120
<i>Gerald Collom (University of New Mexico, USA), Derek Schafer (University of New Mexico, USA), Amanda Bienz (University of New Mexico, USA), Patrick Bridges (University of New Mexico, USA), and Galen Shipman (Los Alamos National Laboratory, USA)</i>	
A Topology- and Load-Aware Design for Neighborhood Allgather	131
<i>Hamed Sharifian (Queen’s University, Canada), Amirhossein Sojoodi (Queen’s University, Canada), and Ahmad Afsahi (Queen’s University, Canada)</i>	

Numerical Libraries

Uncut-GEMMs : Communication-aware matrix multiplication on multi-GPU nodes	143
<i>Petros Anastasiadis (Computing Systems Laboratory, National Technical University of Athens), Nikela Papadopoulou (University of Glasgow), Nectarios Koziris (Computing Systems Laboratory, National Technical University of Athens), and Georgios Goumas (Computing Systems Laboratory, National Technical University of Athens)</i>	
High-Performance FFT Code Generation via MLIR Linalg Dialect and SIMD Micro-Kernels	155
<i>Yifei He (KTH Royal Institute of Technology, Sweden) and Stefano Markidis (KTH Royal Institute of Technology, Sweden)</i>	
Understanding Mixed Precision GEMM with MPGemmFI: Insights into Fault Resilience	166
<i>Bo Fang (Pacific Northwest National Laboratory), Xinyi Li (University of Utah), Harvey Dam (University of Utah), Cheng Tan (Google), Siva Kumar Sastry Hari (NVIDIA), Timothy Tsai (NVIDIA), Ignacio Laguna (Lawrence Livermore National Laboratory), Dingwen Tao (Indiana University), Ganesh Gopalakrishnan (University of Utah), Prashant Nair (University of British Columbia), Kevin Barker (Pacific Northwest National Laboratory), and Ang Li (Pacific Northwest National Laboratory)</i>	

IoT, Cloud, and Data Center 1

Parallelism or Fairness? How to be friendly for SSDs in cloud environments	179
<i>Yang Zhou (Huazhong University of Science and Technology, China), Fang Wang (Huazhong University of Science and Technology, China), Zhan Shi (Huazhong University of Science and Technology, China), and Dan Feng (Huazhong University of Science and Technology, China)</i>	
SLACKVM: Packing Virtual Machines in Oversubscribed Cloud Infrastructures	190
<i>Pierre Jacquet (Inria, France), Thomas Ledoux (IMT Atlantique, France), and Romain Rouvooy (Université de Lille, France)</i>	
RL-Cache: An Efficient Reinforcement Learning based Cache Partitioning Approach for Multi-tenant CDN Services	202
<i>Ranhao Jia (Shanghai Jiao Tong University, China), Zixiao Chen (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), and Hongwen Huang (Tencent Holdings Limited)</i>	

Runtime Optimizations

FCUFS: Core-Level Frequency Tuning for Energy Optimization on Intel Processors	214
<i>Hongjian Zhang (Shanghai Jiao Tong University, China), Akira Nukada (University of Tsukuba, Japan), and Qiucheng Liao (Shanghai Jiao Tong University, China)</i>	

ML-based Dynamic Operator-Level Query Mapping for Stream Processing Systems in Heterogeneous Computing Environments	226
<i>Sejeong Oh (Sogang University, Republic of Korea), Gordon Euhyun Moon (Sogang University, Republic of Korea), and Sungyong Park (Sogang University, Republic of Korea)</i>	
Enabling Practical Transparent Checkpointing for MPI: A Topological Sort Approach	238
<i>Yao Xu (Northeastern University, USA) and Gene Cooperman (Northeastern University, USA)</i>	

