

SC19: International Conference for High Performance Computing, Networking, Storage and Analysis

**Denver, Colorado, USA
17-22 November 2019**

Pages 1-636



**IEEE Catalog Number: CFP19SUP-POD
ISBN: 978-1-6654-8391-9**

**Copyright © 2019, 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:	CFP19SUP-POD
ISBN (Print-On-Demand):	978-1-6654-8391-9
ISBN (Online):	978-1-4503-6229-0
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

A Data-Centric Approach to Extreme-Scale Ab Initio Dissipative Quantum Transport Simulations	1
<i>Alexandros Nikolaos Ziogas, Tal Ben-Nun, Guillermo Indalecio Fernández, Timo Schneider, Mathieu Luisier, Torsten Hoefler</i>	
Fast, Scalable and Accurate Finite-Element Based Ab Initio Calculations Using Mixed Precision Computing: 46 PFLOPS Simulation of a Metallic Dislocation System	14
<i>Sambit Das, Phani Motamarri, Vikram Gavini, Bruno Turcksin, Ying Wai Li, Brent Leback</i>	
GraphM: An Efficient Storage System for High Throughput of Concurrent Graph Processing	25
<i>Jin Zhao, Yu Zhang, Xiaofei Liao, Ligang He, Bingsheng He, Hai Jin, Haikun Liu, Yicheng Chen</i>	
Semantic Query Transformations for Increased Parallelization in Distributed Knowledge Graph Query Processing	39
<i>Hyeongsik Kim, Abhisha Bhattacharyya, Kemafor Anyanwu</i>	
MIQS: Metadata Indexing and Querying Service for Self-Describing File Formats	53
<i>Wei Zhang, Suren Byna, Houjun Tang, Brody Williams, Yong Chen</i>	
Scalable Simulation of Realistic Volume Fraction Red Blood Cell Flows Through Vascular Networks	77
<i>Libin Lu, Matthew J. Morse, Abtin Rahimian, Georg Stadler, Denis Zorin</i>	
Adaptive Neural Network-Based Approximation to Accelerate Eulerian Fluid Simulation	107
<i>Wenqian Dong, Jie Liu, Zhen Xie, Dong Li</i>	
GPU Acceleration of Extreme Scale Pseudo-Spectral Simulations of Turbulence Using Asynchronism	129
<i>Kiran Ravikumar, David Appelhans, P. K. Yeung</i>	
Large-Batch Training for LSTM and Beyond	151
<i>Yang You, Jonathan Hseu, Chris Ying, James Demmel, Kurt Keutzer, Cho-Jui Hsieh</i>	
Channel and Filter Parallelism for Large-Scale CNN Training	167
<i>Nikoli Dryden, Naoya Maruyama, Tim Moon, Tom Benson, Marc Snir, Brian Van Essen</i>	
SparCML: High-Performance Sparse Communication for Machine Learning	187
<i>Cedric Renggli, Saleh Ashkboos, Mehdi Aghagolzadeh, Dan Alistarh, Torsten Hoefler</i>	
Spread-N-Share: Improving Application Performance and Cluster Throughput with Resource-aware Job Placement	202
<i>Xiongchao Tang, Haojie Wang, Xiaosong Ma, Nosayba El-Sayed, Jidong Zhai, Wenguang Chen, Ashraf Aboulnaga</i>	
Swift Machine Learning Model Serving Scheduling: A Region Based Reinforcement Learning Approach	217
<i>Heyang Qin, Syed Zawad, Yanqi Zhou, Lei Yang, Dongfang Zhao, Feng Yan</i>	
Slack Squeeze Coded Computing for Adaptive Straggler Mitigation	240
<i>Krishna Giri Narra, Zhifeng Lin, Mehrdad Kiamari, Salman Avestimehr, Murali Annavaram</i>	
Practical and Efficient Incremental Adaptive Routing for HyperX Networks	256
<i>Nic McDonald, Mikhail Isaev, Adriana Flores, Al Davis, John Kim</i>	

