

SC21: International Conference for High Performance Computing, Networking, Storage and Analysis

**St. Louis, Missouri, USA
14 – 19 November 2021**

Pages 1-731



**IEEE Catalog Number: CFP21SUP-POD
ISBN: 978-1-6654-8390-2**

Printed from e-media with permission by:

Curran Associates, Inc.
57 Morehouse Lane
Red Hook, NY 12571



Some format issues inherent in the e-media version may also appear in this print version.

This work is licensed under a Creative Commons Attribution 4.0 International License.
License details: <http://creativecommons.org/licenses/by/4.0/>.

No changes have been made to the content of these proceedings. There may be changes to pagination, and minor adjustments for aesthetics.

Printed with permission by Curran Associates, Inc. (2022)

For permission requests, please contact the publisher:

The Association for Computing Machinery
1601 Broadway, 10th Floor,
New York, New York 10019, USA

Email: permissions@acm.org

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
Email: curran@proceedings.com
Web: www.proceedings.com

TABLE OF CONTENTS

Anton 3: Twenty Microseconds of Molecular Dynamics Simulation Before Lunch	1
<i>David E. Shawa, Peter J. Adams, Asaph Azaria, Joseph A. Bank, Brannon Batson, Alistair Bell, Michael Bergdorf, Jhanvi Bhatt, J. Adam Butts, Timothy Correia, Robert M. Dirks, Ron O. Dror, Michael P. Eastwood, Bruce Edwards, Amos Even, Peter Feldmann, Michael Fenn, Christopher H. Fenton, Anthony Forte, Joseph Gagliardo, Gennette Gill, Maria Gorlatova, Brian Greskamp, J. P. Grossman, Justin Gullingsrud, Anissa Harper, William Hasenplaugh, Mark Heily, Benjamin Colin Heshmat, Jeremy Hunt, Douglas J. Ierardi, Lev Iserovich, Bryan L. Jackson, Nick P. Johnson, Mollie M. Kirk, John L. Klepeis, Jeffrey S. Kuskin, Kenneth M. Mackenzie, Roy J. Mader, Richard McGowen, Adam McLaughlin, Mark A. Moraes, Mohamed H. Nasrb, Lawrence J. Nociolo, Lief O'Donnell, Andrew Parker, Jon L. Peticolas, Goran Pocina, Cristian Predescu, Terry Quan, John K. Salmon, Carl Schwink, Keun Sup Shim, Naseer Siddiqueb, Jochen Spenglerb, Tamas Szalay, Raymond Tabladillo, Reinhard Tartler, Andrew G. Taube, Michael Theobald, Brian Towlesb, William Vick, Stanley C. Wang, Michael Wazłowski, Madeleine J. Weingarten, John M. Williams, Kevin A. Yuh</i>	
Generalizable Coordination of Large Multiscale Workflows: Challenges and Learnings at Scale	12
<i>Harsh Bhatia, Francesco Di Natale, Joseph Y. Moon, Xiaohua Zhang, Joseph R. Chavez, Fikret Aydin, Chris Stanley, Tomas Ooppelstrup, Chris Neale, Sara Kokkila Schumacher, Dong H. Ahn, Stephen Herbein, Timothy S. Carpenter, Sandrasegaram Gnanakaran, Peer-Timo Bremer, James N. Glosli, Felice C. Lightstone, Helgi I. Ingolfsson</i>	
Online Evolutionary Batch Size Orchestration for Scheduling Deep Learning Workloads in GPU Clusters.....	28
<i>Zhengda Bian, Shenggui Li, Wei Wang, Yang You</i>	
Whale: Efficient One-To-Many Data Partitioning in RDMA-Assisted Distributed Stream Processing Systems	41
<i>Jie Tan, Hanhua Chen, Yonghui Wang, Hai Jin</i>	
Exploiting User Activeness for Data Retention in HPC Systems.....	55
<i>Wei Zhang, Suren Byna, Hyogi Sim, Sangkeun Lee, Sudharshan Vazhkudai, Yong Chen</i>	
Pinpointing Crash-Consistency Bugs in the HPC I/O Stack: A Cross-Layer Approach	69
<i>Jinghan Sun, Jian Huang, Marc Snir</i>	
Characterization and Prediction of Deep Learning Workloads in Large-Scale GPU Datacenters	83
<i>Qinghao Hu, Peng Sun, Shengen Yan, Yonggang Wen, Tianwei Zhang</i>	
Linux vs. Lightweight Multi-Kernels for High Performance Computing: Experiences at Pre-Exascale.....	98
<i>Balazs Gerofi, Kohei Tarumizu, Lei Zhang, Takayuki Okamoto, Masamichi Takagi, Shinji Sumimoto, Yutaka Ishikawa</i>	
Revealing Power, Energy and Thermal Dynamics of a 200PF Pre-Exascale Supercomputer.....	111
<i>Woong Shin, Vladyslav Oles, Ahmad Maroof Karimi, J. Austin Ellis, Feiyi Wang</i>	
KAISA: An Adaptive Second-Order Optimizer Framework for Deep Neural Networks	129
<i>J. Gregory Pauloski, Shivaram Venkataraman, Qi Huang, Kyle Chard, Zhao Zhang, Lei Huang, Ian Foster</i>	