IoT, Cloud, and Data Center 2

Enabling Workload-Driven Elasticity in MPI-based Ensembles	250
<i>Md Rajib Hossen (The University of Texas at Arlington), Vanessa Sochat (Lawrence Livermore National Laboratory), Abhik Sarkar (Lawrence Livermore National Laboratory), Mohammad A. Islam (The University of Texas at Arlington), and Daniel J. Milroy (Lawrence Livermore National Laboratory)</i>	
Geo-Distributed Analytical Streaming Architecture for IoT Platforms	263
<i>MohammadReza HoseinyFarahabady (The University of Sydney Australia) and Albert Y. Zomaya (The University of Sydney Australia)</i>	
Seastar: A Cache-Efficient and Load-Balanced Key-Value Store on Disaggregated Memory	275
<i>Jingwen Du (Huazhong University of Science and Technology), Fang Wang (Huazhong University of Science and Technology), Dan Feng (Huazhong University of Science and Technology), Dexin Zeng (Huazhong University of Science and Technology), and Sheng Yi (Huazhong University of Science and Technology)</i>	

Job Scheduling & Orchestration

HEFTLess: A Bi-Objective Serverless Workflow Batch Orchestration on the Computing Continuum	286
<i>Reza Farahani (University of Klagenfurt), Narges Mehran (Salzburg Research Forschungsgesellschaft mbH), Sashko Ristov (University of Innsbruck), and Radu Prodan (University of Klagenfurt)</i>	
Job Scheduling in High Performance Computing Systems with Disaggregated Memory Resources	297
<i>Jie Li (Texas Tech University, USA), George Michelogiannakis (Lawrence Berkeley National Laboratory, USA), Samuel Maloney (Julich Supercomputing Centre, Forschungszentrum Julich, Germany), Brandon Cook (Lawrence Berkeley National Laboratory, USA), Estela Suarez (Julich Supercomputing Centre, Forschungszentrum Julich, Germany), John Shalf (Lawrence Berkeley National Laboratory, USA), and Yong Chen (Texas Tech University, USA)</i>	

Fully Decentralized Data Distribution for Exascale-HPC: End of the Provider-Demander Matching Puzzle	310
<i>Mingtian Shao (National University of Defense Technology, China), Wenzhe Zhang (National University of Defense Technology, China), Ruibo Wang (National University of Defense Technology, China), Huijun Wu (National University of Defense Technology, China), Yiqin Dai (National University of Defense Technology, China), and Kai Lu (National University of Defense Technology, China)</i>	

Accelerators & In-Network Computing

FT K-means: A High-Performance K-means on GPU with Fault Tolerance	322
<i>Shixun Wu (University of California, Riverside), Yitong Ding (University of California, Riverside), Yujia Zhai (University of California, Riverside), Jinyang Liu (University of Houston, Houston, TX, US), Jiajun Huang (University of California, Riverside), Zizhe Jian (University of California, Riverside), Huangliang Dai (University of California, Riverside), Sheng Di (Argonne National Laboratory), Bryan M. Wong (University of California, Riverside), Zizhong Chen (University of California, Riverside), and Franck Cappello (Argonne National Laboratory)</i>	
ScalFrag: Efficient Tiled-MTTKRP with Adaptive Launching on GPUs	335
<i>Wenqing Lin (China University of Petroleum, Beijing), Hemeng Wang (China University of Petroleum, Beijing), Haodong Deng (China University of Petroleum, Beijing), and Qingxiao Sun (China University of Petroleum, Beijing)</i>	
Leveraging high-performance data transfer to offload data management tasks to SmartNICs	346
<i>Scott Levy (Sandia National Laboratories), Whit Schonbein (Sandia National Laboratories), and Craig Ulmer (Sandia National Laboratories)</i>	

