

2013 International Conference for High Performance Computing, Networking, Storage and Analysis

(SC 2013)

**Denver, Colorado, USA
17-22 November 2013**

Pages 1-555



IEEE Catalog Number: CFP13SUP-POD
ISBN: 978-1-4799-3520-8



Proceedings of SC13

The International Conference for High Performance
Computing, Networking, Storage and Analysis
Denver, Colorado - 17-22 November 2013

Introduction

ACM Gordon Bell Finalist I

[Taking a Quantum Leap in Time to Solution for Simulations of High-Tc Superconductors](#)

Peter Staar, Thomas A. Maier, Raffaele Solca, Gilles Fourestey, Michael Summers, Thomas C. Schulthess

[20 Petaflops Simulation of Protein Suspensions in Crowding Conditions](#)

Massimo Bernaschi, Mauro Bisson, Massimiliano Fatica, Simone Melchionna

[11 PFLOP/s Simulations of Cloud Cavitation Collapse](#)

Diego Rossinelli, Babak Hejazialhosseini, Panagiotis Hadjidoukas, Costas Bekas, Alessandro Curioni, Adam Bertsch, Scott Futral, Steffen Schmidt, Nikolaus Adams, Petros Koumoutsakos

ACM Gordon Bell Finalist II

[The Origin of Mass](#)

Peter Boyle, Michael I. Buchoff, Norman Christ, Taku Izubuchi, Chulwoo Jung, Thomas C. Luu, Robert Mawhinney, Chris Schroeder, Ron Soltz, Pavlos Vranas, Joseph Wasem, Zhongjie Lin, Hantao Yin

[Radiative Signatures of the Relativistic Kelvin-Helmholtz Instability](#)

Michael Bussmann, Heiko Burau, Thomas E. Cowan, Alexander Debus, Axel Huebl, Guido Juckeland, Thomas Kluge, Wolfgang E. Nagel, Richard Pausch, Felix Schmitt, Ulrich Schramm, Joseph Schuchart, Rene Widera

[HACC: Extreme Scaling and Performance Across Diverse Architectures](#)

Salman Habib, Vitali A. Morozov, Nicholas Frontiere, Hal Finkel, Adrian Pope, Katrin Heitmann, Kalyan Kumaran, Venkat Vishwanath, Tom Peterka, Joseph A. Insley, David Daniel, Patricia Fasel, Zarija Lukic

Fault-Tolerant Computing

[ACR: Automatic Checkpoint/Restart for Soft and Hard Error Protection](#)

Xiang Ni, Esteban Meneses, Nikhil Jain, Laxmikant Kale

[SPBC: Leveraging the Characteristics of MPI HPC Applications for Scalable Checkpointing](#)

Thomas Ropars, Tatiana Martsinkevich, Amina Guermouche, André Schiper,
Franck Cappello

[Using Simulation to Explore Distributed Key-Value Stores for Extreme-Scale System Services](#)

Ke Wang, Abhishek Kulkarni, Michael Lang, Dorian Arnold, Ioan Raicu

GPU Programming

[General Transformations for GPU Execution of Tree Traversals](#)

Michael Goldfarb, Youngjoon Jo, Milind Kulkarni

[A Large-Scale Cross-Architecture Evaluation of Thread-Coarsening](#)

Alberto Magni, Christophe Dubach, Michael F.P. O'Boyle

[Semi-Automatic Restructuring of Offloadable Tasks for Many-Core Accelerators](#)

Nishkam Ravi, Yi Yang, Tao Bao, Srimat Chakradhar

Load Balancing

[A Framework for Load Balancing of Tensor Contraction Expressions via Dynamic Task Partitioning](#)

Pai-Wei Lai, Kevin Stock, Samyam Rajbhandari, Sriram Krishnamoorthy, P.
Sadayappan

[Load-Balanced Pipeline Parallelism](#)

Md Kamruzzaman, Steven Swanson, Dean Tullsen

[A Distributed Dynamic Load Balancer for Iterative Applications](#)