Mitigating Network Noise on Dragonfly Networks Through Application-Aware Routing	269
<i>Daniele De Sensi, Salvatore Di Girolamo, Torsten Hoeﬂer</i>	
Topology-Custom UGAL Routing on Dragonfly	301
<i>Md Shafayat Rahman, Saptarshi Bhowmik, Yevgeniy Ryzanskiy, Xin Yuan, Michael Lang</i>	
ComDetective: A Lightweight Communication Detection Tool for Threads.....	316
<i>Muhammad Aditya Sasongko, Milind Chabbi, Palwisha Akhtar, Didem Unat</i>	
Pinpointing Performance Inefficiencies Via Lightweight Variance Profiling.....	337
<i>Pengfei Su, Shuyin Jiao, Milind Chabbi, Xu Liu</i>	
Hatchet: Pruning the Overgrowth in Parallel Profiles	356
<i>Abhinav Bhatele, Stephanie Brink, Todd Gamblin</i>	
Diogenes: Looking for an Honest CPU/GPU Performance Measurement Tool.....	377
<i>Benjamin Welton, Barton P. Miller</i>	
D2P: from Recursive Formulations to Distributed-Memory Codes	397
<i>Nikhil Hegde, Qifan Chang, Milind Kulkarni</i>	
Legate NumPy: Accelerated and Distributed Array Computing.....	419
<i>Michael Bauer, Michael Garland</i>	
Red-Blue Pebbling Revisited: Near Optimal Parallel Matrix-Matrix Multiplication	442
<i>Grzegorz Kwasniewski, Marko Kabic, Maciej Besta, Joost Vandevondele, Raffaele Solcà, Torsten Hoeﬂer</i>	
AutoFFT: A Template-Based FFT Codes Auto-Generation Framework for ARM and X86 CPUs	464
<i>Zhihao Li, Haipeng Jia, Yunquan Zhang, Tun Chen, Liang Yuan, Luning Cao, Xiao Wang</i>	
SLATE: Design of a Modern Distributed and Accelerated Linear Algebra Library.....	479
<i>Mark Gates, Jakub Kurzak, Ali Charara, Asim Yarkhan, Jack Dongarra</i>	
Uncore Power Scavenger: A Runtime for Uncore Power Conservation on HPC Systems.....	497
<i>Neha Gholkar, Frank Mueller, Barry Rountree</i>	
PoDD: Power-Capping Dependent Distributed Applications.....	520
<i>Huazhe Zhang, Henry Hoffmann</i>	
Etalumis: Bringing Probabilistic Programming to Scientific Simulators at Scale.....	543
<i>Atilim Gunes Baydin, Lei Shao, Wahid Bhimji, Lukas Heinrich, Lawrence Meadows, Jialin Liu, Andreas Munk, Saeid Naderiparizi, Bradley Gram-Hansen, Gilles Louppe, Mingfei Ma, Xiaohui Zhao, Philip Torr, Victor Lee, Kyle Cranmer, Prabhat, Frank Wood</i>	
Predicting Faults in High Performance Computing Systems: An In-Depth Survey of the State-of-the-Practice	567
<i>David Jauk, Dai Yang, Martin Schulz</i>	
A Large-Scale Study of MPI Usage in Open-Source HPC Applications.....	580
<i>Ignacio Laguna, Ryan Marshall, Kathryn Mohror, Martin Ruefenacht, Anthony Skjellum, Nawrin Sultana</i>	

