

Guest Editorial: SBAC-PAD 2013

**Guido Araujo¹ · Jean-Luc Gaudiot² ·
Manish Parashar³ · Derek Chiou⁴ ·
José Nelson Amaral⁵ · Chita R. Das⁶**

Received: 13 August 2015 / Accepted: 13 August 2015 / Published online: 9 September 2015
© Springer Science+Business Media New York 2015

Welcome to this special issue, a showcase of an extension of some of the most notable papers which were presented at SBAC-PAD 2013 in Porto de Galinhas, Brazil, in October 2013. SBAC-PAD is an international annual conference, started in 1987, which has continuously presented an overview of new developments, applications, and trends in parallel and distributed computing technologies. SBAC-PAD is open to faculty members, researchers, practitioners, and graduate students around the world. Last year, it was promoted by the Brazilian Computer Society, organized in cooperation with the IEEE Computer Society Technical Committee on Computer Architecture,

✉ Guido Araujo
guido@ic.unicamp.br

Jean-Luc Gaudiot
gaudiot@uci.edu

Manish Parashar
parashar@rutgers.edu

Derek Chiou
derek@utexas.edu

José Nelson Amaral
amaral@cs.ualberta.ca

- 1 University of Campinas, Campinas, Brazil
- 2 University of California, Irvine, USA
- 3 Rutgers University, Rutgers, USA
- 4 University of Texas at Austin, Austin, USA
- 5 University of Alberta, Edmonton, Canada
- 6 Pennsylvania State University, State College, USA

sponsored by CRAY, NEC, CNPq, FINEP and CAPES, co-sponsored by IFIP, and organized by UFPE.

As part of the original selection, the SBAC-PAD 2013 conference Program Committee had the notably difficult task of selecting 28 papers out the 108 submissions we received at the time. The rigorous peer review we sought to achieve was indeed an arduous task for we demanded that no fewer than three independent reviews be obtained for each submission (nearly one in ten actually received four reviews). All “difficult” cases (i.e., when a paper received diverging recommendations) were given additional individual attention by the Track Chairs and the Program Committee co-chairs. It should be noted that SBAC-PAD has now attracted submissions from 23 different countries and that 6 of them actually appear in the proceedings. For this special issue, we, the guest editors, in consultation with the SBAC-PAD 2013 Program Committee, selected the following papers, based upon their overall quality and direct relevance to the spirit of the conference. We then invited the authors to present an extended version of their work. Each was then subjected to a strict peer review and sent to each of the original reviewers in addition to one extra expert in the field. We provided the last stage of quality control and are proud to present to you the following eleven papers. These include a variety of topics, ranging from architecture design to applications and performance evaluation:

- The paper, “A Decomposition-Based Approach for Scalable Many-Field Packet Classification on Multi-core Processors,” by Yun Rock Qu, Shijie Zhou and Viktor Kumar Prasanna presents a decomposition-based packet classification mechanism that supports large rule sets and large packet header fields. Experimental results on state-of-the-art 16-core platforms show that, an overall throughput of 48 Million Packets Per Second (MPPS) and a processing latency of 2000 ns per packet can be achieved for a 32K rule set.
- In the paper, “Fully Optimized Code Block Segmentation Algorithm for LTE-Advanced,” by Karlo Gusso Lenzi, Felipe Augusto Pereira Figueiredo, José Arnaldo Bianco Filho and Fabricio Lira Figueiredo propose a new approach for LTE-A code block segmentation. The experimental results, using standard implementations on DSP and FPGA, show speed-ups of 83x when compared to previous approaches.
- In the paper, “Invasive Compute Balancing for Applications with Shared and Hybrid Parallelization,” Martin Schreiber, Christoph Riesinger, Tobias Neckel, Hans-Joachim Bungartz and Alexander Breuer present a core-distribution scheduler which performs computational power migration by distributing cores according to the requirements of parallel program instances. They validate their technique using artificial workloads as well as realistic Tsunami simulations. Their experiments reveal significantly faster overall execution times and higher hardware utilization than alternative approaches.
- In the paper, “PageRank Computation Using a Multiple Implicitly Restarted Arnoldi Method for Modeling Epidemic Spread,” Zifan Liu, Nahid Emad, Soufian Ben Amor and Michel Lamure propose a parallel multiple approach based on implicitly restarted Arnoldi method (MIRAM) for calculating dominant eigenpair of stochastic matrices derived from large real networks. Their algorithm was tested

on a cluster of clusters Grid5000 using an epidemic spread application. Experiments on very large networks such as Twitter and Yahoo (over 1 billion nodes) showed a speedup of $27\times$ when compared to a sequential solver.

