

2023 IEEE International Conference on Cluster Computing Workshops (CLUSTER Workshops 2023)

**Santa Fe, New Mexico
31 October 2023**



**IEEE Catalog Number: CFP2387K-POD
ISBN: 979-8-3503-7063-8**

**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:	CFP2387K-POD
ISBN (Print-On-Demand):	979-8-3503-7063-8
ISBN (Online):	979-8-3503-7062-1

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 International Conference on Cluster Computing Workshops (CLUSTERWorkshops) **CLUSTERWorkshop 2023**

Table of Contents

Welcome Message from the IEEE Cluster 2023 General Co-Chairs	ix
<i>Scott Pakin (Los Alamos National Laboratory, USA) and Antonio J. Peña (Barcelona Supercomputing Center (BSC), Spain)</i>	
Welcome Message from the IEEE Cluster 2023 Workshops Chairs	x
<i>Hatem Ltaief (King Abdullah University of Science & Technology (KAUST), KSA) and Ulrike Meier Yang (Lawrence Livermore National Laboratory, USA)</i>	
Welcome Message from the IEEE Cluster 2023 Posters Chair	xi
<i>Ahmad Afsahi (Queen's University, Canada)</i>	
Cluster 2023 Committees	xii
<i>Scott Pakin (Los Alamos National Laboratory, USA) and Antonio J. Peña (Barcelona Supercomputing Center (BSC), Spain)</i>	

REX-IO Workshop Papers

Accelerating I/O in Distributed Data Processing Systems with Apache Arrow CHFS	1
<i>Sohei Koyama (University of Tsukuba, Japan), Kohei Hiraga (University of Tsukuba, Japan), and Osamu Tatebe (University of Tsukuba, Japan)</i>	
Does Varying BeeGFS Configuration Affect the I/O Performance of HPC Workloads?	5
<i>Arnav Borkar (BITS Pilani, K. K. Birla Goa Campus, India), Joel Tony (BITS Pilani, K. K. Birla Goa Campus, India), Hari Vamsi K. N (BITS Pilani, K. K. Birla Goa Campus, India), Tushar Barman (BITS Pilani, K. K. Birla Goa Campus, India), Yash Bhisikar (BITS Pilani, K. K. Birla Goa Campus, India), Sreenath T. M. (BITS Pilani, K. K. Birla Goa Campus, India), and Arnab K. Paul (BITS Pilani, K. K. Birla Goa Campus, India)</i>	
DAOS as HPC Storage: Exploring Interfaces	8
<i>Adrian Jackson (The University of Edinburgh, United Kingdom) and Nicolau Manubens (European Centre for Medium-Range Weather Forecasts, Germany)</i>	
I/O-Aware Flushing for HPC Caching Filesystem	11
<i>Osamu Tatebe (University of Tsukuba, Japan), Kohei Hiraga (University of Tsukuba, Japan), and Hiroki Ohtsuji (Fujitsu Limited, Japan)</i>	

Mango-IO: I/O Metrics Consistency Analysis	18
<i>Radita Liem (RWTH Aachen University, Germany), Sebastian Oeste (TU Dresden, Germany), Jay Lofstead (Sandia National Labs, Germany), and Julian Kunkel (Göttingen University/GWDG, Germany)</i>	
An I/O Performance Evaluation of Varying CephFS Striping Patterns	25
<i>Debasmita Biswas (Virginia Tech, USA), Sarah Neuwirth (Johannes Gutenberg University, Germany), Arnab K. Paul (BITS Pilani, K K Birla Goa Campus, India), and Ali R. Butt (Virginia Tech, USA)</i>	

HPCMASPA Workshop Papers

Incorporating Staggered Planned Maintenance Reservations to Improve Performance in Computational Clusters	32
<i>William M. Jones (Coastal Carolina University, USA), Craig S. Walker (Coastal Carolina University, USA), Vivian E. Hafener (Los Alamos National Laboratory, USA), Warren D. Graham (Coastal Carolina University, USA), Nathan A. DeBardleben (Los Alamos National Laboratory, USA), and Steven T. Senator (Los Alamos National Laboratory, USA)</i>	
Autonomy Loops for Monitoring, Operational Data Analytics, Feedback, and Response in HPC Operations	37
<i>Francieli Boito (University of Bordeaux, CNRS, Bordeaux INP, INRIA, LaBRI, France), Jim Brandt (Sandia National Laboratories, USA), Valeria Cardellini (University of Rome Tor Vergata, Italy), Philip Carns (Argonne National Laboratory, USA), Florina M. Ciorba (University of Basel, Switzerland), Hilary Egan (National Renewable Energy Laboratory, USA), Ahmed Eleliemy (University of Basel, Switzerland), Ann Gentile (Sandia National Laboratories, USA), Thomas Gruber (Erlangen National High Performance Computing Center, Germany), Jeff Hanson (Hewlett Packard Enterprise, USA), Utz-Uwe Haus (Hewlett Packard Labs, EMEA Research Lab, Switzerland), Kevin Huck (University of Oregon, USA), Thomas Ilsche (Technische Universität Dresden, Germany), Thomas Jakobsche (University of Basel, Switzerland), Terry Jones (Oak Ridge National Laboratory, US), Sven Karlsson (Technical University of Denmark, Denmark), Abdullah Mueen (University of New Mexico, USA), Michael Ott (Leibniz Supercomputing Centre, Germany), Tapasya Patki (Lawrence Livermore National Laboratory, USA), Ivy Peng (KTH Royal Institute of Technology, Sweden), Krishnan Raghavan (Argonne National Laboratory, USA), Stephen Simms (Lawrence Berkeley National Laboratory, USA), Kathleen Shoga (Lawrence Livermore National Laboratory, USA), Michael Showerman (University of Illinois, USA), Devesh Tiwari (Northeastern University, USA), Torsten Wilde (Hewlett Packard Enterprise, Germany), and Keiji Yamamoto (RIKEN R-CCS, Japan)</i>	

