

Guest Editorial: Computing Frontiers

**Calin Cascaval · Pedro Trancoso ·
Viktor Prasanna**

Published online: 31 January 2013
© Springer Science+Business Media New York 2013

Computing Frontiers is a forum where researchers meet to exchange ideas in all areas of Computer Science. While its roots are in computer systems research, Computing Frontiers has diversified and it is now closer to an interdisciplinary forum. Its main focus is on giving exposure to innovative and visionary work in the broad field of Computer Science and Engineering. The 2011 edition of the conference covered research in a wide range of topics including processor and memory architecture, quantum computing, parallel applications, dynamic binary translation, and fault tolerant design. From all the accepted work, we selected the top four articles for an extended publication in this edition.

In “Emerging Architectures Enable to Boost Massively Parallel Data Mining using Adaptive Sparse Grids,” Alexander Heinecke and Dirk Pflüger present a parallel implementation for the sparse grid algorithms using the vectorization and streaming capabilities of both CPUs and GPUs. The authors target state-of-the-art processors from leading vendors. On a hybrid CPU/GPU system the authors achieve an impressive 188x speedup for single precision operations.

In “Addressing GPU On-chip Shared Memory Bank Conflicts Using Elastic Pipeline,” Chunyang Gou and Georgi N. Gaydadjiev propose optimizations to the GPU pipeline as to avoid the on-chip shared-memory bank conflicts. With this technique the authors are able to improve the performance by up to 21 %.

C. Cascaval (✉)
Qualcomm Research, Santa Clara, CA, USA
e-mail: cascaval@qti.qualcomm.com

P. Trancoso
University of Cyprus, Nicosia, Cyprus

V. Prasanna
University of Southern California, Los Angeles, CA, USA

In “Efficient Stack Distance Computation for a Class of Priority Replacement Policies,” Gianfranco Bilardi, Kattamuri Ekanadham, and Pratap Pattnaik present a comprehensive analysis and a new class of priority replacement policies, the Nearly Static Priorities, which in many cases is close to the OPT replacement policy. The authors evaluate the Min-Tree algorithm that achieves significantly better performance than the traditional Linear Scan algorithm and is amenable to parallelization.

In “Multi-fault Tolerance for Cartesian Data Distributions,” Nawab Ali, Sriram Krishnamoorthy, Mahantesh Halappanavar, and Jeff Daily describe new algorithms that are designed to be tolerant to faults. Such faults are expected to happen more frequently in future large-scale systems. In particular, the authors propose fault-tolerant linear algebra algorithms. These algorithms are tested on a real setup and the results show that faults are tolerated efficiently with little overhead.