

2021 IEEE International Conference on Cluster Computing (CLUSTER 2021)

**Portland, Oregon, USA
7 – 10 September 2021**



**IEEE Catalog Number: CFP21235-POD
ISBN: 978-1-7281-9667-1**

**Copyright © 2021 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:	CFP21235-POD
ISBN (Print-On-Demand):	978-1-7281-9667-1
ISBN (Online):	978-1-7281-9666-4
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

2021 IEEE International Conference on Cluster Computing (CLUSTER) Cluster 2021

Table of Contents

Message from Cluster 2021 General Co-Chairs	xix
Message from Cluster 2021 Program Chair	xx
Cluster 2021 Committees	xxii
Message from EE HPC SOP 2021 Organizers	xxvii
EE HPC SOP 2021 Conference Organization	xxviii
Message from HPCMASPA 2021 Organizers	xxx
HPCMASPA 2021 Conference Organization	xxxi
Message from REX-IO 2021 Organizers	xxxii
REX-IO 2021 Conference Organization	xxxiii
Message from AI4S Workshop Organizers	xxxiv
EA-HPC 2021 Conference Organization	xxxv
Message from EA-HPC Workshop Organizers	xxxvi
Message from the Organizers of FPGA for HPC Workshop 2021	xxxvii

Session 1

Track A: HPC Applications

tcFFT: A Fast Half-Precision FFT Library for NVIDIA Tensor Cores	1
<i>Binrui Li (Shanghai Jiao Tong University, China), Shenggan Cheng (Shanghai Jiao Tong University, China), and James Lin (Shanghai Jiao Tong University, China)</i>	
MiniMod: A Modular Miniapplication Benchmarking Framework for HPC	12
<i>W. Pepper Marts (Sandia National Laboratories, USA), Matthew G. F. Dosanjh (Sandia National Laboratories, USA), Scott Levy (Sandia National Laboratories, USA), Whit Schonbein (Sandia National Laboratories, USA), Ryan E. Grant (Sandia National Laboratories, USA), and Patrick G. Bridges (University of New Mexico, USA)</i>	

Track B: Performance Optimization

- Reproducible Performance Optimization of Complex Applications on the Edge-to-Cloud Continuum .23.....
Daniel Rosendo (University of Rennes, Inria, CNRS, IRISA, France), Alexandru Costan (University of Rennes, Inria, CNRS, IRISA, France), Gabriel Antoniu (University of Rennes, Inria, CNRS, IRISA, France), Matthieu Simonin (University of Rennes, Inria, CNRS, IRISA, France), Jean-Christophe Lombardo (University of Montpellier, Inria, CNRS, LIRMM, France), Alexis Joly (University of Montpellier, Inria, CNRS, LIRMM, France), and Patrick Valduriez (University of Montpellier, Inria, CNRS, LIRMM, France)
- WIRE: Resource-Efficient Scaling with Online Prediction for DAG-Based Workflows .35.....
Bing Xie (Oak Ridge National Laboratory), Qiang Cao (East China Normal University), Mayuresh Kunjir (Qatar Computing Research Institute), Linli Wan (Facebook), Jeff Chase (Duke University), Anirban Mandal (Renaissance Computing Institute), and Mats Rynge (Information Sciences Institute)

Session 2

Track A: AI in practice

- HPC AI500 V2.0: The Methodology, Tools, and Metrics for Benchmarking HPC AI Systems .47.....
Zihan Jiang (University of Chinese Academy of Sciences, China, Institute of Computing Technology, Chinese Academy of Sciences, China), Wanling Gao (Institute of Computing Technology, Chinese Academy of Sciences, China, BenchCouncil (International Open Benchmarking Council)), Fei Tang (University of Chinese Academy of Sciences, China; Institute of Computing Technology, Chinese Academy of Sciences, China), Lei Wang (Institute of Computing Technology, Chinese Academy of Sciences, China), Xingwang Xiong (University of Chinese Academy of Sciences, China; Institute of Computing Technology, Chinese Academy of Sciences, China), Chunjie Luo (Institute of Computing Technology, Chinese Academy of Sciences), Chuanxin Lan (Institute of Computing Technology, Chinese Academy of Sciences), Hongxiao Li (University of Chinese Academy of Sciences, China; Institute of Computing Technology, Chinese Academy of Sciences), and Jianfeng Zhan (University of Chinese Academy of Sciences, China; Institute of Computing Technology, Chinese Academy of Sciences; BenchCouncil (International Open Benchmarking Council))
- RPTCN: Resource Prediction for High-Dynamic Workloads in Clouds Based on Deep Learning .59..
Wenyan Chen (Shenzhen Institute of Advanced Technology, Chinese Academy of Sciences), Chengzhi Lu (Shenzhen Institute of Advanced Technology, Chinese Academy of Sciences; University of Chinese Academy of Sciences), Kejiang Ye (Shenzhen Institute of Advanced Technology, Chinese Academy of Sciences), Yang Wang (Shenzhen Institute of Advanced Technology, Chinese Academy of Sciences), and Cheng-Zhong Xu (State Key Lab of IoTSC, Faculty of Science and Technology, University of Macau)
- READYS: A Reinforcement Learning Based Strategy for Heterogeneous Dynamic Scheduling .70.....
Nathan Grinsztajn (University of Lille, France), Olivier Beaumont (Univ. Bordeaux, France), Emmanuel Jeannot (Univ. Bordeaux, France), and Philippe Preux (University of Lille, France)

