

SC22: International Conference for High Performance Computing, Networking, Storage and Analysis

**Dallas, Texas, USA
13-18 November 2022**

Pages 1-631



**IEEE Catalog Number: CFP22SUP-POD
ISBN: 978-1-6654-5445-2**

**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:	CFP22SUP-POD
ISBN (Print-On-Demand):	978-1-6654-5445-2
ISBN (Online):	978-1-6654-5444-5
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

TABLE OF CONTENTS

Extreme-Scale Many-Against-Many Protein Similarity Search.....	1
<i>Oguz Selvitopi, Saliya Ekanayake, Giulia Guidi, Muaaz G. Awan, Georgios A. Pavlopoulos, Ariful Azad, Nikos Kyrpides, Leonid Olikier, Katherine Yelick, Aydin Buluç</i>	
Reshaping Geostatistical Modeling and Prediction for Extreme-Scale Environmental Applications	13
<i>Qinglei Cao, Sameh Abdulah, Rabab Alomairy, Yu Pei, Pratik Nag, George Bosilca, Jack Dongarra, Marc G. Genton, David E. Keyes, Hatem Ltaief, Ying Sun</i>	
Pushing the Frontier in the Design of Laser-Based Electron Accelerators with Groundbreaking Mesh-Refined Particle-In-Cell Simulations on Exascale-Class Supercomputers.....	25
<i>Luca Fedeli, Axel Huebl, France Boillod-Cerneux, Thomas Clark, Kevin Gott, Conrad Hillairet, Stephan Jaure, Adrien Leblanc, Rémi Lehe, Andrew Myers, Christelle Piechurski, Mitsuhsa Sato, Neil Zaim, Weiqun Zhang, Jean-Luc Vay, Henri Vincenti</i>	
Extreme Scale Earthquake Simulation with Uncertainty Quantification.....	37
<i>Tsuyoshi Ichimura, Kohei Fujita, Ryota Kusakabe, Kentaro Koyama, Sota Murakami, Yuma Kikuchi, Takane Hori, Muneo Hori, Hikaru Inoue, Takafumi Nose, Takahiro Kawashima, Maddegedara Lalith</i>	
2.5 Million-Atom Ab Initio Electronic-Structure Simulation of Complex Metallic Heterostructures with DGDFT.....	48
<i>Wei Hu, Hong An, Zhuoqiang Guo, Qingcai Jiang, Xinming Qin, Junshi Chen, Weile Jia, Chao Yang, Zhaolong Luo, Jielan Li, Wentiao Wu, Guangming Tan, Dongning Jia, Qinglin Lu, Fangfang Liu, Min Tian, Fang Li, Yeqi Huang, Liyi Wang, Sha Liu, Jinlong Yang</i>	
Exaflops Biomedical Knowledge Graph Analytics	61
<i>Ramakrishnan Kannan, Piyush Sao, Hao Lu, Jakub Kurzak, Gundolf Schenk, Yongmei Shi, Seung-Hwan Lim, Sharat Israni, Vijay Thakkar, Guojing Cong, Robert Patton, Sergio E. Baranzini, Richard Vuduc, Thomas Potok</i>	
Scaling Correlated Fragment Molecular Orbital Calculations on Summit	72
<i>Giuseppe M. J. Barca, Calum Snowdon, Jorge L. Galvez Vallejo, Fazeleh Kazemian, Alistair P. Rendell, Mark S. Gordon</i>	
Image Gradient Decomposition for Parallel and Memory-Efficient Ptychographic Reconstruction.....	86
<i>Xiao Wang, Aristeidis Tsaris, Debangshu Mukherjee, Mohamed Wahib, Peng Chen, Mark Oxley, Olga Ovchinnikova, Jacob Hinkle</i>	
P-Massive: A Real-Time Search Engine for a Multi-Terabyte Mass Spectrometry Database	99
<i>Narangerelt Batsoyol, Benjamin Pullman, Mingxun Wang, Nuno Bandeira, Steven Swanson</i>	
Building Blocks for Network-Accelerated Distributed File Systems.....	114
<i>Salvatore Di Girolamo, Daniele De Sensi, Konstantin Taranov, Milos Malesevic, Maciej Besta, Timo Schneider, Severin Kistler, Torsten Hoefler</i>	
HammingMesh: A Network Topology for Large-Scale Deep Learning	128
<i>Torsten Hoefler, Tommaso Bonato, Daniele De Sensi, Salvatore Di Girolamo, Shigang Li, Marco Heddes, Jon Belk, Deepak Goel, Miguel Castro, Steve Scott</i>	
PolarFly: A Cost-Effective and Flexible Low-Diameter Topology	146
<i>Kartik Lakhotia, Maciej Besta, Laura Monroe, Kelly Isham, Patrick Iff, Torsten Hoefler, Fabrizio Petrini</i>	

