

SC16: International Conference for High Performance Computing, Networking, Storage and Analysis

**Salt Lake City, Utah, USA
13-18 November 2016**

Pages 1-513



**IEEE Catalog Number: CFP16SUP-POD
ISBN: 978-1-4673-8816-0**

**Copyright © 2016 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:	CFP16SUP-POD
ISBN (Print-On-Demand):	978-1-4673-8816-0
ISBN (Online):	978-1-4673-8815-3
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

SC16: International Conference for High Performance Computing, Networking, Storage and Analysis

SC 2016

Table of Contents

Organizing Committee	xiii
----------------------------	------

ACM Gordon Bell Finalist I

Towards Green Aviation with Python at Petascale	1
<i>Peter Vincent, Freddie Witherden, Brian Vermeire, Jin Seok Park, and Arvind Iyer</i>	
Modeling Dilute Solutions Using First-Principles Molecular Dynamics: Computing more than a Million Atoms with over a Million Cores	12
<i>Jean-Luc Fattebert, Daniel Osei-Kuffuor, Erik W. Draeger, Tadashi Ogitsu, and William D. Krauss</i>	
Simulations of Below-Ground Dynamics of Fungi: 1.184 Pflops Attained by Automated Generation and Autotuning of Temporal Blocking Codes	23
<i>Takayuki Muranushi, Hideyuki Hotta, Junichiro Makino, Seiya Nishizawa, Hirofumi Tomita, Keigo Nitadori, Masaki Iwasawa, Natsuki Hosono, Yutaka Maruyama, Hikaru Inoue, Hisashi Yashiro, and Yoshifumi Nakamura</i>	

ACM Gordon Bell Finalist II

Extreme-Scale Phase Field Simulations of Coarsening Dynamics on the Sunway TaihuLight Supercomputer	34
<i>Jian Zhang, Chunbao Zhou, Yangang Wang, Lili Ju, Qiang Du, Xuebin Chi, Dongsheng Xu, Dexun Chen, Yong Liu, and Zhao Liu</i>	
A Highly Effective Global Surface Wave Numerical Simulation with Ultra-High Resolution	46
<i>Fangli Qiao, Wei Zhao, Xunqiang Yin, Xiaomeng Huang, Xin Liu, Qi Shu, Guansuo Wang, Zhenya Song, Xinfang Li, Haixing Liu, Guangwen Yang, and Yeli Yuan</i>	

10M-Core Scalable Fully-Implicit Solver for Nonhydrostatic Atmospheric Dynamics	57
<i>Chao Yang, Wei Xue, Haohuan Fu, Hongtao You, Xinliang Wang, Yulong Ao, Fangfang Liu, Lin Gan, Ping Xu, Lanning Wang, Guangwen Yang, and Weimin Zheng</i>	

Molecular Dynamics Simulation

The Vectorization of the Tersoff Multi-body Potential: An Exercise in Performance Portability	69
<i>Markus Höhnerbach, Ahmed E. Ismail, and Ahmed E. Ismail</i>	
Increasing Molecular Dynamics Simulation Rates with an 8-Fold Increase in Electrical Power Efficiency	82
<i>W. Michael Brown, Andrey Semin, Michael Hebenstreit, Sergey Khvostov, Karthik Raman, and Steven J. Plimpton</i>	
Enhanced MPSM3 for Applications to Quantum Biological Simulations	96
<i>A. Pozdneev, V. Weber, T. Laino, C. Bekas, and A. Curioni</i>	

