

SC17: International Conference for High Performance Computing, Networking, Storage and Analysis

**Denver, Colorado, USA
12 – 17 November 2017**



**IEEE Catalog Number: CFP17SUP-POD
ISBN: 978-1-4503-5114-0**

**Copyright © 2017, The 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:	CFP17SUP-POD
ISBN (Print-On-Demand):	978-1-4503-5114-0
ISBN (Online):	978-1-4503-5114-0
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



ACM Gordon Bell Finalists

1. Redesigning CAM-SE for Petascale Climate Modeling Performance on Sunway TaihuLight. Haohuan Fu, Junfeng Liao, and Nan Ding (Tsinghua University); Xiaohui Duan (Shandong University); Lin Gan (Tsinghua University); Yishuang Liang (Beijing Normal University); Xinliang Wang (Tsinghua University); Jinzhe Yang (Imperial College, London); Yan Zheng (National Research Center of Parallel Computer Engineering and Technology, China); Weiguo Liu (Shandong University); Lanning Wang (Beijing Normal University); and Guangwen Yang (Tsinghua University)...1

2. 15-Pflops Nonlinear Earthquake Simulation on Sunway TaihuLight: Enabling Depiction of Realistic 10 Hz Scenarios. Haohuan Fu, Conghui He, and Bingwei Chen (Tsinghua University); Zekun Yin (Shandong University); Zhenguo Zhang (Southern University of Science and Technology, China); Wenqiang Zhang (University of Science and Technology of China); Tingjian Zhang (Shandong University); Wei Xue (Tsinghua University); Weiguo Liu (Shandong University); Wanwang Yin (National Research Center of Parallel Computer Engineering and Technology, China); Guangwen Yang (Tsinghua University); and Xiaofei Chen (Southern University of Science and Technology, China)....13

3. Massively Parallel 3D Image Reconstruction. Xiao Wang (Purdue University), Amit Sabne (Microsoft), Putt Sakdhnagool (Purdue University), Sherman J. Kisner (High Performance Imaging LLC), and Charles A. Bouman and Samuel P. Midkiff (Purdue University)....25

Technical Papers

4. LocoFS: A Loosely-Coupled Metadata Service for Distributed File Systems. Siyang Li, Youyou Lu, and Jiwu Shu (Tsinghua University); Yang Hu (University of Texas, Dallas); and Tao Li (University of Florida)....37

5. TagIt: An Integrated Indexing and Search Service for File Systems. Hyogi Sim (Virginia Tech, Oak Ridge National Laboratory); Youngjae Kim (Sogang University); Sudharshan S. Vazhkudai, Geoffroy R. Vallee, and Seung-Hwan Lim (Oak Ridge National Laboratory); and Ali R. Butt (Virginia Tech)....49

6. A Configurable Rule-Based Classful Token Bucket Filter Network Request Scheduler for the Lustre File System. Yingjin Qian, Xi Li, and Shuichi Ihara (DataDirect Networks) and Lingfang Zeng, Jürgen Kaiser, Tim Süß, and André Brinkmann (Johannes Gutenberg University Mainz)....61

7. Deep Learning at 15PF: Supervised and Semi-Supervised Classification for Scientific Data. Thorsten Kurth (Lawrence Berkeley National Laboratory), Jian Zhang (Stanford University), Nadathur Satish (Intel Corporation), Evan Racah (Lawrence Berkeley National Laboratory), Ioannis Mitliagkas (Stanford University), Md Mostofa Patwary (Intel Corporation), Tareq Malas (Lawrence Berkeley National Laboratory), Narayanan Sundaram (Intel Corporation), Wahid Bhimji (Lawrence Berkeley National Laboratory), Mikhail E. Smorkalov (Intel Corporation), Jack Deslippe (Lawrence Berkeley National Laboratory), Mikhail Shiryayev and Srinivas Sridharan (Intel Corporation), Mr Prabhat (Lawrence Berkeley National Laboratory), and Pradeep Dubey (Intel Corporation)....73

8. Understanding Error Propagation in Deep Learning Neural Network (DNN) Accelerators and Applications. Guanpeng Li (University of British Columbia); Siva Hari, Michael Sullivan, and Timothy Tsai (NVIDIA); Karthik Pattabiraman (University of British Columbia); and Joel Emer and Stephen Keckler (NVIDIA)....84

9. Scaling Deep Learning on GPU and Knights Landing Clusters. Yang You (University of California, Berkeley); Aydin Buluc (Lawrence Berkeley National Laboratory; University of California, Berkeley); and James Demmel (University of California, Berkeley)....96

