

2022 IEEE International Conference on Cluster Computing (CLUSTER 2022)

**Heidelberg, Germany
6 – 9 September 2022**



**IEEE Catalog Number: CFP22235-POD
ISBN: 978-1-6654-9857-9**

**Copyright © 2022 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:	CFP22235-POD
ISBN (Print-On-Demand):	978-1-6654-9857-9
ISBN (Online):	978-1-6654-9856-2
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

2022 IEEE International Conference on Cluster Computing (CLUSTER) **CLUSTER 2022**

Table of Contents

Welcome Message from the IEEE Cluster 2022 General Co-Chairs	xvi
Welcome Message from the Cluster 2022 Program Chairs	xviii
IEEE Cluster 2022 Committees	xix
Welcome Message from HPCEuropeLatAm 2022	xxv
HPCEuropeLatAm 2022 Committees	xxvi
Welcome Message from EA-HPC 2022 Organizers	xxvii
EA-HPC 2022 Committees	xxviii
Welcome Message from REX-IO 2022	xxix
REX-IO 2022 Committees	xxx
Welcome Message from HPCMASPA 2022	xxxi
HPCMASPA 2022 Committees	xxxii

Papers

Networking and Security

SKV: A SmartNIC-Offloaded Distributed Key-Value Store	1
<i>Shangyi Sun (Fudan University, China), Rui Zhang (Fudan University, China), Ming Yan (Fudan University, China), and Jie Wu (Fudan University, China)</i>	
Bring the BitCODE — Moving Compute and Data in Distributed Heterogeneous Systems	12
<i>Wenbin Lu (Stony Brook University, USA), Luis E. Peña (Arm Research, USA), Pavel Shamis (Arm Research, USA), Valentin Churavy (MIT, USA), Barbara Chapman (Stony Brook University, USA), and Steve Poole (Los Alamos National Laboratory, USA)</i>	
Exploring Light-Weight Cryptography for Efficient and Secure Lossy Data Compression	23
<i>Ruiwen Shan (Clemson University, USA), Sheng Di (Argonne National Laboratory, USA), Jon C. Calhoun (Clemson University, USA), and Frank Cappello (Argonne National Laboratory, USA; University of Illinois at Urbana-Champaign, USA)</i>	

Scheduling and Multi-tenancy

What does Inter-Cluster Job Submission and Execution Behavior Reveal to Us?	35
<i>Tirthak Patel (Northeastern University, USA), Devesh Tiwari (Northeastern University, USA), Raj Kettimuthu (Argonne National Laboratory, USA), William Allcock (Argonne National Laboratory, USA), Paul Rich (Argonne National Laboratory, USA), and Zhengchun Liu (Argonne National Laboratory, USA)</i>	
MRSch: Multi-Resource Scheduling for HPC	47
<i>Boyang Li (Illinois Institute of Technology, USA), Yuping Fan (Illinois Institute of Technology, USA), Matthew Dearing (Illinois Institute of Technology, USA), Zhiling Lan (Illinois Institute of Technology, USA), Paul Rich (Argonne National Laboratory, USA), William Allcock (Argonne National Laboratory, USA), and Michael Papka (Argonne National Laboratory, USA; Northern Illinois University, USA)</i>	
Matching-Based Scheduling of Asynchronous Data Processing Workflows on the Computing Continuum	58
<i>Narges Mehran (Alpen-Adria-Universität Klagenfurt, Austria), Zahra Najafabadi Samani (Alpen-Adria-Universität Klagenfurt, Austria), Dragi Kimovski (Alpen-Adria-Universität Klagenfurt, Austria), and Radu Prodan (Alpen-Adria-Universität Klagenfurt, Austria)</i>	