State-of-the-Practice: Advanced Applications Development

Development Effort Estimation in HPC	107
<i>Sandra Wienke, Julian Miller, Martin Schulz, and Matthias S. Müller</i>	
MetaMorph: A Library Framework for Interoperable Kernels on Multi- and Many-Core Clusters	119
<i>Ahmed E. Helal, Virginia Tech, Paul Sathre, and Wu-chun Feng</i>	
TrueNorth Ecosystem for Brain-Inspired Computing: Scalable Systems, Software, and Applications	130
<i>Jun Sawada, Filipp Akopyan, Andrew S. Cassidy, Brian Taba, Michael V. Debole, Pallab Datta, Rodrigo Alvarez-Icaza, Arnon Amir, John V. Arthur, Alexander Andreopoulos, Rathinakumar Appuswamy, Heinz Baier, Davis Barch, David J. Berg, Carmelo di Nolfo, Steven K. Esser, Myron Flickner, Thomas A. Horvath, Bryan L. Jackson, Jeff Kusnitz, Scott Lekuch, Michael Mastro, Timothy Melano, Paul A. Merolla, Steven E. Millman, Tapan K. Nayak, Norm Pass, Hartmut E. Penner, William P. Risk, Kai Schleupen, Benjamin Shaw, Hayley Wu, Brian Giera, Adam T. Moody, Nathan Mundhenk, Brian C. Van Essen, Eric X. Wang, David P. Widemann, Qing Wu, William E. Murphy, Jamie K. Infantolino, James A. Ross, Dale R. Shires, Manuel M. Vindiola, Raju Namburu, and Dharmendra S. Modha</i>	

Systems and Networks I

Scheduling-Aware Routing for Supercomputers	142
<i>Jens Domke and Torsten Hoefler</i>	
Evaluating HPC Networks via Simulation of Parallel Workloads	154
<i>Nikhil Jain, Abhinav Bhatele, Sam White, Todd Gamblin, and Laxmikant V. Kale</i>	
Flexfly: Enabling a Reconfigurable Dragonfly through Silicon Photonics	166
<i>Ke Wen, Payman Samadi, Sébastien Rumley, Christine P. Chen, Yiwen Shen, Meisam Bahadori, Keren Bergman, and Jeremiah Wilke</i>	

Numerical Algorithms I

PFEAST: A High Performance Sparse Eigenvalue Solver Using Distributed-Memory Linear Solvers	178
<i>James Kestyn, Vasileios Kalantzis, Eric Polizzi, and Yousef Saad</i>	
Block Iterative Methods and Recycling for Improved Scalability of Linear Solvers	190
<i>Pierre Jolivet and Pierre-Henri Tournier</i>	
Scalable Non-blocking Preconditioned Conjugate Gradient Methods	204
<i>Paul R. Eller and William Gropp</i>	

Resilience and Error Handling

Pinpointing Scale-Dependent Integer Overflow Bugs in Large-Scale Parallel Applications	216
<i>Ignacio Laguna and Martin Schulz</i>	
Compiler-Directed Lightweight Checkpointing for Fine-Grained Guaranteed Soft Error Recovery	228
<i>Qingrui Liu, Changhee Jung, Dongyoon Lee, and Devesh Tiwari</i>	
Understanding Error Propagation in GPGPU Applications	240
<i>Guanpeng Li, Karthik Pattabiraman, Chen-Yang Cher, and Pradip Bose</i>	

Scientific Data Management and Visualization

Simulation and Performance Analysis of the ECMWF Tape Library System	252
<i>Markus Mäscher, Lars Nagel, Tim Süß, André Brinkmann, and Lennart Sorth</i>	
Real-Time Synthesis of Compression Algorithms for Scientific Data	264
<i>Martin Burtscher, Hari Mukka, Annie Yang, and Farbod Hesaaraki</i>	
Performance Modeling of In Situ Rendering	276
<i>Matthew Larsen, Cyrus Harrison, James Kress, David Pugmire, Jeremy S. Meredith, and Hank Childs</i>	

Topics in Distributed Computing

HARP: Predictive Transfer Optimization Based on Historical Analysis and Real-Time Probing	288
<i>Engin Arslan, Kemal Guner, and Tevfik Kosar</i>	
SERF: Efficient Scheduling for Fast Deep Neural Network Serving via Judicious Parallelism	300
<i>Feng Yan, Olatunji Ruwase, Yuxiong He, and Evgenia Smirni</i>	