10. Egeria: A Framework for Auto-Construction of HPC Advising Tools through Multi-Layered Natural Language Processing. Hui Guan, Xipeng Shen, and Hamid Krim (North Carolina State University)....108

11. DataRaceBench: A Benchmark Suite for Systematic Evaluation of Data Race Detection Tools. Chunhua Liao, Pei-Hung Lin, Joshua Asplund, Markus Schordan, and Ian Karlin (Lawrence Livermore National Laboratory)....122

12. Optimizing the Query Performance of Block Index Through Data Analysis and I/O Modeling. Tzuhsien Wu, Jerry Chou, and Hao Shyng (National Tsing Hua University, Taiwan); Bin Dong (Lawrence Berkeley National Laboratory); Scott Klasky (University of Tennessee); and Kesheng Wu (Lawrence Berkeley National Laboratory)....136

13. Sympiler: Transforming Sparse Matrix Codes by Decoupling Symbolic Analysis. Kazem Cheshmi (Rutgers University), Shoaib Kamil (Adobe Research), Michelle Mills Strout (University of Arizona), and Maryam Mehri Dehnavi (Rutgers University)....146

14. Control Replication: Compiling Implicit Parallelism to Efficient SPMD with Logical Regions. Elliott Slaughter (Stanford University, SLAC National Accelerator Laboratory); Wonchan Lee (Stanford University); Sean Treichler (Stanford University, NVIDIA); Wen Zhang (Stanford University); Michael Bauer (NVIDIA); Galen Shipman and Patrick McCormick (Los Alamos National Laboratory); and Alex Aiken (Stanford University)....159

15. Optimizing Geometric Multigrid Method Computation Using a DSL Approach. Vinay Vasista and Kumudha Narasimhan (Indian Institute of Science); Siddharth Bhat (International Institute of Information Technology, Hyderabad); and Uday Bondhugula (Indian Institute of Science)....171

16. Efficient Process Mapping in Geo-Distributed Cloud Data Centers. Amelie Chi Zhou (Shenzhen University), Yifan Gong (TuSimple), Bingsheng He (National University of Singapore), and Jidong Zhai (Tsinghua University)....184

17. Topology-Aware GPU Scheduling for Learning Workloads in Cloud Environments. Marcelo Amaral, Jorda Polo, and David Carrera (Barcelona Supercomputing Center) and Seetharami Seelam and Malgorzata Steinder (IBM)....196

18. Probabilistic Guarantees of Execution Duration for Amazon Spot Instances. Rich Wolski (University of California, Santa Barbara); John Brevik (California State University, Long Beach); Ryan Chard (Argonne National Laboratory); and Kyle Chard (University of Chicago)....208

19. A Framework for Scalable Biophysics-Based Image Analysis. Amir Gholami and Andreas Mang (University of Texas), Klaudius Scheufele (University of Stuttgart), Christos Davatzikos (University of Pennsylvania), Miriam Mehl (University of Stuttgart), and George Biros (University of Texas)....219

20. Galactos: Computing the 3-pt Anisotropic Correlation for 2 Billion Galaxies. Brian Friesen (Lawrence Berkeley National Laboratory), Md. Mostofa Ali Patwary (Intel Corporation), Brian Austin (Lawrence Berkeley National Laboratory), Nadathur Satish (Intel Corporation), Zachary Slepian (Lawrence Berkeley National Laboratory), Narayanan Sundaram (Intel Corporation), Deborah Bard (Lawrence Berkeley National Laboratory), Daniel Eisenstein (Harvard University), Jack Deslippe (Lawrence Berkeley National Laboratory), Pradeep Dubey (Intel Corporation), and Mr Prabhat (Lawrence Berkeley National Laboratory)....232

21. Extreme Scale Multi-Physics Simulations of the Tsunamigenic 2004 Sumatra Megathrust Earthquake. Carsten Uphoff, Sebastian Rettenberger, and

Michael Bader (Technical University Munich) and Stephanie Wollherr, Thomas Ulrich, Elizabeth H. Madden, and Alice-Agnes Gabriel (Ludwig Maximilian University of Munich)....243

22. GPU Triggered Networking for Intra-Kernel Communications. Michael LeBeane (University of Texas, Advanced Micro Devices Inc); Khaled Hamidouche and Brad Benton (Advanced Micro Devices Inc); Mauricio Breternitz (University Institute of Lisbon); Steven K. Reinhardt (Microsoft); and Lizy K. John (University of Texas)....259