Preparation and Optimization of a Diverse Workload for a Large-Scale Heterogeneous System	594
<i>Ian Karlin, Yoonho Park, Bronis R. De Supinski, Peng Wang, Bert Still, David Beckingsale, Robert Blake, Tong Chen, Guojing Cong, Carlos Costa, Johann Dahm, Giacomo Domeniconi, Thomas Epperly, Aaron Fisher, Sara Kokkila Schumacher, Steven Langer, Hai Le, Eun Kyung Lee, Naoya Maruyama, Xinyu Que, David Richards, Bjorn Sjogreen, Jonathan Wong, Carol Woodward, Ulrike Yang, Xiaohua Zhang, Bob Anderson, David Appelhans, Levi Barnes, Peter Barnes, Sorin Bastea, David Boehme, Jamie A. Bramwell, Jim Brase, Jose Brunheroto, Barry Chen, Charway R. Cooper, Tony Degroot, Rob Falgout, Todd Gamblin, David Gardner, James Glosli, John Gunnels, Max Katz, Tzanio Kolev, I-Feng W. Kuo, Matthew P. Legendre, Ruipeng Li, Pei-Hung Lin, Shelby Lockhart, Kathleen McCandless, Claudia Misale, Jaime Moreno, Rob Neely, Jarom Nelson, Rao Nimmakayala, Kathryn O'Brien, Kevin O'Brien, Ramesh Pankajakshan, Roger Pearce, Slaven Peles, Phil Regier, Steve Rennich, Martin Schulz, Howard Scott, James Sexton, Kathleen Shoga, Shiv Sundram, Guillaume Thomas-Collignon, Brian Van Essen, Alexey Voronin, Bob Walkup, Lu Wang, Chris Ward, Hui-Fang Wen, Dan White, Christopher Young, Cyril Zeller, Ed Zywick</i>	
Significantly Improving Lossy Compression Quality Based on an Optimized Hybrid Prediction Model	611
<i>Xin Liang, Sheng Di, Sihuan Li, Dingwen Tao, Bogdan Nicolae, Zizhong Chen, Franck Cappello</i>	
Moment Representation in the Lattice Boltzmann Method on Massively Parallel Hardware	637
<i>Madhurima Vardhan, John Gounley, Luiz Hegele, Erik W. Draeger, Amanda Randles</i>	
Slim Graph: Practical Lossy Graph Compression for Approximate Graph Processing, Storage, and Analytics.....	658
<i>Maciej Besta, Simon Weber, Lukas Gianinazzi, Robert Gerstenberger, Andrey Ivanov, Yishai Oltchik, Torsten Hoefler</i>	
PruneTrain: Fast Neural Network Training by Dynamic Sparse Model Reconfiguration	683
<i>Sangkug Lym, Esha Choukse, Siavash Zangeneh, Wei Wen, Sujay Sanghavi, Mattan Erez</i>	
Scalable Reinforcement-Learning-Based Neural Architecture Search for Cancer Deep Learning Research	696
<i>Prasanna Balaprakash, Romain Egele, Misha Salim, Stefan Wild, Venkatram Vishwanath, Fangfang Xia, Tom Brettin, Rick Stevens</i>	
BSTC: A Novel Binarized-Soft-Tensor-Core Design for Accelerating Bit-Based Approximated Neural Nets.....	729
<i>Ang Li, Tong Geng, Tianqi Wang, Martin Herbordt, Shuaiwen Leon Song, Kevin Barker</i>	
An Evaluation of the CORAL Interconnects	759
<i>Christopher Zimmer, Scott Atchley, Ramesh Pankajakshan, Brian E. Smith, Ian Karlin, Matthew L. Leininger, Adam Bertsch, Brian S. Ryujin, Jason Burmark, Andre Walker-Loudo, M. A. Clark, Olga Pearce</i>	
HyperX Topology: First At-Scale Implementation and Comparison to the Fat-Tree	777
<i>Jens Domke, Satoshi Matsuoka, Ivan R. Ivanov, Yuki Tsushima, Tomoya Yuki, Akihiro Nomura, Shin'Ichi Miura, Nic McDonald, Dennis L. Floyd, Nicolas Dubé</i>	
Bandwidth Steering in HPC Using Silicon Nanophotonics.....	800
<i>George Michelogiannakis, Yiwen Shen, Min Yee Teh, Xiang Meng, Benjamin Aivazi, Taylor Groves, John Shalf, Madeleine Glick, Manya Ghobadi, Larry Dennison, Keren Bergman</i>	