Resilience

Failure Detection and Propagation in HPC systems	312
<i>George Bosilca, Aurelien Bouteiller, Amina Guermouche, Thomas Herault, Yves Robert, Pierre Sens, and Jack Dongarra</i>	
Improving Application Resilience to Memory Errors with Lightweight Compression	323
<i>Scott Levy, Kurt B. Ferreira, and Patrick G. Bridges</i>	
FlipBack: Automatic Targeted Protection against Silent Data Corruption	335
<i>Xiang Ni and Laxmikant V. Kale</i>	

Tensor and Graph Algorithms

Graph Colouring as a Challenge Problem for Dynamic Graph Processing on Distributed Systems	347
<i>Scott Sallinen, Keita Iwabuchi, Suraj Poudel, Maya Gokhale, Matei Ripeanu, and Roger Pearce</i>	
An Exploration of Optimization Algorithms for High Performance Tensor Completion	359
<i>Shaden Smith, Jongsoo Park, and George Karypis</i>	
An Efficient and Scalable Algorithmic Method for Generating Large-Scale Random Graphs	372
<i>Maksudul Alam, Maleq Khan, Anil Vullikanti, and Madhav Marathe</i>	

Performance Measurement and Analysis

Understanding Performance Interference in Next-Generation HPC Systems	384
<i>Oscar H. Mondragon, Patrick G. Bridges, Scott Levy, Kurt B. Ferreira, and Patrick Widener</i>	
Reliable and Efficient Performance Monitoring in Linux	396
<i>Maria Dimakopoulou, Stéphane Eranian, Nectarios Koziris, and Nicholas Bambos</i>	

Evaluating and Optimizing OpenCL Kernels for High Performance Computing with FPGAs	409
<i>Hamid Reza Zohouri, Naoya Maruyamay, Aaron Smith, Motohiko Matsuda, and Satoshi Matsuoka</i>	

Systems and Networks II

Enhancing InfiniBand with OpenFlow-Style SDN Capability	421
<i>Jason Lee, Zhou Tong, Karthik Achalkar, Xin Yuan, and Michael Lang</i>	
Designing MPI Library with On-Demand Paging (ODP) of InfiniBand: Challenges and Benefits	433
<i>Mingzhe Li, Khaled Hamidouche, Xiaoyi Lu, Hari Subramoni, Jie Zhang, and Dhabaleswar K. Panda</i>	
The Mont-Blanc Prototype: An Alternative Approach for HPC Systems	444
<i>Nikola Rajovic, Alejandro Rico, Filippo Mantovani, Daniel Ruiz, Josep Oriol Vilarrubi, Constantino Gomez, Luna Backes, Diego Nieto, Harald Servat, Xavier Martorell, Jesus Labarta, Eduard Ayguade, Chris Adeniyi-Jones, Said Derradji, Herve Gloaguen, Piero Lanucara, Nico Sanna, Jean-Francois Mehaut, Kevin Pouget, Brice Videau, Eric Boyer, Momme Allalen, Axel Auweter, David Brayford, Daniele Tafani, Volker Weinberg, Dirk Brömmel, René Halver, Jan H. Meinke, Ramon Beivide, Mariano Benito, Enrique Vallejo, Mateo Valero, and Alex Ramirez</i>	

Compilation for Enhanced Parallelism

PIPES: A Language and Compiler for Task-Based Programming on Distributed-Memory Clusters	456
<i>Martin Kong, Louis-Noël Pouchet, P. Sadayappan, and Vivek Sarkar</i>	
A Domain-Specific Compiler for a Parallel Multiresolution Adaptive Numerical Simulation Environment	468
<i>Samyam Rajbhandari, Jinsung Kim, Sriram Krishnamoorthy, Louis-Noel Pouchet, Fabrice Rastello, Robert J. Harrison, and P. Sadayappan</i>	
Automating Wavefront Parallelization for Sparse Matrix Computations	480
<i>Anand Venkat, Mahdi Soltan Mohammadi, Jongsoo Park, Hongbo Rong, Rajkishore Barik, Michelle Mills Strout, and Mary Hall</i>	

