

SC23: International Conference for High Performance Computing, Networking, Storage and Analysis

**Denver, Colorado, USA
11-17 November 2023**

Pages 1-696



**IEEE Catalog Number: CFP23SUP-POD
ISBN: 979-8-3503-7663-0**

**Copyright © 2023, Association for Computing Machinery (ACM)
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:	CFP23SUP-POD
ISBN (Print-On-Demand):	979-8-3503-7663-0
ISBN (Online):	979-8-4007-0109-2
ISSN:	2167-4329

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

Session: ACM Gordon Bell Finalists

Large-Scale Materials Modeling at Quantum Accuracy: Ab Initio Simulations of Quasicrystals and Interacting Extended Defects in Metallic Alloys...1

Authors: Sambit Das,Bikash Kanungo,Vishal Subramanian,Gourab Panigrahi,Phani Motamarri,David Rogers,Paul Zimmerman,Vikram Gavini

Scaling the Leading Accuracy of Deep Equivariant Models to Biomolecular Simulations of Realistic Size...13

Authors: Boris Kozinsky,Albert Musaelian,Anders Johansson,Simon Batzner

Exascale Multiphysics Nuclear Reactor Simulations for Advanced Designs...25

Authors: Elia Merzari,Steven Hamilton,Thomas Evans,Misun Min,Paul Fischer,Stefan Kerkemeier,Jun Fang,Paul Romano,Yu-Hsiang Lan,Malachi Phillips,Elliott Biondo,Katherine Royston,Tim Warburton,Noel Chalmers,Thilina Rathnayake

Towards Exascale Computation for Turbomachinery Flows...36

Authors: Yuhang Fu,Weiqi Shen,Jiahuan Cui,Yao Zheng,Guangwen Yang,Zhao Liu,Jifa Zhang,Tingwei Ji,Fangfang Xie,Xiaojing Lv,Hanyue Liu,Xu Liu,Xiyang Liu,Xiaoyu Song,Guocheng Tao,Yan Yan,Paul Tucker,Steven Miller,Shirui Luo,Seid Koric,Weimin Zheng

Exploring the Ultimate Regime of Turbulent Rayleigh–Bénard Convection through Unprecedented Spectral-Element Simulations...48

Authors: Niclas Jansson,Martin Karp,Adalberto Perez,Timofey Mukha,Yi Ju,Jiahui Liu,Szilárd Páll,Erwin Laure,Tino Weinkauff,Jörg Schumacher,Philipp Schlatter, Stefano Markidis

Scaling the “Memory Wall” for Multi-Dimensional Seismic Processing with Algebraic Compression on Cerebras CS-2 Systems...57

Authors: Hatem Ltaief,Yuxi Hong,Leighton Wilson,Mathias Jacquelin,Matteo Ravasi,David Elliot Keyes

Session: ACM Gordon Bell Climate Modelling Finalists

The Simple Cloud-Resolving E3SM Atmosphere Model Running on the Frontier Exascale System...69

Authors: Mark Taylor,Peter M. Caldwell,Luca Bertagna,Conrad Clevenger,Aaron Donahue,James Foucar,Oksana Guba,Benjamin Hillman,Noel Keen,Jayesh Krishna,Matthew Norman,Sarat Sreepathi,Christopher Terai,James B. White,Andrew G Salinger,Renata B McCoy,Lai-yung Ruby Leung,David C. Bader,Danqing Wu

Big Data Assimilation: Real-Time 30-Second-Refresh Heavy Rain Forecast Using Fugaku during Tokyo Olympics and Paralympics...80

Authors: Takemasa Miyoshi,Arata Amemiya,Shigenori Otsuka,Yasumitsu Maejima,James Taylor,Takumi Honda,Hirofumi Tomita,Seiya Nishizawa,Kenta Sueki,Tsuyoshi Yamaura,Yutaka Ishikawa,Shinsuke Satoh,Tomoo Ushio,Kana Koike,Atsuya Uno

Establishing a Modeling System in 3-km Horizontal Resolution for Global Atmospheric Circulation Triggered by Submarine Volcanic Eruptions with 200 Billion Smoothed Particles Hydrodynamics...90