- In the paper, “Cluster Cache Monitor: Leveraging the Proximity Data in CMP,” Guohong Li, Olivier Temam, Zhenyu Liu, Sanchuan Guo and Dongsheng Wang investigate the cost of L1 miss latencies caused by longer average distance between nodes in large multi-core architectures. By observing that there is a high probability that the target data of an L1 miss resides in the L1 cache of a neighbor node, they organize cores into clusters and propose a Cluster Cache Monitor (CCM) to detect whether an L1 miss can be served by them. When using a 64-node multi-core and the SPLASH-2 and PARSEC benchmarks, the authors achieved a 15% reduction in execution time and a 14% decrease in energy consumption.
- In the paper, “BPLG: A Tuned Butterfly Processing Library for GPU Architectures,” Jacobo Lobeiras, Margarita Amor and Ramón Doallo propose a library to easy (simplify?) the task of programming GPUs. The parametrization approach adopted by their library considerably simplifies the design of algorithms for a variety of GPU architectures. Implementation of a set of orthogonal signal transforms and an algorithm to solve tridiagonal equations systems showed competitive performance for two NVIDIA GPU architectures.
- The paper, “List Scheduling in Embedded Systems Under Memory Constraints,” by Paul-Antoine Arras, Didier Fuin, Emmanuel Jeannot, Arthur Stoutchinin and Samuel Thibault expands list-scheduling heuristics with static priorities to introduce the notion of memory. Their approach reduces memory footprint when compared to classical heuristics, and shows that through adequate adjustment of task priorities and judicious resort to insertion-based policy, speedups up to 20% can be achieved.
- In the paper, “A Hardware/Software Approach for Database Query Acceleration with FPGAs,” Bharat Sukhwani, Mathew Thoennes, Hong Min, Parijat Dube, Bernard Brezzo, Sameh Asaad and Donna Dillenberger describe an FPGA-based composable accelerator that offloads analytics queries from the host CPU during the execution of online transactions from real-time analytics workloads. Experimental results reveal performance improvement and query-specific accelerator customization without requiring FPGA-reconfiguration.
- In the paper, “An Autotuning Engine for the 3D Fast Wavelet Transform on Clusters with Hybrid CPU + GPU Platforms,” Gregorio Bernabé, Javier Cuenca and Domingo Giménez present an optimization method to run the 3D-Fast Wavelet Transform (3D-FWT) on combined CPU + GPU systems. Their mechanism detects the different computing components of the system, and executes kernels in the appropriate component. Different parallel programming paradigms are combined to fully exploit the performance capacity of the various computational elements of the system, resulting in an important reduction in the compression time of long video sequences.
- In the paper, “The Scalability of Disjoint Data Structures on a New Hardware Transactional Memory System,” Gong Su and Stephen Heisig analyze the scalability of disjoint data structures in hardware transactional memory systems. To achieve that they propose an order preserving hashed structure and a concurrent

open addressing hash table that restricts memory conflicts to actual contention situations, thus enabling the efficient usage of transactional memory techniques. Their experiments show near linear scalability of insertion and deletion operations for up to 96 CPUs in an IBM zEnterprise EC12 server.

- In the paper, “Extending Summation Precision for Network Reduction Operations,” George Michalogiannakis, Xiaoye S. Li, David H. Bailey and John Shalf investigate the accumulation of rounding errors due to compressed representations, an increasing problem when scaling modern HPC systems. To address this problem they propose a fixed-point representation of double precision variables that enable arbitrarily large summations without error and exact and reproducible results.

We would like to thank Alex Nicolau, the Editor-in-Chief of the International Journal on Parallel Programming, for giving us a chance to guest edit this Special Issue, the Program Committee of SBAC-PAD 2013, for their kind and expert work in evaluating these papers, the SBAC-PAD Steering Committee for their expert guidance, and of course the authors for their participation and for making this outstanding issue possible.