Fluid Dynamics

Granularity and the Cost of Error Recovery in Resilient AMR Scientific Applications	492
<i>Anshu Dubey, Hajime Fujita, Daniel T. Graves, Andrew Chien, and Devesh Tiwari</i>	

Extreme Scale Plasma Turbulence Simulations on Top Supercomputers Worldwide	502
<i>William Tang, Bei Wang, Stephane Ethier, Grzegorz Kwasniewski, Torsten Hoefler, Khaled Z. Ibrahim, Kamesh Madduri, Samuel Williams, Leonid Oliker, Carlos Rosales-Fernandez, and Tim Williams</i>	

A Parallel Arbitrary-Order Accurate AMR Algorithm for the Scalar Advection-Diffusion Equation	514
<i>Arash Bakhtiari, Dhairya Malhotra, Amir Raoofy, Miriam Mehl, Hans-Joachim Bungartz, and George Biros</i>	

Performance Tools

MUSA: A Multi-level Simulation Approach for Next-Generation HPC Machines	526
<i>Thomas Grass, César Allande, Adrià Armejach, Alejandro Rico, Eduard Ayguadé, Jesus Labarta, Mateo Valero, Marc Casas, and Miquel Moreto</i>	

A Machine Learning Framework for Performance Coverage Analysis of Proxy Applications	538
<i>Tanzima Z. Islam, Jayaraman J. Thiagarajan, Abhinav Bhatele, Martin Schulz, and Todd Gamblin</i>	

Caliper: Performance Introspection for HPC Software Stacks	550
<i>David Boehme, Todd Gamblin, David Beckingsale, Peer-Timo Bremer, Alfredo Gimenez, Matthew LeGendre, Olga Pearce, and Martin Schulz</i>	

Storage Systems

Exploring the Potentials of Parallel Garbage Collection in SSDs for Enterprise Storage Systems	561
<i>Narges Shahidi, Mahmut T. Kandemir, Mohammad Arjomand, Chita R. Das, Myoungsoo Jung, and Anand Sivasubramaniam</i>	

Týr: Blob Storage Meets Built-In Transactions	573
<i>Pierre Matri, Alexandru Costan, Gabriel Antoniu, Jesús Montes, and María S. Pérez</i>	

DAOS and Friends: A Proposal for an Exascale Storage System	585
<i>Jay Lofstead, Ivo Jimenez, Carlos Maltzahn, Quincey Koziol, John Bent, and Eric Barton</i>	

Accelerator Programming Tools

Translating OpenMP Device Constructs to OpenCL Using Unnecessary Data Transfer Elimination	597
<i>Junghyun Kim, Yong-Jun Lee, Jungho Park, and Jaejin Lee</i>	

dCUDA: Hardware Supported Overlap of Computation and Communication	609
<i>Tobias Gysi, Jeremia Bär, and Torsten Hoefler</i>	

Daino: A High-Level Framework for Parallel and Efficient AMR on GPUs	621
<i>Mohamed Wahib, Naoya Maruyama, and Takayuki Aoki</i>	

Memory and Power

Optimizing Memory Efficiency for Deep Convolutional Neural Networks on GPUs	633
<i>Chao Li, Yi Yang, Min Feng, Srimat Chakradhar, and Huiyang Zhou</i>	
Unprotected Computing: A Large-Scale Study of DRAM Raw Error Rate on a Supercomputer	645
<i>Leonardo Bautista-Gomez, Ferad Zyulkyarov, Osman Unsal, and Simon McIntosh-Smith</i>	
A Data Driven Scheduling Approach for Power Management on HPC Systems	656
<i>Sean Wallace, Xu Yang, Venkatram Vishwanath, William E. Allcock, Susan Coghlan, Michael E. Papka, and Zhiling Lan</i>	