Authors: Shenghong Huang,Junshi Chen,Ziyu Zhang,Xiaoyu Hao,Jun Gu,Hong An,Chun Zhao,Yan Hu,Zhanming Wang,Longkui Chen,Yifan Luo,Jineng Yao,Yi Zhang,Yang Zhao,Zhihao Wang,Dongning Jia,Zhao Jin,Changming Song,Xisheng Luo,Xiaobin He,Dexun Chen

Technical Papers

Session: Extreme-Scale Applications

69.7-PFlops Extreme Scale Earthquake Simulation with Crossing Multi-Faults and Topography on Sunway...102

Authors: Wubing Wan, Lin Gan, Wenqiang Wang, Zekun Yin, Haodong Tian, Zhenguo Zhang, YINUO Wang, Mengyuan Hua, Xiaohui Liu, Shengye Xiang, Zhongqiu He, Zijia Wang, Ping Gao, Xiaohui Duan, Weiguo Liu, Wei Xue, Haohuan Fu, Guangwen Yang, Xiaofei Chen, Zeyu Song, Yaojian Chen, Xin Liu, Wei Zhang

Large-Scale Simulation of Structural Dynamics Computing on GPU Clusters...118

Authors: Yumeng Shi, Ningming Nie, Jue Wang, Kehao Lin, Chunbao Zhou, Shigang Li, Kehan Yao, Shunde Li, Yangde Feng, Yan Zeng, Fang Liu, Yangang Wang, Yue Gao

ANT-MOC: Scalable Neutral Particle Transport Using 3D Method of Characteristics on Multi-GPU Systems...132

Authors: Shunde Li, Zongguo Wang, Lingkun Bu, Jue Wang, Zhikuang Xin, Shigang Li, Yangang Wang, Yangde Feng, Peng Shi, Yun Hu, Xuebin Chi

Session: Global Task Parallelism

Legate Sparse: Distributed Sparse Computing in Python...145

Authors: Rohan Yadav, Wonchan Lee, Melih Elibol, Manolis Papadakis, Taylor Lee-Patti, Michael Garland, Alex Aiken, Fredrik Kjolstad, Michael Bauer

Itoyori: Reconciling Global Address Space and Global Fork-Join Task Parallelism...159

Authors: Shumpei Shiina, Kenjiro Taura

Automated Mapping of Task-Based Programs onto Distributed and Heterogeneous Machines...174

Authors: Thiago S. F. X. Teixeira, Alexandra Henzinger, Rohan Yadav, Alex Aiken

Session: Graph Algorithms in HPC

Efficient Maximal Biclique Enumeration on GPUs...188

Authors: Zhe Pan, Shuibing He, Xu Li, Xuechen Zhang, Rui Wang, Gang Chen

A GPU Algorithm for Detecting Strongly Connected Components...202

Authors: Ghadeer Alabandi, William Sands, George Biros, Martin Burtscher

Peek: A Prune-Centric Approach for K Shortest Path Computation...215

Authors: Wang Feng, Shiyang Chen, Hang Liu, Yuede Ji

Session: Sustainable Computing

Toward Sustainable HPC: Carbon Footprint Estimation and Environmental Implications of HPC Systems...229

Authors: Baolin Li, Rohan Basu Roy, Daniel Wang, Siddharth Samsi, Vijay Gadepally, Devesh Tiwari

Clover: Toward Sustainable AI with Carbon-Aware Machine Learning Inference Service...244

Authors: Baolin Li, Siddharth Samsi, Vijay Gadepally, Devesh Tiwari

GreenNFV: Energy-Efficient Network Function Virtualization with Service Level Agreement Constraints...259
Authors: Md S. Q. Zulkar Nine, Tefvik Kosar, Muhammed Fatih Bulut, Jinho Hwang

Session: Graph Frameworks and Databases

The Graph Database Interface: Scaling Online Transactional and Analytical Graph Workloads to Hundreds of Thousands of Cores...271

Authors: Maciej Besta, Robert Gerstenberger, Marc Fischer, Michal Podstawski, Nils Blach, Berke Egeli, Georgy Mitenkov, Wojciech Chlapek, Marek Michalewicz, Hubert Niewiadomski, Juergen Mueller, Torsten Hoefler