Harshitha Menon, Laxmikant Kale

MPI Performance and Debugging

[Distributed Wait State Tracking for Runtime MPI Deadlock Detection](#)

Tobias Hilbrich, Bronis R. de Supinski, Wolfgang E. Nagel, Joachim Protze,
Christel Baier, Matthias S. Mueller

[Globalizing Selectively: Shared-Memory Efficiency with Address-Space Separation](#)

Nilesh Mahajan, Uday Pitambare, Arun Chauhan

[Hybrid MPI: Efficient Message Passing for Multi-Core Systems](#)

Andrew Friedley, Greg Bronevetsky, Torsten Hoefler, Andrew Lumsdaine

Memory Hierarchy

[Performance Evaluation of Intel Transactional Synchronization Extensions for High Performance Computing](#)

Richard Yoo, Christopher Hughes, Konrad Lai, Ravi Rajwar

[Location-Aware Cache Management for Many-Core Processors with Deep Cache Hierarchy](#)

Jongsoo Park, Richard M. Yoo, Daya S. Khudia, Christopher J. Hughes,
Daehyun Kim

[Practical Nonvolatile Multilevel-Cell Phase Change Memory](#)   

Doe Hyun Yoon, Jichuan Chang, Robert S. Schreiber, Norman P. Jouppi

Memory Resilience

[Feng Shui of Supercomputer Memory: Positional Effects in DRAM and SRAM Faults](#)   

Vilas Sridharan, Jon Stearley, Nathan DeBardeleben, Sean Blanchard,
Sudhanva Gurumurthi

[Exploring DRAM Organizations for Energy-Efficient and Resilient Exascale Memories](#)   

Bharan Giridhar, Michael Cieslak, Deepankar Duggal, Ronald Dreslinski, Hsing
Min Chen, Robert Patti, Betina Hold, Chaitali Chakrabarti, Trevor Mudge, David
Blaauw




[Low-Power, Low-Storage-Overhead Chipkill Correct via Multi-Line Error Correction](#)   

Xun Jian, Henry Duwe, John Sartori, Vilas Sridharan, Rakesh Kumar

Optimizing Numerical Code

[AUGEM: Automatically Generate High Performance Dense Linear Algebra Kernels on x86 CPUs](#)   

Qian Wang, Xianyi Zhang, Yunquan Zhang, Qing Yi

[Accelerating Sparse Matrix-Vector Multiplication on GPUs using Bit-Representation-Optimized Schemes](#)   

Wai Teng Tang, Wen Jun Tan, Rajarshi Ray, Yi Wen Wong, Weiguang Chen,
Shyh-hao Kuo, Rick Siow Mong Goh, Stephen John Turner, Weng-Fai Wong

[Precimonious: Tuning Assistant for Floating-Point Precision](#)   

Cindy Rubio-González, Cuong Nguyen, Hong Diep Nguyen, James Demmel,
William Kahan, Koushik Sen, David H. Bailey, Costin Iancu, David Hough

Parallel Performance Tools

[A Data-Centric Profiler for Parallel Programs](#)   

Xu Liu, John Mellor-Crummey

[On the Usefulness of Object Tracking Techniques in Performance Analysis](#)   

German Llort, Harald Servat, Juan Gonzalez, Judit Gimenez, Jesus Labarta

[Detection of False Sharing Using Machine Learning](#)   

Sanath Jayasena, Saman Amarasinghe, Asanka Abeyweera, Gayashan
Amarasinghe, Himeshi De Silva, Sunimal Rathnayake, Xiaoqiao Meng, Yanbin
Liu

Parallel Programming Models and Compilation

[Parallelizing the Execution of Sequential Scripts](#)    

Zhao Zhang, Daniel S. Katz, Timothy G. Armstrong, Justin M. Wozniak, Ian Foster

[Deterministic Scale-Free Pipeline Parallelism with Hyperqueues](#) 