Numerical Algorithms, Part II

GreenLA: Green Linear Algebra Software for GPU-accelerated Heterogeneous Computing	667
<i>Jieyang Chen, Li Tan, Panruo Wu, Dingwen Tao, Hongbo Li, Xin Liang, Sihuan Li, Rong Ge, Laxmi Bhuyan, and Zizhong Chen</i>	
Merge-Based Parallel Sparse Matrix-Vector Multiplication	678
<i>Duane Merrill and Michael Garland</i>	
Strassen's Algorithm Reloaded	690
<i>Jianyu Huang, Tyler M. Smith, Greg M. Henry, and Robert A. van de Geijn</i>	

Data Analytics

Optimal Execution of Co-analysis for Large-Scale Molecular Dynamics Simulations	702
<i>Preeti Malakar, Venkatram Vishwanath, Christopher Knight, Todd Munson, and Michael E. Papka</i>	
ScaleMine: Scalable Parallel Frequent Subgraph Mining in a Single Large Graph	716
<i>Ehab Abdelhamid, Ibrahim Abdelaziz, Panos Kalnis, Zuhair Khayyat, and Fuad Jamour</i>	
Efficient Delaunay Tessellation through K-D Tree Decomposition	728
<i>Dmitriy Morozov and Tom Peterka</i>	

Performance Analysis of Network Systems

A PCIe Congestion-Aware Performance Model for Densely Populated Accelerator Servers	739
<i>Maxime Martinasso, Grzegorz Kwasniewski, Sadaf R. Alam, Thomas C. Schulthess, and Torsten Hoefler</i>	
Watch Out for the Bully! Job Interference Study on Dragonfly Network	750
<i>Xu Yang, John Jenkins, Misbah Mubarak, Robert B. Ross, and Zhiling Lan</i>	
Measuring and Understanding Throughput of Network Topologies	761
<i>Sangeetha Abdu Jyothi, Ankit Singla, P. Brighten Godfrey, and Alexandra Kolla</i>	

Combinatorial and Multigrid Algorithms

Designing Scalable b -MATCHING Algorithms on Distributed Memory Multiprocessors by Approximation	773
<i>Arif Khan, Alex Pothan, Md. Mostofa Ali Patwary, Mahantesh Halappanavar, Nadathur Rajagopalan Satish, Narayanan Sundaram, and Pradeep Dubey</i>	
A Parallel Algorithm for Finding All Pairs κ -Mismatch Maximal Common Substrings	784
<i>Sriram P. Chockalingam, Sharma V. Thankachan, and Srinivas Aluru</i>	
Accelerating Lattice QCD Multigrid on GPUs Using Fine-Grained Parallelization	795
<i>M. A. Clark, Bálint Joó, Alexei Strelchenko, Michael Cheng, Arjun Gambhir, and Richard. C. Brower</i>	

File Systems and I/O

An Ephemeral Burst-Buffer File System for Scientific Applications	807
<i>Teng Wang, Kathryn Mohror, Adam Moody, Kento Sato, and Weikuan Yu</i>	
Server-Side Log Data Analytics for I/O Workload Characterization and Coordination on Large Shared Storage Systems	819
<i>Yang Liu, Raghul Gunasekaran, Xiaosong Ma, and Sudharshan S. Vazhkudai</i>	
G-Store: High-Performance Graph Store for Trillion-Edge Processing	830
<i>Pradeep Kumar and H. Howie Huang</i>	

Inverse Problems and Quantum Circuits

Distributed-Memory Large Deformation Diffeomorphic 3D Image Registration	842
<i>Andreas Mang, Amir Gholami, and George Biros</i>	
ZNN i : Maximizing the Inference Throughput of 3D Convolutional Networks on CPUs and GPUs	854
<i>Aleksandar Zlateski, Kisuk Lee, and H. Sebastian Seung</i>	

High Performance Emulation of Quantum Circuits	866
<i>Thomas Häner, Damian S. Steiger, Mikhail Smelyanskiy, and Matthias Troyer</i>	