Workflows

DaYu: Optimizing Distributed Scientific Workflows by Decoding Dataflow Semantics and Dynamics	357
<i>Meng Tang (Illinois Institute of Technology), Jaime Cernuda (Illinois Institution of Technology), Jie Ye (Illinois Institution of Technology), Luanzheng Guo (Pacific Northwest National Laboratory (PNNL)), Nathan R. Tallent (Pacific Northwest National Laboratory (PNNL)), Anthony Kougkas (Illinois Institute of Technology), and Xian-He Sun (Illinois Institute of Technology)</i>	
Sizey: Memory-Efficient Execution of Scientific Workflow Tasks	370
<i>Jonathan Bader (Technische Universität Berlin, Germany), Fabian Skalski (Technische Universität Berlin, Germany), Fabian Lehmann (Humboldt-Universität zu Berlin, Germany), Dominik Scheinert (Technische Universität Berlin, Germany), Jonathan Will (Technische Universität Berlin), Lauritz Thamsen (University of Glasgow, United Kingdom), and Odej Kao (Technische Universität Berlin)</i>	

Phase-based Data Placement Optimization in Heterogeneous Memory	382
<i>Jannis Klinckenberg (Chair for High Performance Computing, RWTH Aachen University), Clément Foyer (Université de Reims Champagne-Ardenne, CEA, LRC DIGIT, LICIS), Pierre Clouzet (Inria, Univ. Bordeaux, LaBRI), Brice Goglin (Inria, Univ. Bordeaux, LaBRI), Emmanuel Jeannot (Inria, Univ. Bordeaux, LaBRI), Christian Terboven (Chair for High Performance Computing, RWTH Aachen University), and Anara Kozhokanova (Chair for High Performance Computing, RWTH Aachen University)</i>	

Applications

Xphase3d: Memory-Distributed Phase Retrieval for Reconstructing Large-Scale 3D Density Maps of Biological Macromolecules	394
<i>Wenyang Zhao (RIKEN Center for Computational Science), Osamu Miyashita (RIKEN Center for Computational Science), Miki Nakano (RIKEN Center for Computational Science), and Florence Tama (Nagoya University, RIKEN Center for Computational Science)</i>	
Accuracy-Efficiency Optimization for Multi-Stage Small Object Detection in Surveillance Video with Collaborative Frame Sampling	403
<i>Chunhong Du (Tianjin University, China), Shanjiang Tang (Tianjin University, China), Song Meng (Tianjin University, China), Jiekai Gou (Tianjin University, China), Ce Yu (Tianjin University), Yusen Li (Nankai University, China), Hao Fu (National Supercomputing Center of Tianjin, China), Ye Tian (Tianjin University, China), and Ding Yuan (Tianjin University, China)</i>	
Modernizing an Operational Real-time Tsunami Simulator to Support Diverse Hardware Platforms	414
<i>Keichi Takahashi (Tohoku University, Japan), Takashi Abe (Tohoku University, Japan), Akihiro Musa (NEC Corporation, Japan), Yoshihiko Sato (NEC Solution Innovators, Japan), Yoichi Shimomura (NEC Solution Innovators, Japan), Hiroyuki Takizawa (Tohoku University, Japan), and Shunichi Koshimura (Tohoku University, Japan)</i>	

Storage & I/O

I/O Behind the Scenes: Bandwidth Requirements of HPC Applications With Asynchronous I/O	426
<i>Ahmad Tarraf (Technical University of Darmstadt, Germany), Javier Fernandez Muñoz (University Carlos III of Madrid, Spain), David E. Singh (University Carlos III of Madrid, Spain), Taylan Özden (Technical University of Darmstadt, Germany), Jesus Carretero (University Carlos III of Madrid, Spain), and Felix Wolf (Technical University of Darmstadt, Germany)</i>	
FINCHFS: Design of Ad-Hoc File System for I/O Heavy HPC Workloads	440
<i>Sohei Koyama (University of Tsukuba), Kohei Hiraga (University of Tsukuba), and Osamu Tatebe (University of Tsukuba)</i>	
A High-Performance and Fast-Recovery Scheme for Secure Non-Volatile Memory Systems	451
<i>Yujie Shi (Huazhong University of Science and Technology, China), Yu Hua (Huazhong University of Science and Technology, China), and Jianming Huang (Huazhong University of Science and Technology, China)</i>	

Author Index 465