Hans Vandierendonck, Kallia Chronaki, Dimitrios S. Nikolopoulos

[Compiling Affine Loop Nests for Distributed-Memory Parallel Architectures](#) 

Uday Bondhugula

Performance Analysis of Applications at Large Scale

[Tera-Scale 1D FFT with Low-Communication Algorithm on Intel Xeon Phi Coprocessors](#) 

Jongsoo Park, Ganesh Bikshandi, Karthikeyan Vaidyanathan, Ping Tak Peter Tang, Pradeep Dubey, Daehyun Kim

[A Framework for Hybrid Parallel Flow Simulations with a Trillion Cells in Complex Geometries](#) 

Christian Godenschwager, Florian Schornbaum, Martin Bauer, Harald Köstler, Ulrich Rüde

[A New Routing Scheme for Jellyfish and its Performance with HPC Workloads](#) 

Xin Yuan, Santosh Mahapatra, Wickus Nienaber, Scott Pakin, Michael Lang

Performance Management of HPC Systems

[Enabling Fair Pricing on HPC Systems with Node Sharing](#) 

Alex D. Breslow, Ananta Tiwari, Martin Schulz, Laura Carrington, Lingjia Tang, Jason Mars

[ACIC: Automatic Cloud I/O Configurator for HPC Applications](#) 

Mingliang Liu, Ye Jin, Jidong Zhai, Yan Zhai, Qianqian Shi, Xiaosong Ma, Wenguang Chen

[COCA: Online Distributed Resource Management for Cost Minimization and Carbon Neutrality in Data Centers](#) 

Shaolei Ren, Yuxiong He

System-wide Application Performance Assessments

[Supercomputing with Commodity CPUs: Are Mobile SoCs Ready for HPC?](#) 

Nikola Rajovic, Paul M. Carpenter, Isaac Gelado, Nikola Puzovic, Alex Ramirez, Mateo Valero

[There Goes the Neighborhood: Performance Degradation due to Nearby Jobs](#) 

Abhinav Bhatele, Kathryn Mohror, Steven H. Langer, Katherine E. Isaacs



[CooMR: Cross-Task Coordination for Efficient Data Management in MapReduce Programs](#) 

Xiaobing Li, Yandong Wang, Yizheng Jiao, Cong Xu, Weikuan Yu

Tools for Scalable Analysis

[Effective Sampling-Driven Performance Tools for GPU-Accelerated Supercomputers](#)  

Milind Chabbi, Karthik Murthy, Michael Fagan, John Mellor-Crummey

[Rethinking Algorithm-Based Fault Tolerance with a Cooperative Software-Hardware Approach](#)  

Dong Li, Zizhong Chen, Panruo Wu, Jeffrey S. Vetter

[Using Automated Performance Modeling to Find Scalability Bugs in Complex Codes](#)  

Alexandru Calotoiu, Torsten Hoefler, Marius Poke, Felix Wolf



Data Management in the Cloud

[Efficient Data Partitioning Model for Heterogeneous Graphs in the Cloud](#)  

Kisung Lee, Ling Liu

[SDQuery DSI: Integrating Data Management Support with a Wide Area Data Transfer Protocol](#)  

Yu Su, Yi Wang, Gagan Agrawal, Rajkumar Kettimuthu

[Design and Performance Evaluation of NUMA-Aware RDMA-Based End-to-End Data Transfer Systems](#)  

Yufei Ren, Tan Li, Dantong Yu, Shudong Jin, Thomas Robertazzi

Graph Partitioning and Data Clustering

[Scalable Parallel OPTICS Data Clustering Using Graph Algorithmic Techniques](#)  

Md. Mostofa Ali Patwary, Diana Palsetia, Ankit Agrawal, Wei-keng Liao, Fredrik Manne, Alok Choudhary

[Scalable Matrix Computations on Large Scale-Free Graphs Using 2D Graph Partitioning](#)  

Erik G. Boman, Karen D. Devine, Sivasankaran Rajamanickam

[Scalable Parallel Graph Partitioning](#)  

