

SC18: International Conference for High Performance Computing, Networking, Storage and Analysis

**Dallas, Texas, USA
11 – 16 November 2018**



**IEEE Catalog Number: CFP18SUP-POD
ISBN: 978-1-5386-8385-9**

**Copyright © 2018 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:	CFP18SUP-POD
ISBN (Print-On-Demand):	978-1-5386-8385-9
ISBN (Online):	978-1-5386-8384-2

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

SC18: The International Conference for High Performance Computing, Networking, Storage, and Analysis SC 2018

Table of Contents

Data and Storage

SP-Cache: Load-Balanced, Redundancy-Free Cluster Caching with Selective Partition .1.....	
<i>Yinghao Yu (Hong Kong University of Science and Technology), Renfei Huang (Hong Kong University of Science and Technology), Wei Wang (Hong Kong University of Science and Technology), Jun Zhang (Hong Kong University of Science and Technology), and Khaled Ben Letaief (Hong Kong University of Science and Technology)</i>	
BESPOKV: Application Tailored Scale-Out Key-Value Stores .14.....	
<i>Ali Anwar (IBM Research–Almaden), Yue Cheng (George Mason University), Hai Huang (IBM Research–T.J. Watson), Jingoo Han (Virginia Tech), Hyogi Sim (Oak Ridge National Laboratory), Dongyoon Lee (Virginia Tech), Fred Douglass (Perspecta Labs), and Ali R. Butt (Virginia Tech)</i>	
Scaling Embedded In-Situ Indexing with DeltaFS .30.....	
<i>Qing Zheng (Carnegie Mellon University), Charles D. Cranor (Carnegie Mellon University), Danhao Guo (Carnegie Mellon University), Gregory R. Ganger (Carnegie Mellon University), George Amvrosiadis (Carnegie Mellon University), Garth A. Gibson (Carnegie Mellon University), Bradley W. Settlemyer (Los Alamos National Laboratory), Gary Grider (Los Alamos National Laboratory), and Fan Guo (Los Alamos National Laboratory)</i>	

Next-Generation Networking

Exploiting Idle Resources in a High-Radix Switch for Supplemental Storage .45.....	
<i>Matthias A. Blumrich (NVIDIA Corporation), Nan Jiang (NVIDIA Corporation), and Larry R. Dennison (NVIDIA Corporation)</i>	

Fine-Grained, Multi-Domain Network Resource Abstraction as a Fundamental Primitive to Enable High-Performance, Collaborative Data Sciences .58.....	
<i>Qiao Xiang (Tongji University), J. Jensen Zhang (Tongji University),</i>	
<i>X. Tony Wang (Tongji University), Y. Jace Liu (Tongji University),</i>	
<i>Chin Guok (Lawrence Berkeley National Laboratory), Franck Le (IBM T.J.</i>	
<i>Watson Research Center), John MacAuley (Lawrence Berkeley National</i>	
<i>Laboratory), Harvey Newman (California Institute of Technology), and</i>	
<i>Y. Richard Yang (Tongji University)</i>	
Light-Weight Protocols for Wire-Speed Ordering .71.....	
<i>Hans Eberle (NVIDIA Research) and Larry Dennison (NVIDIA Research)</i>	

Resilience

GPU Age-Aware Scheduling to Improve the Reliability of Leadership Jobs on Titan .83.....	
<i>Christopher Zimmer (Oak Ridge National Laboratory), Don Maxwell (Oak</i>	
<i>Ridge National Laboratory), Stephen McNally (Oak Ridge National</i>	
<i>Laboratory), Scott Atchley (Oak Ridge National Laboratory), and</i>	
<i>Sudharshan S. Vazhkudai (Oak Ridge National Laboratory)</i>	
FlipTracker: Understanding Natural Error Resilience in HPC Applications .94.....	
<i>Luanzheng Guo (EECS, UC Merced), Dong Li (EECS, UC Merced), Ignacio</i>	
<i>Laguna (Lawrence Livermore National Laboratory), and Martin Schulz</i>	
<i>(Technical University of Munich)</i>	
Doomsday: Predicting Which Node Will Fail When on Supercomputers .108.....	
<i>Anwsha Das (North Carolina State University), Frank Mueller (North</i>	
<i>Carolina State University), Paul Hargrove (Lawrence Berkeley National</i>	
<i>Laboratory), Eric Roman (Lawrence Berkeley National Laboratory), and</i>	
<i>Scott Baden (Lawrence Berkeley National Laboratory)</i>	