MPI

Spark Meets MPI: Towards High-Performance Communication Framework for Spark Using MPI ...	71
<i>Kinan Al-Attar (The Ohio State University), Aamir Shafi (The Ohio State University), Mustafa Abduljabbar (The Ohio State University), Hari Subramoni (The Ohio State University), and Dhabaleswar K. Panda (The Ohio State University)</i>	
Deadlock Detection for MPI Programs Based on Refined Match-Sets	82
<i>Shushan Li (Hebei University, China), Meng Wang (Hebei University, China), and Hong Zhang (Hebei University, China)</i>	
A Framework for Hierarchical Single-Copy MPI Collectives on Multicore Nodes	94
<i>George Katevenis (Institute of Computer Science (ICS), Foundation for Research and Technology – Hellas (FORTH), Greece), Manolis Ploumidis (Institute of Computer Science (ICS), Foundation for Research and Technology – Hellas (FORTH), Greece), and Manolis Marazakis (Institute of Computer Science (ICS), Foundation for Research and Technology – Hellas (FORTH), Greece)</i>	

Runtimes

Pythia: An Oracle to Guide Runtime System Decisions	106
<i>Alexis Colin (Télécom SudParis, Institut Polytechnique de Paris, France), François Trahay (Télécom SudParis, Institut Polytechnique de Paris, France), and Denis Conan (Télécom SudParis, Institut Polytechnique de Paris, France)</i>	

Pushing the Boundaries of Small Tasks: Scalable Low-Overhead Data-Flow Programming in TTG .	117
<i>Joseph Schuchart (The University of Tennessee, USA), Poornima Nookala (Stony Brook University, USA), Thomas Herault (University of Tennessee, USA), Edward F. Valeev (Virginia Polytechnic Institute and State University, USA), and George Bosilca (University of Tennessee, USA)</i>	
Distributed Continuation Stealing is More Scalable than You Might Think	129
<i>Shumpei Shiina (The University of Tokyo, Japan) and Kenjiro Taura (The University of Tokyo, Japan)</i>	

MPI Collectives

Fast(er) Construction of Round-Optimal n-Block Broadcast Schedules	142
<i>Jesper Larsson Träff (TU Wien, Austria)</i>	
Lossy all-to-all Exchange for Accelerating Parallel 3-D FFTs on Hybrid Architectures with GPUs	152
<i>Sébastien Cayrols (University of Tennessee Knoxville, USA), Jiali Li (University of Tennessee Knoxville, USA), George Bosilca (University of Tennessee Knoxville, USA), Stanimire Tomov (University of Tennessee Knoxville, USA), Alan Ayala (University of Tennessee Knoxville, USA), and Jack Dongarra (University of Tennessee Knoxville, USA; Oak Ridge National Laboratory, USA)</i>	
ACCLAiM: Advancing the Practicality of MPI Collective Communication Autotuning Using Machine Learning	161
<i>Michael Wilkins (Northwestern University), Yanfei Guo (Argonne National Laboratory), Rajeev Thakur (Argonne National Laboratory), Peter Dinda (Northwestern University), and Nikos Hardavellas (Northwestern University)</i>	

Serverless and Virtual Networks

Call Scheduling to Reduce Response Time of a FaaS System	172
<i>Paweł Żuk (University of Warsaw, Poland), Bartłomiej Przybylski (University of Warsaw, Poland), and Krzysztof Rządca (University of Warsaw, Poland)</i>	
FaaS: Optimize Makespan of Serverless Workflows in Federated Commercial FaaS	183
<i>Sashko Ristov (University of Innsbruck, Austria) and Philipp Gritsch (University of Innsbruck, Austria)</i>	
Last-Mile Matters: Mitigating the Tail Latency of Virtualized Networks with Multipath Data Plane	195
<i>Dian Shen (Southeast University, China), Yi Zhai (Southeast University, China), Fang Dong (Southeast University, China), and Junzhou Luo (Southeast University, China)</i>	