Shad Kirmani, Padma Raghavan

Inter-Node Communication

[Channel Reservation Protocol for Over-Subscribed Channels and Destinations](#)  

George Michelogiannakis, Nan Jiang, Daniel U. Becker, William J. Dally

[Enabling Highly-Scalable Remote Memory Access Programming with MPI-3 One Sided](#)  

Robert Gerstenberger, Maciej Besta, Torsten Hoefler

[MVAPICH-PRISM: A Proxy-Based Communication Framework Using InfiniBand and SCIF for Intel MIC Clusters](#)  

Sreeram Potluri, Devendar Bureddy, Khaled Hamidouche, Akshay Venkatesh, Krishna Kandalla, Hari Subramoni, Dhableswar K. Panda

Cloud Resource Management and Scheduling

[Exploring Portfolio Scheduling for Long-Term Execution of Scientific Workloads in IaaS Clouds](#)  

Kefeng Deng, Junqiang Song, Kaijun Ren, Alexandru Iosup

[Cost-Effective Cloud HPC Resource Provisioning by Building Semi-Elastic Virtual Clusters](#)

Shuangcheng Niu, Jidong Zhai, Xiaosong Ma, Xiongchao Tang, Wenguang Chen

[Exploiting Application Dynamism and Cloud Elasticity for Continuous Dataflows](#)

Alok Gautam Kumbhare, Yogesh Simmhan, Viktor K. Prasanna

Energy Management

[A "Cool" Way of Improving the Reliability of HPC Machines](#)

Osman Sarood, Esteban Meneses, Laxmikant Kale

[Coordinated Energy Management in Heterogeneous Processors](#)

Indrani Paul, Vignesh Ravi, Srilatha Manne, Manish Arora, Sudhakar Yalamanchili

[Integrating Dynamic Pricing of Electricity into Energy Aware Scheduling for HPC Systems](#)

Xu Yang, Zhou Zhou, Sean Wallace, Zhiling Lan, Wei Tang, Susan Coghlan, Michael E. Papka

Extreme-Scale Applications

[Petascale Direct Numerical Simulation of Turbulent Channel Flow on up to 786K Cores](#)

Myoungkyu Lee, Nicholas Malaya, Robert D. Moser

[Solving the Compressible Navier-Stokes Equations on up to 1.97 Million Cores and 4.1 Trillion Grid Points](#)

Ivan Bermejo-Moreno, Julien Bodart, Johan Larsson, Blaise Barney, Joseph Nichols, Steve Jones

[Petascale WRF Simulation of Hurricane Sandy: Deployment of NCSA's Cray XE6 Blue Waters](#)

Peter Johnsen, Mark Straka, Melvyn Shapiro, Alan Norton, Thomas Galarneau

Fault Tolerance and Migration in the Cloud

[Optimization of Cloud Task Processing with Checkpoint-Restart Mechanism](#)

Sheng Di, Yves Robert, Frederic Vivien, Derrick Kondo, Cho-Li Wang, Franck Cappello

[Scalable Virtual Machine Deployment Using VM Image Caches](#)

Kaveh Razavi, Thilo Kielmann

[Guide-Copy: Fast and Silent Migration of Virtual Machine for Datacenters](#)

Jihun Kim, Dongju Chae, Jangwoo Kim, Jong Kim

IO Tuning

[Characterization and Modeling of PIDX Parallel I/O for Performance Optimization](#)

Sidharth Kumar, Avishek Saha, Venkatram Vishwanath, Philip Carns, John A. Schmidt, Robert Latham, Giorgio Scorzelli, Hemanth Kolla, Robert Ross, Jackie Chen, Michael E. Papka, Ray Grout, Valerio Pascucci

[Taming Parallel I/O Complexity with Auto-Tuning](#)

Babak Behzad, Huong Vu Thanh Luu, Joseph Huchette, Surendra Byna, Prabhat Mr., Ruth Aydt, Quincey Koziol, Marc Snir

