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

Kobe, Japan
24-27 September 2024

IEEE

IEEE Catalog Number: CFP2487K-POD
ISBN: 979-8-3503-8346-1

****** 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:	CFP2487K-POD
ISBN (Print-On-Demand):	979-8-3503-8346-1
ISBN (Online):	979-8-3503-8345-4

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

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

CLUSTERW 2024

Table of Contents

Welcome Message from the IEEE CLUSTER 2024 General Co-Chairs ..... **xi**
Welcome Message from the IEEE CLUSTER 2024 Workshop Chair ..... **xii**
Welcome Message from Sustainable HPC SOP Workshop ..... **xiii**
Welcome Message from REX-IO Workshop ..... **xiv**
Welcome Message from LLMxHPC Workshop ..... **xv**
Welcome Message from the IEEE CLUSTER 2024 Posters Chair ..... **xvi**
CLUSTER 2024 Committees ..... **xvii**

Sustainable HPC SOP 2024 Workshop Papers

PowerSched - managing power consumption in overprovisioned systems 1
Christian Simmendinger (Hewlett Packard Enterprise, Germany), Marcel Marquardt (Hewlett Packard Enterprise, Germany), Jan Mäder (Hewlett Packard Enterprise, Germany), and Ralf Schneider (High Performance Computing Center Stuttgart, Germany)

"How-to" Guide for Transitioning from Air to Liquid-cooled High Performance Computing Systems 9
Dave Martinez (Sandia National Laboratory), David Sickinger (National Renewable Energy Laboratory), and Aaron Andersen (National Renewable Energy Laboratory)

Optimizing Idle Power of HPC Systems: Practical Insights and Methods 19
Thomas Ilsche (Technische Universität Dresden, Germany), Sebastian Schrader (Technische Universität Dresden, Germany), and Robert Schöne (Technische Universität Dresden, Germany)

Calculating User-Centric Carbon Footprints for HPC 26
Christian Wassermann (IT Center, RWTH Aachen University, Aachen, Germany), Mario Bielert (ZIH, CIDS, TU Dresden, Dresden, Germany), Gert Vanberg (IT Center, RWTH Aachen University, Aachen, Germany), Daniel Hackenberg (ZIH, CIDS, TU Dresden, Dresden, Germany), Christian Terboven (IT Center, RWTH Aachen University, Aachen, Germany), and Matthias S. Müller (IT Center, RWTH Aachen University, Aachen, Germany)

Evolving Large Scale HPC Monitoring & Analysis to Track Modern Dynamic Environments 36
Kathleen Shoga (Lawrence Livermore National Laboratory, USA), Jim Brandt (Sandia National Laboratories, USA), Benjamin Schwaller (Sandia National Laboratories, USA), and Thomas Tucker (Open Grid Computing, USA)

Microgrid Integration with High Performance Computing Systems for Microreactor Operation 44
Matthew Anderson (Idaho National Laboratory) and Matthew Sgambati (Idaho National Laboratory)

Power-Efficiency Variation on A64FX Supercomputers and its Application to System Operation..... 55
Tomoya Kusaba (The University of Electro-Communications), Awaki Yusuke (The University of Electro-Communications), Kohei Yoshida (The University of Electro-Communications), Shinobu Miwa (The University of Electro-Communications), Hayato Yamaki (The University of Electro-Communications), Tonoshiro Hanawa (The University of Tokyo), and Hiroki Honda (The University of Electro-Communications)

Towards Improving Resource Allocation for Multi-Tenant HPC Systems: An Exploratory HPC Cluster Utilization Case Study 66
Robert Keßler (University of Cologne, Germany), Simon Volpert (University of Cologne, Germany), and Stefan Wesner (University of Cologne, Germany)

16 Years of SPEC Power: An Analysis of x86 Energy Efficiency Trends 76
Hannes Tröpgen (Technische Universität Dresden, Germany), Robert Schöne (Technische Universität Dresden, Germany), Thomas Ilsche (Technische Universität Dresden, Germany), and Daniel Hackenberg (Technische Universität Dresden, Germany)

Advanced Visualization of Power, Temperature, and Energy Metrics in HPE Cray EX Systems 81
Lavanya L (Hewlett Packard Enterprise, India) and Stefan Ceballos (Hewlett Packard Enterprise, USA)