Tensor Processing Primitives: A Programming Abstraction for Efficiency and Portability in Deep Learning Workloads	143
<i>Evangelos Georganas, Dhiraj Kalamkar, Sasikanth Avancha, Menachem Adelman, Cristina Anderson, Alexander Breuer, Jeremy Bruestle, Narendra Chaudhary, Abhisek Kundu, Denise Kutnick, Frank Laub, Vasimuddin Md, Sanchit Misra, Ramanarayan Mohanty, Hans Pabst, Barukh Ziv, Alexander Heinecke</i>	
Enable Simultaneous DNN Services Based on Deterministic Operator Overlap and Precise Latency Prediction	159
<i>Weihao Cui, Han Zhao, Quan Chen, Ningxin Zheng, Jingwen Leng, Jieru Zhao, Zhuo Song, Tao Ma, Yong Yang, Chao Li, Minyi Guo</i>	
Distributed Quantum Computing with QMPI	174
<i>Thomas Häner, Damian S. Steiger, Torsten Hoefler, Matthias Troyer</i>	
BAASH: Lightweight, Efficient, and Reliable Blockchain-As-A-Service for HPC Systems	189
<i>Abdullah Al Mamun, Feng Yan, Dongfang Zhao</i>	
Representation of Women in HPC Conferences.....	204
<i>Eitan Frachtenberg, Rhody D. Kaner</i>	
Preparing an Incompressible-Flow Fluid Dynamics Code for Exascale-Class Wind Energy Simulations.....	216
<i>Paul Mullaney, Ruipeng Li, Stephen Thomas, Shreyas Ananthan, Ashesh Sharma, Jon S. Rood, Alan B. Williams, Michael A. Sprague</i>	
Symplectic Structure-Preserving Particle-In-Cell Whole-Volume Simulation of Tokamak Plasmas to 111.3 Trillion Particles and 25.7 Billion Grids.....	229
<i>Jianyuan Xiao, Junshi Chen, Jiangshan Zheng, Hong An, Shenghong Huang, Chao Yang, Fang Li, Ziyu Zhang, Yeqi Huang, Wenting Han, Xin Liu, Dexun Chen, Zixi Liu, Ge Zhuang, Jiale Chen, Guoqiang Li, Xuan Sun, Qiang Chen</i>	
Scalable Adaptive PDE Solvers in Arbitrary Domains	242
<i>Saurabh Kumar, Masado Ishii, Milinda Fernando, Boshun Gao, Kendrick Tan, Ming-Chen Hsu, Adarsh Krishnamurthy, Hari Sundar, Baskar Ganapathysubramanian</i>	
A Next-Generation Discontinuous Galerkin Fluid Dynamics Solver with Application to High-Resolution Lung Airflow Simulations.....	260
<i>Martin Kronbichler, Niklas Fehn, Peter Munch, Maximilian Bergbauer, Karl-Robert Wichmann, Carolin Geitner, Momme Allalen, Martin Schulz, Wolfgang A. Wall</i>	
Understanding, Predicting and Scheduling Serverless Workloads Under Partial Interference	274
<i>Laiping Zhao, Yanan Yang, Yiming Li, Xian Zhou, Keqiu Li</i>	
The Hidden Cost of the Edge: A Performance Comparison of Edge and Cloud Latencies.....	288
<i>Ahmed Ali-Eldin, Bin Wang, Prashant Shenoy</i>	
RIBBON: Cost-Effective and QoS-Aware Deep Learning Model Inference Using a Diverse Pool of Cloud Computing Instances	303
<i>Baolin Li, Rohan Basu Roy, Tirthak Patel, Vijay Gadepally, Karen Gettings, Devesh Tiwari</i>	
E.T.: Re-Thinking Self-Attention for Transformer Models on GPUs	319
<i>Shiyang Chen, Shaoyi Huang, Santosh Pandey, Bingbing Li, Guang R. Gao, Long Zheng, Caiwen Ding, Hang Liu</i>	
Parallel Construction of Module Networks	333
<i>Ankit Srivastava, Sriram P. Chockalingam, Maneesha Aluru, Srinivas Aluru</i>	