Accelerating DNN Architecture Search at Scale using Selective Weight Transfer .82.....
Hongyuan Liu (William & Mary, USA), Bogdan Nicolae (Argonne National Laboratory, USA), Sheng Di (Argonne National Laboratory, USA), Franck Cappello (Argonne National Laboratory, USA), and Adwait Jog (William & Mary, USA)

SAP-SGD: Accelerating Distributed Parallel Training with High Communication Efficiency on Heterogeneous Clusters .94.....
Jing Cao (University of Science and Technology of China, China), Zongwei Zhu (University of Science and Technology of China, China), and Xuehai Zhou (University of Science and Technology of China, China)

2PGraph: Accelerating GNN Training over Large Graphs on GPU Clusters .103.....
Lizhi Zhang (National University of Defence Technology, China), Zhiquan Lai (National University of Defence Technology, China), Shengwei Li (National University of Defence Technology, China), Yu Tang (National University of Defence Technology, China), Feng Liu (National University of Defence Technology, China), and Dongsheng Li (National University of Defence Technology, China)

Track B: Storage and I/O

HFlow: A Dynamic and Elastic Multi-layered I/O Forwarder .114.....
Jaime Cernuda (Illinois Institute of Technology), Hariharan Devarajan (Illinois Institute of Technology), Luke Logan (Illinois Institute of Technology), Keith Bateman (Illinois Institute of Technology), Neeraj Rajesh (Illinois Institute of Technology), Jie Ye (Illinois Institute of Technology), Anthony Kougkas (Illinois Institute of Technology), and Xian-He Sun (Illinois Institute of Technology)

Building a Fast and Efficient LSM-Tree Store by Integrating Local Storage with Cloud Storage .125.....
Peng Xu (Huazhong University of Science and Technology, China), Nannan Zhao (Northwestern Polytechnical University, China), Jiguang Wan (Huazhong University of Science and Technology, China), Wei Liu (PingCAP, China), Shuning Chen (PingCAP, China), Yuanhui Zhou (Huazhong University of Science and Technology, China), Hadeel Albahar (Virginia Tech, USA), Hanyang Liu (PingCAP, China), Liu Tang (PingCAP, China), and Changsheng Xie (Huazhong University of Science and Technology, China)

Virtual Log-Structured Storage for High-Performance Streaming .135.....
Ovidiu-Cristian Marcu (University of Luxembourg, Luxembourg), Alexandru Costan (University of Rennes, Inria, CNRS, IRISA, France), Bogdan Nicolae (Argonne National Laboratory, USA), and Gabriel Antoniu (University of Rennes, Inria, CNRS, IRISA, France)

RISE: Reducing I/O Contention in Staging-Based Extreme-Scale In-Situ Workflows .146.....
Pradeep Subedi (University of Utah, USA), Philip E. Davis (University of Utah, USA), and Manish Parashar (University of Utah, USA)

Lazy-WL: A Wear-Aware Load Balanced Data Redistribution Method for Efficient SSD Array Scaling .157.....	
<i>Hanchen Guo (Shanghai Jiao Tong University, Shanghai, China), Zhehan Lin (Shanghai Jiao Tong University, Shanghai, China), Yunfei Gu (Shanghai Jiao Tong University, Shanghai, China), Chentao Wu (Shanghai Jiao Tong University, Shanghai, China, Sichuan Research Institute, Shanghai Jiao Tong University, Sichuan, China), Li Jiang (Shanghai Jiao Tong University, Shanghai, China, Shanghai Qi Zhi Institute, Shanghai, China), Jie Li (Shanghai Jiao Tong University, Shanghai, China), Guangtao Xue (Shanghai Jiao Tong University, Shanghai, China), and Minyi Guo (Shanghai Jiao Tong University, Shanghai, China)</i>	
Streamlining Distributed Deep Learning I/O with ad hoc File Systems .169.....	
<i>Frederic Schimmelpfennig (Johannes Gutenberg University Mainz, Germany), Marc-André Vef (Johannes Gutenberg University Mainz, Germany), Reza Salkhordeh (Johannes Gutenberg University Mainz, Germany), Alberto Miranda (Barcelona Supercomputing Center, Spain), Ramon Nou (Barcelona Supercomputing Center, Spain), and André Brinkmann (Johannes Gutenberg University Mainz, Germany)</i>	

Session 3

Best Paper Candidates

Accelerating GPU Message Communication for Autonomous Navigation Systems .181.....	
<i>Hao Wu (TuSimple), Jiangming Jin (TuSimple), Jidong Zhai (Tsinghua University), Yifan Gong (TuSimple), and Wei Liu (TuSimple)</i>	
csTuner: Scalable Auto-Tuning Framework for Complex Stencil Computation on GPUs .192.....	
<i>Qingxiao Sun (Beihang University, China), Yi Liu (Beihang University, China), Hailong Yang (Beihang University, China), Zhonghui Jiang (Beihang University, China), Xiaoyan Liu (Beihang University, China), Ming Dun (Beihang University, China), Zhongzhi Luan (Beihang University, China), and Depei Qian (Beihang University, China)</i>	