[Toward Millions of File System IOPS on Low-Cost Commodity Hardware](#)

Da Zheng, Randal Burns, Alexander S. Szalay

Physical Frontiers

[Physics-Based Seismic Hazard Analysis on Petascale Heterogeneous Supercomputers](#)

Yifeng Cui, Efekan Poyraz, Kim B. Olsen, Jun Zhou, Kyle Withers, Scott Callaghan, Jeff Larkin, Clark C. Guest, Dong Ju Choi, Amit Chourasia, Zheqiang Shi, Steven M. Day, Philip J. Maechling, Thomas H. Jordan

[A Scalable Parallel Algorithm for Dynamic Range-Limited N-Tuple Computation in Many-Body Molecular Dynamics Simulation](#)

Manaschai Kunaseth, Rajiv K. Kalia, Aiichiro Nakano, Ken-ichi Nomura, Priya Vashishta

[2HOT: An Improved Parallel Hashed Oct-Tree N-body Algorithm for Cosmological Simulation](#)

Michael S. Warren

Optimizing Data Movement

[SIDR: Structure-Aware Intelligent Data Routing in Hadoop](#)

Joe Buck, Noah Watkins, Greg Levin, Adam Crume, Kleoni Ioannidou, Scott Brandt, Carlos Maltzahn, Neoklis Polyzotis, Aaron Torres

[Using Cross-Layer Adaptations for Dynamic Data Management in Large Scale Coupled Scientific Workflows](#)

Tong Jin, Fan Zhang, Qian Sun, Hoang Bui, Manish Parashar, Hongfeng Yu, Scott Klasky, Norbert Podhorszki, Hasan Abbasi

[Exploring the Future of Out-Of-Core Computing with Compute-Local Non-Volatile Memory](#)

Myoungsoo Jung, Ellis H. Wilson III, Wonil Choi, John Shalf, Hasan Metin Aktulga, Chao Yang, Erik Saule, Umit V. Catalyurek, Mahmut Kandemir

In-Situ Data Analytics and Reduction

[Assessing the Effects of Data Compression in Simulations Using Physically Motivated Metrics](#)

Daniel E. Laney, Steven H. Langer, Christopher R. Weber, Peter G. Lindstrom, Al Wegener


Matrix Computations

[Parallel Reduction to Hessenberg Form with Algorithm-Based Fault Tolerance](#) 

Yulu Jia, George Bosilca, Piotr Luszczek, Jack Dongarra


[A Computationally Efficient Algorithm for the 2D Covariance Method](#) 

Oded Green, Yitzhak Birk


[An Improved Parallel Singular Value Algorithm and Its Implementation for Multicore Hardware](#) 

Azzam Haidar, Jakub Kurzak, Piotr Luszczek

Sorting and Graph Algorithms

[Distributed-Memory Parallel Algorithms for Generating Massive Scale-Free Networks Using Preferential Attachment Model](#) 

Maksudul Alam, Maleq Khan, Madhav Marathe

[On Fast Parallel Detection of Strongly Connected Components \(SCC\) in Small-World Graphs](#) 

Sungpack Hong, Nicole C. Rodia, Kunle Olukotun


[Algorithms for High-Throughput Disk-to-Disk Sorting](#) 

Hari Sundar, Dhairya Malhotra, Karl Schulz

Application Performance Characterization

[An Early Performance Evaluation of Many Integrated Core Based SGI Rackable Computing System](#) 

Subhash Saini, Haoqiang Jin, Dennis Jespersen, Huiyu Feng, Jahed Djomehri, William Arasin, Robert Hood, Piyush Mehrotra, Rupak Biswas

[Predicting Application Performance Using Supervised Learning on Communication Features](#) 

Nikhil Jain, Abhinav Bhatele, Michael Robson, Todd Gamblin, Laxmikant Kale

[Investigating Applications Portability with the Uintah DAG-Based Runtime System on PetaScale Supercomputers](#) 

Qingyu Meng, Alan Humphrey, John Schmidt, Martin Berzins