Chimera: Efficiently Training Large-Scale Neural Networks with Bidirectional Pipelines.....	347
<i>Shigang Li, Torsten Hoefler</i>	
Bootstrapping In-Situ Workflow Auto-Tuning Via Combining Performance Models of Component Applications.....	361
<i>Tong Shu, Yanfei Guo, Justin Wozniak, Xiaoning Ding, Ian Foster, Tahsin Kurc</i>	
Meeting the Real-Time Challenges of Ground-Based Telescopes Using Low-Rank Matrix Computations.....	376
<i>Hatem Ltaief, Jesse Cranney, Damien Gratadour, Yuxi Hong, Laurent Gatineau, David Keyes</i>	
Closing the “Quantum Supremacy” Gap: Achieving Real-Time Simulation of a Random Quantum Circuit Using a New Sunway Supercomputer	390
<i>Yong Liu, Xin Liu, Fang Li, Haohuan Fu, Yuling Yang, Jiawei Song, Pengpeng Zhao, Zhen Wang, Dajia Peng, Huarong Chen, Chu Guo, Heliang Huang, Wenzhao Wu, Dexun Chen</i>	
AgEBO-Tabular: Joint Neural Architecture and Hyperparameter Search with Autotuned Data-Parallel Training for Tabular Data.....	402
<i>Romain Égelé, Prasanna Balaprakash, Isabelle Guyon, Venkatram Vishwanath, Fangfang Xia, Rick Stevens, Zhengying Liu</i>	
Non-Recurring Engineering (NRE) Best Practices: A Case Study with the NERSC/NVIDIA OpenMP Contract.....	418
<i>Christopher S. Daley, Annemarie Southwell, Rahulkumar Gayatri, Scott Biersdorff, Craig Toepfer, Güray Özen, Nicholas J. Wright</i>	
Minimizing Privilege for Building HPC Containers	432
<i>Reid Priedhorsky, R. Shane Canon, Timothy Randles, Andrew J. Younge</i>	
Systematically Inferring I/O Performance Variability by Examining Repetitive Job Behavior.....	445
<i>Emily Costa, Tirthak Patel, Benjamin Schwaller, Jim M. Brandt, Devesh Tiwari</i>	
SEEC: Stochastic Escape Express Channel.....	460
<i>Mayank Parasar, Natalie Enright Jerger, Paul V. Gratz, Joshua San Miguel, Tushar Krishna</i>	
Flare: Flexible In-Network Allreduce	474
<i>Daniele De Sensi, Salvatore Di Girolamo, Saleh Ashkboos, Shigang Li, Torsten Hoefler</i>	
HatRPC: Hint-Accelerated Thrift RPC Over RDMA.....	489
<i>Tianxi Li, Haiyang Shi, Xiaoyi Lu</i>	
APNN-TC: Accelerating Arbitrary Precision Neural Networks on Ampere GPU Tensor Cores.....	504
<i>Boyuan Feng, Yuke Wang, Tong Gena, Ang Li, Yufei Ding</i>	
Scalable Edge-Based Hyperdimensional Learning System with Brain-Like Neural Adaptation	518
<i>Zhuowen Zou, Yeseong Kim, Farhad Imani, Haleh Alimohamadi, Rosario Cammarota, Mohsen Imani</i>	
Dr. Top-k: Delegate-Centric Top-k on GPUs.....	533
<i>Anil Gaihre, Da Zheng, Scott Weitze, Lingda Li, Shuaiwen Leon Song, Caiwen Ding, Xiaoye S. Li, Hang Liu</i>	
Billion Atom Molecular Dynamics Simulations of Carbon at Extreme Conditions and Experimental Time and Length Scales	547
<i>Kien Nguyen-cong, Jonathan T. Willman, Stan G. Moore, Anatoly B. Belonoshko, Rahulkumar Gayatri, Evan Weinberg, Mitchell A. Wood, Aidan P. Thompson, Ivan I. Oleynik</i>	