REX-IO 2024 Workshop Papers

Enabling High-Throughput Parallel I/O in Particle-in-Cell Monte Carlo Simulations with openPMD and Darshan I/O Monitoring 86
Jeremy Johnathan Williams (KTH Royal Institute of Technology, Sweden), Daniel Medeiros (KTH Royal Institute of Technology, Sweden), Stefan Costea (LeCAD, University of Ljubljana, Slovenia), David Tskhakaya (Institute of Plasma Physics of the CAS, Czech Republic), Franz Poeschel (Helmholtz-Zentrum Dresden-Rossendorf, Germany), René Widera (Helmholtz-Zentrum Dresden-Rossendorf, Germany), Axel Huebl (Lawrence Berkeley National Laboratory, USA), Scott Klasky (Oak Ridge National Laboratory, USA), Norbert Podhorszki (Oak Ridge National Laboratory, USA), Leon Kos (LeCAD, University of Ljubljana, Slovenia), Ales Podolnik (Institute of Plasma Physics of the CAS, Czech Republic), Jakub Hromadka (Institute of Plasma Physics of the CAS, Czech Republic), Tapish Narwal (Helmholtz-Zentrum Dresden-Rossendorf, Germany), Klaus Steiniger (Helmholtz-Zentrum Dresden-Rossendorf, Germany), Michael Bussmann (Helmholtz-Zentrum Dresden-Rossendorf, Germany), Erwin Laure (Max Planck Computing and Data Facility, Germany), and Stefano Markidis (KTH Royal Institute of Technology, Sweden)

Understanding Highly Configurable Storage for Diverse Workloads 96
Olga Kogiou (Florida State University), Hariharan Devarajan (Lawrence Livermore National Laboratory), Chen Wang (Lawrence Livermore National Laboratory), Weikuan Yu (Florida State University), and Kathryn Mohror (Lawrence Livermore National Laboratory)

Object-Centric Data Management in HPC Workflows - A Case Study 104
Chen Wang (Lawrence Livermore National Laboratory), Houjun Tang (Lawrence Berkeley National Laboratory), Jean Luca Bez (Lawrence Berkeley National Laboratory), and Suren Byna (The Ohio State University)

Studying the Effects of Asynchronous I/O on HPC I/O Patterns 109
Arnav Gupta (BITS Pilani, K. K. Birla Goa Campus, India), Druva Dhakshinamoorthy (BITS Pilani, K. K. Birla Goa Campus, India), and Arnab K. Paul (BITS Pilani, K. K. Birla Goa Campus, India)

Challenges in Understanding Metadata Performance: A Case of Metadata Analysis Using Score-P 113
Boris Kosmynin (RWTH Aachen University) and Radita Liem (RWTH Aachen University, IT Center)

LLMxHPC 2024 Workshop Papers

RAPID: A Rapid Automatic Parallelizer for Immense Deep Neural Networks 118
Chong Li (Huawei Technologies France S.A.S.U), Thibaut Tachon (Huawei Technologies France S.A.S.U), and Haoran Wang (Huawei Technologies Co., Ltd.)

Automatic Parallelization with CodeT5+: A Model for Generating OpenMP Directives 127
Soratouch Pornmaneerattanatri (Division of Information Science, Nara Institute of Science and Technology), Keichi Takahashi (Cyberscience Center, Tohoku University), Yutaro Kashiwa (Division of Information Science, Nara Institute of Science and Technology), Kohei Ichikawa (Division of Information Science, Nara Institute of Science and Technology), and Hajimu Iida (Division of Information Science, Nara Institute of Science and Technology)

LASSI: An LLM-based Automated Self-Correcting Pipeline for Translating Parallel Scientific Codes 136
Matthew Dearing (University of Illinois Chicago, USA), Yiheng Tao (University of Illinois Chicago, USA), Xingfu Wu (Argonne National Laboratory, USA), Zhiling Lan (University of Illinois Chicago, USA), and Valerie Taylor (Argonne National Laboratory, USA)

Posters

Design, Implementation and Deployment of Sunrise Integrated Health Care Cluster 144
Tao Yu (Tongji University), Zhifeng Gu (Sunrise medical tech co. ltd), Yizheng Sun (Sunrise medical tech co. ltd), and Xiaofei Wang (Tsinghua University, China)