Parallel Sessions

Track A: Applications

Octo-Tiger's New Hydro Module and Performance using HPX+CUDA on ORNL's Summit .204....	
<i>Patrick Diehl (Louisiana State University, USA), Gregor Daiß (University of Stuttgart, Germany), Dominic Marcello (Louisiana State University, USA), Kevin Huck (University of Oregon, USA), Sagiv Shiber (Louisiana State University, USA), Hartmut Kaiser (Louisiana State University, USA), Juhan Frank (Louisiana State University, USA), Geoffrey C. Clayton (Louisiana State University, USA), and Dirk Pflüger (University of Stuttgart, Germany)</i>	
Pipelined Preconditioned s-Step Conjugate Gradient Methods for Distributed Memory Systems .215	
<i>Manasi Tiwari (Indian Institute of Science, India) and Sathish Vadhiyar (Indian Institute of Science, India)</i>	

Thrifty Label Propagation: Fast Connected Components for Skewed-Degree Graphs .226.....
*Mohsen Koochi Esfahani (Queen’s University Belfast, United Kingdom),
Peter Kilpatrick (Queen’s University Belfast, United Kingdom), and
Hans Vandierendonck (Queen’s University Belfast, United Kingdom)*

Track B: Workloads

Optimizing Distributed Load Balancing for Workloads with Time-Varying Imbalance .238.....
*Jonathan Lifflander (Sandia National Laboratories, USA), Nicole
Lemaster Slattengren (Sandia National Laboratories, USA), Philippe P.
Pebay (NexGen Analytics, USA), Phil Miller (Intense Computing, USA),
Francesco Rizzi (NexGen Analytics, USA), and Matthew T. Bettencourt
(NexGen Analytics, USA)*

Distributed Work Stealing at Scale via Matchmaking .250.....
*Hrushit Parikh (Georgia Institute of Technology), Vinit Deodhar
(Georgia Institute of Technology), Ada Gavrilovska (Georgia Institute
of Technology), and Santosh Pande (Georgia Institute of Technology)*

Bellamy: Reusing Performance Models for Distributed Dataflow Jobs Across Contexts .261.....
*Dominik Scheinert (TU Berlin, Germany), Lauritz Thamsen (TU Berlin,
Germany), Houkun Zhu (TU Berlin, Germany), Jonathan Will (TU Berlin,
Germany), Alexander Acker (TU Berlin, Germany), Thorsten Wittkopp (TU
Berlin, Germany), and Odej Kao (TU Berlin, Germany)*

Session 4

Track A: Compression and data reduction

CSwap: A Self-Tuning Compression Framework for Accelerating Tensor Swapping in GPUs .271...
*Ping Chen (Zhejiang University, China), Shuibing He (Zhejiang
University, Illinois Institute of Technolog, China), Xuechen Zhang
(Washington State University, USA), Shuaiben Chen (Zhejiang
University, China), Peiyi Hong (Zhejiang University, China), Yanlong
Yin (Zhejiang Lab, China), Xian-He Sun (Illinois Institute of
Technology, USA), and Gang Chen (Zhejiang University, China)*

Optimizing Error-Bounded Lossy Compression for Scientific Data on GPUs .283.....
*Jiannan Tian (Washington State University, USA), Sheng Di (Argonne
National Laboratory, USA), Xiaodong Yu (Argonne National Laboratory,
USA), Cody Rivera (The University of Alabama, USA), Kai Zhao
(University of California, USA), Sian Jin (Washington State
University, USA), Yunhe Feng (University of Washington, USA), Xin
Liang (Missouri University of Science and Technology, USA), Dingwen
Tao (Washington State University, USA), and Franck Cappello (Argonne
National Laboratory, USA)*

Exploring Autoencoder-Based Error-Bounded Compression for Scientific Data .294.....
Jinyang Liu (University of California, Riverside, USA), Sheng Di (Argonne National Laboratory, USA), Kai Zhao (University of California, Riverside, USA), Sian Jin (Washington State University, USA), Dingwen Tao (Washington State University, USA), Xin Liang (Missouri University of Science and Technology, USA), Zizhong Chen (University of California, Riverside, USA), and Franck Cappello (Argonne National Laboratory, USA, University of Illinois, USA)

cuZ-Checker: A GPU-Based Ultra-Fast Assessment System for Lossy Compressions .307.....
Xiaodong Yu (Argonne National Laboratory), Sheng Di (Argonne National Laboratory), Ali Murat Gok (Cerebras Systems), Dingwen Tao (Washington State University), and Franck Cappello (Argonne National Laboratory)

DPZ: Improving Lossy Compression Ratio with Information Retrieval on Scientific Data .320.....
Jialing Zhang (University of Massachusetts Lowell, USA), Jiayi Chen (University of Massachusetts Lowell, USA), Xiaoyan Zhuo (University of Massachusetts Lowell, USA), Aekyeung Moon (ETRI, South Korea), and Seung Woo Son (University of Massachusetts Lowell, USA)