Enabling Large-Scale Correlated Electronic Structure Calculations: Scaling the RI-MP2 Method on Summit	559
<i>Giuseppe M. J. Barca, Jorge L. Galvez Vallejo, David L. Poole, Melisa Alkan, Ryan Stocks, Alistair P. Rendell, Mark S. Gordon</i>	
Accelerating All-Electron <i>Ab initio</i> Simulation of Raman Spectra for Biological Systems.....	573
<i>Honghui Shang, Fang Li, Yunquan Zhang, Ying Liu, Libo Zhang, Mingchuan Wu, Yangjun Wu, Di Wei, Huimin Cui, Xin Liu, Fei Wang, Yuxi Ye, Yingxiang Gao, Shuang Ni, Xin Chen, Dexun Chen</i>	
LMFF: Efficient and Scalable Layered Materials Force Field on Heterogeneous Many-Core Processors.....	587
<i>Ping Gao, Xiaohui Duan, Jiaxu Guo, Jin Wang, Zhenya Song, Lizhen Cui, Xiangxu Meng, Xin Liu, Wusheng Zhang, Ming Ma, Guohui Li, Dexun Chen, Haohuan Fu, Wei Xue, Weiguo Liu, Guangwen Yang</i>	
Hardware Acceleration of Tensor-Structured Multilevel Ewald Summation Method on MDGRAPE-4A, a Special-Purpose Computer System for Molecular Dynamics Simulations	603
<i>Gentaro Morimoto, Yohei M. Koyama, Hao Zhang, Teruhisa S. Komatsu, Yousuke Ohno, Keigo Nishida, Itta Ohmura, Hiroshi Koyama, Makoto Taijit</i>	
Accelerating Bandwidth-Bound Deep Learning Inference with Main-Memory Accelerators	618
<i>Benjamin Y. Cho, Jeageun Jung, Mattan Erez</i>	
LCCG: A Locality-Centric Hardware Accelerator for High Throughput of Concurrent Graph Processing.....	632
<i>Jin Zhao, Yu Zhang, Xiaofei Liao, Liaang He, Binzheng He, Haikun Liu, Hai Jin</i>	
Simurgh: A Fully Decentralized and Secure NVMM User Space File System.....	647
<i>Nafiseh Moti, Frederic Schimmelpfennig, Reza Salkhordeh, David Klopp, Toni Cortes, Ulrich Rückert, André Brinkmann</i>	
Lunule: An Agile and Judicious Metadata Load Balancer for CephFS.....	663
<i>Yiduo Wang, Cheng Li, Xinyang Shao, Youxu Chen, Feng Yan, Yinlong Xu</i>	
DeltaFS: A Scalable No-Ground-Truth Filesystem for Massively-Parallel Computing.....	677
<i>Qing Zheng, Charles D. Cranor, Gregory R. Ganger, Garth A. Gibson, George Amvrosiadis, Bradley W. Settlemyer, Gary A. Grider</i>	
Distributed Multigrid Neural Solvers on Megavoxel Domains.....	692
<i>Aditya Balu, Sergio Botelho, Biswajit Khara, Vinay Rao, Soumik Sarkar, Chinmay Hegde, Adarsh Krishnamurthy, Santi Adavani, Baskar Ganapathysubramanian</i>	
A 400 Trillion-Grid Vlasov Simulation on Fugaku Supercomputer: Large-Scale Distribution of Cosmic Relic Neutrinos in a Six-Dimensional Phase Space	707
<i>Kohji Yoshikawa, Satoshi Tanaka, Naoki Yoshida</i>	
EIGA: Elastic and Scalable Dynamic Graph Analysis.....	718
<i>Kasimir Gabert, Kaan Sancak, M. Yusuf Ozkaya, Ali Pinar, Umit V. Catalyurek</i>	
Krill: A Compiler and Runtime System for Concurrent Graph Processing.....	732
<i>Hongzheng Chen, Minghua Shen, Nong Xiao, Yutong Lu</i>	
Pilgrim: Scalable and (near) Lossless MPI Tracing	746
<i>Chen Wang, Pavan Balaji, Marc Snir</i>	