23. Gravel: Fine-Grain GPU-Initiated Network Messages. Marc S. Orr (University of Wisconsin); Shuai Che and Bradford M. Beckmann (Advanced Micro Devices Inc); Mark Oskin (Advanced Micro Devices Inc, University of Washington); Steven K. Reinhardt (Microsoft); and David A. Wood (University of Wisconsin, Advanced Micro Devices Inc)....271

24. Toward Standardized Near-Data Processing with Unrestricted Data Placement for GPUs. Gwangsun Kim (ARM Ltd), Niladrish Chatterjee and Mike O'Connor (NVIDIA), and Kevin Hsieh (Carnegie Mellon University)....283

25. Understanding Object-Level Memory Access Patterns Across the Spectrum. Xu Ji (Tsinghua University, Qatar Computing Research Institute); Chao Wang (Oak Ridge National Laboratory); Nosayba El-Sayed (Massachusetts Institute of Technology, Qatar Computing Research Institute); Xiaosong Ma (Qatar Computing Research Institute); Youngjae Kim (Sogang University); Sudharshan S. Vazhkudai (Oak Ridge National Laboratory); Wei Xue (Tsinghua University); and Daniel Sanchez (Massachusetts Institute of Technology)....295

26. Exploring and Analyzing the Real Impact of Modern On-Package Memory on HPC Scientific Kernels. Ang Li (Pacific Northwest National Laboratory); Weifeng Liu (University of Copenhagen, Norwegian University of Science and Technology); Mads R. B. Kristensen and Brian Vinter (University of Copenhagen); Hao Wang and Kaixi Hou (Virginia Tech); and Andres Marquez and Shuaiwen Leon Song (Pacific Northwest National Laboratory)....307

27. Large-Scale Adaptive Mesh Simulations Through Non-Volatile Byte-Addressable Memory. Bao Nguyen, Hua Tan, and Xuechen Zhang (Washington State University, Vancouver)....321

28. Experimental and Analytical Study of Xeon Phi Reliability. Daniel Oliveira (Federal University of Rio Grande do Sul); Laercio Pilla (Federal University of Santa Catarina); Nathan DeBardeleben, Sean Blanchard, and Heather Quinn (Los Alamos National Laboratory); Israel Koren (University of Massachusetts); and Philippe Navaux and Paolo Rech (Federal University of Rio Grande do Sul)....333

29. REFINE: Realistic Fault Injection via Compiler-Based Instrumentation for Accuracy, Portability and Speed. Giorgis Georgakoudis (Queen's University Belfast); Ignacio Laguna (Lawrence Livermore National Laboratory); Dimitrios S. Nikolopoulos (Queen's University Belfast); and Martin Schulz (Lawrence Livermore National Laboratory, Technical University Munich)....345

30. Correcting Soft Errors Online in Fast Fourier Transform. Xin Liang, Jieyang Chen, Dingwen Tao, Sihuan Li, Panruo Wu, Hongbo Li, Kaiming Ouyang, and Yuanlai Liu (University of California, Riverside); Fengguang Song (Indiana University-Purdue University Indianapolis); and Zizhong Chen (University of California, Riverside)....359

31. Performance Modeling under Resource Constraints Using Deep Transfer Learning. Aniruddha Marathe, Rushil Anirudh, Nikhil Jain, Abhinav Bhatele, Jayaraman Thiagarajan, Bhavya Kailkhura, Jae-Seung Yeom, Barry Rountree, and Todd Gamblin (Lawrence Livermore National Laboratory)....371

32. Obtaining Dynamic Scheduling Policies with Simulation and Machine Learning. Danilo Carastan-Santos (Federal University of ABC, Santo André, Brazil; University of Grenoble) and Raphael Y. de Camargo (Federal University of ABC, Santo André, Brazil)....383

33. 0.5 Petabyte Simulation of a 45-Qubit Quantum Circuit. Thomas Häner and Damian S. Steiger (ETH Zurich)....396

34. Representative Paths Analysis. Nathan Tallent, Darren Kerbyson, and Adolfo Hoisie (Pacific Northwest National Laboratory)....406

35. ScrubJay: Deriving Knowledge from the Disarray of HPC Performance Data. Alfredo Gimenez (Lawrence Livermore National Laboratory; University of California, Davis); Todd Gamblin and Abhinav Bhatele (Lawrence Livermore National Laboratory); Chad Wood (University of Oregon); Kathleen Shoga, Aniruddha Marathe, and Peer-Timo Bremer (Lawrence Livermore National Laboratory); Bernd Hamann (University of California, Davis); and Martin Schulz (Technical University Munich, Lawrence Livermore National Laboratory)....418