An optimization pass for training speed-up and strategy search in 3D parallelism 146
Ryubu Hosoki (Tokyo Institute of Technology, Japan), Kento Sato (RIKEN Center for Computational Science, Japan), Toshio Endo (Tokyo Institute of Technology, Japan), Julien Bigot (CEA, France), and Edouard Audit (CEA, France)

Beyond Training: A Zero-Shot Framework to Neural Architecture and Accelerator Co-Exploration 148
Wei Fu (University of Science and Technology of China), Wenqi Lou (University of Science and Technology of China), Lei Gong (University of Science and Technology of China), Chao Wang (University of Science and Technology of China), and Xuehai Zhou (University of Science and Technology of China)

Implementing Fast Modal Filtering of SCALE-DG 150
Xuanzhengbo Ren (Nagoya University, Graduate School of Informatics, Japan), Yuta Kawai (RIKEN Center for Computational Science, Japan), Hirofumi Tomita (RIKEN Center for Computational Science, Japan), Seiya Nishizawa (RIKEN Center for Computational Science, Japan), Takahiro Katagiri (Nagoya University, Information Technology Center, Japan), Masatoshi Kawai (Nagoya University, Information Technology Center, Japan), Tetsuya Hoshino (Nagoya University, Information Technology Center, Japan), and Toru Nagai (Nagoya University, Information Technology Center, Japan)

Enhancing Large Scale Brain Simulation with Optimized Parallel Algorithms on Fugaku Supercomputer 152
Zhe Sun (Juntendo University), Mitsuhisa Sato (Juntendo University), Shigeki Aoki (Juntendo University), Ryutaro Himeno (Juntendo University), and Tianxiang Lyu (Juntendo University)

Innovative Computational Science by Integration of Simulation/Data/Learning on Heterogeneous Supercomputers 154
Kengo Nakajima (The University of Tokyo), Takashi Furumura (The University of Tokyo), France Boillod-Cerneux (CEA), Edoardo Di Napoli (Forschungszentrum Jülich GmbH), Estela Suarez (Forschungszentrum Jülich GmbH), Takashi Arakawa (CliMTech/University of Tokyo), Shinji Sumimoto (The University of Tokyo), and Hisashi Yashiro (NIES)

Neko: A Modern, Portable, and Scalable Framework for Extreme-Scale Computational Fluid Dynamics 156
Niclas Jansson (KTH Royal Institute of Technology), Martin Karp (KTH Royal Institute of Technology), Stefano Markidis (KTH Royal Institute of Technology), and Philipp Schlatter (Friedrich-Alexander-Universität, Erlangen-Nürnberg and KTH Royal Institute of Technology)

vBoost: A Lock-free Distributed Index based on vEB Tree for Disaggregated Memory 158
Yuting Li (University of Science and Technology of China, China), Yun Xu (University of Science and Technology of China, China), Pengcheng Wang (Huawei, China), Yonghui Xu (Huawei, China), and Weiguang Wang (Huawei, China)

Communication Optimization for Distributed GCN Training on ABCI Supercomputer 160
Chen Zhuang (Tokyo Institute of Technology), Peng Chen (National Institute of Advanced Industrial Science and Technology), Xin Liu (National Institute of Advanced Industrial Science and Technology), Toshio Endo (Tokyo Institute of Technology), Satoshi Matsuoka (RIKEN Center for Computational Science), and Mohamed Wahib (RIKEN Center for Computational Science)

Optimizing STAR Aligner for High Throughput Computing in the Cloud 162
Piotr Kica (Sano Centre for Computational Medicine, Poland), Sabina Lichołai (Sano Centre for Computational Medicine, Poland), Michał Orzechowski (Sano Centre for Computational Medicine), and Maciej Malawski (Sano Centre for Computational Medicine, Poland)

A Lossless-Ethernet-based interconnect for FPGA clusters toward FTQC 164
Yoshito Higa (Kumamoto University, Japan) and Yasunori Osana (Kumamoto University, Japan)