Applications

Towards Virtual Certification of Gas Turbine Engines with Performance-Portable Simulations.....	206
<i>Gihan R. Mudalige (University of Warwick, UK), Istvan Z. Reguly (Pazmany Peter Catholic University, Hungary), Arun Prabhakar (University of Warwick, UK), Dario Amirante (University of Surrey, UK), Leigh Lapworth (Rolls-Royce plc., UK), and Stephen A. Jarvis (University of Birmingham, UK)</i>	
Hybrid Analysis of Fusion Data for Online Understanding of Complex Science on Extreme Scale Computers	218
<i>Eric Suchyta (Oak Ridge National Laboratory, USA), Jong Youl Choi (Oak Ridge National Laboratory, USA), Seung-Hoe Ku (Princeton Plasma Physics Laboratory, USA), David Pugmire (Oak Ridge National Laboratory, USA), Ana Gainaru (Oak Ridge National Laboratory, USA), Kevin Huck (University of Oregon, USA), Ralph Kube (Princeton Plasma Physics Laboratory, USA), Aaron Scheinberg (Jubilee Development), Frédéric Suter (Oak Ridge National Laboratory, USA), Choongseock Chang (Princeton Plasma Physics Laboratory, USA), Todd Munson (Argonne National Laboratory, USA), Norbert Podhorzski (Oak Ridge National Laboratory, USA), and Scott Klasky (Oak Ridge National Laboratory, USA)</i>	
High Performance Adaptive Physics Refinement to Enable Large-Scale Tracking of Cancer Cell Trajectory	230
<i>Daniel F. Puleri (Duke University, USA), Sayan Roychowdhury (Duke University, USA), Peter Balogh (New Jersey Institute of Technology, USA), John Gounley (Oak Ridge National Laboratory, USA), Erik W. Draeger (Lawrence Livermore National Laboratory, USA), Jeff Ames (Duke University, USA), Adebayo Adebisi (Duke University, USA), Simbarashe Chidyagwai (Duke University, USA), Benjamín Hernández (Oak Ridge National Laboratory, USA), Seyong Lee (Oak Ridge National Laboratory, USA), Shirley V. Moore (University of Texas at El Paso, USA), Jeffrey S. Vetter (Oak Ridge National Laboratory, USA), and Amanda Randles (Duke University, USA)</i>	

I/O

Extracting and Characterizing I/O Behavior of HPC Workloads	243
<i>Hariharan Devarajan (Lawrence Livermore National Laboratory, USA) and Kathryn Mohror (Lawrence Livermore National Laboratory, USA)</i>	
Be SMART, Save I/O: A Probabilistic Approach to Avoid Uncorrectable Errors in Storage Systems	256
<i>Md Arifuzzaman (University of Nevada, USA), Masudul Bhuiyan (CISPA, Germany), Mehmet Gumus (University of Nevada, USA), and Engin Arslan (University of Nevada, USA)</i>	
The Role of Storage Target Allocation in Applications' I/O Performance with BeeGFS	267
<i>Francieli Boito (Univ. Bordeaux, France), Guillaume Pallez (Univ. Bordeaux, France), and Luan Teylo (Univ. Bordeaux, France)</i>	

Best Paper Nominees

ecoHMEM: Improving Object Placement Methodology for Hybrid Memory Systems in HPC	278
<i>Marc Jordà (Barcelona Supercomputing Center), Siddharth Rai (Barcelona Supercomputing Center), Eduard Ayguadé (Barcelona Supercomputing Center), Jesús Labarta (Barcelona Supercomputing Center), and Antonio J. Peña (Barcelona Supercomputing Center)</i>	
Efficient Hierarchical State Vector Simulation of Quantum Circuits via Acyclic Graph Partitioning	289
<i>Bo Fang (Pacific Northwest National Laboratories, USA), M. Yusuf Özkaya (Georgia Institute of Technology, USA), Ang Li (Pacific Northwest National Laboratories, USA), Ümit V. Çatalyürek (Georgia Institute of Technology, USA), and Sriram Krishnamoorthy (Pacific Northwest National Laboratories, USA)</i>	