36. Charliecloud: Unprivileged Containers for User-Defined Software Stacks in HPC. Reid Priedhorsky and Tim Randles (Los Alamos National Laboratory)....430

37. Securing HPC: Development of a Low Cost, Open Source, Multi-Factor Authentication Infrastructure. W. Cyrus Proctor, Patrick Storm, Matthew R. Hanlon, and Nathaniel Mendoza (Texas Advanced Computing Center, University of Texas)....440

38. Embracing a New Era of Highly Efficient and Productive Quantum Monte Carlo Simulations. Amrita Mathuriya (Intel Corporation), Ye Luo (Argonne

National Laboratory), Raymond C. Clay III (Sandia National Laboratories), Anouar Benali (Argonne National Laboratory), Luke Shulenburg (Sandia National Laboratories), and Jeongnim Kim (Intel Corporation)....451

39. An Efficient MPI/OpenMP Parallelization of the Hartree-Fock Method for the Second Generation of Intel Xeon Phi Processor. Vladimir Mironov (Lomonosov Moscow State University), Alexander Moskovsky (RSC Technologies), Kristopher Keipert (Iowa State University), Michael D'mello (Intel Corporation), Mark Gordon (Iowa State University), and Yuri Alexeev (Argonne National Laboratory)....463

40. Efficient and Scalable Calculation of Complex Band Structure Using Sakurai-Sugiura Method. Shigeru Iwase, Yasunori Futamura, Akira Imakura, Tetsuya Sakurai, and Tomoya Ono (University of Tsukuba)....475

41. Towards Fine-Grained Dynamic Tuning of HPC Applications on Modern Multi-Core Architectures. Mohammed Sourouri (Norwegian University of Science and Technology), Espen Birger Raknes (Aker BP ASA), Nico Reissmann (Norwegian University of Science and Technology), Johannes Langguth (Simula Research Laboratory), Daniel Hackenberg and Robert Schöne (Technical University Dresden), and Per Gunnar Kjeldsberg (Norwegian University of Science and Technology)....487

42. CAPES: Unsupervised Storage Performance Tuning Using Neural Network-Based Deep Reinforcement Learning. Yan Li, Kenneth Chang, Oceane Bel, Ethan L. Miller, and Darrell D. E. Long (University of California, Santa Cruz)....499

43. Input-Aware Auto-Tuning of Compute-Bound HPC Kernels. Philippe Tillet and David Cox (Harvard University)....513

44. Failures in Large Scale Systems: Long-Term Measurement, Analysis, and Implications. Saurabh Gupta (Intel Corporation), Tirthak Patel (Northeastern University), Christian Engelmann (Oak Ridge National Laboratory), and Devesh Tiwari (Northeastern University)....525

45. GUIDE: A Scalable Information Directory Service to Collect, Federate, and Analyze Logs for Operational Insights into a Leadership HPC Facility. Sudharshan S. Vazhkudai and Ross Miller (Oak Ridge National Laboratory); Devesh Tiwari (Northeastern University); and Christopher Zimmer, Feiyi Wang, Sarp Oral, Raghul Gunasekaran, and Deryl Steinert (Oak Ridge National Laboratory)....537

46. Scientific User Behavior and Data-Sharing Trends in a Petascale File System. Seung-Hwan Lim, Hyogi Sim, Raghul Gunasekaran, and Sudharshan S. Vazhkudai (Oak Ridge National Laboratory)....549

47. Scaling Betweenness Centrality Using Communication-Efficient Sparse Matrix Multiplication. Edgar Solomonik (University of Illinois), Maciej Besta (ETH

Zurich), Flavio Vella (Sapienza University of Rome), and Torsten Hoefler (ETH Zurich)....561

48. Distributed Southwell: An Iterative Method with Low Communication Costs. Jordi Wolfson-Pou and Edmond Chow (Georgia Institute of Technology)....575

49. Tessellating Stencils. Liang Yuan, Yunquan Zhang, Peng Guo, and Shan Huang (Chinese Academy of Sciences)....588

50. Predicting the Performance Impact of Different Fat-Tree Configurations. Nikhil Jain, Abhinav Bhatele, Louis Howell, David Boehme, Ian Karlin, and Edgar A. Leon (Lawrence Livermore National Laboratory); Misbah Mubarak (Argonne National Laboratory); Noah Wolfe (Rensselaer Polytechnic Institute); and Todd Gamblin and Matthew Leininger (Lawrence Livermore National Laboratory)....601