Post-Route Power Estimation: a Case Study of RIKEN-CGRA 166
Chenlin Shi (The University of Electro-Communications, Japan), Boma Adhi (RIKEN Center for Computational Science (R-CCS), Japan), Shinobu Miwa (The University of Electro-Communications, Japan), and Kentaro Sano (RIKEN Center for Computational Science (R-CCS), Japan)

Scalable Connection of Qubits to Quantum Error Correction Systems using Ethernet 168
Jan-Erik R. Wichmann (RIKEN Center for Computational Science) and Kentaro Sano (RIKEN Center for Computational Science)

Workload Analytics of LLMs Training on ABCI 170
Yusuke Tanimura (National Institute of Advanced Industrial Science and Technology), Naoki Onishi (National Institute of Advanced Industrial Science and Technology), and Shinichiro Takizawa (National Institute of Advanced Industrial Science and Technology)

Evaluating MPI Performance on SGX and Gramine 172
Shinobu Miwa (The University of Electro-Communications), Hayato Yamaki (The University of Electro-Communications), and Hiroki Honda (The University of Electro-Communications)

Investigating Nvidia GPU Architecture Trends via Microbenchmarks 174
Lingqi Zhang (RIKEN Center for Computational Science, Japan), Ryan Barton (Tokyo Institute of Technology, Japan), Peng Chen (National Institute of Advanced Industrial Science and Technology, Japan), Xiao Wang (Oak Ridge National Laboratory, United States of America), Toshio Endo (Tokyo Institute of Technology, Japan), Satoshi Matsuoka (RIKEN Center for Computational Science, Japan), and Mohamed Wahib (RIKEN Center for Computational Science, Japan)

Leveraging Portals4 Microbenchmarks to Enhance GASPI Performance on BXI Networks 176
Niklas Bartelheimer (Johannes Gutenberg University Mainz, Germany) and Sarah Neuwirth (Johannes Gutenberg University Mainz, Germany)

Evaluation of Vectorization Methods on Arm SVE Using the Exo Language 178
Rin Iwai (Toyohashi University of Technology, Japan), Emil Vatai (R-CCS, Riken, Japan), Jens Domke (R-CCS, Riken, Japan), and Yukinori Sato (Toyohashi University of Technology, Japan)

Introduction of WHEEL: An analysis workflow tool for industrial users and its use case on supercomputer Fugaku 180
Tomohiro Kawanabe (RIKEN Center for Computational Science, Japan), Naoyuki Sogo (Longtail software LLC, Japan), and Kenji Ono (Research Institute for Information Technology, Kyushu University, Japan)

Heterogeneous Application Coupling Library for Center-Wide QC-HPC Hybrid Computing 182
Shinji Sumimoto (The University of Tokyo), Kazuya Yamazaki (The University of Tokyo), Yao Hu (The University of Tokyo), and Kengo Nakajima (The University of Tokyo/RIKEN CCS)

Preliminary Performance Evaluation of Grace-Hopper GH200 184
Toshihiro Hanawa (The University of Tokyo, Japan), Kengo Nakajima (The University of Tokyo, Japan), Yohei Miki (The University of Tokyo, Japan), Takashi Shimokawabe (The University of Tokyo, Japan), Kazuya Yamazaki (The University of Tokyo, Japan), Shinji Sumimoto (The University of Tokyo, Japan), Osamu Tatebe (University of Tsukuba, Japan), Taisuke Boku (University of Tsukuba, Japan), Daisuke Takahashi (University of Tsukuba, Japan), Akira Nukada (University of Tsukuba, Japan), Norihisa Fujita (University of Tsukuba, Japan), Ryohei Kobayashi (University of Tsukuba, Japan), Hiroto Tadano (University of Tsukuba, Japan), and Akira Naruse (NVIDIA, Japan)

Performance Insights into Supporting Kokkos Views in the Kokkos Comm MPI Library 186
C. Nicole Avans (Sandia National Laboratories, USA), Jan Ciesko (Sandia National Laboratories, USA), Carl Pearson (Sandia National Laboratories, USA), Evan Drake Suggs (Tennessee Technological University, USA), Stephen L. Olivier (Sandia National Laboratories, USA), and Anthony Skjellum (Tennessee Technological University, USA)

