Guest Editorial

**Alba Melo · Jean-Luc Gaudiot · Luiz DeRose ·
Kunle Olukotun · Albert Zomaya**

Published online: 18 June 2013

© Springer Science+Business Media New York 2013

Welcome to this special issue, a showcase of some of the most notable papers which were presented at SBAC-PAD 2011 in Vitória, Brazil, in October 2011. 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

A. Melo

Department of Computer Science, University of Brasilia (UnB),
Module 18, Room 3A, Campus UnB - ICC Norte, Sub-solo, Asa Norte,
Brasilia, DF CEP 70910-900, Brazil
e-mail: albam@ic.unb.br

J.-L. Gaudiot (✉)

Department of Electrical Engineering and Computer Science,
University of California, Engineering Hall 4424, Irvine, CA 92697-2625, USA
e-mail: gaudiot@uci.edu
<http://pascal.eng.uci.edu/people/gaudiot.html>

L. DeRose

Department of Research and Development, Cray Inc.,
Cray Plaza, 380 Jackson Street, St. Paul, MN 55101, USA
e-mail: ldr@cray.com

K. Olukotun

Department of Electrical Engineering, Stanford University,
Gates Hall 3A, Room 302, Stanford, CA 94305-9030, USA
e-mail: kunle@stanford.edu

A. Zomaya

School of Information Technologies, The University of Sydney,
Building J12, Sydney, NSW 2006, USA
e-mail: albert.zomaya@sydney.edu.au

Society Technical Committee on Computer Architecture, sponsored by CRAY, NEC, CNPq, FINEP and CAPES, co-sponsored by IFIP, and organized by FUES.

As part of the original selection, the SBAC-PAD 2011 conference Program Committee had the notably difficult task of selecting 26 papers out the 76 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 3 independent reviews be obtained for each submission (54 % actually received 4 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 13 foreign countries and that 7 of them actually appear in the proceedings. For this special issue, we, the guest editors, in consultation with the SBAC-PAD 2011 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 ten papers. These include a variety of topics, ranging from architecture design to applications and performance evaluation:

- The paper by Avilez–Gonzalez, Piernas and Gonzalez–Ferez presents the design and implementation of an enhanced type of Object Storage Device (OSD), called OSD+ device, which supports both data and directory objects, and a metadata cluster based on it. When compared to Lustre, this proposal presents impressive improvements of 60–80 %.
- Sun and Kaeli explore how to leverage aggressive software-based value prediction techniques on a GPU to accelerate programs that lack inherent data parallelism. The experimental results show that, despite the overhead incurred due to software speculation and the communication overhead between the CPU and GPU, this approach is able to obtain up to 6.5x speedup on a selected set of kernels taken from the SPEC CPU2006, PARSEC and Sequoia benchmark suites.
- Bahi and Eisenbeis present the basic ideas of their algorithm for rematerialization in register allocation with reverse computing. To evaluate their algorithm, the authors use the memory demanding LQCD (Lattice Quantum Chromodynamics) application to demonstrate that important gains of up to 33 % on register pressure can be obtained.
- Schneider and Linnert investigate and discuss different data structures in order to fit to planning based resource management with advance reservations. As a result the benefits of using lists of resource allocation are exposed. Simulation results show a better runtime and higher reservation success rate compared with the currently favored approach of a slotted time and the more sophisticated approach based on AVL trees.
- Rosas, Morajko, Jorba and Cesar propose a methodology for improving performance of data-intensive applications based on performing multiple data partitions prior to the execution, and ordering the data chunks according to their processing times during the application execution. Also, the data gathered by a monitoring

tool is used to dynamically tune the parameters of the application. The proposed methodology was initially tested using the bioinformatics tool BLAST, obtaining results with up to a 40 % of performance improvement.

- Tomic, Cristal, Unsal and Valero propose dynamic runtime testing, a methodology for rapid development and accurate error detection in architectural cycle-accurate simulators. The paper shows that dynamic testing can help the developers of architectural simulators to get a reliable and accurate verification of functional correctness, being able to reduce the simulator modification time from 12–18 person-months to 3–4 person-months.
- Borin, Breternitz Jr., Wu and Araújo describe techniques for microcode compression that can achieve significant area and power savings. It also proposes a streamlined architecture that enables high throughput within the constraints of a high performance CPU. The experimental results for microcode compression on several commercial CPU designs were able to obtain compression ratios ranging from 50 to 62 %.
- Arunagiri, Kwok, Teller, Portillo and Seelam propose FAIRIO, a cycle-based I/O scheduling algorithm that provides differentiated service to workloads concurrently accessing a consolidated RAID storage system. The experimental results show that, for a broad range of workload request types, sizes, and access characteristics, the proposed algorithm provides differentiated storage throughput that is within 10 % of being perfectly proportional to workload weights, with little or no degradation of aggregate throughput.
- Waliullah and Stenstrom analyze the performance of hardware transactional memory systems, investigate the source of the performance weaknesses and propose solutions to reduce these weaknesses. To this end, they propose a cache-miss classification model and present two techniques for the removal of data conflicts. Experimental results demonstrate a significant improvement in performance and energy consumption.
- Finally, Ma, Xia and Prasanna investigate data parallelism for belief propagation in acyclic factor graphs on multicore systems. The authors propose a complete belief propagation algorithm that is able to perform exact inference in factor graphs. The experimental results show that the proposed algorithms exhibit good scalability using a representative set of factor graphs.

We would like to thank Alex Nicolau, the Editor-in-Chief of the International Journal on Parallel Programming, for giving us a chance to host this Special Issue, the Program Committee of SBAC-PAD 2011, 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.