Biology Applications

Extreme Scale De Novo Metagenome Assembly .122.....	
<i>Evangelos Georganas (Parallel Computing Lab, Intel Corp), Rob Egan</i>	
<i>(Lawrence Berkeley National Laboratory), Steven Hofmeyr (Computational</i>	
<i>Research Division), Eugene Goltsman (Lawrence Berkeley National</i>	
<i>Laboratory), Bill Arndt (National Energy Research Scientific Computing</i>	
<i>Center), Andrew Tritt (National Energy Research Scientific Computing</i>	
<i>Center), Aydin Buluç (Computational Research Division), Leonid Olikier</i>	
<i>(Computational Research Division), and Katherine Yelick (Computational</i>	
<i>Research Division)</i>	
Optimizing High Performance Distributed Memory Parallel Hash Tables for DNA k-mer Counting .135.....	
<i>Tony C. Pan (Georgia Institute of Technology), Sanchit Misra (Intel</i>	
<i>Corporation), and Srinivas Aluru (Georgia Institute of Technology)</i>	

Redesigning LAMMPS for Peta-Scale and Hundred-Billion-Atom Simulation on Sunway TaihuLight .148.....	
<i>Xiaohui Duan (Shandong University), Ping Gao (Shandong University), Tingjian Zhang (Shandong University), Meng Zhang (Shandong University), Weiguo Liu (Shandong University), Wusheng Zhang (Tsinghua University), Wei Xue (Tsinghua University), Haohuan Fu (Tsinghua University), Lin Gan (Tsinghua University), Dexun Chen (Tsinghua University), Xiangxu Meng (Shandong University), and Guangwen Yang (Tsinghua University)</i>	

Large-Scale Algorithms

Large-Scale Hierarchical k-means for Heterogeneous Many-Core Supercomputers .160.....	
<i>Liandeng Li (Tsinghua University), Teng Yu (University of St Andrews), Wenlai Zhao (Tsinghua University), Haohuan Fu (Tsinghua University), Chenyu Wang (University of St Andrews), Li Tan (Beijing Technology and Business University), Guangwen Yang (Tsinghua University), and John Thomson (University of St Andrews)</i>	
TriCore: Parallel Triangle Counting on GPUs .171.....	
<i>Yang Hu (The George Washington University), Hang Liu (University of Massachusetts Lowell), and H. Howie Huang (The George Washington University)</i>	
Distributed-Memory Hierarchical Compression of Dense SPD Matrices .183.....	
<i>Chenhan D. Yu (The University of Texas at Austin), Severin Reiz (Technical University of Munich), and George Biros (The University of Texas at Austin)</i>	

Performance and Energy Analysis

A Parallelism Profiler with What-If Analyses for OpenMP Programs .198.....	
<i>Nader Boushehrinejadmoradi (Rutgers University), Adarsh Yoga (Rutgers University), and Santosh Nagarakatte (Rutgers University)</i>	
Energy Efficiency Modeling of Parallel Applications .212.....	
<i>Mark Endrei (The University of Queensland), Chao Jin (The University of Queensland), Minh Ngoc Dinh (The University of Queensland), David Abramson (The University of Queensland), Heidi Poxon (Cray Inc.), Luiz DeRose (Cray Inc.), and Bronis R. de Supinski (Lawrence Livermore National Laboratory)</i>	
HPL and DGEMM Performance Variability on the Xeon Platinum 8160 Processor .225.....	
<i>John D. McCalpin (University of Texas at Austin)</i>	

Algorithms on Sparse Data

HiCOO: Hierarchical Storage of Sparse Tensors .238.....	
<i>Jiajia Li (Georgia Institute of Technology), Jimeng Sun (Georgia Institute of Technology), and Richard Vuduc (Georgia Institute of Technology)</i>	