Combining Hard and Soft Constraints in Quantum Constraint-Satisfaction Systems	161
<i>Ellis Wilson, Frank Mueller, Scott Pakin</i>	
Large-Scale Simulation of Quantum Computational Chemistry on a New Sunway Supercomputer	175
<i>Honghui Shang, Li Shen, Yi Fan, Zhiqian Xu, Chu Guo, Jie Liu, Wenhao Zhou, Huan Ma, Rongfen Lin, Yuling Yang, Fang Li, Zhuoya Wang, Yunquan Zhang, Zhenyu Li</i>	
Charter: Identifying the Most-Critical Gate Operations in Quantum Circuits Via Amplified Gate Reversibility	189
<i>Tirthak Patel, Daniel Silver, Devesh Tiwari</i>	
A Taxonomy of Error Sources in HPC I/O Machine Learning Models	205
<i>Mihailo Isakov, Mikaela Currier, Eliakin Del Rosario, Sandeep Madireddy, Prasanna Balaprakash, Philip Carns, Robert B. Ross, Glenn K. Lockwood, Michel A. Kinsky</i>	
Mitigating Silent Data Corruptions in HPC Applications Across Multiple Program Inputs	219
<i>Yafan Huang, Shengjian Guo, Sheng Di, Guanpeng Li, Franck Cappello</i>	
Optimizing Random Access to Hierarchically-Compressed Data on GPU	233
<i>Feng Zhang, Yihua Hu, Haipeng Ding, Zhiming Yao, Zhewei Wei, Xiao Zhang, Xiaoyong Du</i>	
Scaling Graph 500 SSSP to 140 Trillion Edges with Over 40 Million Cores.....	248
<i>Yuanwei Wang, Huanqi Cao, Zixuan Ma, Wanwang Yin, Wenguang Chen</i>	
Study of Workload Interference with Intelligent Routing on Dragonfly	263
<i>Yao Kang, Xin Wang, Zhiling Lan</i>	
SERVIZ: A Shared in Situ Visualization Service	277
<i>Srinivasan Ramesh, Hank Childs, Allen Malony</i>	
DayDream: Executing Dynamic Scientific Workflows on Serverless Platforms with Hot Starts	291
<i>Rohan Basu Roy, Tirthak Patel, Devesh Tiwari</i>	
LabStor: A Modular and Extensible Platform for Developing High-Performance, Customized I/O Stacks in Userspace	309
<i>Luke Logan, Jaime Cernuda Garcia, Jay Lofstead, Xian-He Sun, Anthony Kougkas</i>	
Towards Scalable Resource Management for Supercomputers	324
<i>Yiqin Dai, Yong Dong, Kai Lu, Ruibo Wang, Wei Zhang, Juan Chen, Mingtian Shao, Zheng Wang</i>	
Deinsum: Practically I/O Optimal Multi-Linear Algebra	339
<i>Alexandros Nikolaos Ziogas, Grzegorz Kwasniewski, Tal Ben-Nun, Timo Schneider, Torsten Hoefler</i>	
Addressing Irregular Patterns of Matrix Computations on GPUs and Their Impact on Applications Powered by Sparse Direct Solvers.....	354
<i>Ahmad Abdelfattah, Pieter Ghysels, Wajih Boukaram, Stanimire Tomov, Xiaoye Sherry Li, Jack Dongarra</i>	
Accelerating Elliptic Curve Digital Signature Algorithms on GPUs.....	368
<i>Zonghao Feng, Qipeng Xie, Qiong Luo, Yujie Chen, Haoxuan Li, Huizhong Li, Qiang Yan</i>	
CA3DMM: A New Algorithm Based on a Unified View of Parallel Matrix Multiplication	381
<i>Hua Huang, Edmond Chow</i>	