O(1) Communication for Distributed SGD Through Two-Level Gradient Averaging .332.....
Subhadeep Bhattacharya (Florida State University), Weikuan Yu (Florida State University), Fahim Tahmid Chowdhury (Florida State University), and Kathryn Mohror (Lawrence Livermore National Laboratory)

Track B: Systems support for parallel applications

Distributed Computation of Persistent Homology from Partitioned Big Data .344.....
Nicholas O. Malott (University of Cincinnati, USA), Rishi R. Verma (University of Cincinnati, USA), Rohit P. Singh (University of Cincinnati, USA), and Philip A. Wilsey (University of Cincinnati, USA)

FineQuery: Fine-Grained Query Processing on CPU-GPU Integrated Architectures .355.....
Dalin Wang (Renmin University of China, China), Feng Zhang (Renmin University of China, China), Weitao Wan (Renmin University of China, China), Hourun Li (Renmin University of China, China), and Xiaoyong Du (Renmin University of China, China)

Packet Forwarding Cache of Commodity Switches for Parallel Computers .366.....
Shoichi Hirasawa (National Institute of Informatics; The Graduate University for Advanced Studies), Hayato Yamaki (The University of Electro-Communications), and Michihiro Koibuchi (National Institute of Informatics; The Graduate University for Advanced Studies)

Two-Chains: High Performance Framework for Function Injection and Execution .377.....
Megan Grodowitz (Arm Research, USA), Luis E. Peña (Arm Research, USA), Curtis Dunham (Arm Research, USA), Dong Zhong (The University of Tennessee, USA), Pavel Shamis (Arm Research, USA), and Steve Poole (Los Alamos National Laboratory, USA)

HNGraph: Parallel Graph Processing in Hybrid Memory Based NUMA Systems .388.....	
	<i>Wei Liu (Huazhong University of Science and Technology, China), Haikun Liu (Huazhong University of Science and Technology, China), Xiaofei Liao (Huazhong University of Science and Technology, China), Hai Jin (Huazhong University of Science and Technology, China), and Yu Zhang (Huazhong University of Science and Technology, China)</i>
Modeling the Linux page Cache for Accurate Simulation of Data-Intensive Applications .398.....	
	<i>Hoang-Dung Do (Concordia University, Canada), Valérie Hayot-Sasson (Concordia University, Canada), Rafael Ferreira da Silva (University of Southern California, USA), Christopher Steele (Concordia University, Canada), Henri Casanova (University of Hawai'i at Manoa, USA), and Tristan Glatard (Concordia University, Canada)</i>

Session 5

Track A: Impacts of Errors on Applications

Characterizing Impacts of Storage Faults on HPC Applications: A Methodology and Insights .409..	
	<i>Bo Fang (Pacific Northwest National Laboratory, USA), Daoce Wang (Washington State University, USA), Sian Jin (Washington State University, USA), Quincey Koziol (Lawrence Berkeley National Laboratory, USA), Zhao Zhang (Texas Advanced Computing Center, USA), Qiang Guan (Kent State University, USA), Suren Byna (Lawrence Berkeley National Laboratory, USA), Sriram Krishnamoorthy (Pacific Northwest National Laboratory, USA), and Dingwen Tao (Washington State University, USA)</i>
Understanding the Effects of DRAM Correctable Error Logging at Scale .421.....	
	<i>Kurt Ferreira (Sandia National Laboratories), Scott Levy (Sandia National Laboratories), Victor Kuhns (Sandia National Laboratories), Nathan DeBardleben (Los Alamos National Laboratory), and Sean Blanchard (Los Alamos National Laboratory)</i>

Track B: Supporting Applications

Tackling Cold Start of Serverless Applications by Efficient and Adaptive Container Runtime Reusing .433.....	
	<i>Kun Suo (Kennesaw State University), Junggab Son (Kennesaw State University), Dazhao Cheng (University of North Carolina at Charlotte), Wei Chen (Nvidia), and Sabur Baidya (University of Louisville)</i>
Reusability First: Toward FAIR Workflows .444.....	
	<i>Matthew Wolf (Oak Ridge National Laboratory), Jeremy Logan (Oak Ridge National Laboratory), Kshitij Mehta (Oak Ridge National Laboratory), Daniel Jacobson (Oak Ridge National Laboratory, University of Tennessee), Mikaela Cashman (Oak Ridge National Laboratory), Angelica M. Walker (University of Tennessee), Greg Eisenhauer (Georgia Institute of Technology), Patrick Widener (Sandia National Laboratories), and Ashley Cliff (University of Tennessee)</i>