GPCNeT: Designing a Benchmark Suite for Inducing and Measuring Contention in HPC Networks	825
<i>Sudheer Chunduri, Taylor Groves, Peter Mendygral, Brian Austin, Jacob Balma, Krishna Kandalla, Kalyan Kumaran, Glenn Lockwood, Scott Parker, Steven Warren, Nathan Wichmann, Nicholas Wright</i>	
Understanding Congestion in High Performance Interconnection Networks Using Sampling	858
<i>Philip Taffet, John Mellor-Crummey</i>	
TriEC: Tripartite Graph Based Erasure Coding NIC Offload.....	882
<i>Haiyang Shi, Xiaoyi Lu</i>	
A Constraint-Based Approach to Automatic Data Partitioning for Distributed Memory Execution	916
<i>Wonchan Lee, Manolis Papadakis, Elliott Slaughter, Alex Aiken</i>	
Understanding Priority-Based Scheduling of Graph Algorithms on a Shared-Memory Platform.....	940
<i>Serif Yesil, Azin Heidarshenas, Adam Morrison, Josep Torrellas</i>	
Almost Deterministic Work Stealing	954
<i>Shumpei Shiina, Kenjiro Taura</i>	
Conflict-Free Symmetric Sparse Matrix-Vector Multiplication on Multicore Architectures.....	970
<i>Athena Elafrou, Georgios Goumas, Nectarios Koziris</i>	
An Efficient Mixed-Mode Representation of Sparse Tensors	985
<i>Israt Nisa, Jiajia Li, Aravind Sukumaran-Rajam, Prasant Singh Rawat, Sriram Krishnamoorthy, P. Sadayappan</i>	
Regularizing Irregularly Sparse Point-To-Point Communications	1010
<i>Oguz Selvitopi, Cevdet Aykanat</i>	
Compiler Assisted Hybrid Implicit and Explicit GPU Memory Management Under Unified Address Space	1024
<i>Lingda Li, Barbara Chapman</i>	
Exploiting Reuse and Vectorization in Blocked Stencil Computations on CPUs and GPUs.....	1040
<i>Tuowen Zhao, Protonu Basu, Samuel Williams, Mary Hall, Hans Johansen</i>	
A Versatile Software Systolic Execution Model for GPU Memory-Bound Kernels	1084
<i>Peng Chen, Mohamed Wahib, Shinichiro Takizawa, Ryousei Takano, Satoshi Matsuoka</i>	
INCA: In-Network Compute Assistance	1165
<i>Whit Schonbein, Ryan E. Grant, Matthew G. F. Dosanjh, Dorian Arnold</i>	
Near-Memory Data Transformation for Efficient Sparse Matrix Multi-Vector Multiplication	1178
<i>Daichi Fujiki, Niladrish Chatterjee, Donghyuk Lee, Mike O'Connor</i>	
Network-Accelerated Non-Contiguous Memory Transfers.....	1195
<i>Salvatore Di Girolamo, Konstantin Taranov, Andreas Kurth, Michael Schaffner, Timo Schneider, Jakub Beránek, Maciej Besta, Luca Benini, Duncan Roweth, Torsten Hoefler</i>	