Symmetric Block-Cyclic Distribution: Fewer Communications Leads to Faster Dense Cholesky Factorization.....	396
<i>Olivier Beaumont, Philippe Duchon, Lionel Eyraud-Dubois, Julien Langou, Mathieu V�rit�</i>	
Scalable Distributed High-Order Stencil Computations.....	411
<i>Mathias Jacquelin, Mauricio Araya-Polo, Jie Meng</i>	
Memory Optimizations in an Array Language	424
<i>Philip Munksgaard, Troels Henriksen, Ponnuswamy Sadayappan, Cosmin Oancea</i>	
Vectorizing Sparse Matrix Computations with Partially-Strided Codelets.....	439
<i>Kazem Cheshmi, Zachary Cetinic, Maryam Mehri Dehnavi</i>	
Finding Inputs that Trigger Floating-Point Exceptions in GPUs Via Bayesian Optimization	454
<i>Ignacio Laguna, Ganesh Gopalakrishnan</i>	
Mapping Out the HPC Dependency Chaos	468
<i>Farid Zakaria, Thomas R. W. Scogland, Todd Gamblin, Carlos Maltzahn</i>	
Using Answer Set Programming for HPC Dependency Solving.....	480
<i>Todd Gamblin, Massimiliano Culpo, Gregory Becker, Sergei Shudler</i>	
SPATL: Salient Parameter Aggregation and Transfer Learning for Heterogeneous Federated Learning	495
<i>Sixing Yu, Phuong Nguyen, Waqwoya Abebe, Wei Qian, Ali Anwar, Ali Jannesari</i>	
Efficient Quantized Sparse Matrix Operations on Tensor Cores	509
<i>Shigang Li, Kazuki Osawa, Torsten Hoefler</i>	
LightSeq2: Accelerated Training for Transformer-Based Models on GPUs.....	524
<i>Xiaohui Wang, Yang Wei, Ying Xiong, Guyue Huang, Xian Qian, Yufei Ding, Mingxuan Wang, Lei Li</i>	
CoGNN: Efficient Scheduling for Concurrent GNN Training on GPUs.....	538
<i>Qingxiao Sun, Yi Liu, Hailong Yang, Ruizhe Zhang, Ming Dun, Mingzhen Li, Xiaoyan Liu, Wencong Xiao, Yong Li, Zhongzhi Luan, Depei Qian</i>	
Using Unused: Non-Invasive Dynamic FaaS Infrastructure with HPC-Whisk	553
<i>Bartlomiej Przybylski, Maciej Pawlik, Pawel Zuk, Bartlomiej Lagosz, Maciej Malawski, Krzysztof Rzacza</i>	
Canary: Fault-Tolerant FaaS for Stateful Time-Sensitive Applications.....	568
<i>Moiz Arif, Kevin Assogba, M. Mustafa Rafique</i>	
SFS: Smart OS Scheduling for Serverless Functions	584
<i>Yuqi Fu, Li Liu, Haoliang Wang, Yue Cheng, Songqing Chen</i>	
ProbGraph: High-Performance and High-Accuracy Graph Mining with Probabilistic Set Representations	600
<i>Maciej Besta, Cesare Miglioli, Paolo Sylos Labini, Jakub Tetek, Patrick Iff, Raghavendra Kanakagiri, Saleh Ashkboos, Kacper Janda, Michal Podstawski, Grzegorz Kwasniewski, Niels Gleinig, Flavio Vella, Onur Muthu, Torsten Hoefler</i>	
Blaze: Fast Graph Processing on Fast SSDs.....	617
<i>Juno Kim, Steven Swanson</i>	