Hybrid, Scalable, Trace-Driven Performance Modeling of GPGPUs	759
<i>Yehia Arafa, Abdel-Hameed Badawy, Ammar Elwazir, Atanu Barai, Ali Eker, Gopinath Chennupati, Nandakishore Santhi, Stephan Eidenbenz</i>	
G-SEPM: Building an Accurate and Efficient Soft Error Prediction Model for GPGPUs	774
<i>Hengshan Yue, Xiaohui Wei, Guangli Li, Jianpeng Zhao, Nan Jiang, Jingweijia Tan</i>	
Single-Node Partitioned-Memory for Huge Graph Analytics: Cost and Performance Trade-Offs	788
<i>Sayan Ghosh, Nathan R. Tallent, Marco Minutoli, Mahantesh Halappanavar, Ramesh Peri, Ananth Kalyanaraman</i>	
Accelerating Applications Using Edge Tensor Processing Units	804
<i>Kuan-Chieh Hsu, Hung-Wei Tseng</i>	
Enabling and Scaling the HPCG Benchmark on the Newest Generation Sunway Supercomputer with 42 Million Heterogeneous Cores	820
<i>Qianchao Zhu, Hao Luo, Chao Yang, Mingshuo Ding, Wanwang Yin, Xinhui Yuan</i>	
Efficient Large-Scale Language Model Training on GPU Clusters Using Megatron-LM	833
<i>Deepak Narayanan, Mohammad Shoeybi, Jared Casper, Patrick Legresley, Mostofa Patwary, Vijay Korthikanti, Dmitri Vainbrand, Prethvi Kashinkunti, Julie Bernauer, Bryan Catanzaro, Amar Phanishayee, Matei Zaharia</i>	
ZeRO-Infinity: Breaking the GPU Memory Wall for Extreme Scale Deep Learning	847
<i>Samyam Rajbhandari, Olatunji Ruwase, Jeff Rasley, Shaden Smith, Yuxiong He</i>	
Extreme-Scale <i>Ab initio</i> Quantum Raman Spectra Simulations on the Leadership HPC System in China	862
<i>Honghui Shang, Fang Li, Yunquan Zhang, Libo Zhang, You Fu, Yingxiang Gao, Yangjun Wu, Xiaohui Duan, Rongfen Lin, Xin Liu, Ying Liu, Dexun Chen</i>	
FedAT: A High-Performance and Communication-Efficient Federated Learning System with Asynchronous Tiers	875
<i>Zheng Chai, Yujing Chen, Ali Anwar, Liang Zhao, Yue Cheng, Huzefa Rangwala</i>	
Reverse-Mode Automatic Differentiation and Optimization of GPU Kernels Via Enzyme	892
<i>William S. Moses, Valentin Churavy, Ludger Paehler, Jan Hüchelheim, Sri Hari Krishna Narayanan, Michel Schanen, Johannes Doerfert</i>	
Overcoming Barriers to Scalability in Variational Quantum Monte Carlo	910
<i>Tianchen Zhao, Saibal De, Brian Chen, James Stokes, Shravan Veerapaneni</i>	
3D Acoustic-Elastic Coupling with Gravity: The Dynamics of the 2018 Palu, Sulawesi Earthquake and Tsunami	922
<i>Lukas Krenz, Carsten Uphoff, Thomas Ulrich, Alice-Agnes Gabriel, Lauren S. Abrahams, Eric M. Dunham, Michael Bader</i>	
In-Depth Analyses of Unified Virtual Memory System for GPU Accelerated Computing	939
<i>Tyler Allen, Rong Ge</i>	
Paths to OpenMP in the Kernel	953
<i>Jiacheng Ma, Wenyi Wang, Aaron Nelson, Michael Cuevas, Brian Homerding, Conghao Liu, Zhen Huang, Simone Campanoni, Kyle Hale, Peter Dinda</i>	
Index Launches: Scalable, Flexible Representation of Parallel Task Groups	968
<i>Rupanshu Soi, Michael Bauer, Sean Treichler, Manolis Papadakis, Wonchan Lee, Patrick McCormick, Alex Aiken, Elliott Slaughter</i>	