Deep Learning

AutoPipe: A Fast Pipeline Parallelism Approach with Balanced Partitioning and Micro-Batch Slicing	301
<i>Weijie Liu (National University Of Defense Technology, China), Zhiquan Lai (National University Of Defense Technology, China), Shengwei Li (National University Of Defense Technology, China), Yabo Duan (National University Of Defense Technology, China), Keshi Ge (National University Of Defense Technology, China), and Dongsheng Li (National University Of Defense Technology, China)</i>	
HPH: Hybrid Parallelism on Heterogeneous Clusters for Accelerating Large-Scale DNNs Training	313
<i>Yabo Duan (National University of Defence Technology, China), Zhiquan Lai (National University of Defence Technology, China), Shengwei Li (National University of Defence Technology, China), Weijie Liu (National University of Defence Technology, China), Keshi Ge (National University of Defence Technology, China), Peng Liang (National University of Defence Technology, China), and Dongsheng Li (National University of Defence Technology, China)</i>	
HVAC: Removing I/O Bottleneck for Large-Scale Deep Learning Applications	324
<i>Azwaiz Khan (Oak Ridge National Laboratory, USA), Arnab K. Paul (BITS Pilani, K K Birla Goa Campus, India), Christopher Zimmer (Oak Ridge National Laboratory, USA), Sarp Oral (Oak Ridge National Laboratory, USA), Sajal Dash (Oak Ridge National Laboratory, USA), Scott Atchley (Oak Ridge National Laboratory, USA), and Feiyi Wang (Oak Ridge National Laboratory, USA)</i>	

Node Technologies

Enabling Dynamic Virtual Frequency Scaling for Virtual Machines in the Cloud	336
<i>Emile Cadorel (Inria / Univ. Lille, France) and Romain Rouvoy (Univ. Lille / Inria, France)</i>	

The Cost of Flexibility: Embedded Versus Discrete Routers in CGRAs for HPC	347
<i>Boma Adhi (Center for Computational Science (R-CCS), RIKEN, Japan), Carlos Cortes (Center for Computational Science (R-CCS), RIKEN, Japan), Yiyu Tan (Iwate University, Japan), Takuya Kojima (Center for Computational Science (R-CCS), RIKEN, Japan; The University of Tokyo, Japan), Artur Podobas (KTH Royal Institute of Technology, Sweden), and Kentaro Sano (Center for Computational Science (R-CCS), RIKEN, Japan)</i>	
SVAGC: Garbage Collection with a Scalable Virtual Address Swapping Technique	357
<i>Ismail Ataie (Florida State University, USA) and Weikuan Yu (Florida State University, USA)</i>	

Operations and ML Training Strategies

ALBADross: Active Learning Based Anomaly Diagnosis for Production HPC Systems	369
<i>Burak Aksar (Boston University, USA), Efe Sencan (Boston University, USA), Benjamin Schwaller (Sandia National Laboratories, USA), Omar Aaziz (Sandia National Laboratories, USA), Vitus J. Leung (Sandia National Laboratories, USA), Jim Brandt (Sandia National Laboratories, USA), Brian Kulis (Boston University, USA), and Ayse K. Coskun (Boston University, USA)</i>	
HPC Storage Service Autotuning Using Variational-Autoencoder-Guided Asynchronous Bayesian Optimization	381
<i>Matthieu Dorier (Argonne National Laboratory, USA), Romain Egele (Argonne National Laboratory, USA; Université Paris-Saclay, France), Prasanna Balaprakash (Argonne National Laboratory, USA), Jaehoon Koo (Argonne National Laboratory, USA), Sandeep Madireddy (Argonne National Laboratory, USA), Srinivasan Ramesh (University of Oregon, USA), Allen D. Malony (University of Oregon, USA), and Rob Ross (Argonne National Laboratory, USA)</i>	
fairDMS: Rapid Model Training by Data and Model Reuse	394
<i>Ahsan Ali (Argonne National laboratory, USA), Hemant Sharma (Argonne National laboratory, USA), Rajkumar Kettimuthu (Argonne National laboratory, USA), Peter Kenesei (Argonne National laboratory, USA), Dennis Trujillo (Argonne National laboratory, USA), Antonino Miceli (Argonne National laboratory, USA), Ian Foster (Argonne National laboratory, USA; University of Chicago, USA), Ryan Coffee (SLAC National Accelerator Laboratory, USA), Jana Thayer (SLAC National Accelerator Laboratory, USA), and Zhengchun Liu (Argonne National laboratory, USA; University of Chicago, USA)</i>	

Distributed Memory Applications

Painless Transposition of Reproducible Distributed Environments with NixOS Compose	406
<i>Quentin Guilloteau (Univ. Grenoble Alpes, France), Jonathan Bleuzen (Univ. Grenoble Alpes, France), Millian Poquet (Univ. Grenoble Alpes, France), and Olivier Richard (Univ. Grenoble Alpes, France)</i>	