GraphFly: Efficient Asynchronous Streaming Graphs Processing Via Dependency-Flow	632
<i>Dan Chen, Chuangyi Gui, Yi Zhang, Hai Jin, Long Zheng, Yu Huang, Xiaofei Liao</i>	
DeepSpeed- Inference: Enabling Efficient Inference of Transformer Models at Unprecedented Scale	646
<i>Reza Yazdani Aminabadi, Samyam Rajbhandari, Ammar Ahmad Awan, Cheng Li, Du Li, Elton Zheng, Olatunji Ruwase, Shaden Smith, Minjia Zhang, Jeff Rasley, Yuxiong He</i>	
HyLo: A Hybrid Low-Rank Natural Gradient Descent Method.....	661
<i>Baorun Mu, Saeed Soori, Bugra Can, Mert Gürbüzbalaban, Maryam Mehri Dehnavi</i>	
AI for Quantum Mechanics: High Performance Quantum Many-Body Simulations Via Deep Learning	677
<i>Xuncheng Zhao, Mingfan Li, Qian Xiao, Junshi Chen, Fei Wang, Li Shen, Meijia Zhao, Wenhao Wu, Hong An, Lixin He, Xiao Liang</i>	
UniQ: A Unified Programming Model for Efficient Quantum Circuit Simulation.....	692
<i>Chen Zhang, Haojie Wang, Zixuan Ma, Lei Xie, Zeyu Song, Jidong Zhai</i>	
Scalable Irregular Parallelism with GPUs: Getting CPUs Out of the Way.....	708
<i>Yuxin Chen, Benjamin Brock, Serban Porumbescu, Aydin Buluç, Katherine Yelick, John D. Owens</i>	
Parla: A Python Orchestration System for Heterogeneous Architectures	724
<i>Hochan Lee, William Ruys, Ian Henriksen, Arthur Peters, Yineng Yan, Sean Stephens, Bozhi You, Henrique Fingler, Martin Burtcher, Milos Gligoric, Karl Schulz, Keshav Pingali, Christopher J. Rossbach, Mattan Erez, George Biros</i>	
VSGM: View-Based GPU-Accelerated Subgraph Matching on Large Graphs.....	739
<i>Guanxian Jiang, Qihui Zhou, Tatiana Jin, Boyang Li, Yunjian Zhao, Yichao Li, James Cheng</i>	
STMatch: Accelerating Graph Pattern Matching on GPU with Stack-Based Loop Optimizations	754
<i>Yihua Wei, Peng Jiang</i>	
WholeGraph: A Fast Graph Neural Network Training Framework with Multi-GPU Distributed Shared Memory Architecture.....	767
<i>Dongxu Yang, Junhong Liu, Jiaying Qi, Junjie Lai</i>	
SeqDLM: A Sequencer-Based Distributed Lock Manager for Efficient Shared File Access in a Parallel File System.....	781
<i>Qi Chen, Shaonan Ma, Kang Chen, Teng Ma, Xin Liu, Dexun Chen, Yongwei Wu, Zuoning Chen</i>	
MetaWBC: POSIX-Compliant Metadata Write-Back Caching for Distributed File Systems.....	795
<i>Yingjin Qian, Wen Cheng, Lingfang Zeng, Marc-André Vef, Oleg Drokin, Andreas Dilger, Shuichi Ihara, Wusheng Zhang, Yang Wang, André Brinkmann</i>	
GUFI: Fast, Secure File System Metadata Search for Both Privileged and Unprivileged Users	815
<i>Dominic Manno, Jason Lee, Prajwal Challa, Qing Zheng, David Bonnie, Gary Grider, Bradley Settlemyer</i>	
AD for an Array Language with Nested Parallelism	829
<i>Robert Schenck, Ola Ronning, Troels Henriksen, Cosmin E. Oancea</i>	
SpDISTAL: Compiling Distributed Sparse Tensor Computations	844
<i>Rohan Yadav, Alex Aiken, Fredrik Kjolstad</i>	

Scalable Automatic Differentiation of Multiple Parallel Paradigms Through Compiler Augmentation	859
<i>William S. Moses, Sri Hari Krishna Narayanan, Ludger Paehler, Valentin Churavy, Michel Schanen, Jan Hückelheim, Johannes Doerfert, Paul Hovland</i>	
Accelerating Parallel Write Via Deeply Integrating Predictive Lossy Compression with HDF5	877
<i>Sian Jin, Dingwen Tao, Houjun Tang, Sheng Di, Suren Byna, Zarija Lukic, Franck Cappello</i>	
Dynamic Quality Metric Oriented Error Bounded Lossy Compression for Scientific Datasets.....	892
<i>Jinyang Liu, Sheng Di, Kai Zhao, Xin Liang, Zizhong Chen, Franck Cappello</i>	
VGraph: Memory-Efficient Multicore Graph Processing for Traversal-Centric Algorithms	907
<i>Menghan Jia, Yiming Zhang, Xinbiao Gan, Dongsheng Li, Erci Xu, Ruibo Wang, Kai Lu</i>	
Boosting Performance Optimization with Interactive Data Movement Visualization.....	921
<i>Philipp Schaad, Tal Ben-Nun, Torsten Hoefler</i>	
Not All GPUs Are Created Equal: Characterizing Variability in Large-Scale, Accelerator-Rich Systems.....	937
<i>Prasoon Sinha, Akhil Guliani, Rutwik Jain, Brandon Tran, Matthew D. Sinclair, Shivaram Venkataraman</i>	
AlphaSparse: Generating High Performance SpMV Codes Directly from Sparse Matrices	952
<i>Zhen Du, Jiajia Li, Yinshan Wang, Xueqi Li, Guangming Tan, Ninghui Sun</i>	
Approximate Computing Through the Lens of Uncertainty Quantification	967
<i>Konstantinos Parasyris, James Diffenderfer, Harshitha Menon, Ignacio Laguna, Jackson Vanover, Ryan Vogt, Daniel Osei-Kuffuor</i>	
Positive-Phase Temperature Scaling for Quantum-Assisted Boltzmann Machine Training.....	981
<i>Jose P. Pinilla, Steven J. E. Wilton</i>	
QoS-Aware Irregular Collaborative Inference for Improving Throughput of DNN Services	993
<i>Kaihua Fu, Jiuchen Shi, Quan Chen, Ningxin Zheng, Wei Zhang, Deze Zeng, Minyi Guo</i>	
EL-Rec: Efficient Large-Scale Recommendation Model Training Via Tensor-Train Embedding Table	1007
<i>Zheng Wang, Yuke Wang, Boyuan Feng, Dheevatsa Mudigere, Bharath Muthiah, Yufei Ding</i>	
STRONGHOLD: Fast and Affordable Billion-Scale Deep Learning Model Training	1021
<i>Xiaoyang Sun, Wei Wang, Shenghao Qiu, Renyu Yang, Songfang Huang, Jie Xu, Zheng Wang</i>	
HGL: Accelerating Heterogeneous GNN Training with Holistic Representation and Optimization.....	1038
<i>Yuntao Gui, Yidi Wu, Han Yang, Tatiana Jin, Boyang Li, Qihui Zhou, James Cheng, Fan Yu</i>	
Productive Performance Engineering for Weather and Climate Modeling with Python	1053
<i>Tal Ben-Nun, Linus Groner, Florian Deconinck, Tobias Wicky, Eddie Davis, Johann Dahm, Oliver D. Elbert, Rhea George, Jeremy McGibbon, Lukas Trümper, Elynn Wu, Oliver Fuhrer, Thomas Schulthess, Torsten Hoefler</i>	
Optimization of Full-Core Reactor Simulations on Summit	1067
<i>Misun Min, Yu-Hsiang Lan, Paul Fischer, Elia Merzari, Stefan Kerkemeier, Malachi Phillips, Thilina Rathnayake, April Novak, Derek Gaston, Noel Chalmers, Tim Warburton</i>	
A GPU-Accelerated AMR Solver for Gravitational Wave Propagation	1078
<i>Milinda Fernando, David Neilsen, Eric Hirschmann, Yosef Zlochower, Hari Sundar, Omar Ghattas, George Biros</i>	