TriPoll: Computing Surveys of Triangles in Massive-Scale Temporal Graphs with Metadata.....	982
<i>Trevor Steil, Tahsin Reza, Keita Iwabuchi, Benjamin W. Priest, Geoffrey Sanders, Roger Pearce</i>	
Discovering and Balancing Fundamental Cycles in Large Signed Graphs	996
<i>Ghadeer Alabandi, Jelena Tesic, Lucas Rusnak, Martin Burtscher</i>	
cuTS: Scaling Subgraph Isomorphism on Distributed Multi-GPU Systems Using Trie Based Data Structure	1010
<i>Lizhi Xiang, Arif Khan, Edoardo Serra, Mahantesh Halappanavar, Aravind Sukumaran-Rajam</i>	
Accelerating Large Scale De Novo Metagenome Assembly Using GPUs.....	1023
<i>Muaaz Gul Awan, Steven Hofmeyr, Rob Egan, Nan Ding, Aydin Buluc, Jack Deslippe, Leonid Oliker, Katherine Yelick</i>	
On the Parallel I/O Optimality of Linear Algebra Kernels: Near-Optimal Matrix Factorizations	1034
<i>Grzegorz Kwasniewski, Marko Kabic, Tal Ben-Nun, Alexandros Nikolaos Ziogas, Jens Eirik Saethre, André Gaillard, Timo Schneider, Maciej Besta, Anton Kozhevnikov, Joost Vandevondele, Torsten Hoefler</i>	
STM-Multifrontal QR: Streaming Task Mapping Multifrontal QR Factorization Empowered by GCN	1049
<i>Shengle Lin, Wangdong Yang, Haotian Wang, Qinyun Tsai, Kenli Li</i>	
LIBSHALOM: Optimizing Small and Irregular-Shaped Matrix Multiplications on ARMv8 Multi-Cores.....	1063
<i>Weiling Yang, Jianbin Fang, Dezun Dong, Xing Su, Zheng Wang</i>	
TensorKMC: Kinetic Monte Carlo Simulation of 50 Trillion Atoms Driven by Deep Learning on a New Generation of Sunway Supercomputer	1078
<i>Honghui Shang, Xin Chen, Xingyu Gao, Rongfen Lin, Lifang Wang, Fang Li, Qian Xiao, Lei Xu, Qiang Sun, Leilei Zhu, Fei Wang, Yunquan Zhang, Haifeng Song</i>	
High-Throughput Virtual Screening of Small Molecule Inhibitors for SARS-CoV-2 Protein Targets with Deep Fusion Models.....	1092
<i>Garrett A. Stevenson, Derek Jones, Hyojin Kim, W. F. Drew Bennett, Brian J. Bennion, Monica Borucki, Feliza Bourguet, Aidan Epstein, Magdalena Franco, Brooke Harmon, Stewart He, Max P. Katz, Daniel Kirshner, Victoria Lao, Edmond Y. Lau, Jacky Lo, Kevin McLoughlin, Richard Mosesso, Deepa K. Muruges, Oscar A. Negrete, Edwin A. Saada, Brent Segelke, Maxwell Stefan, Marisa W. Torres, Dina Weilhammer, Sergio Wong, Yue Yang, Adam Zemla, Xiaohua Zhang, Fangqiang Zhu, Felice C. Lightstone, Jonathan E. Allen</i>	
High Performance Uncertainty Quantification with Parallelized Multilevel Markov Chain Monte Carlo.....	1107
<i>Linus Seelinger, Anne Reinartz, Leonhard Rannabauer, Michael Bader, Peter Bastian, Robert Scheichl</i>	
DistGNN: Scalable Distributed Training for Large-Scale Graph Neural Networks.....	1120
<i>Vasimuddin Md, Sanchit Misra, Guixiang Ma, Ramnarayan Mohanty, Evangelos Georganas, Alexander Heinecke, Dhiraj Kalamkar, Nesreen K. Ahmed, Sasikanth Avancha</i>	
Efficient Scaling of Dynamic Graph Neural Networks	1134
<i>Venkatesan T. Chakaravarthy, Shivmaran S. Pandian, Saurabh Raje, Yogish Sabharwal, Toyotaro Suzumura, Shashanka Ubaru</i>	