Session 6

Track A: Understanding system interactions with applications

- Thinking More About RDMA Memory Semantics .456.....
Teng Ma (Alibaba), Kang Chen (Tsinghua University), Shaonan Ma (Tsinghua University), Zhuo Song (Alibaba), and Yongwei Wu (Tsinghua University)
- Monitoring Large Scale Supercomputers: A Case Study with the Lassen Supercomputer .468.....
Tapasya Patki (Lawrence Livermore National Laboratory), Adam Bertsch (Lawrence Livermore National Laboratory), Ian Karlin (Lawrence Livermore National Laboratory), Dong H. Ahn (Lawrence Livermore National Laboratory), Brian Van Essen (Lawrence Livermore National Laboratory), Barry Rountree (Lawrence Livermore National Laboratory), Bronis R. de Supinski (Lawrence Livermore National Laboratory), and Nathan Besaw (IBM, USA)
- Robustness Analysis of Loop-Free Floating-Point Programs via Symbolic Automatic Differentiation .481.....
Arnab Das (University of Utah, USA), Tanmay Tirpankar (University of Utah, USA), Ganesh Gopalakrishnan (University of Utah, USA), and Sriram Krishnamoorthy (Google, USA)
- Understanding Soft Error Sensitivity of Deep Learning Models and Frameworks through Checkpoint Alteration .492.....
Elois Rojas (Costa Rica Institute of Technology, National University of Costa Rica), Diego Pérez (Costa Rica Institute of Technology), Jon C. Calhoun (Clemson University), Leonardo Bautista Gomez (Barcelona Supercomputing Center), Terry Jones (Oak Ridge National Laboratory), and Esteban Meneses (Costa Rica National High Technology Center, Costa Rica National High Technology Center)

Track B: Communication Optimization

- On-the-Fly, Robust Translation of MPI Libraries .504.....
Edgar A. León (Lawrence Livermore National Laboratory, USA), Marc Joos (CEA, DAM, DIF, France), Nathan Hanford (Lawrence Livermore National Laboratory, USA), Adrien Cotte (AS+ Groupe Eolen, France), Tony Delforge (AS+ Groupe Eolen, France), François Diakhaté (CEA, DAM, DIF, France), Vincent Ducrot (AS+ Groupe Eolen, France), Ian Karlin (Lawrence Livermore National Laboratory, USA), and Marc Pérache (CEA, DAM, DIF, France)
- Daps: A Dynamic Asynchronous Progress Stealing Model for MPI Communication .516.....
Kaiming Ouyang (University of California, Riverside, USA), Min Si (Argonne National Laboratory, USA), Astushi Hori (National Institute of Informatics), Zizhong Chen (University of California, Riverside, USA), and Pavan Balaji (Facebook, USA)
- Combining One-Sided Communications with Task-Based Programming Models .528.....
Kevin Sala (Barcelona Supercomputing Center (BSC), Spain), Sandra Macià (Barcelona Supercomputing Center (BSC), Spain), and Vicenç Beltran (Barcelona Supercomputing Center (BSC), Spain)

Optimizing Barrier Synchronization on ARMv8 Many-Core Architectures .542.....
*Wanrong Gao (National University of Defense Technology, China),
Jianbin Fang (National University of Defense Technology, China), Chun
Huang (National University of Defense Technology, China), Chuanfu Xu
(National University of Defense Technology, China), and Zheng Wang
(University of Leeds, United Kingdom)*

Energy Efficient HPC State of the Practice Workshop (EE HPC SOP 2021)

Evaluation of SPEC CPU and SPEC OMP on the A64FX .553.....
*Yuetsu Kodama (RIKEN, Japan), Masaaki Kondo (RIKEN, Japan), and
Mitsuhisa Sato (RIKEN, Japan)*

Energy Efficiency Aspects of the AMD Zen 2 Architecture .562.....
*Robert Schöne (Technische Universität Dresden, Germany), Thomas Ilsche
(Technische Universität Dresden, Germany), Mario Bielert (Technische
Universität Dresden, Germany), Markus Velten (Technische Universität
Dresden, Germany), Markus Schmidl (Technische Universität Dresden,
Germany), and Daniel Hackenberg (Technische Universität Dresden,
Germany)*

Explicit Uncore Frequency Scaling for Energy Optimisation Policies with EAR in Intel
Architectures .572.....
*Julita Corbalan (Barcelona Supercomputing Center, Spain), Oriol Vidal
(Barcelona Supercomputing Center, Spain), Lluís Alonso (Barcelona
Supercomputing Center, Spain), and Jordi Aneas (Barcelona
Supercomputing Center, Spain)*

FIRESTARTER 2: Dynamic Code Generation for Processor Stress Tests .582.....
*Robert Schöne (Technische Universität Dresden, Germany), Markus
Schmidl (Technische Universität Dresden, Germany), Mario Bielert
(Technische Universität Dresden, Germany), and Daniel Hackenberg
(Technische Universität Dresden, Germany)*

Cooling the Data Center: Design of a Mechanical Controls Owner Project Requirements (OPR)
Template .591.....
*Stefan Robila (Montclair State University, USA), David Grant (Oak
Ridge National Laboratory, USA), Chris DePrater (Lawrence Livermore
National Laboratory, USA), Vali Sorell (Microsoft, USA), Terry Rodgers
(JLL, USA), David Martinez (Sandia National Laboratory, USA), and
Shlomo Novotny (Chillydyne, USA)*

A Conceptual Framework for HPC Operational Data Analytics .596.....
*Alessio Netti (Leibniz Supercomputing Centre), Woong Shin (Oak Ridge
National Laboratory), Michael Ott (Leibniz Supercomputing Centre),
Torsten Wilde (Hewlett Packard Enterprise), and Natalie Bates (Energy
Efficient HPC Working Group)*