A Massively Parallel Infrastructure for Adaptive Multiscale Simulations: Modeling RAS Initiation Pathway for Cancer	1209
<i>Francesco Di Natale, Chris Neale, Liam Stanton, Thomas R. W. Scogland, Yue Yang, Carlos Costa, Sandrasegaram Gnanakaran, Felice C. Lightstone, Harsh Bhatia, Sara Kokkila Schumacher, Xiaohua Zhang, Gautham Dharuman, Claudia Misale, Changhoan Kim, Dwight V. Nissley, Peer-Timo Bremer, Helgi I. Ingólfsson, Timothy S. Carpenter, Tomas Ooppelstrup, Shiv Sundram, Michael P. Surh, Lars Schneidenbach, Bruce D'Amora, Fred Streitz, James N. Glosli</i>	
CARE: Compiler-Assisted Recovery from Soft Failures	1225
<i>Chao Chen, Greg Eisenhauer, Santosh Pande, Qiang Guan</i>	
Code Generation for Massively Parallel Phase-Field Simulations	1248
<i>Martin Bauer, Johannes Hötzer, Dominik Ernst, Julian Hammer, Marco Seiz, Henrik Hierl, Jan Hönig, Harald Köstler, Gerhard Wellein, Britta Nestler, Ulrich Rüde</i>	
Local-Global Merge Tree Computation with Local Exchanges	1280
<i>Arnur Nigmatov, Dmitriy Morozov</i>	
Solving PDEs in Space-Time: 4D Tree-Based Adaptivity, Mesh-Free and Matrix-Free Approaches	1293
<i>Masado Ishii, Milinda Fernando, Kumar Saurabh, Biswajit Khara, Baskar Ganapathysubramanian, Hari Sundar</i>	
From Piz Daint to the Stars: Simulation of Stellar Mergers Using High-Level Abstractions.....	1315
<i>Gregor Daiß, Parsa Amini, John Biddiscombe, Patrick Diehl, Juhan Frank, Kevin Huck, Hartmut Kaiser, Dominic Marcello, David Pfander, Dirk Pflüger</i>	
End-To-end I/O Portfolio for the Summit Supercomputing Ecosystem	1352
<i>Sarp Oral, Sudharshan S. Vazhkudai, Feiyi Wang, Christopher Zimmer, Christopher Brumgard, Jesse Hanley, George Markomanolis, Ross Miller, Dustin Leverman, Scott Atchley, Veronica Vergara Larrea</i>	
From Facility to Application Sensor Data: Modular, Continuous and Holistic Monitoring with DCDB.....	1366
<i>Alessio Netti, Micha Müller, Axel Auweter, Carla Guillen, Michael Ott, Daniele Tafani, Martin Schulz</i>	
Revisiting I/O Behavior in Large-Scale Storage Systems: The Expected and the Unexpected.....	1393
<i>Tirthak Patel, Suren Byna, Glenn K. Lockwood, Devesh Tiwari</i>	
SW_GROMACS: Accelerate GROMACS on Sunway TaihuLight.....	1406
<i>Tingjian Zhang, Yuxuan Li, Ping Gao, Qi Shao, Mingshan Shao, Meng Zhang, Jinxiao Zhang, Xiaohui Duan, Zhao Liu, Lin Gan, Haohuan Fu, Wei Xue, Weiguo Liu, Guangwen Yang</i>	
Fully Integrated FPGA Molecular Dynamics Simulations	1420
<i>Chen Yang, Tong Geng, Tianqi Wang, Rushi Patel, Qingqing Xiong, Ahmed Sanaullah, Chunshu Wu, Jiayi Sheng, Charles Lin, Vipin Sachdeva, Woody Sherman, Martin Herbordt</i>	
OpenKMC: A KMC Design for Hundred-Billion-Atom Simulation Using Millions of Cores on Sunway Taihulight.....	1451
<i>Kun Li, Honghui Shang, Yunquan Zhang, Shigang Li, Baodong Wu, Dong Wang, Libo Zhang, Fang Li, Dexun Chen, Zhiqiang Wei</i>	
BinFI: An Efficient Fault Injector for Safety-Critical Machine Learning Systems	1467
<i>Zitao Chen, Guanpeng Li, Karthik Pattabiraman, Nathan Debardeleben</i>	