Efficient Tensor Core-Based GPU Kernels for Structured Sparsity Under Reduced Precision	1147
<i>Zhaodong Chen, Zheng Qu, Liu Liu, Yufei Ding, Yuan Xie</i>	
Arithmetic-Intensity-Guided Fault Tolerance for Neural Network Inference on GPUs.....	1160
<i>Jack Kosaian, K. V. Rashmi</i>	
FastZ: Accelerating Gapped Whole Genome Alignment on GPUs.....	1175
<i>Sree Charan Gundabolu, T. N. Vijaykumar, Mithuna Thottethodi</i>	
PEPPA-X: Finding Program Test Inputs to Bound Silent Data Corruption Vulnerability in HPC Applications.....	1188
<i>Md Hasanur Rahman, Aabid Shamji, Shengjian Guo, Guanpeng Li</i>	
Cuttlefish: Library for Achieving Energy Efficiency in Multicore Parallel Programs	1202
<i>Sunil Kumar, Akshat Gupta, Vivek Kumar, Sridutt Bhalachandra</i>	
Temporal Vectorization for Stencils.....	1216
<i>Liang Yuan, Hang Cao, Yunquan Zhang, Kun Li, Pengqi Lu, Yue Yue</i>	
PAGANI: A Parallel Adaptive GPU Algorithm for Numerical Integration	1230
<i>Ioannis Sakiotis, Kamesh Arumugam, Marc Paterno, Desh Ranjan, Balsša Terzic, Mohammad Zubair</i>	
Reducing Redundancy in Data Organization and Arithmetic Calculation for Stencil Computations	1243
<i>Kun Li, Liang Yuan, Yunquan Zhang, Yue Yue</i>	
CAKE: Matrix Multiplication Using Constant-Bandwidth Blocks	1258
<i>H. T. Kung, Vikas Natesh, Andrew Sabot</i>	
HPAC: Evaluating Approximate Computing Techniques on HPC OpenMP Applications.....	1271
<i>Konstantinos Parasyris, Giorgis Georgakoudis, Harshitha Menon, James Diffenderfer, Ignacio Laguna, Daniel Osei-Kuffuor, Markus Schordan</i>	
Accelerating XOR-Based Erasure Coding Using Program Optimization Techniques	1285
<i>Yuya Uezato</i>	
Error-Controlled, Progressive, and Adaptable Retrieval of Scientific Data with Multilevel Decomposition.....	1300
<i>Xin Liang, Qian Gong, Jieyang Chen, Ben Whitney, Lipeng Wan, Qing Liu, David Pugmire, Rick Archibald, Norbert Podhorszki, Scott Klasky</i>	
LogECMem: Coupling Erasure-Coded In-Memory Key-Value Stores with Parity Logging.....	1314
<i>Liangfeng Cheng, Yuchong Hu, Zhaokang Ke, Jia Xu, Qiaori Yao, Dan Feng, Weichun Wang, Wei Chen</i>	
Scalable FBP Decomposition for Cone-Beam CT Reconstruction	1328
<i>Peng Chen, Mohamed Wahib, Xiao Wang, Takahiro Hirofuchi, Hirotaka Ogawa, Ander Biguri, Richard Boardman, Thomas Blumensath, Satoshi Matsuoka</i>	
Online Optimization of File Transfers in High-Speed Networks	1344
<i>Md Arifuzzaman, Engin Arslan</i>	
Hardware-Supported Remote Persistence for Distributed Persistent Memory	1357
<i>Zhuohui Duan, Haodi Lu, Haikun Liu, Xiaofei Liao, Hai Jin, Yu Zhang, Song Wu</i>	
Clairvoyant Prefetching for Distributed Machine Learning I/O.....	1372
<i>Nikoli Dryden, Roman Böhringer, Tal Ben-Nun, Torsten Hoefler</i>	