Distributed Memory Sparse Inverse Covariance Matrix Estimation on High-Performance Computing Architectures .253.....	
<i>Aryan Eftekhari (Universita della Svizzera italiana), Matthias Bollhöfer (TU Braunschweig), and Olaf Schenk (Universita della Svizzera italiana)</i>	
PruneJuice: Pruning Trillion-edge Graphs to a Precise Pattern-Matching Solution .265.....	
<i>Tahsin Reza (Lawrence Livermore National Laboratory), Matei Ripeanu (University of British Columbia), Nicolas Tripoul (University of British Columbia), Geoffrey Sanders (Lawrence Livermore National Laboratory), and Roger Pearce (Lawrence Livermore National Laboratory)</i>	

Performance Optimization Studies

Many-Core Graph Workload Analysis .282.....	
<i>Stijn Eyerman (Intel Corporation), Wim Heirman (Intel Corporation), Kristof Du Bois (Intel Corporation), Joshua B. Fryman (Intel Corporation), and Ibrahim Hur (Intel Corporation)</i>	
Lessons Learned from Analyzing Dynamic Promotion for User-Level Threading .293.....	
<i>Shintaro Iwasaki (The University of Tokyo), Abdelhalim Amer (Argonne National Laboratory), Kenjiro Taura (The University of Tokyo), and Pavan Balaji (Argonne National Laboratory)</i>	
Topology-Aware Space-Shared Co-Analysis of Large-Scale Molecular Dynamics Simulations .305.....	
<i>Preeti Malakar (Argonne National Laboratory), Todd Munson (Argonne National Laboratory), Christopher Knight (Argonne National Laboratory), Venkatram Vishwanath (Argonne National Laboratory), and Michael E. Papka (Argonne National Laboratory)</i>	

Resource Management and Interference

RM-Replay: A High-Fidelity Tuning, Optimization and Exploration Tool for Resource Management .320.....	
<i>Maxime Martinasso (ETH Zurich), Miguel Gila (ETH Zurich), Mauro Bianco (ETH Zurich), Sadaf R. Alam (ETH Zurich), Colin McMurtrie (ETH Zurich), and Thomas C. Schulthess (ETH Zurich)</i>	
Evaluation of an Interference-free Node Allocation Policy on Fat-tree Clusters .333.....	
<i>Samuel D. Pollard (University of Oregon), Nikhil Jain (Lawrence Livermore National Laboratory), Stephen Herbein (Lawrence Livermore National Laboratory), and Abhinav Bhatele (Lawrence Livermore National Laboratory)</i>	
Mitigating Inter-Job Interference Using Adaptive Flow-Aware Routing .346.....	
<i>Staci A. Smith (University of Arizona), Clara E. Cromei (University of Arizona), David K. Lowenthal (University of Arizona), Jens Domke (Tokyo Institute of Technology), Nikhil Jain (Lawrence Livermore National Laboratory), Jayaraman J. Thiagarajan (Lawrence Livermore National Laboratory), and Abhinav Bhatele (Lawrence Livermore National Laboratory)</i>	

MPI Optimization and Characterization

- Cooperative Rendezvous Protocols for Improved Performance and Overlap .361.....
S. Chakraborty (The Ohio State University), M. Bayatpour (The Ohio State University), J. Hashmi (The Ohio State University), H. Subramoni (The Ohio State University), and D. K. Panda (The Ohio State University)
- Framework for Scalable Intra-Node Collective Operations using Shared Memory .374.....
Surabhi Jain (Intel Corporation), Rashid Kaleem (Intel Corporation), Marc Gamell Balmana (Intel Corporation), Akhil Langer (Intel Corporation), Dmitry Durnov (Intel Corporation), Alexander Sannikov (Intel Corporation), and Maria Garzaran (Intel Corporation)
- Characterization of MPI Usage on a Production Supercomputer .386.....
Sudheer Chunduri (Argonne National Laboratory), Scott Parker (Argonne National Laboratory), Pavan Balaji (Argonne National Laboratory), Kevin Harms (Argonne National Laboratory), and Kalyan Kumaran (Argonne National Laboratory)