Workshop on Monitoring and Analysis for HPC Systems Plus Applications (HPCMASPA 2021)

- An Execution Fingerprint Dictionary for HPC Application Recognition .604.....
Thomas Jakobsche (University of Basel), Nicolas Lachiche (University of Strasbourg), Aurélien Cavelan (University of Basel), and Florina Ciorba (University of Basel)
- An Integrated Job Monitor, Analyzer and Predictor .609.....
Ashish Pal (IIT Kanpur) and Preeti Malakar (IIT Kanpur)
- Backfilling HPC Jobs with a Multimodal-Aware Predictor .618.....
Kenneth Lamar (University of Central Florida), Alexander Goponenko (University of Central Florida), Christina Peterson (University of Central Florida), Benjamin A. Allan (Sandia National Laboratories), Jim M. Brandt (Sandia National Laboratories), and Damian Dechev (University of Central Florida)
- Sequence-RTG: Efficient and Production-Ready Pattern Mining in System Log Messages .623.....
Louise Harding (CNRS, IN2P3 Computing Center, France), Fabien Wernli (CNRS, IN2P3 Computing Center, France), and Frédéric Suter (CNRS, IN2P3 Computing Center, France)
- The Challenge of Disproportionate Importance of Temporal Features in Predicting HPC Power Consumption .632.....
Chengcheng Li (The University of Tennessee), Ahmad M. Karimi (Oak Ridge National Laboratory), Woong Shin (Oak Ridge National Laboratory), Hairong Qi (The University of Tennessee), and Feiyi Wang (Oak Ridge National Laboratory)
- Dynamic and Adaptive Monitoring and Analysis for Many-Task Ensemble Computing .637.....
Shantenu Jha (Rutgers University) and Allen D. Malony (University of Oregon)

Workshop on Re-envisioning Extreme-Scale I/O for Emerging Hybrid HPC Workloads (REX-IO 2021)

Session I: I/O Optimization Technique

- A Scalability Study of Data Exchange in HPC Multi-component Workflows .642.....
Jie Yin (National Institute of Informatics, Japan), Atsushi Hori (National Institute of Informatics, Japan), Balazs Gerofi (RIKEN Center for Computational Science, Japan), and Yutaka Ishikawa (National Institute of Informatics, Japan)
- The Case for Storage Optimization Decoupling in Deep Learning Frameworks .649.....
Ricardo Macedo (INESC TEC & University of Minho), Cláudia Correia (INESC TEC & University of Minho), Marco Dantas (INESC TEC & University of Minho), Cláudia Brito (INESC TEC & University of Minho), Weijia Xu (Texas Advanced Computing Center), Yusuke Tanimura (National Institute of Advanced Industrial Science and Technology), Jason Haga (National Institute of Advanced Industrial Science and Technology), and João Paulo (INESC TEC & University of Minho)

MONARCH: Hierarchical Storage Management for Deep Learning Frameworks .657.....
Marco Dantas (INESC TEC & University of Minho), Diogo Leitão (INESC TEC & University of Minho), Cláudia Correia (INESC TEC & University of Minho), Ricardo Macedo (INESC TEC & University of Minho), Weijia Xu (Texas Advanced Computing Center), and João Paulo (INESC TEC & University of Minho)

Session II: I/O Libraries and Evaluation Methodologies

pMEMCPY: A Simple, Lightweight, and Portable I/O Library for Storing Data in Persistent Memory .664.....
Luke Logan (Illinois Institute of Technology and Sandia National Labs), Jay Lofstead (Illinois Institute of Technology and Sandia National Labs), Scott Levy (Illinois Institute of Technology and Sandia National Labs), Patrick Widener (Illinois Institute of Technology and Sandia National Labs), Xian-He Sun (Illinois Institute of Technology and Sandia National Labs), and Anthony Kougkas (Illinois Institute of Technology and Sandia National Labs)

Parallel I/O Evaluation Techniques and Emerging HPC Workloads: A Perspective .671.....
Sarah Neuwirth (Goethe-University Frankfurt, Germany) and Arnab K. Paul (Oak Ridge National Laboratory, USA)

Workshop on Artificial Intelligence and Machine Learning for Scientific Applications (AI4S 2021)

Special Function Neural Network (SFNN) Models .680.....
Yuzhen Liu (Argonne National Laboratory, USA) and Oana Marin (Argonne National Laboratory, USA)

AMR-Net: Convolutional Neural Networks for Multi-resolution Steady Flow Prediction .686.....
Yuuichi Asahi (Japan atomic energy agency, Japan), Sora Hatayama (The University of Tokyo, Japan), Takashi Shimokawabe (The University of Tokyo, Japan), Naoyuki Onodera (Japan Atomic Energy Agency, Japan), Yuta Hasegawa (Japan Atomic Energy Agency, Japan), and Yasuhiro Idomura (Japan Atomic Energy Agency, Japan)

A Deep Learning-Based Particle-in-Cell Method for Plasma Simulations .692.....
Xavier Aguilar (KTH Royal Institute of Technology, Sweden) and Stefano Markidis (KTH Royal Institute of Technology, Sweden)