Toward providing root privilege to flagship HPC users with thin-hypervisor 188
Takaaki Fukai (National Institute of Advanced Industrial Science and Technology, Japan), Manami Mori (Tokyo Metropolitan University / National Institute of Advanced Industrial Science and Technology, Japan), Keiji Yamamoto (RIKEN Center for Computational Science, Japan), Takahiro Hirofuchi (National Institute of Advanced Industrial Science and Technology, Japan), and Takuya Asaka (Tokyo Metropolitan University, Japan)

Cheetah: An Efficient Deterministic Concurrency Control Scheme with Non-visible Write Elimination and Re-designed Garbage Collection 190
Haowen Li (Keio University, Japan), Rina Onishi (Keio University, Japan), and Hideyuki Kawashima (Keio University, Japan)

Using SYCLomatic to migrate CUDA code to oneAPI adapting NVIDIA GPU 192
Wentao Liang (University of Tsukuba, Japan), Norihisa Fujita (University of Tsukuba, Japan), Ryohei Kobayashi (University of Tsukuba, Japan), and Taisuke Boku (University of Tsukuba, Japan)

Preliminary Evaluation of Kyokko for Inter-FPGA Communication Framework CIRCUS 194
Kaito Kitazume (Master's Program in Computer Science, Degree Programs in Systems and Information Engineering, Graduate School of Science and Technology, University of Tsukuba), Norihisa Fujita (Center for Computational Sciences, University of Tsukuba), Ryohei Kobayashi (Center for Computational Sciences, University of Tsukuba), and Taisuke Boku (Center for Computational Sciences, University of Tsukuba)

Asynchronous I/O Optimization for X-ray Imaging via GPUDirect Storage 196
Du Wu (Tokyo Institute of Technology / RIKEN-CCS), Peng Chen (National Institute of Advanced Industrial Science and Technology), Yiyu Tan (Iwate University), Yusuke Tanimura (National Institute of Advanced Industrial Science and Technology), Toshio Endo (Tokyo Institute of Technology), Satoshi Matsuoka (RIKEN-CCS), and Mohamed Wahib (RIKEN-CCS)

FDPVirt: Flexible Data Placement SSD Emulator 198
Joonyeop Park (Seoul National University, South Korea), Haeram Kim (Seoul National University, South Korea), Jiwon Ha (Seoul National University, South Korea), Hyungsoo Jung (Seoul National University, South Korea), and Hyeonsang Eom (Seoul National University, South Korea)

Enhanced Simulation and Analysis of Air Pollutants Using Multi-Platform HPC and In-Situ Visualization 200
Chongke Bi (Tianjin University, China), Fumiyoshi Shoji (RIKEN R-CCS, Japan), Kenji Ono (Kyushu University, Japan), Naohisa Sakamoto (Kobe University, Japan), Jorji Nonaka (RIKEN R-CCS, Japan), Honggang Yin (China Meteorological Administration, China), and Wenjuan Cui (Chinese Academy of Sciences, China)

On the Building of a Common In-Situ Visualization Environment for Arm A64FX Supercomputers 202
Jorji Nonaka (RIKEN R-CCS, Japan), Daichi Obinata (Fujitsu Limited, Japan), Hiroyuki Ito (Ryoyu Systems, Japan), Shunji Uno (JAXA, Japan), Takanori Haga (JAXA, Japan), Atsushi Toyoda (Intelligent Light, Japan), Naohisa Sakamoto (Kobe University, Japan), Masahiro Nakao (RIKEN R-CCS, Japan), Hitoshi Murai (RIKEN R-CCS, Japan), Keiji Yamamoto (RIKEN R-CCS, Japan), Masaaki Terai (RIKEN R-CCS, Japan), Tomohiro Kawanabe (RIKEN R-CCS, Japan), Toshihiko Kai (RIKEN R-CCS, Japan), Manabu Motokawa (JAXA, Japan), Atsushi Fujino (JAXA, Japan), Naoyuki Fujita (JAXA, Japan), Seiji Tsutsumi (JAXA, Japan), and Fumiyoshi Shoji (RIKEN R-CCS, Japan)

Refining Compaction Offloading I/O Stack for LSM-based Key-Value Stores with SPDK 204
Honghyeon Yoo (Sogang University), Hongsu Byun (Sogang University), and Sungyong Park (Sogang University)

Author Index **207**