Non-Volatile Memory

- Runtime Data Management on Non-Volatile Memory-based Heterogeneous Memory for Task-Parallel Programs....
401
Kai Wu (University of California, Merced), Jie Ren (University of California, Merced), and Dong Li (University of California, Merced)
- DRAGON: Breaking GPU Memory Capacity Limits with Direct NVM Access .414.....
Pak Markthub (Tokyo Institute of Technology), Mehmet E. Belviranli (Oak Ridge National Laboratory), Seyong Lee (Oak Ridge National Laboratory), Jeffrey S. Vetter (Oak Ridge National Laboratory), and Satoshi Matsuoka (Oak Ridge National Laboratory)
- Siena: Exploring the Design Space of Heterogeneous Memory Systems .427.....
Ivy B. Peng (Oak Ridge National Laboratory) and Jeffrey S. Vetter (Oak Ridge National Laboratory)

Task-Based Programming

- Dynamic Tracing: Memoization of Task Graphs for Dynamic Task-Based Runtimes .441.....
Wonchan Lee (Stanford University), Elliott Slaughter (SLAC National Accelerator Laboratory), Michael Bauer (NVIDIA), Sean Treichler (NVIDIA), Todd Warszawski (Stanford University), Michael Garland (NVIDIA), and Alex Aiken (Stanford University)
- Runtime-Assisted Cache Coherence Deactivation in Task Parallel Programs .454.....
Paul Caheny (miBarcelona Supercomputing Center), Lluc Alvarez (Barcelona Supercomputing Center), Mateo Valero (Barcelona Supercomputing Center), Miquel Moretó (Barcelona Supercomputing Center), and Marc Casas (Barcelona Supercomputing Center)
- A Divide and Conquer Algorithm for DAG Scheduling under Power Constraints .466.....
Gökalep Demirci (University of Chicago), Ivana Marincic (University of Chicago), and Henry Hoffmann (University of Chicago)

Clouds and Distributed Computing

- A Reference Architecture for Datacenter Scheduling: Design, Validation, and Experiments .478.....
Georgios Andreadis (Delft University of Technology), Laurens Versluis (Vrije Universiteit Amsterdam), Fabian Mastenbroek (Delft University of Technology), and Alexandru Iosup (Delft University of Technology)
- Dynamically Negotiating Capacity Between On-demand and Batch Clusters .493.....
Feng Liu (University of Minnesota), Kate Keahey (Argonne National Laboratory), Pierre Riteau (University of Chicago), and Jon Weissman (University of Minnesota)
- A Lightweight Model for Right-Sizing Master-Worker Applications .504.....
Nathaniel Kremer-Herman (University of Notre Dame), Benjamin Tovar (University of Notre Dame), and Douglas Thain (University of Notre Dame)

Physics and Tensor Applications

- Simulating the Wenchuan Earthquake with Accurate Surface Topography on Sunway TaihuLight .517.....
Bingwei Chen (Tsinghua University), Haohuan Fu (Tsinghua University), Yanwen Wei (Tsinghua University), Conghui He (Tsinghua University), Wenqiang Zhang (University of Science and Technology of China), Yuxuan Li (Tsinghua University), Wubin Wan (National Supercomputing Center in Wuxi), Wei Zhang (National Supercomputing Center in Wuxi), Lin Gan (Tsinghua University), Wei Zhang (Southern University of Science and Technology), Zhenguo Zhang (Southern University of Science and Technology), Guangwen Yang (Tsinghua University), and Xiaofei Chen (Southern University of Science and Technology)
- Accelerating Quantum Chemistry with Vectorized and Batched Integrals .529.....
Hua Huang (Georgia Institute of Technology) and Edmond Chow (Georgia Institute of Technology)
- High-Performance Dense Tucker Decomposition on GPU Clusters .543.....
Jee Choi (IBM T. J. Watson Research Center), Xing Liu (IBM T. J. Watson Research Center), and Venkatesan Chakaravarthy (IBM India Research Lab)

Resilience II

- Lessons Learned from Memory Errors Observed Over the Lifetime of Cielo .554.....
Scott Levy (Sandia National Laboratories), Kurt B. Ferreira (Sandia National Laboratories), Nathan DeBardeleben (Los Alamos National Laboratory), Taniya Siddiqua (Advanced Micro Devices, Inc.), Vilas Sridharan (Advanced Micro Devices, Inc.), and Elisabeth Baseman (Los Alamos National Laboratory)
- Partial Redundancy in HPC Systems with Non-Uniform Node Reliabilities .566.....
Zaeem Hussain (University of Pittsburgh), Taieb Znati (University of Pittsburgh), and Rami Melhem (University of Pittsburgh)