ndzip-gpu: Efficient Lossless Compression of Scientific Floating-Point Data on GPUs	1386
<i>Fabian Knorr, Peter Thoman, Thomas Fahringer</i>	
Resilient Error-Bounded Lossy Compressor for Data Transfer	1399
<i>Sihuan Li, Sheng Di, Kai Zhao, Xin Liang, Zizhong Chen, Franck Cappello</i>	
Productivity, Portability, Performance: Data-Centric Python	1413
<i>Alexandros Nikolaos Ziogas, Timo Schneider, Tal Ben-Nun, Alexandru Calotoiu, Tiziano De Matteis, Johannes De Fine Licht, Luca Lavarini, Torsten Hoefler</i>	
Empirical Evaluation of Circuit Approximations on Noisy Quantum Devices	1428
<i>Ellis Wilson, Frank Mueller, Lindsay Bassman, Costin Iancu</i>	
SV-Sim: Scalable PGAS-Based State Vector Simulation of Quantum Circuits.....	1440
<i>Ang Li, Bo Fang, Christopher Granade, Guen Prawiroatmodjo, Bettina Heim, Martin Roetteler, Sriram Krishnamoorthy</i>	
SW_Qsim: A Minimize-Memory Quantum Simulator with High-Performance on a New Sunway Supercomputer.....	1454
<i>Fang Li, Xin Liu, Yong Liu, Pengpeng Zhao, Yuling Yang, Honghui Shang, Weizhe Sun, Zhen Wang, Enming Dong, Dexun Chen</i>	
MAPA: Multi-Accelerator Pattern Allocation Policy for Multi-Tenant GPU Servers.....	1467
<i>Kiran Ranganath, Joshua D. Suetterlein, Joseph B. Manzano, Shuaiwen Leon Song, Daniel Wong</i>	

Author Index