Manycore Architectures

Elastic Multi-resource Fairness: Balancing Fairness and Efficiency in Coupled CPU-GPU Architectures	875
<i>Shanjiang Tang, BingSheng He, Shuhao Zhang, and Zhaojie Niu</i>	
DCA: A DRAM-cache-Aware DRAM Controller	887
<i>Cheng-Chieh Huang, Vijay Nagarajan, and Arpit Joshi</i>	
Enabling Efficient Preemption for SIMT Architectures with Lightweight Context Switching	898
<i>Zhen Lin, Lars Nyland, and Huiyang Zhou</i>	

State-of-the-Practice: System Characterization and Design

Characterizing Parallel Scientific Applications on Commodity Clusters: An Empirical Study of a Tapered Fat-Tree	909
<i>Edgar A. León, Ian Karlin, Abhinav Bhatele, Steven H. Langer, Chris Chembreau, Louis H. Howell, Trent D'Hooge, and Matthew L. Leininger</i>	
Performance Analysis, Design Considerations, and Applications of Extreme-Scale In Situ Infrastructures	921
<i>Utkarsh Ayachit, Andrew Bauer, Earl P. N. Duque, Greg Eisenhauer, Nicola Ferrier, Junmin Gu, Kenneth E. Jansen, Burlen Loring, Zarija Lukic, Suresh Menon, Dmitriy Morozov, Patrick O'Leary, Reetesh Ranjan, Michel Rasquin, Christopher P. Stone, Venkat Vishwanath, Gunther H. Weber, Brad Whitlock, Matthew Wolf, K. John Wu, and E. Wes Bethel</i>	

Task-Oriented Runtimes

Extended Task Queuing: Active Messages for Heterogeneous Systems	933
<i>Michael LeBeane, Brandon Potter, Abhisek Pan, Alexandru Dutu, Vinay Agarwala, Wonchan Lee, Deepak Majeti, Bibek Ghimire, Eric Van Tassell, Samuel Wasmundt, Brad Benton, Mauricio Breternitz, Michael L. Chu, Mithuna Thottethodi, Lizy K. John, and Steven K. Reinhardt</i>	
Perilla: Metadata-Based Optimizations of an Asynchronous Runtime for Adaptive Mesh Refinement	945
<i>Tan Nguyen, Didem Unat, Weiqun Zhang, Ann Almgren, Nufail Farooqi, and John Shalf</i>	

Accelerating Science

High-Frequency Nonlinear Earthquake Simulations on Petascale Heterogeneous Supercomputers	957
<i>Daniel Roten, Yifeng Cui, Kim B. Olsen, Steven M. Day, Kyle Withers, William H. Savran, Peng Wang, and Dawei Mu</i>	
Refactoring and Optimizing the Community Atmosphere Model (CAM) on the Sunway TaihuLight Supercomputer	969
<i>Haohuan Fu, Junfeng Liao, Wei Xue, Lanning Wang, Dexun Chen, Long Gu, Jinxu Xu, Nan Ding, Xinliang Wang, Conghui He, Shizhen Xu, Yishuang Liang, Jiarui Fang, Yuanchao Xu, Weijie Zheng, Jingheng Xu, Zhen Zheng, Wanjing Wei, Xu Ji, He Zhang, Bingwei Chen, Kaiwei Li, Xiaomeng Huang, Wenguang Chen, and Guangwen Yang</i>	
LIBXSMM: Accelerating Small Matrix Multiplications by Runtime Code Generation	981
<i>Alexander Heinecke, Greg Henry, Maxwell Hutchinson, and Hans Pabst</i>	

Clouds & Job Scheduling

Transient Guarantees: Maximizing the Value of Idle Cloud Capacity	992
<i>Supreeth Shastri, Amr Rizk, and David Irwin</i>	
Multi-resource Fair Sharing for Datacenter Jobs with Placement Constraints	1003
<i>Wei Wang, Baochun Li, Ben Liang, and Jun Li</i>	
A Multi-faceted Approach to Job Placement for Improved Performance on Extreme-Scale Systems	1015
<i>Christopher Zimmer, Saurabh Gupta, Scott Atchley, Sudharshan S. Vazhkudai, and Carl Albing</i>	

Author Index