Evaluating and Accelerating High-Fidelity Error Injection for HPC .577.....	
<i>Chun-Kai Chang (University of Texas at Austin), Sangkug Lym (University of Texas at Austin), Nicholas Kelly (University of Texas at Austin), Michael B. Sullivan (University of Texas at Austin), and Mattan Erez (University of Texas at Austin)</i>	

Arithmetic and Optimization

Associative Instruction Reordering to Alleviate Register Pressure .590.....	
<i>Prashant Singh Rawat (The Ohio State University), Aravind Sukumaran-Rajam (The Ohio State University), Atanas Rountev (The Ohio State University), Fabrice Rastello (University Grenoble Alpes), Louis-Noël Pouchet (Colorado State University), and P. Sadayappan (The Ohio State University)</i>	
Harnessing GPU Tensor Cores for Fast FP16 Arithmetic to Speed up Mixed-Precision Iterative Refinement Solvers .603.....	
<i>Azzam Haidar (University of Tennessee), Stanimire Tomov (University of Tennessee), Jack Dongarra (University of Tennessee), and Nicholas J. Higham (University of Manchester)</i>	
ADAPT: Algorithmic Differentiation Applied to Floating-Point Precision Tuning .614.....	
<i>Harshitha Menon (Lawrence Livermore National Laboratory), Michael O. Lam (James Madison University), Daniel Osei-Kuffuor (Lawrence Livermore National Laboratory), Markus Schordan (Lawrence Livermore National Laboratory), Scott Lloyd (Lawrence Livermore National Laboratory), Kathryn Mohror (Lawrence Livermore National Laboratory), and Jeffrey Hittinger (Lawrence Livermore National Laboratory)</i>	

Gordon Bell Prize Finalist #1

A Fast Scalable Implicit Solver for Nonlinear Time-Evolution Earthquake City Problem on Low-Ordered Unstructured Finite Elements with Artificial Intelligence and Transprecision Computing .627.....	
<i>Tsuyoshi Ichimura (missing), Kohei Fujita (The University of Tokyo), Takuma Yamaguchi (The University of Tokyo), Akira Naruse (NVIDIA Corporation), Jack C. Wells (Oak Ridge National Laboratory), Thomas C. Schulthess (Swiss National Supercomputing Centre), Tjerk P. Straatsma (Oak Ridge National Laboratory), Christopher J. Zimmer (Oak Ridge National Laboratory), Maxime Martinasso (Swiss National Supercomputing Centre), Kengo Nakajima (The University of Tokyo), Muneo Hori (The University of Tokyo), and Lalith Maddegadara (The University of Tokyo)</i>	
167-PFlops Deep Learning for Electron Microscopy: From Learning Physics to Atomic Manipulation .638.....	
<i>Robert M. Patton (Oak Ridge National Laboratory), J. Travis Johnston (Oak Ridge National Laboratory), Steven R. Young (Oak Ridge National Laboratory), Catherine D. Schuman (Oak Ridge National Laboratory), Don D. March (Oak Ridge National Laboratory), Thomas E. Potok (Oak Ridge National Laboratory), Derek C. Rose (Oak Ridge National Laboratory), Seung-Hwan Lim (Oak Ridge National Laboratory), Thomas P. Karnowski (Oak Ridge National Laboratory), Maxim A. Ziatdinov (Oak Ridge National Laboratory), and Sergei V. Kalinin (Oak Ridge National Laboratory)</i>	

Exascale Deep Learning for Climate Analytics .649.....	<i>Thorsten Kurth (Lawrence Berkeley National Laboratory), Sean Treichler (NVIDIA), Joshua Romero (NVIDIA), Mayur Mudigonda (Lawrence Berkeley National Laboratory), Nathan Luehr (NVIDIA), Everett Phillips (NVIDIA), Ankur Mahesh (Lawrence Berkeley National Laboratory), Michael Matheson (Oak Ridge National Laboratory), Jack Deslippe (Lawrence Berkeley National Laboratory), Massimiliano Fatica (NVIDIA), Prabhat (Lawrence Berkeley National Laboratory), and Michael Houston (NVIDIA)</i>
--	---