51. A Comparative Study of SDN and Adaptive Routing on Dragonfly Networks. Peyman Faizian, Md Atiqul Mollah, Zhou Tong, and Xin Yuan (Florida State University) and Michael Lang (Los Alamos National Laboratory)....614

52. Run-to-Run Variability on Xeon Phi Based Cray XC Systems. Sudheer Chunduri, Kevin Harms, Scott Parker, and Vitali Morozov (Argonne National Laboratory); Samuel Oshin and Naveen Cherukuri (Intel Corporation); and Kalyan Kumaran (Argonne National Laboratory)....625

53. Geometry-Oblivious FMM for Compressing Dense SPD Matrices. Chenhan D. Yu and James Levitt (University of Texas), Severin Reiz (Technical University Munich), and George Biros (University of Texas)....638

54. Low Communication FMM-Accelerated FFT on GPUs. Cris Cecka (NVIDIA)....652

55. Designing Vector-Friendly Compact BLAS and LAPACK Kernels. Kyungjoo Kim (Sandia National Laboratories); Timothy B. Costa (Intel Corporation); Mehmet Deveci, Andrew M. Bradley, and Simon D. Hammond (Sandia National Laboratories); Murat E. Guney, Sarah Knepper, and Shane Story (Intel Corporation); and Sivasankaran Rajamanickam (Sandia National Laboratories)....663

56. Transactional NVM Cache with High Performance and Crash Consistency. Qingsong Wei and Chundong Wang (Data Storage Institute); Cheng Chen (Data Storage Institute, National University of Singapore); and Yechao Yang, Jun Yang, and Mingdi Xue (Data Storage Institute)....675

57. PapyrusKV: A High-Performance Parallel Key-Value Store for Distributed NVM Architectures. Jungwon Kim, Seyong Lee, and Jeffrey Vetter (Oak Ridge National Laboratory)....687

- 58. Unimem: Runtime Data Management on Non-Volatile Memory-Based Heterogeneous Main Memory.** Kai Wu, Yingchao Huang, and Dong Li (University of California, Merced)....701
- 59. sPIN: High-Performance Streaming Processing in the Network.** Torsten Hoeffler, Salvatore Di Girolamo, and Konstantin Taranov (ETH Zurich) and Ryan Grant and Ronald Brightwell (Sandia National Laboratories)....715
- 60. Leveraging Near Data Processing for High-Performance Checkpoint/Restart.** Abhinav Agrawal (North Carolina State University), Gabriel H. Loh (Advanced Micro Devices Inc), and James Tuck (North Carolina State University)....731
- 61. Melissa: Large Scale In Transit Global Sensitivity Analysis Avoiding Intermediate Files.** Théophile Terraz (French Institute for Research in Computer Science and Automation (INRIA)); Alejandro Ribes, Yvan Fournier, and Bertrand Iooss (EDF France); and Bruno Raffin (French Institute for Research in Computer Science and Automation (INRIA))....743
- 62. Why Is MPI So Slow? Analyzing the Fundamental Limits in Implementing MPI-3.1.** Ken Raffenetti, Abdelhalim Amer, and Lena Oden (Argonne National Laboratory); Charles Archer, Wesley Bland, and Hajime Fujita (Intel Corporation); Yanfei Guo (Argonne National Laboratory); Tomislav Janjusic (Mellanox Technologies); Dmitry Durnov and Michael Blocksome (Intel Corporation); Min Si and Sangmin Seo (Argonne National Laboratory); Akhil Langer and Gengbin Zheng (Intel Corporation); Masamichi Takagi (RIKEN); Paul Coffman (Argonne National Laboratory); Jithin Jose, Sayantan Sur, Alexander Sannikov, Sergey Oblomov, and Michael Chuvelev (Intel Corporation); Masayuki Hatanaka (RIKEN); Xin Zhao (Mellanox Technologies); Paul Fischer and Thilina Rathnayake (University of Illinois); Matt Otten (Cornell University); and Misun Min and Pavan Balaji (Argonne National Laboratory)....757
- 63. ParaStack: Efficient Hang Detection for MPI Programs at Large Scale.** Hongbo Li, Zizhong Chen, and Rajiv Gupta (University of California, Riverside)....769
- 64. Scalable Reduction Collectives with Data Partitioning-Based Multi-Leader Design.** Mohammadreza Bayatpour, Sourav Chakraborty, Hari Subramoni, Xiaoyi Lu, and Dhabaleswar K. Panda (Ohio State University)....781