J. Parallel Distrib. Comput. 120 (2018) 195

Contents lists available at ScienceDirect

J. Parallel Distrib. Comput.

journal homepage: www.elsevier.com/locate/jpdc

Editorial

Special issue on Computer Architecture and High Performance Computing

This special issue is focused on Computer Architecture and High Performance Computing. It also includes extended papers presented at SBAC-PAD 2016, 28th International Symposium on Computer Architecture and High Performance Computing, which took place in Los Angeles, USA, from October 26–28, 2016.

All submitted papers to this special issue were rigorously reviewed by at least three expert reviewers, and further carefully evaluated by the guest editors. After the review process, only 9 papers were finally accepted for publication.

Below, we provide an overview of the papers appearing in this volume.

In "Janus: Diagnostics and Reconfiguration of Data Parallel Programs", the authors present the design and implementation of Janus, a tool that automates the reconfiguration of Spark applications. It leverages logs from previous executions as input, enforces configurable adjustment policies over the collected statistics and makes its decisions taking into account communication behaviors specific of the application evaluated, showing gains of up to 1.9x in the scenarios considered.

The work entitled "An Experimental Evaluation of a Parallel Simulated Annealing Approach for the 0–1 Multidimensional Knapsack Problem" focuses on the proposal of a parallel simulated annealing algorithm (SA) using GPGPU. The results achieved by the parallel SA were compared to other reference works and showed that GPGPU is effective on the task of obtaining better quality solutions in reduced execution time when compared to sequential programs.

In "Aspen-Based Performance and Energy Modeling Frameworks", the authors propose and evaluate two energy estimation techniques: ACEE (Algorithmic and Categorical Energy Estimation), which uses a combination of analytical and empirical modeling techniques; and AEEM (Aspen's Embedded Energy Estimation), a system-level analytical energy estimation technique, that incorporate Aspen domain specific language for performance modeling.

In "MR-Advisor: A Comprehensive Tuning, Profiling, and Prediction Tool for MapReduce Execution Frameworks on HPC Clusters", the MR-Advisor tool is proposed and described in detail.

It also presents the MR-Advisor generalization to provide performance optimizations for Hadoop, Spark, and RDMA-enhanced Hadoop MapReduce designs over different file systems such as HDFS, Lustre, and Tachyon.

The paper "A Scalable Algorithm for Simulating the Structural Plasticity of the Brain" proposes a scalable approximation algorithm for Model of Structural Plasticity that presents the complexity of $O(n \log^2 n)$. It also shows an MPI-based parallel implementation of the algorithm that can simulate the structural plasticity of up to 10^9 neurons – four orders of magnitude more than the naive $O(n^2)$ version.

The proposal presented in "DITVA: Dynamic Inter-Thread Vectorization Architecture" leverages the implicit Data Level Parallelism that exists across threads on SPMD applications. The experimental evaluation of the DITVA architecture on the SPMD applications from the PARSEC and Rodinia OpenMP benchmarks shows that it presents a higher performance when compared to a 4- thread 4-issue SMT architecture with AVX instructions, while fetching and issuing 51% fewer instructions, and achieving an overall 24% energy reduction.

In "Exploiting Social Network Graph Characteristics for Efficient BFS on Heterogeneous Chips", three approaches to perform BFS on different heterogeneous chips (a multicore CPU and an integrated GPU) are evaluated. The paper also identifies how to take advantage of the features of social network graphs, that are a particular example of highly connected graphs, as well as the drawbacks of the three proposed heterogeneous approaches. Through exhaustive evaluation it is shown that the proposed implementations can be faster when compared with the best of only one device baselines and with other related heterogenous approach.

The work "Dark-Silicon Aware Design Space Exploration" introduces a less conservative dark silicon estimate based on chip components power density and technological process, and a technique that performs the design space exploration aware of the dark silicon constraints. In addition to the design, implementation, and validation, the proposal is integrated to the Multiexplorer platform design tool so that users can design a computing platform, simulate the performance, analyze physical estimates, and perform design space exploration looking for alternative designs that mitigate the dark silicon.

In "A novel hierarchical architecture for wireless network-on-chip", a combination of several topologies are investigated to develop an efficient hierarchical structure for the architecture of Wireless Network-on-Chip – NoC. The performance of considered hierarchical structures are compared under different traffic patterns. Finally, by using the Analytic Hierarchy Process technique, a new hierarchical wireless NoC is proposed. Based on the results of simulations, the proposed hierarchical structure has better efficiency than a wired mesh topology in NoC.

Finally, we thank all the reviewers for their critical and expert comments and the authors for their high-quality contributions. We hope that this volume will be of timely value to the readers.

Guest Editors
Lúcia Drummond
Edson Borin

https://doi.org/10.1016/j.jpdc.2018.07.016