Large Scale System Deployments

The Design, Deployment, and Evaluation of the CORAL Pre-Exascale Systems .661.....	<i>Sudharshan S. Vazhkudai (Oak Ridge National Laboratory), Bronis R. de Supinski (Lawrence Livermore National Laboratory), Arthur S. Bland (Oak Ridge National Laboratory), Al Geist (Oak Ridge National Laboratory), James Sexton (IBM), Jim Kahle (IBM), Christopher J. Zimmer (Oak Ridge National Laboratory), Scott Atchley (Oak Ridge National Laboratory), Sarp Oral (Oak Ridge National Laboratory), Don E. Maxwell (Oak Ridge National Laboratory), Veronica G. Vergara Larrea (Oak Ridge National Laboratory), Adam Bertsch (Lawrence Livermore National Laboratory), Robin Goldstone (Lawrence Livermore National Laboratory), Wayne Joubert (Oak Ridge National Laboratory), Chris Chambreau (Lawrence Livermore National Laboratory), David Appelhans (IBM), Robert Blackmore (IBM), Ben Casses (Lawrence Livermore National Laboratory), George Chochia (IBM), Gene Davison (IBM), Matthew A. Ezell (Oak Ridge National Laboratory), Tom Gooding (IBM), Elsa Gonsiorowski (Lawrence Livermore National Laboratory), Leopold Grinberg (IBM), Bill Hanson (IBM), Bill Hartner (IBM), Ian Karlin (Lawrence Livermore National Laboratory), Matthew L. Leininger (Lawrence Livermore National Laboratory), Dustin Leverman (Oak Ridge National Laboratory), Chris Marroquin (IBM), Adam Moody (Lawrence Livermore National Laboratory), Martin Ohmacht (IBM), Ramesh Pankajakshan (Lawrence Livermore National Laboratory), Fernando Pizzano (IBM), James H. Rogers (Oak Ridge National Laboratory), Bryan Rosenburg (IBM), Drew Schmidt (Oak Ridge National Laboratory), Mallikarjun Shankar (Oak Ridge National Laboratory), Feiyi Wang (Oak Ridge National Laboratory), Py Watson (Lawrence Livermore National Laboratory), Bob Walkup (IBM), Lance D. Weems (Lawrence Livermore National Laboratory), and Junqi Yin (Oak Ridge National Laboratory)</i>
--	---

Best Practices and Lessons from Deploying and Operating a Sustained-Petascale System: The Blue Waters Experience .673.....	<i>Gregory H. Bauer (University of Illinois), Brett Bode (University of Illinois), Jeremy Enos (University of Illinois), William T. Kramer (University of Illinois), Scott Lathrop (University of Illinois), Celso L. Mendes (University of Illinois), and Roberto R. Sisneros (University of Illinois)</i>
--	---

Performance Evaluation of a Vector Supercomputer SX-Aurora TSUBASA .685.....	<i>Kazuhiko Komatsu (Tohoku University), Shintaro Momose (Tohoku University, NEC Corporation), Yoko Isobe (Tohoku University, NEC Corporation), Osamu Watanabe (Tohoku University, NEC Corporation), Akihiro Musa (Tohoku University, NEC Corporation), Mitsuo Yokokawa (Kobe University), Toshikazu Aoyama (NEC Corporation), Masayuki Sato (Tohoku University), and Hiroaki Kobayashi (Tohoku University)</i>
--	---

Gordon Bell Prize Finalist #2

Simulating the Weak Death of the Neutron in a Femtoscale Universe with Near-Exascale Computing .697.....

Evan Berkowitz (Institut für Kernphysik and Institute for Advanced Simulation), M.A. Clark (NVIDIA Corporation), Arjun Gambhir (Lawrence Livermore National Laboratory; Lawrence Berkeley National Laboratory; University of California, Berkeley), Ken McElvain (University of California, Berkeley; Lawrence Berkeley National Laboratory), Amy Nicholson (University of North Carolina), Enrico Rinaldi (Brookhaven National Laboratory; Lawrence Berkeley National Laboratory), Pavlos Vranas (Lawrence Livermore National Laboratory; Lawrence Berkeley National Laboratory;), André Walker-Loud (Lawrence Berkeley National Laboratory; Lawrence Livermore National Laboratory), Chia Cheng Chang (Lawrence Berkeley National Laboratory), Bálint Joó (Thomas Jefferson National Accelerator Facility), Thorsten Kurth (Lawrence Berkeley National Laboratory), and Kostas Orginos (The College of William & Mary)

ShenTu: Processing Multi-Trillion Edge Graphs on Millions of Cores in Seconds .706.....