Integrating Process, Control-flow, and Data Resiliency Layers Using a Hybrid Fenix/Kokkos Approach	418
<i>Matthew Whitlock (Sandia National Laboratories, USA; Georgia Institute of Technology, USA), Nicolas Morales (Sandia National Laboratories, USA), George Bosilca (University of Tennessee, Knoxville, USA), Aurelien Bouteiller (University of Tennessee, Knoxville, USA), Bogdan Nicolae (Argonne National Laboratory, USA), Keita Teranishi (Sandia National Laboratories, USA), Elisabeth Giem (University of California, Riverside, Sandia National Laboratories, USA), and Vivek Sarkar (Georgia Institute of Technology, USA)</i>	
Fast Dynamic Updates and Dynamic SpGEMM on MPI-Distributed Graphs	429
<i>Alexander van der Grinten (Humboldt-Universität zu Berlin, Germany), Geert Custers (Delft University of Technology, Netherlands), Duy Le Thanh (Humboldt-Universität zu Berlin, Germany), and Henning Meyerhenke (Humboldt-Universität zu Berlin, Germany)</i>	

Tensors and Linear Algebra

BALA-CPD: BALanced and Asynchronous Distributed Tensor Decomposition	440
<i>Zheng Miao (Hangzhou Dianzi University), Jiajia Li (North Carolina State University), Jon C. Calhoun (Clemson University), and Rong Ge (Clemson University)</i>	
Optimizing Irregular-Shaped Matrix-Matrix Multiplication on Multi-core DSPs	451
<i>Shangfei Yin (National University of Defense Technology, China), Qinglin Wang (National University of Defense Technology, China), Ruochen Hao (National University of Defense Technology, China), Tianyang Zhou (National University of Defense Technology, China), Songzhu Mei (National University of Defense Technology, China), and Jie Liu (National University of Defense Technology, China)</i>	
Optimizations of H-Matrix-Vector Multiplication for Modern Multi-core Processors	462
<i>Tetsuya Hoshino (The University of Tokyo, Tokyo), Akihiro Ida (Center for Earth Information Science and Technology Japan Agency for Marine-Earth Science and Technology, Japan), and Toshihiro Hanawa (The University of Tokyo, Japan)</i>	

Shared Memory

Recursive Multi-section on the Fly: Shared-Memory Streaming Algorithms for Hierarchical Graph Partitioning and Process Mapping	473
<i>Marcelo Fonseca Faraj (Heidelberg University, Germany) and Christian Schulz (Heidelberg University, Germany)</i>	
MemGaze: Rapid and Effective Load-Level Memory Trace Analysis	484
<i>Ozgur O. Kilic (Pacific Northwest National Laboratory, USA), Nathan R. Tallent (Pacific Northwest National Laboratory, USA), Yasodha Suriyakumar (Portland State University, USA), Chenhao Xie (Pacific Northwest National Laboratory, USA), Andrés Marquez (Pacific Northwest National Laboratory, USA), and Stephane Eranian (Google Inc.)</i>	