From Correctable Memory Errors to Uncorrectable Memory Errors: What Error Bits Tell.....	1093
<i>Cong Li, Yu Zhang, Jialei Wang, Hang Chen, Xian Liu, Tai Huang, Liang Peng, Shen Zhou, Lixin Wang, Shijian Ge</i>	
Lessons Learned on MPI+Threads Communication	1107
<i>Rohit Zambre, Aparna Chandramowlishwaran</i>	
Climbing the Summit and Pushing the Frontier of Mixed Precision Benchmarks at Extreme Scale	1123
<i>Hao Lu, Michael Matheson, Vladyslav Oles, Austin Ellis, Wayne Joubert, Feiyi Wang</i>	
Scalable Deep Learning-Based Microarchitecture Simulation on GPUs	1138
<i>Santosh Pandey, Lingda Li, Thomas Flynn, Adolfy Hoisie, Hang Liu</i>	
TD-NUCA: Runtime Driven Management of NUCA Caches in Task Dataflow Programming Models.....	1153
<i>Paul Caheny, Lluc Alvarez, Marc Casas, Miquel Moreto</i>	
ReSemble: Reinforced Ensemble Framework for Data Prefetching	1168
<i>Pengmiao Zhang, Rajgopal Kannan, Ajitesh Srivastava, Anant V. Nori, Viktor K. Prasanna</i>	
W-Cycle SVD: A Multilevel Algorithm for Batched SVD on GPUs	1182
<i>Junmin Xiao, Yunfei Pang, Qing Xue, Chaoyang Shui, Ke Meng, Hui Ma, Mingyi Li, Xiaoyang Zhang, Guangming Tan</i>	
Scalable Linear Time Dense Direct Solver for 3-D Problems Without Trailing Sub-Matrix Dependencies.....	1198
<i>Qianxiang Ma, Sameer Deshmukh, Rio Yokota</i>	
Solving Linear Systems on a GPU with Hierarchically Off-Diagonal Low-Rank Approximations.....	1210
<i>Chao Chen, Per-Gunnar Martinsson</i>	
Graph Neural Networks Based Memory Inefficiency Detection Using Selective Sampling.....	1225
<i>Pengcheng Li, Yixin Guo, Yingwei Luo, Xiaolin Wang, Zhenlin Wang, Xu Liu</i>	
Predicting Reuse Interval for Optimized Web Caching: An LSTM-Based Machine Learning Approach	1239
<i>Pengcheng Li, Yixin Guo, Yongbin Gu</i>	
Out of Hypervisor (OoH): Efficient Dirty Page Tracking in Userspace Using Hardware Virtualization Features.....	1254
<i>Stella Bitchebe, Alain Tchana</i>	

Author Index