Graph3PO: A Temporal Graph Data Processing Method for Latency QoS Guarantee in Object Cloud Storage System...292

Authors: Wang Zhang, Zhan Shi, Ziyi Liao, Yiling Li, Yu Du, Yutong Wu, Fang Wang, Dan Feng

MBFGraph: An SSD-Based External Graph System for Evolving Graphs...309

Authors: Chun-Yi Liu, Wonil Choi, Soheil Khadirsharbiyani, Mahmut Kandemir

Session: Resource Management

Mirage: Toward Low-interruption Services on Batch GPU Clusters with Reinforcement Learning...323

Authors: Qiyang Ding, Pengfei Zheng, Shreyas Kudari, Shivaram Venkataraman, Zhao Zhang

Prodigy: Toward Unsupervised Anomaly Detection in Production HPC Systems...337

Authors: Burak Aksar, Efe Sencan, Benjamin Schwaller, Omar Aaziz, Vitus J. Leung, Jim Brandt, Brian Kulis, Manuel Egele, Ayse K. Coskun

DPS: Adaptive Power Management for Overprovisioned Systems...351

Authors: Jianru Ding, Henry Hoffmann

Session: GPU Middleware and System Software

Scalable Tuning of (OpenMP) GPU Applications via Kernel Record and Replay...366

Authors: Konstantinos Parasyris, Giorgis Georgakoudis, Esteban Rangel, Ignacio Laguna, Johannes Doerfert

HPAC-Offload: Accelerating HPC Applications with Portable Approximate Computing on the GPU...381

Authors: Zane Fink, Konstantinos Parasyris, Giorgis Georgakoudis, Harshitha Menon

Interference-Aware Multiplexing for Deep Learning in GPU Clusters: A Middleware Approach...396

Authors: Wenyan Chen, Zizhao Mo, Huanle Xu, Kejiang Ye, Chengzhong Xu

Session: High Performance for Graph Operations

Bringing Order to Sparsity: A Sparse Matrix Reordering Study on Multicore CPUs...410

Authors: James D. Trotter, Sinan Ekmek, Johannes Langguth, Tugba Torun, Emre Düzakin, Aleksandar Ilic, Didem Unat

GraphSet: High Performance Graph Mining through Equivalent Set Transformations...423

Authors: Tianhui Shi, Jidong Zhai, Haojie Wang, Qiqian Chen, Mingshu Zhai, Zixu Hao, Haoyu Yang, Wenguang Chen

Space Efficient Sequence Alignment for SRAM-Based Computing: X-Drop on the Graphcore IPU...438
Authors: Luk Burchard,Max Xiaohang Zhao,Johannes Langguth,Ayd,Giulia Guidi

Session: Message Passing Innovations

Optimizing MPI Collectives on Shared Memory Multi-Cores...454
Authors: Jintao Peng,Jianbin Fang,Jie Liu,Min Xie,Yi Dai,Bo Yang,Shengguo Li,Zheng Wang

Embracing Irregular Parallelism in HPC with YGM...469
Authors: Trevor Steil,Tahsin Reza,Benjamin Priest,Roger Pearce

HEAR: Homomorphically Encrypted Allreduce...482
Authors: Marcin Chrapek,Mikhail Khalilov,Torsten Hoefler

Session: Training Graph Neural Networks

BLAD: Adaptive Load Balanced Scheduling and Operator Overlap Pipeline for Accelerating the Dynamic GNN Training...500
Authors: Kaihua Fu,Quan Chen,Yuzhuo Yang, Jiuchen Shi,Chao Li,Minyi Guo

TANGO: Re-Thinking Quantization for Graph Neural Network Training on GPUs...513
Authors: Shiyang Chen, Da Zheng, Caiwen Ding, Chengying Huan, Yuede Ji, Hang Liu

DistTGL: Distributed Memory-Based Temporal Graph Neural Network Training...528
Authors: Hongkuan Zhou, Da Zheng, Xiang Song, George Karypis, Viktor Prasanna

Session: Applications in Materials Science and Biology

Portable and Scalable All-Electron Quantum Perturbation Simulations on Exascale Supercomputers...540
Authors: Zhikun Wu, Yangjun Wu, Ying Liu, Honghui Shang, Yingxiang Gao, Zhongcheng Zhang, Yuyang Zhang, Yingchi Long, Xiaobing Feng, Huiming Cui