Hybrid Workflow of Simulation and Deep Learning on HPC: A Case Study for Material Behavior Determination .698.....
Li Zhong (High Performance Computing Center Stuttgart, Germany), Dennis Hoppe (High Performance Computing Center Stuttgart, Germany), Naweiluo Zhou (High Performance Computing Center Stuttgart, Germany), and Oleksandr Shcherbakov (High Performance Computing Center Stuttgart, Germany)

Higgs Boson Classification: Brain-Inspired BCPNN Learning with StreamBrain .705.....
Martin Svedin (KTH Royal Institute of Technology, Sweden), Artur Podobas (KTH Royal Institute of Technology, Sweden), Steven W. D. Chien (KTH Royal Institute of Technology, Sweden), and Stefano Markidis (KTH Royal Institute of Technology, Sweden)

Embracing Arm for High Performance Computing Workshop (EAHPC 2021)

- A64FX Performance: Experience on Ookami .711.....
Md Abdullah Shahneous Bari (Stony Brook University, USA), Barbara Chapman (Stony Brook University, USA), Anthony Curtis (Stony Brook University, USA), Robert J. Harrison (Stony Brook University, USA), Eva Siegmann (Stony Brook University, USA), Nikolay A. Simakov (University at Buffalo, USA), and Matthew D. Jones (University at Buffalo, USA)
- Early Evaluation of Fugaku A64FX Architecture using Climate Workloads .719.....
Sarat Sreepathi (Oak Ridge National Laboratory, USA) and Mark Taylor (Sandia National Laboratory, USA)
- Performance Evaluation and Analysis of A64FX Many-Core Processor for the Fiber Miniapp Suite .728.....
Miwako Tsuji (RIKEN, Japan) and Mitsuhsa Sato (RIKEN, Japan)
- A64FX - Your Compiler You Must Decide! .736.....
Jens Domke (RIKEN Center for Computational Science (R-CCS), Japan)
- Cluster of Emerging Technology: Evaluation of a Production HPC System Based on A64FX .741....
Fabio Banchelli (Barcelona Supercomputing Center, Spain), Kilian Peiro (Barcelona Supercomputing Center, Spain), Guillem Ramirez-Gargallo (Barcelona Supercomputing Center, Spain), Joan Vinyals (Barcelona Supercomputing Center, Spain), David Vicente (Barcelona Supercomputing Center, Spain), Marta Garcia-Gasulla (Barcelona Supercomputing Center, Spain), and Filippo Mantovani (Barcelona Supercomputing Center, Spain)
- Sequences of Sparse Matrix-Vector Multiplication on Fugaku's A64FX Processors .751.....
Jérôme Gurhem (USR 3441 - Maison de la Simulation, CNRS, France), Maxence Vandromme (USR 3441 - Maison de la Simulation, CNRS, France), Miwako Tsuji (RIKEN Center for Computational Science, Japan), Serge G. Petiton (Univ. Lille\ CNRS, France), and Mitsuhsa Sato (RIKEN Center for Computational Science, Japan)

FPGA for HPC Workshop 2021 (HPC FPGA 2021)

- From Domain-Specific Languages to Memory-Optimized Accelerators for Fluid Dynamics .759.....
Karl F. A. Friebel (Technische Universität Dresden, Germany), Stephanie Soldavini (Politecnico di Milano, Italy), Gerald Hempel (Technische Universität Dresden, Germany), Christian Pilato (Politecnico di Milano, Italy), and Jeronimo Castrillon (Technische Universität Dresden, Germany)
- Accelerating Advection for Atmospheric Modelling on Xilinx and Intel FPGAs .767.....
Nick Brown (EPCC at the University of Edinburgh, UK)
- Optimisation of an FPGA Credit Default Swap Engine by Embracing Dataflow Techniques .775....
Nick Brown (EPCC at the University of Edinburgh, UK), Mark Klaisoongnoen (EPCC at the University of Edinburgh, UK), and Oliver Thomson Brown (EPCC at the University of Edinburgh, UK)

TIGRA: A Tightly Integrated Generic RISC-V Accelerator Interface .779.....	779
<i>Brad Green (Clemson University, United States), Dillon Todd (Clemson University, United States), Jon C. Calhoun (Clemson University, United States), and Melissa C. Smith (Clemson University, United States)</i>	
HBM2 Memory System for HPC Applications on an FPGA .783.....	783
<i>Norihisa Fujita (University of Tsukuba, Japan), Ryohei Kobayashi (University of Tsukuba, Japan), Yoshiki Yamaguchi (University of Tsukuba, Japan), and Taisuke Boku (University of Tsukuba, Japan)</i>	
A Memory Bandwidth Improvement with Memory Space Partitioning for Single-Precision Floating-Point FFT on Stratix 10 FPGA .787.....	787
<i>Takaaki Miyajima (Meiji University, Japan) and Kentaro Sano (RIKEN, Japan)</i>	
An FPGA-Based Storage Control with load Balancing .791.....	791
<i>Naoya Umezu (University of Tsukuba, Japan), Yoshiki Yamaguchi (University of Tsukuba, Japan), and Taisuke Boku (University of Tsukuba, Japan)</i>	