Poster Papers

I/O Characterization and Performance Evaluation of Large-Scale Storage Architectures for Heterogeneous Workloads	44
<i>Olga Kogiou (Florida State University, USA), Hariharan Devarajan (Lawrence Livermore National Laboratory, USA), Chen Wang (Lawrence Livermore National Laboratory, USA), Weikuan Yu (Florida State University, USA), and Kathryn Mohror (Lawrence Livermore National Laboratory, USA)</i>	
Few-Shot HPC Application Runtime Prediction	46
<i>Si Chen (Emory University, USA), Simon Garcia De Gonzalo (Sandia National Laboratories, USA), and Avani Wildani (Emory University, USA)</i>	
An Efficient and Accurate Compression Ratio Estimation Model for SZx	48
<i>Arham Khan (University of Chicago), Sheng Di (Argonne National Laboratory), Kai Zhao (University of Alabama), Jinyang Liu (University of California Riverside), Kyle Chard (University of Chicago), Ian Foster (University of Chicago), and Franck Cappello (Argonne National Laboratory; University of California Riverside)</i>	
VINARCH: A Visual Analytics Interactive Tool for Neural Network Archaeology	50
<i>Seoyoung An (University of Tennessee, USA), Georgia Channing (University of Tennessee, USA), Catherine Schuman (University of Tennessee, USA), and Michela Taufer (University of Tennessee, USA)</i>	
Latency and Bandwidth Microbenchmarks of Six US Department of Energy Systems in the Top500..	52
<i>Carl Pearson (Sandia National Laboratories, USA), Christopher M. Siefert (Sandia National Laboratories, USA), Stephen L. Olivier (Sandia National Laboratories, USA), Andrey Prokopenko (Oak Ridge National Laboratories, USA), Timothy J. Fuller (Sandia National Laboratories, USA), and Jonathan J. Hu (Sandia National Laboratories, USA)</i>	
Efficient Particle Tracing for Scalable Kinetic Plasma Simulation Analysis	54
<i>Nigel Tan (University of Tennessee, USA), Scott V. Luedtke (Los Alamos National Laboratory, USA), Michela Taufer (University of Tennessee, USA), and Brian Albright (Los Alamos National Laboratory, USA)</i>	
Accelerating Distributed ML Training via Selective Synchronization (Poster Abstract)	56
<i>Sahil Tyagi (Indiana University Bloomington, USA) and Martin Swamy (Indiana University Bloomington, USA)</i>	
Performance Improvement by Enhancing Spatial Parallelism on FPGA for HPC Applications	58
<i>Yuka Sano (University of Tsukuba, Japan), Taisuke Boku (University of Tsukuba, Japan), Mitsuhisa Sato (Juntendo University; Center for Computational Science, RIKEN, Japan), Miwako Tsuji (Center for Computational Science, RIKEN, Japan), Norihisa Fujita (University of Tsukuba, Japan), and Ryohei Kobayashi (University of Tsukuba, Japan)</i>	
Accelerating Hyperdimensional Classifier with SYCL	60
<i>Zheming Jin (Oak Ridge National Laboratory, USA) and Jeffrey S. Vetter (Oak Ridge National Laboratory, USA)</i>	

Mappings and Patterns to Improve the Triangular Matrix Product on Distributed Systems	62
<i>Inmaculada Santamaria-Valenzuela (Universidad de Valladolid, Spain), Rocío Carratalá-Sáez (Universidad de Valladolid, Spain), Yuri Torres (Universidad de Valladolid, Spain), Diego R. Llanos (Universidad de Valladolid, Spain), and Arturo Gonzalez-Escribano (Universidad de Valladolid, Spain)</i>	
OpenMP Offloading to DPU	64
<i>Muhammad Usman (Barcelona Supercomputing Center, Spain), Sergio Iserte (Barcelona Supercomputing Center, Spain), Roger Ferrer (Barcelona Supercomputing Center, Spain), and Antonio J. Peña (Barcelona Supercomputing Center, Spain)</i>	
Performance Insights into Device-Initiated RMA Using Kokkos Remote Spaces	66
<i>Daniel Mishler (Sandia National Laboratories, USA), Jan Ciesko (Sandia National Laboratories, USA), Stephen Olivier (Sandia National Laboratories, USA), and George Bosilca (University of Tennessee, USA)</i>	
A Lightweight Network Traffic Prediction Method for SmartNICs	68
<i>Whit Schonbein (Sandia National Laboratories, USA), Tinotenda Matsika (Queen’s University, Canada), and Ryan E. Grant (Queen’s University, Canada)</i>	
On the Multi-dimensional Acceleration of Stochastic Blockmodeling for Community Detection	70
<i>Frank Wanye (Virginia Tech, USA) and Wu-chun Feng (Virginia Tech, USA)</i>	
Author Index	73