Enhancing Adaptive Physics Refinement Simulations through the Addition of Realistic Red Blood Cell Counts...554
Authors: Sayan Roychowdhury, Samreen T. Mahmud, Aristotle Martin, Peter Balogh, Daniel F. Puleri, John Gounley, Erik W. Draeger, Amanda Randles

NNQS-Transformer: An Efficient and Scalable Neural Network Quantum States Approach for Ab Initio Quantum Chemistry...568
Authors: Yangjun Wu, Chu Guo, Yi Fan, Pengyu Zhou, Honghui Shang

Session: Data Compression

cuSZp: An Ultra-Fast GPU Error-Bounded Lossy Compression Framework with Optimized End-to-End Performance...582
Authors: Yafan Huang, Sheng Di, Xiaodong Yu, Guanpeng Li, Franck Cappello

AMRIC: A Novel In Situ Lossy Compression Framework for Efficient I/O in Adaptive Mesh Refinement Applications...596

Authors: Daoce Wang, Jesus Pulido, Pascal Grosset, Jiannan Tian, Sian Jin, Houjun Tang, Jean Sexton, Sheng Di, Kai Zhao, Bo Fang, Zarija Lukić, Franck Cappello, James Ahrens, Dingwen Tao

ADT-FSE: A New Encoder for SZ...612

Authors: Tao Lu, Yu Zhong, Zibin Sun, Xiang Chen, You Zhou, Fei Wu, Ying Yang, Yunxin Huang, Yafei Yang

Session: Handling Hardware Faults

Understanding the Effects of Permanent Faults in GPU's Parallelism Management and Control Units...626

Authors: Juan David Guerrero Balaguera, Josie Esteban Rodriguez Condia, Fernando Fernandes Dos Santos, Matteo Sonza Reorda, Paolo Rech

Design Considerations and Analysis of Multi-Level Erasure Coding in Large-Scale Data Centers...641

Authors: Meng Wang, Jiajun Mao, Rajdeep Rana, John Bent, Serkay Olmez, Anjus George, Garrett Wilson Ransom, Jun Li, Haryadi S. Gunawi

Unity ECC: Unified Memory Protection Against Bit and Chip Errors...654

Authors: Dongwhee Kim, Jaeyoon Lee, Wonyeong Jung, Michael Sullivan, Jungrae Kim

Session: Linear Algebra I

Optimizing High-Performance Linpack for Exascale Accelerated Architectures...670

Authors: Noel Chalmers, Jakub Kurzak, Damon Mcdougall, Paul Bauman

Unified Communication Optimization Strategies for Sparse Triangular Solver on CPU and GPU Clusters...682

Authors: Yang Liu, Nan Ding, Piyush Sao, Samuel Williams, Xiaoye Sherry Li

PanguLU: A Scalable Regular Two-Dimensional Block-Cyclic Sparse Direct Solver on Distributed Heterogeneous Systems...697

Authors: Xu Fu, Bingbin Zhang, Tengcheng Wang, Wenhao Li, Yuechen Lu, Enxin Yi, Jianqi Zhao, Xiaohan Geng, Fangying Li, Jingwen Zhang, Zhou Jin, Weifeng Liu

Session: Exascale Computing

Frontier: Exploring Exascale...712

Authors: Scott Atchley, Christopher Zimmer, John Lange, David Bernholdt, Veronica Melesse Vergara, Thomas Beck, Michael Brim, Reuben Budiardja, Sunita Chandrasekaran, Markus Eisenbach, Thomas Evans, Matthew Ezell, Nicholas Frontiere, Antigoni Georgiadou, Joe Glenski, Philipp Grete, Steven Hamilton, John Holmen, Axel Huebl, Daniel Jacobson, Wayne Joubert, Kim McMahon, Elia Merzari, Stan Moore, Andrew Myers, Stephen Nichols, Sarp Oral, Thomas Papatheodore, Danny Perez, David M. Rogers, Evan Schneider, Jean-Luc Vay, P. K. Yeung