Heng Lin (Tsinghua University; Fma Technology), Xiaowei Zhu (Tsinghua University; Qatar Computing Research Institute), Bowen Yu (Tsinghua University), Xiongchao Tang (Tsinghua University; Qatar Computing Research Institute), Wei Xue (Tsinghua University), Wenguang Chen (Tsinghua University), Lufei Zhang (State Key Laboratory of Mathematical Engineering and Advanced Computing), Torsten Hoeftler (ETH Zurich), Xiaosong Ma (Qatar Computing Research Institute), Xin Liu (National Research Centre of Parallel Computer Engineering and Technology), Weimin Zheng (Tsinghua University), and Jingfang Xu (Beijing Sogou Technology Development Co., Ltd.)

Attacking the Opioid Epidemic: Determining the Epistatic and Pleiotropic Genetic Architectures for Chronic Pain and Opioid Addiction .717.....

Wayne Joubert (Oak Ridge National Laboratory), Deborah Weighill (Oak Ridge National Laboratory; University of Tennessee), David Kainer (Oak Ridge National Laboratory), Sharlee Climer (University of Missouri-St. Louis), Amy Justice (Yale University/Department of Veterans Affairs), Kjersten Fagnan (DOE Joint Genome Institute), and Daniel Jacobson (Oak Ridge National Laboratory)

Graph Algorithms and Systems

iSpan: Parallel Identification of Strongly Connected Components with Spanning Trees .731.....

Yuede Ji (George Washington University), Hang Liu (University of Massachusetts, Lowell), and H. Howie Huang (George Washington University)

Adaptive Anonymization of Data using b-Edge Cover .743.....

Arif Khan (Pacific Northwest National Laboratory), Krzysztof Choromanski (Google Brain Robotics, New York), Alex Pothen (Purdue University), S. M. Ferdous (Purdue University), Mahantesh Halappanavar (Pacific Northwest National Laboratory), and Antonino Tumeo (Pacific Northwest National Laboratory)

faimGraph: High Performance Management of Fully-Dynamic Graphs Under Tight Memory Constraints on the GPU .754.....

Martin Winter (Graz University of Technology, Austria), Daniel Mlakar (Graz University of Technology, Austria), Rhaleb Zayer (missingMax Planck Institute for Informatics), Hans-Peter Seidel (Max Planck Institute for Informatics), and Markus Steinberger (Max Planck Institute for Informatics)

Programming Systems Tools

Dynamic Data Race Detection for OpenMP Programs .767.....

Yizi Gu (Rice University) and John Mellor-Crummey (Rice University)

ParSy: Inspection and Transformation of Sparse Matrix Computations for Parallelism .779.....

Kazem Cheshmi (University of Toronto), Shoaib Kamil (Adobe Research), Michelle Mills Strout (University of Arizona), and Maryam Mehri Dehnavi (University of Toronto)

Detecting MPI Usage Anomalies via Partial Program Symbolic Execution .794.....

Fangke Ye (Georgia Institute of Technology), Jisheng Zhao (Georgia Institute of Technology), and Vivek Sarkar (Georgia Institute of Technology)

Deep Learning

Exploring Flexible Communications for Streamlining DNN Ensemble Training Pipelines .807.....

Randall Pittman (North Carolina State University), Hui Guan (North Carolina State University), Xipeng Shen (North Carolina State University), Seung-Hwan Lim (Oak Ridge National Laboratory), and Robert M. Patton (Oak Ridge National Laboratory)

CosmoFlow: Using Deep Learning to Learn the Universe at Scale .819.....