Poster Short Papers

CVFCC: CV-Based Framework for Container Consolidation in Cloud Data Centers .795.....	795
<i>Yuting Li (University of Science and Technology of China, China), Yun Xu (University of Science and Technology of China, China), and Xuehai Zhou (University of Science and Technology of China, China)</i>	
A Dynamic Power Capping Library for HPC Applications .797.....	797
<i>Sahil Sharma (Illinois Institute of Technology), Zhiling Lan (Illinois Institute of Technology), Xingfu Wu (Argonne National Laboratory), and Valerie Taylor (Argonne National Laboratory)</i>	
SDIS: A PB-Level Seismic Data Index System with ML Methods .799.....	799
<i>Shaoheng Luo (Beihang University, China), Lei Wang (Beihang University, China), Yufeng Liu (Beihang University, China), Changhai Zhao (China National Petroleum Corporation, China), and Xudong Zhang (China National Petroleum Corporation, China)</i>	
Malleability Implementation in a MPI Iterative Method .801.....	801
<i>Iker Martín-Álvarez (Universitat Jaume I, Spain), José I. Aliaga (Universitat Jaume I, Spain), María Isabel Castillo (Universitat Jaume I, Spain), Rafael Mayo (Universitat Jaume I, Spain), and Sergio Iserte (Universitat Jaume I, Spain)</i>	
Computational Storage to Increase the Analysis Capability of Tier-2 HEP Data Sites .803.....	803
<i>Chen Zou (University of Chicago, USA), Andrew A. Chien (University of Chicago, USA), Robert Gardner (University of Chicago, USA), and Ilija Vukotic (University of Chicago, USA)</i>	
NUMA-Aware I/O System Call Steering .805.....	805
<i>Chan-Gyu Lee (Konkuk University, Korea) and Hyun-Wook Jin (Konkuk University, Korea)</i>	

A Roadmap to Robust Science for High-throughput Applications: The Developers' Perspective	807
<i>Michela Taufer (The University of Tennessee), Ewa Deelman (University of Southern California), Rafael Ferreira da Silva (University of Southern California), Trilce Estrada (University of New Mexico), Mary Hall (University of Utah), and Miron Livny (U. Wisconsin–Madison)</i>	
A Transfer Learning Scheme for Time Series Forecasting using Facebook Prophet	809
<i>Menuka Warushavithana (Colorado State University, USA), Saptashwa Mitra (Colorado State University, USA), Mazdak Arabi (Colorado State University, USA), Jay Breidt (Colorado State University, USA), Sangmi Pallickara (Colorado State University, USA), and Shrideep Pallickara (Colorado State University, USA)</i>	
Exploring Node Connection Modes in Multi-rail Fat-Tree	811
<i>Yuyang Wang (National University of Defense Technology, China), Fei Lei (National University of Defense Technology, China), and Dezun Dong (National University of Defense Technology, China)</i>	
RELAR: A Reinforcement Learning Framework for Adaptive Routing in Network-on-Chips	813
<i>Changhong Wang (National University of Defense Technology, China), Dezun Dong (National University of Defense Technology, China), Zicong Wang (National University of Defense Technology, China), Xiaoyun Zhang (National University of Defense Technology, China), and Zhenyu Zhao (National University of Defense Technology, China)</i>	
A Generative Approach to Visualizing Satellite Data	815
<i>Saptashwa Mitra (Colorado State University, USA), Daniel Rammer (Colorado State University, USA), Shrideep Pallickara (Colorado State University, USA), and Sangmi Pallickara (Colorado State University, USA)</i>	
Load Balancing Policies for Nested Fork-Join	817
<i>Lukas Reitz (University of Kassel, Germany)</i>	
Supporting Elastic Compaction of LSM-Tree with a FaaS Cluster	819
<i>Xiaoliang Wang (University of Science and Technology of China, China), Jianchuan Li (University of Science and Technology of China, China), Peiquan Jin (University of Science and Technology of China, China), Kuankuan Guo (ByteDance Inc., China), Yuanjin Lin (ByteDance Inc., China), and Ming Zhao (ByteDance Inc., China)</i>	
Automatic Parallelisation of Structured Mesh Computations with SYCL	821
<i>Gábor Dániel Balogh (Pázmány Péter Catholic University, Hungary) and István Reguly (Pázmány Péter Catholic University, Hungary)</i>	
Halcyon: Unified HPC Center Operations	823
<i>Kevin D. Colby (Purdue University) and Shawn Rice (Purdue University)</i>	
CASQ: Accelerate Distributed Deep Learning with Sketch-Based Gradient Quantization	825
<i>Keshi Ge (National University of Defense Technology, China), Yiming Zhang (National University of Defense Technology, China), Yongquan Fu (National University of Defense Technology, China), Zhiquan Lai (National University of Defense Technology, China), Xiaoge Deng (National University of Defense Technology, China), and DongSheng Li (National University of Defense Technology, China)</i>	
Toward a Comprehensive Benchmark Suite for Evaluating GASPI in HPC Environments	827
<i>Sarah Neuwirth (Goethe-University Frankfurt, Germany)</i>	
Incorporating Fault-Tolerance Awareness into System-Level Modeling and Simulation	829
<i>Trokon Johnson (University of Florida, USA) and Herman Lam (The University of Florida, USA)</i>	
Author Index	831