Poster Papers

Empirical Study on the GPU-Accelerated HPL Performance: Effects of PCIe Communication	496
<i>Jieun Choi (Korea Institute of Science and Technology Information, Republic of Korea), Yosang Jeong (Korea Institute of Science and Technology Information, Republic of Korea), Ji-Hoon Kang (Korea Institute of Science and Technology Information, Republic of Korea), Gibeom Gu (Korea Institute of Science and Technology Information, Republic of Korea), and Hoon Ryu (Korea Institute of Science and Technology Information, Republic of Korea)</i>	
H2M: Towards Heuristics for Heterogeneous Memory	498
<i>Clément Foyer (Inria, Univ. Bordeaux, LaBRI, France), Brice Goglin (Inria, Univ. Bordeaux, LaBRI, France), Emmanuel Jeannot (Inria, Univ. Bordeaux, LaBRI, France), Jannis Klinkenberg (RWTH Aachen University, Germany), Anara Kozhokanova (RWTH Aachen University, Germany), and Christian Terboven (RWTH Aachen University, Germany)</i>	
An Analysis of Performance Variability on Dragonfly+ Topology	500
<i>Majid Salimi Beni (University of Salerno, Italy) and Biagio Cosenza (University of Salerno, Italy)</i>	
An Asynchronous Parallel Algorithm to Improve the Scalability of Finite Element Solvers	502
<i>Zhuo Tian (Laboratory of Parallel Software and Computational Science Institute of Software Chinese Academy of Sciences, China) and Changyou Zhang (Laboratory of Parallel Software and Computational Science Institute of Software Chinese Academy of Sciences, China)</i>	
An Efficient Sparse CNNs Accelerator on FPGA	504
<i>Yonghua Zhang (Beihang University, China), Hongxu Jiang (Beihang University, China), Xiaobin Li (Beihang University, China), Haojie Wang (Tsinghua University, China), Dong Dong (Beihang University, China), and Yongxiang Cao (Beihang University, China)</i>	

Workshops

HPC for International Collaboration Between Europe and Latin America: HPCEuropeLatAm2022

A Performance Evaluation of Adaptive MPI for a Particle-in-Cell Code	506
<i>Christian Asch (National High Technology Center, Costa Rica), Diego Jiménez (National High Technology Center, Costa Rica), Markus Rampp (Max Planck Computing and Data Facility, Germany), Erwin Laure (Max Planck Computing and Data Facility, Germany), and Esteban Meneses (National High Technology Center and Costa Rica Institute of Technology, Costa Rica)</i>	

Scalable Architectures to Support Sustainable Advanced Information Technologies	512
<i>Oscar Carrillo (Univ Lyon, CPE, INSA Lyon, France), Carlos Jaime Barrios Hernández (Universidad Industrial de Santander, Colombia), Frédéric Le Mouël (Univ Lyon, INSA Lyon, Inria, France), Harold E. Castro Barrera (Universidad de los Andes, Colombia), Yves Denneulin (Université de Grenoble-Alpes, France), José Tiberio Hernández (Universidad de los Andes, Colombia), Fernando Jiménez Vargas (Universidad de los Andes, Colombia), Lola Xiomara Bautista Roza (Universidad Industrial de Santander, Colombia), Claudia Roncancio (Université de Grenoble Alpes, France), and Michel Riveill (Université Côte d'Azur, France)</i>	
Early Experiences of Noise-Sensitivity Performance Analysis of a Distributed Deep Learning Framework	516
<i>Elvis Rojas (Costa Rica National High Technology; National University of Costa Rica), Michael Knobloch (Jülich Supercomputing Centre, Germany), Nour Daoud (Jülich Supercomputing Centre, Germany), Esteban Meneses (Costa Rica National High Technology; Costa Rica National High Technology Center), and Bernd Mohr (Jülich Supercomputing Centre, Germany)</i>	
Learning Tenant Behavior and Evolutionary Approach for Demand Response in Colocation Datacenters	523
<i>Jonathan Muraña (Universidad de la República, Uruguay), Santiago Iturriaga (Universidad de la República, Uruguay), and Sergio Nesmachnow (Universidad de la República, Uruguay)</i>	
Impact of Containerization on Low-Cost Post Moore Computing Architectures	528
<i>Pablo Josue Rojas Yepes (Universidad Industrial de Santander, Colombia), Carlos Jaime Barrios Hernandez (Universidad Industrial de Santander, Colombia), and Luiz Angelo Steffanel (Université de Reims Champagne Ardenne, France)</i>	
Automatic Vehicle Counting area Creation Based on Vehicle Deep Learning Detection and DBSCAN	535
<i>Gerardo Alvarez Piña (Universidad Autónoma de Guadalajara, Mexico), E. Ulises Moya-Sánchez (Coordinación general de innovación and Universidad Autónoma de Guadalajara, Mexico), Abraham Sánchez (Coordinación general de innovación, Gobierno de Jalisco), and Ulises Cortés (Barcelona Supercomputing Center/UPC, Spain)</i>	