Experiences Readying Applications for Exascale...728

Authors: Nicholas Malaya, Bronson Messer, Joseph Glenski, Antigoni Georgiadou, Justin Lietz, Kalyana Gottiparthi, Marc Day, Jackie Chen, Jon Rood, Lucas Esclapez, James White Iii, Gustav R. Jansen, Nicholas Curtis, Stephen Nichols, Jakub Kurzak, Noel Chalmers, Chip Freitag, Paul Bauman, Alessandro Fanfarillo, Reuben D. Budiardja, Thomas Papatheodore, Nicholas Frontiere, Damon Mcdougall, Matthew Norman, Sarat Sreepathi, Philip Roth, Dmytro Bykov, Noah Wolfe, Paul Mullowney, Markus Eisenbach, Marc T. Henry De Frahan, Wayne Joubert

5 ExaFlop/s HPL-MxP Benchmark with Linear Scalability on the 40-Million-Core Sunway Supercomputer...741

Authors: Rongfen Lin,Xinhui Yuan,Wei Xue,Wanwang Yin,Jienan Yao,Junda Shi,Qiang Sun,Chaobo Song,Fei Wang

Session: Training in HPC Machine Learning

EasyScale: Elastic Training with Consistent Accuracy and Improved Utilization on GPUs...754

Authors: Mingzhen Li,Wencong Xiao,Hailong Yang,Biao Sun,Hanyu Zhao,Shiru Ren,Zhongzhi Luan,Xianyan Jia,Yi Liu,Yong Li,Wei Lin,Depei Qian

Hanayo: Harnessing Wave-Like Pipeline Parallelism for Enhanced Large Model Training Efficiency...769

Authors: Ziming Liu,Shenggan Cheng,Haotian Zhou,Yang You

High Throughput Training of Deep Surrogates from Large Ensemble Runs...782

Authors: Lucas Thibaut Meyer,Marc Schouler,Robert Alexander Caulk,Alejandro Ribes,Bruno Raffin

Session: Data Coordination

Data Flow Lifecycles for Optimizing Workflow Coordination...798

Authors: Hyungro Lee,Luanzheng Guo,Meng Tang,Jesun Firoz,Nathan Tallent,Anthony Kougkas,Xian-He Sun

Accelerating Communications in Federated Applications with Transparent Object Proxies...813

Authors: J. Gregory Pauloski,Valerie Hayot-Sasson,Logan Ward,Nathaniel Hudson,Charlie Sabino,Matt Baughman,Kyle Chard,Ian Foster

A Quantitative Approach for Adopting Disaggregated Memory in HPC Systems...829

Authors: Jacob Wahlgren,Gabin Schieffer,Maya Gokhale,Ivy Peng

Session: Quantum Computing

GRAPHINE: Enhanced Neutral Atom Quantum Computing Using Application-Specific Rydberg Atom Arrangement...843

Authors: Tirthak Patel,Daniel Silver,Devesh Tiwari

Mitigating Coupling Map Constrained Correlated Measurement Errors on Quantum Devices...858

Authors: Alan Robertson,Shuaiwen Song

Experimental Evaluation of Xanadu X8 Photonic Quantum Computer: Error Measurement, Characterization, and Implications...871

Authors: Aditya Ranjan,Tirthak Patel,Harshitta Gandhi,Daniel Silver,William Cutler,Devesh Tiwari

Session: Tensor Computation

Automatic Generation of Distributed-Memory Mappings for Tensor Computations...884

Authors: Martin Kong,Raneem Abu Yosef,Atanas Rountev,P. Sadayappan

Application Performance Modeling via Tensor Completion...898

Authors: Edward Hutter,Edgar Solomonik

High-Performance and Programmable Attentional Graph Neural Networks with Global Tensor Formulations...912

Authors: Maciej Besta,Pawel Renc,Robert Gerstenberger,Paolo Sylos Labini,Alexandros Ziogas,Tiancheng Chen,Lukas Gianinazzi,Florian Scheidl,Kalman Szenes,Armon Carigiet,Patrick Iff,Grzegorz Kwasniewski,Raghavendra Kanakagiri,Chio Ge,Sammy Jaeger,Jaroslav Was,Flavio Vella,Torsten Hoefler

Session: Topics in Cloud Computing

Rethinking Deployment for Serverless Functions: A Performance-First Perspective...930