*Amrita Mathuriya (Intel Corporation), Deborah Bard (Lawrence Berkeley National Laboratory), Peter Mendygral (Cray Inc.), Lawrence Meadows (Intel Corporation), James Arnemann (U.C. Berkeley), Lei Shao (Intel Corporation), Siyu He (Flatiron Institute), Tuomas Kärnä (Intel Corporation), Diana Moise (Cray Inc.), Simon J. Pennycook (Intel Corporation), Kristyn Maschhoff (Cray Inc.), Jason Sewall (Intel Corporation), Nalini Kumar (Intel Corporation), Shirley Ho (Flatiron Institute), Michael F. Ringenburt (Cray Inc.), * Prabhat (Lawrence Berkeley National Laboratory), and Victor Lee (Intel Corporation)*

Anatomy of High-Performance Deep Learning Convolutions on SIMD Architectures .830.....

Evangelos Georganas (Intel Corporation), Sasikanth Avancha (Intel Corporation), Kunal Banerjee (Intel Corporation), Dhiraj Kalamkar (Intel Corporation), Greg Henry (Intel Corporation), Hans Pabst (Intel Corporation), and Alexander Heinecke (Intel Corporation)

Resilience III: GPUs

- Optimizing Software-Directed Instruction Replication for GPU Error Detection .842.....
Abdulrahman Mahmoud (University of Illinois at Urbana-Champaign), Siva Kumar Sastry Hari (NVIDIA), Michael B. Sullivan (NVIDIA), Timothy Tsai (NVIDIA), and Stephen W. Keckler (NVIDIA)
- Fault Tolerant One-sided Matrix Decompositions on Heterogeneous Systems with GPUs .854.....
Jieyang Chen (University of California, Riverside), Hongbo Li (University of California, Riverside), Sihuan Li (University of California, Riverside), Xin Liang (University of California, Riverside), Panruo Wu (University of Houston), Dingwen Tao (University of Alabama), Kaiming Ouyang (University of California, Riverside), Yuanlai Liu (University of California, Riverside), Kai Zhao (University of California, Riverside), Qiang Guan (Kent State University), and Zizhong Chen (University of California, Riverside)
- PRISM: Predicting Resilience of GPU Applications Using Statistical Methods .866.....
Charu Kalra (Northeastern University), Fritz Previlon (Northeastern University), Xiangyu Li (Northeastern University), Norman Rubin (NVIDIA), and David Kaeli (Northeastern University)

Astrophysics Applications

- Phase Asynchronous AMR Execution for Productive and Performant Astrophysical Flows .880.....
Muhammad Nufail Farooqi (Koc University), Tan Nguyen (Lawrence Berkeley National Laboratory), Weiqun Zhang (Lawrence Berkeley National Laboratory), Ann S. Almgren (Lawrence Berkeley National Laboratory), John Shalf (Lawrence Berkeley National Laboratory), and Didem Unat (Koc University)
- Computing Planetary Interior Normal Modes with a Highly Parallel Polynomial Filtering Eigensolver .894.....
Jia Shi (Rice University), Ruipeng Li (Lawrence Livermore National Laboratory), Yuanzhe Xi (Emory University), Yousef Saad (University of Minnesota), and Maarten V. de Hoop (Rice University)

File Systems: Data Movement and Provenance

- Dac-Man: Data Change Management for Scientific Datasets on HPC systems .907.....
Devarshi Ghoshal (Lawrence Berkeley National Laboratory), Lavanya Ramakrishnan (Lawrence Berkeley National Laboratory), and Deborah Agarwal (Lawrence Berkeley National Laboratory)
- Stacker: An Autonomic Data Movement Engine for Extreme-Scale Data Staging-Based In-Situ Workflows .920
Pradeep Subedi (Rutgers University), Philip Davis (Rutgers University), Shaohua Duan (Rutgers University), Scott Klasky (Oak Ridge National Laboratory), Hemanth Kolla (Sandia National Laboratories), and Manish Parashar (Rutgers University)

A Year in the Life of a Parallel File System .931.....	
<i>Glenn K. Lockwood (Lawrence Berkeley National Laboratory), Shane Snyder (Argonne National Laboratory), Teng Wang (Lawrence Berkeley National Laboratory), Suren Byna (Lawrence Berkeley National Laboratory), Philip Carns (Argonne National Laboratory), and Nicholas J. Wright (Lawrence Berkeley National Laboratory)</i>	

Author Index 945.	
--------------------------------	--