Embracing Arm for High Performance Computing Workshop: EAHPC-2022

On Using Linux Kernel Huge Pages with FLASH, an Astrophysical Simulation Code	539
<i>Alan C. Calder (Stony Brook University, USA), Catherine Feldman (Stony Brook University, USA), Eva Siegmann (Stony Brook University, USA), John Dey (Stony Brook University, USA), Anthony Curtis (Stony Brook University, USA), Smeet Chheda (Stony Brook University, USA), and Robert J. Harrison (Stony Brook University, USA)</i>	

Performance of an Astrophysical Radiation Hydrodynamics Code Under Scalable Vector Extension Optimization	545
<i>Dennis C. Smolarski (Santa Clara University, USA), F. Douglas Swesty (Stony Brook University, USA), and Alan C. Calder (Stony Brook University, USA)</i>	
Productivity Meets Performance: Julia on A64FX	549
<i>Mosè Giordano (UCL, United Kingdom), Milan Klöwer (University of Oxford, United Kingdom), and Valentin Churavy (Massachusetts Institute of Technology, USA)</i>	
Assessing the State of Autovectorization Support Based on SVE	556
<i>Bine Brank (Forschungszentrum Jülich, Germany) and Dirk Pleiter (KTH Royal Institute of Technology, Germany)</i>	
Performance Analysis of a State Vector Quantum Circuit Simulation on A64FX Processor	563
<i>Miwako Tsuji (RIKEN Center for Computational Science, Japan) and Mitsuhsa Sato (RIKEN Center for Computational Science, Japan)</i>	

The 2nd Workshop on Re-Envisioning Extreme-Scale I/O for Emerging Hybrid HPC Workloads (REX-IO)

Protecting Metadata Servers From Harm Through Application-Level I/O Control	573
<i>Ricardo Macedo (TACC & UT Austin), Mariana Miranda (INESC TEC & University of Minho), Yusuke Tanimura (AIST), Jason Haga (AIST), Amit Ruhela (TACC & UT Austin), Stephen Lien Harrell (TACC & UT Austin), Richard Todd Evans (Intel), and João Paulo (INESC TEC & University of Minho)</i>	
A Comprehensive I/O Knowledge Cycle for Modular and Automated HPC Workload Analysis	581
<i>Zhaobin Zhu (Goethe University Frankfurt, Germany), Sarah Neuwirth (Goethe University Frankfurt, Germany; Juelich Supercomputing Centre, Germany), and Thomas Lippert (Goethe University Frankfurt, Germany; Juelich Supercomputing Centre, Germany)</i>	
Assessment of the I/O and Storage Subsystem in Modular Supercomputing Architectures	589
<i>Sarah Neuwirth (Goethe University Frankfurt, Germany)</i>	

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

Towards Real-Time Classification of HPC Workloads via Out-of-Band Telemetry	597
<i>Steven Presser (University of Stuttgart, Germany)</i>	
Shasta Log Aggregation, Monitoring and Alerting in HPC Environments with Grafana Loki and ServiceNow	602
<i>Elizabeth Bautista (Lawrence Berkeley National Lab, USA), Nitin Sukhija (Slippery Rock University of Pennsylvania, USA), and Siqi Deng (Lawrence Berkeley National Lab, USA)</i>	
Bridging the Gap Between Application Performance Analysis and System Monitoring	611
<i>Thomas Ilsche (Technische Universität Dresden, Germany), Mario Bielert (Technische Universität Dresden, Germany), and Christian von Elm (Technische Universität Dresden, Germany)</i>	

IncProf: Efficient Source-Oriented Phase Identification for Application Behavior Understanding	616
<i>Omar Aaziz (Sandia National Laboratories, USA), Mohammad Al-Tahat (New Mexico State University, USA), Strahinja Trecakov (New Mexico State University, USA), and Jonathan Cook (New Mexico State University, USA)</i>	
LDMS Darshan Connector: For Run Time Diagnosis of HPC Application I/O Performance	626
<i>Sara Walton (Sandia National Laboratories, USA), Omar Aaziz (Sandia National Laboratories, USA), Ana Luisa V. Solórzano (Northeastern University, USA), and Ben Schwaller (Sandia National Laboratories, USA)</i>	
Author Index	635