
Preface

Claudio L. Amorim

Programa de Engenharia de Sistemas e Computação,
COPPE, Centro de Tecnologia – H318,
Universidade Federal do Rio de Janeiro,
21941-972 Rio de Janeiro, RJ, Brazil
E-mail: amorim@cos.ufrj.br

Biographical notes: Claudio L. Amorim is an Associate Professor in the Systems and Computer Engineering Program at COPPE-Federal University of Rio de Janeiro. His research interests include parallel and distributed systems, and mobile networking. He received his PhD Degree in Computer Science from Imperial College, London/UK. He is a member of the Brazilian-SBC and of the IEEE Computer Societies.

This special issue of the journal contains five selected papers from the *17th Symposium on Computer Architecture and High-Performance Computing (SBAC-PAD)*, which was held in Rio de Janeiro, Brazil, on 24–27 October 2005. The SBAC-PAD series of symposia is sponsored by the Brazilian Computer Society and was started in 1987 to discuss interdisciplinary research focused on high performance that span computer architecture, systems software, parallel and distributed algorithms, and large-scale applications. Currently, SBAC-PAD is a well-recognised international symposium co-sponsored by the IEEE Technical Committees on Computer Architecture and Scalable Computing (TCCA and TCSC).

In 2005, the symposium received 83 submissions from 18 countries in the Americas, Europe, and Asia, of which 29 papers were selected for publication in the proceedings based on at least three reviews. For this journal issue, five best papers were selected and the authors were asked to extend them significantly. These papers discuss recent results on several aspects of high-performance computing, including techniques to reduce microprocessor simulation time, advanced Grid management systems, virtualisation software, and OLAP query processing in database clusters.

Using cycle-accurate microprocessor simulation for evaluating processor design is often prohibitively time-consuming. The first paper by Luo et al. proposes a set of statistical metrics to assess the proposed techniques to reduce simulation time, including clustering algorithms, distance measures, and phase classification features in representative sampling for microprocessor simulation. In addition, the authors introduce a new phase identification scheme based on the reuse distance distribution that outperformed the Basic Block Vector technique for many SPEC CPU2000 benchmark programs.

Grid management systems rely on advance reservation mechanisms, which allocate resources to admitted jobs even before they are actually required. The paper by Burchard et al. examines the impact of failures to the resource management of admitted but not yet started jobs. While

current grid management approaches use estimation of the downtime to remap reservations to other resources, the authors propose a novel dynamic load-based approach that remaps affected jobs accordingly and yet shows superior performance.

In a typical computational Grid, the management of available resources is performed by a central Resource Management System (RMS). However, the design complexity of RMS to efficiently deal with widely different Grid applications has led Nascimento et al. to propose the use of autonomous or autonomic applications. An autonomic application is managed by its own Application Management System (AMS) that utilises the grid according to the resources available and the application's characteristics. The paper by Nascimento et al. assesses the feasibility of an MPI-based implementation of an AMS to adapt efficiently the runtime execution of bag-of-tasks and parameters sweep applications in the context of grid-enabled MPI implementations.

The paper by Pinter et al. investigates the problem of improving performance of virtualisation software to enable the efficient running of many virtual machines. Specifically, they proposed and evaluate a new method for handling and using dynamic memory extension for Linux guests in a virtualisation system so that more virtual machines can utilise a virtualisation platform by dynamically enlarging and shrinking the size of the virtual machines.

On-line Analytical Processing (OLAP) applications require high-performance database support to achieve good response time. Furtado et al. present a new partitioning scheme called Adaptive Hybrid Partitioning (AHP) that combines virtual partitioning with physical partitions to implement intra-query parallelism for OLAP query processing. Experimental results using a 32-node shared-nothing cluster running PostgreSQL DBMS for representative TPC-H queries showed that AHP achieved linear and super-linear speedups with a significant reduction on disk space requirements.