

Guest Editorial

Special Issue: Selected Papers from Super Computing 2012

Jeffrey S. Vetter^{a,b,*,‡} and Padma Raghavan^{c,‡}

^a *Georgia Institute of Technology, Atlanta, GA, USA*

^b *Oak Ridge National Laboratory, Oak Ridge, TN, USA*

^c *Pennsylvania State University, University Park, PA, USA*

As the chairs for the SC12 Technical Papers, we were surprised and delighted to have the highest number of technical paper submissions ever to the SC conference series. Authors submitted 472 paper for review, and the 175 member program committee accepted 100 of these papers for presentation during the conference – an acceptance rate of 21%.

From these 100 papers, the committee also nominated exceptional papers as finalists for Best Paper and Best Student Paper, numbering six and nine papers, respectively, with two papers in both categories. The authors of several of these finalist papers were offered the opportunity to publish extended versions of their papers in this special issue; ten authors accepted.

Here, we provide a brief introduction to these ten papers.

“Compiler-directed file layout optimization for hierarchical storage systems” by Wei Ding, Yuanrui Zhang, Mahmut Kandemir and Seung Woo Son. This paper describes a new perspective on optimizing file layouts for hierarchical storage caches using compiler-driven analysis. The fully automated method shows improvements on 16 I/O intensive workloads, increasing performance by 23.7% on average.

“Efficient and reliable network tomography in heterogeneous networks using BitTorrent broadcasts and clustering algorithms” by Kiril Dichev, Fergal Reid and Alexey Lastovetsky. This paper describes a new solution to the problem of managing end-to-end bandwidth in BitTorrent-style data transfers. Their method

counts the fragments exchanged during a broadcast, and then uses this information to reconstruct logical clusters of nodes, and any bottlenecks between the logical clusters.

“A divide and conquer strategy for scaling weather simulations with multiple regions of interest” by Preeti Malakar, Thomas George, Sameer Kumar, Rashmi Mittal, Vijay Natarajan, Yogish Sabharwal, Vaibhav Saxena and Sathish S. Vadhiyar. In this paper, the authors describe their efforts to improve the performance of nested simulations for weather predictions. These efforts use performance prediction in order to help drive processor allocation and topology-aware mapping onto Blue Gene system, with results improving performance up to 33%.

“MPI runtime error detection with MUST: Advances in deadlock detection” by Tobias Hilbrich, Joachim Protze, Martin Schulz, Bronis R. de Supinski and Matthias S. Müller. This manuscript presents improvements to the Marmot Umpire Scalable Tool (MUST) that detects errors in MPI applications. The improvements include higher tool scaling, testing new MPI constructs, and handling more complex real-world applications.

“Characterizing and mitigating work time inflation in task parallel programs” by Stephen L. Olivier, Bronis R. de Supinski, Martin Schulz and Jan F. Prins. This paper describes an examination of several problems causing poor performance in task parallel programs. In an attempt to improve performance, the paper proposes a locality framework for OpenMP, which shows an 3X performance improvement, when compared to Intel OpenMP task scheduler.

*Corresponding author. E-mail: vetter@ornl.gov.

‡Both authors are SC12 Technical Papers Co-chairs.

“Direction-optimizing breadth-first search” by Scott Beamer, Krste Asanovic and David Patterson. In this paper, the authors present a new, hybrid top-down, bottom-up algorithm for performing breadth-first search effectively over low-diameter and scale free graphs.

“McrEngine: A scalable checkpointing system using data-aware aggregation and compression” by Tanzima Zerine Islam, Kathryn Mohror, Saurabh Bagch, Adam Moody, Bronis R. de Supinski and Rudolf Eigenmann. This paper describes a scalable checkpoint-restart system, mcrEngine, that uses data aggregation and compression to reduce overheads by up to 87% and compressibility up to 115%.

“Efficient backprojection-based synthetic aperture radar computation with many-core processors” by Jongsoo Park, Ping Tak Peter Tang, Mikhail Smelyanskiy, Daehyun Kim and Thomas Benson. This paper describes an important image reconstruction method, synthetic aperture radar via backprojection, and the techniques that they used to port it to a Intel Knights Corner manycore platform. The optimizations include

strength reduction, locality enhancements, and exposing parallelism at the thread and instruction level.

“A framework for low-communication 1-D FFT” by Ping Tak Peter Tang, Jongsoo Park, Daehyun Kim and Vladimir Petrov. The authors introduce a new algorithm for a distributed 1-D FFT, where communication costs are reduced. The implementation on large-scale problems can be as twice as fast as leading FFT libraries.

“Containment domains: A scalable, efficient and flexible resilience scheme for exascale systems” by Jinsuk Chung, Ikhwan Lee, Michael Sullivan, Jee Ho Ryoo, Dong Wan Kim, Doe Hyun Yoony, Larry Kaplanz and Mattan Erez. In this paper, the authors describe and evaluate the concept of containment domains, which is a programming construct that allows applications to express resiliency requirements to the underlying programming system. The evaluation of containment domains shows that it is superior to both checkpoint restart and redundant execution techniques.

Enjoy reading these papers!



Hindawi

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

