

## Guest editorial

---

# Special issue: SC13 – The International Conference for High Performance Computing, Networking, Storage and Analysis

**Abstract.** The technical papers program for SC13 received 449 submissions of which 90 were selected for the program giving an acceptance rate of 20%. A rigorous peer review process, including author rebuttals and a 1.5 day face-to-face program committee meeting ensured that selected papers were the very best in our field. One of the tasks at the face-to-face meeting was also to select finalists for the best paper award, from which one is selected by a committee during the conference. To further highlight their achievement of being selected as the very top tier of all the submitted papers to SC13, the authors of these finalist papers were offered the opportunity to publish extended versions of their papers in this special issue journal; all eight authors accepted.

### 1. Introduction

We provide a brief introduction to the eight papers included in this special issue below.

“Enabling fair pricing on high performance computer systems with node sharing”, by Alex D. Breslow, Ananta Tiwari, Martin Schulz, Laura Carrington, Lingjia Tang and Jason Mars describes a runtime system that enables node sharing by applications with fair pricing – that is, charging for resources – by accurately and automatically determining the inter-application interference.

“Enabling highly-scalable remote memory access programming with MPI-3 One Sided”, by Robert Gerstenberger, Maciej Besta and Torsten Hoefler describes a high-performance implementation of the new one-sided (e.g., put, get and accumulate) in MPI-3 that takes advantage of RDMA hardware in modern networks. The implementation is atop Cray’s DMAPP and performance measurements show this MPI-3 implementation to provide comparable or faster performance to the implementations of UPC and CAF provided by Cray.

“Coordinated energy management in heterogeneous processors”, by Indrani Paul, Vignesh Ravi, Srilatha Manne, Manish Arora and Sudhakar Yalamanchili de-

scribes an approach for adjusting CPU and GPU frequencies, taking into account the sensitivity of the performance of the application to frequency changes. Significant improvements in energy use are shown with a small average loss in performance.

“2HOT: An improved parallel hashed oct-tree N-body algorithm for cosmological simulation”, by Michael S. Warren details improvements to an adaptive treecode for N-body problems, with scaling to 256 K processors and results from applications with 69 billion particles.

“Exploring the future of out-of-core computing with compute-local non-volatile memory”, by Myoungsoo Jung, Ellis H. Wilson III, Wonil Choi, John Shalf, Hasan Metin Aktulga, Chao Yang, Erik Saule, Umit V. Catalyurek and Mahmut Kandemir discusses the use of non-volatile memory near the processors for HPC applications that use out-of-core algorithms (or that use large distributed memory systems to supply sufficient memory). By considering both hardware and software issues, they demonstrate near optimal performance in simulation.

“Assessing the effects of data compression in simulations using physically motivated metrics”, by Daniel Laney, Steven Langer, Christopher Weber, Peter Lindstrom and Al Wegener examines the use of lossy com-

pression in physics simulations, using compression techniques developed for floating point data. The evaluation is in terms of application-relevant metrics such as energy and the location of fronts, and is considered for three codes.

“Scalable domain decomposition preconditioners for heterogeneous elliptic problems”, by Pierre Jolivet, Frédéric Hecht, Frédéric Nataf and Christophe Prud’homme describes a scalable implementation of a domain decomposition preconditioner, with particular attention to the coarse grid problem. The approach is based on a rigorous mathematical analysis and makes use of new features in MPI-3. Results from several examples and upto 16 K threads are presented.





“The Science DMZ: A network design pattern for data-intensive science”, by Eli Dart, Lauren Rotman, Brian Tierney, Mary Hester and Jason Zurawski describes an approach to end-to-end network design and

analysis for data-intensive science workflows based on network design patterns. Particular attention is paid to the consequences of packet loss. Case studies from universities, supercomputer centers, and research laboratories are discussed in detail.

We hope you enjoy these papers!

William Gropp  
*SC13 General Chair*  
*University of Illinois at Urbana-Champaign*  
*USA*

Satoshi Matsuoka  
*SC13 Technical Program Chair*  
*Tokyo Institute of Technology*  
*Japan*



# Hindawi

Submit your manuscripts at  
<http://www.hindawi.com>