Authors: Yiming Li,Laiping Zhao,Yanan Yang,Wenyu Qu

FISCO-BCOS: An Enterprise-Grade Permissioned Blockchain System with High-Performance...945

Authors: Huizhong Li,Yujie Chen,Xiang Shi,Xingqiang Bai,Nan Mo,Wenlin Li,Rui Guo,Zhang Wang,Yi Sun

SYnergy: Fine-Grained Energy-Efficient Heterogeneous Computing for Scalable Energy Saving...962

Authors: Kaijie Fan,Marco D'Antonio,Lorenzo Carpentieri,Biagio Cosenza,Federico Ficarelli,Daniele Cesarini

Session: Architecture-Specific Optimization

Optimizing Direct Convolutions on ARM Multi-Cores...976

Authors: Pengyu Wang,Weiling Yang,Jianbin Fang,Dezun Dong,Chun Huang,Peng Zhang,Tao Tang,Zheng Wang

Calculon: a Methodology and Tool for High-Level Codesign of Systems and Large Language Models...990

Authors: Mikhail Isaev,Nic Mcdonald,Larry Dennison,Richard Vuduc

VENOM: A Vectorized N:M Format for Unleashing the Power of Sparse Tensor Cores...1004

Authors: Roberto L. Castro,Andrei Ivanov,Diego Andrade,Tal Ben-Nun,Basilio B. Fraguera,Torsten Hoefler

Session: Linear Algebra II

DASP: Specific Dense Matrix Multiply-Accumulate Units Accelerated General Sparse Matrix-Vector Multiplication...1020

Authors: Yuechen Lu,Weifeng Liu

High-Performance SVD Partial Spectrum Computation...1035

Authors: David Keyes,Hatem Ltaief,Yuji Nakatsukasa,Dalal Sukkari

ReFloat: Low-Cost Floating-Point Processing in ReRAM for Accelerating Iterative Linear Solvers...1047

Authors: Linghao Song,Fan Chen,Hai Li,Yiran Chen

Session: Algorithms on GPUs

Parallel Top-K Algorithms on GPU: A Comprehensive Study and New Methods...1062

Authors: Jingrong Zhang,Akira Naruse,Xipeng Li,Yong Wang

A High-Performance MST Implementation for GPUs...1075

Authors: Alex Fallin, Andres Gonzalez, Jarim Seo, Martin Burtscher

Adaptive Workload-Balanced Scheduling Strategy for Global Ocean Data Assimilation on Massive GPUs...1089
Authors: Junmin Xiao, Chaoyang Shui, Di Cai, Kangyu Wang, Yunfei Pang, Mingyi Li, Hui Ma, Guangming Tan

Session: Applications of Machine Learning

Rapid Simulations of Atmospheric Data Assimilation of Hourly-Scale Phenomena with Modern Neural Networks...1105

Authors: Yiyuan Li, Xiting Ju, Yi Xiao, Qilong Jia, Yongxiao Zhou, Simeng Qian, Rongfen Lin, Bin Yang, Shupeng Shi, Xin Liu, Jie Gao, Zhen Wang, Sha Liu, Jian Tan, Xuan Wang, Zhengding Hu, Limin Yan, Wei Xue

Breaking Boundaries: Distributed Domain Decomposition with Scalable Physics-Informed Neural PDE Solvers...1118

Authors: Arthur Feeney, Zitong Li, Ramin Bostanabad, Aparna Chandramowliswaran

FORGE: Pre-Training Open Foundation Models for Science...1133

Authors: Junqi Yin, Sajal Dash, Feiyi Wang, Mallikarjun Shankar

Session: Data Centers and Large Distributed Systems

Cloud Computing to Enable Wearable-Driven Longitudinal Hemodynamic Maps...1146

Authors: Cyrus Tanade, Emily Rakestraw, William Ladd, Erik Draeger, Amanda Randles

Optimizing Reconfigurable Optical Datacenters: The Power of Randomization...1160

Authors: Marcin Bienkowski, David Fuchssteiner, Stefan Schmid

Leveraging the Compute Power of Two HPC Systems for Higher-Dimensional Grid-Based Simulations with the Widely-Distributed Sparse Grid Combination Technique...1171