Assessing the Impact of Timing Errors on HPC Applications.....	1490
<i>Chun-Kai Chang, Wenqi Yin, Mattan Erez</i>	
FT-ISort: Efficient Fault Tolerance for Introsort	1509
<i>Sihuan Li, Hongbo Li, Xin Liang, Jieyang Chen, Elisabeth Giem, Kaiming Ouyang, Kai Zhao, Sheng Di, Franck Cappello, Zizhong Chen</i>	
Distributed Enhanced Suffix Arrays: Efficient Algorithms for Construction and Querying	1526
<i>Patrick Flick, Srinivas Aluru</i>	
Scalable Generation of Graphs for Benchmarking HPC Community-Detection Algorithms.....	1543
<i>George M. Slota, Jonathan W. Berry, Simon D. Hammond, Stephen L. Olivier, Cynthia A. Phillips, Sivasankaran Rajamanickam</i>	
Analytical Cache Modeling and Tilesize Optimization for Tensor Contractions	1557
<i>Rui Li, Aravind Sukumaran-Rajam, Richard Veras, Tze Meng Low, Fabrice Rastello, Atanas Rountev, P. Sadayappan</i>	
SSD Failures in the Field: Symptoms, Causes, and Prediction Models	1570
<i>Jacob Alter, Ji Xue, Alma Dimnaku, Evgenia Smirni</i>	
An Early Evaluation of Intel's Optane DC Persistent Memory Module and Its Impact on High-Performance Scientific Applications	1584
<i>Michèle Weiland, Holger Brunst, Tiago Quintino, Nick Johnson, Olivier Iffrig, Simon Smart, Christian Herold, Antonino Bonanni, Adrian Jackson, Mark Parsons</i>	
Performance Optimality Or Reproducibility: That is the Question	1603
<i>Tapasya Patki, Jayaraman J. Thiagarajan, Alexis Ayala, Tanzima Z. Islam</i>	
Optimizing the Data Movement in Quantum Transport Simulations Via Data-Centric Parallel Programming	1633
<i>Alexandros Nikolaos Ziogas, Tal Ben-Nun, Guillermo Indalecio Fernández, Timo Schneider, Mathieu Luisier, Torsten Hoefler</i>	
Parallel Transport Time-Dependent Density Functional Theory Calculations with Hybrid Functional on Summit	1650
<i>Weile Jia, Lin-Wang Wang, Lin Lin</i>	
Full-State Quantum Circuit Simulation by Using Data Compression	1673
<i>Xin-Chuan Wu, Sheng Di, Emma Maitreyee Dasgupta, Franck Cappello, Hal Finkel, Yuri Alexeev, Frederic T. Chong</i>	
Stateful Dataflow Multigraphs: A Data-Centric Model for Performance Portability on Heterogeneous Architectures	1697
<i>Tal Ben-Nun, Johannes De Fine Licht, Alexandros N. Ziogas, Timo Schneider, Torsten Hoefler</i>	
Streaming Message Interface: High-Performance Distributed Memory Programming on Reconfigurable Hardware.....	1711
<i>Tiziano De Matteis, Johannes De Fine Licht, Jakub Beránek, Torsten Hoefler</i>	
High Performance Monte Carlo Simulation of Ising Model on TPU Clusters	1744
<i>Kun Yang, Yi-Fan Chen, Georgios Roumpos, Chris Colby, John Anderson</i>	
IFDK: A Scalable Framework for Instant High-Resolution Image Reconstruction.....	1759
<i>Peng Chen, Mohamed Wahib, Shinichiro Takizawa, Ryousei Takano, Satoshi Matsuoka</i>	

MemXCT: Memory-Centric X-Ray CT Reconstruction with Massive Parallelization	1783
<i>Mert Hidayetoglu, Tekin Biçer, Simon Garcia De Gonzalo, Bin Ren, Doga Gürsoy, Rajkumar Kettimuthu, Ian T. Foster, Wen-Mei W. Hwu</i>	
Consensus Equilibrium Framework for Super-Resolution and Extreme-Scale CT Reconstruction	1839
<i>Xiao Wang, Venkatesh Sridhar, Zahra Ronaghi, Rollin Thomas, Jack Deslippe, Dilworth Parkinson, Gregory T. Buzzard, Samuel P. Midkiff, Charles A. Bouman, Simon K. Warfield</i>	
Addressing Data Resiliency for Staging Based Scientific Workflows	1862
<i>Shaohua Duan, Pradeep Subedi, Philip E. Davis, Manish Parashar</i>	
LPCC: Hierarchical Persistent Client Caching for Lustre	1884
<i>Yingjin Qian, Xi Li, Shuichi Ihara, Andreas Dilger, Carlos Thomaz, Shilong Wang, Wen Cheng, Chunyan Li, Lingfang Zeng, Fang Wang, Dan Feng, Tim Süß, André Brinkmann</i>	
Replication is More Efficient than You Think	1898
<i>Anne Benoit, Thomas Herault, Valentin Le Fèvre, Yves Robert</i>	

Author Index