Authors: Theresa Pollinger, Alexander Van Craen, Christoph Niethammer, Marcel Breyer, Dirk Pfluger

Session: Fault Tolerance and FPGA Codesign

Structural Coding: A Low-Cost Scheme to Protect CNNs from Large-Granularity Memory Faults...1186

Authors: Ali Asgari Khoshouyeh, Florian Geissler, Syed Qutub, Michael Paulitsch, Prashant Nair, Karthik Pattabiraman

Demystifying and Mitigating Cross-Layer Deficiencies of Soft Error Protection in Instruction Duplication...1203

Authors: Zhengyang He, Yafan Huang, Hui Xu, Dingwen Tao, Guanpeng Li

Co-Design Hardware and Algorithm for Vector Search...1217

Authors: Wenqi Jiang, Shigang Li, Yu Zhu, Johannes De Fine Licht, Zhenhao He, Runbin Shi, Cedric Renggli, Shuai Zhang, Theodoros Rekatsinas, Torsten Hoefler, Gustavo Alonso

Session: Code Optimization

FuzzyFlow: Leveraging Dataflow to Find and Squash Program Optimization Bugs...1233

Authors: Philipp Schaad, Timo Schneider, Tal Ben-Nun, Alexandru Calotoiu, Alexandros Nikolaos Ziogas, Torsten Hoefler

Runtime Composition of Iterations for Fusing Loop-Carried Sparse Dependence...1248

Authors: Kazem Cheshmi,Michelle Strout,Maryam Mehri Dehnavi

TrivialSpy: Identifying Software Triviality via Fine-Grained and Dataflow-Based Value Profiling...1263

Authors: Xin You,Hailong Yang,Kelun Lei,Zhongzhi Luan,Depei Qian

Session: Graph Analytics

Phases, Modalities, Spatial and Temporal Locality: Domain Specific ML Prefetcher for Accelerating Graph Analytics...1277

Authors: Pengmiao Zhang,Rajgopal Kannan,Viktor K. Prasanna

Choosing the Best Parallelization and Implementation Styles for Graph Analytics Codes: Lessons Learned from 1106 Programs...1293

Authors: Yiqian Liu,Noushin Azami,Avery Vanausdal,Martin Burtscher

DGAP: Efficient Dynamic Graph Analysis on Persistent Memory...1308

Authors: Abdullah Al Raqibul Islam,Dong Dai

Session: High Performance I/O

I/O in WRF: A Case Study in Modern Parallel I/O Techniques...1322

Authors: Zanhua Huang,Kaiyuan Hou,Ankit Agrawal,Alok Choudhary,Robert Ross,Wei-Keng Liao

Fine-Grained Policy-Driven I/O Sharing for Burst Buffers...1335

Authors: Ed Karrels,Lei Huang,Yuhong Kan,Ishank Arora,Yinzhi Wang,Daniel S. Katz,William Gropp,Zhao Zhang

Xfast: Extreme File Attribute Stat Acceleration for Lustre...1347

Authors: Yingjin Qian,Wen Cheng,Lingfang Zeng,Xi Li,Marc-Andre Vef,Andreas Dilger,Siyao Lai,Shuichi Ihara,Yong Fan,Andre Brinkmann

Session: Molecular Dynamics Applications and Accelerators

Enhance the Strong Scaling of LAMMPS on Fugaku...1364

Authors: Jianxiong Li,Tong Zhao,Zhuoqiang Guo,Shunchen Shi,Lijun Liu,Guangming Tan,Weile Jia,Guojun Yuan,Zhan Wang

FASDA: An FPGA-Aided, Scalable and Distributed Accelerator for Range-Limited Molecular Dynamics...1377

Authors: Chunshu Wu,Tong Geng,Anqi Guo,Sahan Bandara,Pouya Haghi,Chuan Liu,Ang Li,Martin Herbordt

Enabling Real World Scale Structural Superlubricity All-Atom Simulation on the Next-Generation Sunway Supercomputer...1391

Authors: Xiaohui Duan,Jin Wang,Ping Gao,Ming Ma,Lin Gan,Xin Liu,Haohuan Fu,Wei Xue,Dexun Chen,Guangwen Yang,Weiguo Liu