Melo, A., Carretero, J., Stenstrom, P., Ranka, S. Ayguade, E. (2019). Trends on heterogeneous and innovative hardware and software systems. *Journal of Parallel and Distributed Computing*, 133, pp. 362-364.

DOI: 10.1016/j.jpdc.2019.08.001

Trends on Heterogeneous and Innovative Hardware and Software Systems

Alba Melo
University of Brasilia, Brazil
Jesus Carretero
Universidad Carlos III de Madrid, Spain
Per Stenstrom
Chalmers University of Technology, Sweden
Sanjay Ranka
Eduard Ayguade
Universitat Politecnica de Catalunya, Spain

1 Introduction

To cope with the increasing scale and heterogeneity, in the last years, there has been a clear move towards heterogeneous hardware and software systems. Those systems have been incorporated to hybrid High Performance Computing (HPC) and cloud computing infrastructures, so that combining the benefits of several technologies is possible [6]. These solutions are more and more important due to the increasing synergies between cloud computing and data intensive applications. This convergence requires cyberinfrastructures that must be powerful in a broad sense (computation, storage, I/O capacity, communications, etc). In this scenario, the hardware and software designer is faced with the need of innovative techniques to manage efficiently such complex systems.

The purpose of this special issue is to collect recent trends in heterogeneous and innovative systems, including computer architecture, languages and compilers, algorithms, and applications. The best papers of The International Symposium on Computer Architecture and High Performance Computing 2017 (SBAC-PAD 2017), an annual international conference series that started in 1987, were invited to send extended versions; in addition there was an open call for papers. After the peer review process, 10 papers were accepted for publication.

2 Special issue contents

This special issue contains 10 selected papers that cover a variety of topics in parallel and distributed heterogeneous systems, including heterogeneous programming environments, cloud and publish/subscribe systems, parallel applications, asymmetric multicores and heterogeneous embedded systems.

The paper *Data-flow Analysis and Optimization for Data Coherence in Heterogeneous Architectures* by Sousa, Pereira, Pereira and Araujo [8] proposes

an innovative technique to allocate shared buffers between host and devices in OpenCL, that aims to minimize the number of data coherence operations. The experimental results showed speed-ups of up to 5.25x on an ARM Mali-T880 and up to 8.87x on an NVIDIA GPU Pascal Titan X.

Perez, Stafford, Beivide, Mateo, Turel, Ayguadé and Martorell propose in *Auto-tuned OpenCL kernel co-execution in OmpSs for heterogeneous systems* [7] a novel extension to the OmpSs programming model to allow the co-execution of a single OpenCL kernel in several devices, including the Auto-Tune algorithm that provides adaptive load balancing strategies. Experimental results reveal that the co-execution of single kernels on all the devices in the node is beneficial in terms of performance and energy consumption, and that the proposed scheduling algorithm gives the best overall results.

Paper *LXCloudFT: Towards high availability, fault tolerant Cloud System based Linux Containers* [5] by Louati, Abbes, Cérin and Jemni addresses the issue of replication and contributes with a novel replication model to provide fault-tolerance. This topic is becoming critical with the increasing scale of systems, which makes faults to become a frequent occurrence, making availability a challenge. The proposed LXCloudFT fault tolerant Cloud system is composed of LXCloud-CR, a Checkpoint Restart model, and GC-CR, a garbage collector component that eliminates old snapshots of containers. Although LXCloudFT is designed, originally, for scientic applications, authors also want to adapt it to serve stateless loosely coupled applications such as web applications. Large-scale experiments on Grid'5000 are shown in the paper.

In *VCube-PS: A Causal Broadcast Topic-based Publish/Subscribe System* [3], Araujo, Arantes, Duarte, Rodrigues and Sens present VCube-PS, a topic-based Publish/Subscribe system built on the top of a virtual hypercube-like topology. Membership information and published messages are broadcast to subscribers (members) of a topic group over dynamically built spanning trees rooted at the publisher, showing that, for a given topic, the delivery of published messages respects the causal order. The results presented in the paper confirm the efficiency of VCube-PS in terms of scalability, latency, number and size of messages.

Predicting resource usage in virtualized systems is a challenge currently. The paper *An Intelligent Regressive Ensemble Approach for Predicting Resource Usage in Cloud Computing* [4] by Kaur, Bala and Chana proposes a REAP approach which integrates feature selection and resource usage prediction techniques to achieve high performance. The effectiveness of the proposed approach is evaluated in a real cloud environment by conducting a series of experiments. The experimental results in the paper show that the proposed approach outperforms the existing models by significantly improving the accuracy rate and reducing the execution time. The results are further validted by comparing the existing Learning Automata (LA) based ensemble approach with the proposed approach on the basis of error rate.

Chronaki, Casas, Rico, Valero, Moreto, Badia and Ayguadé perform in *On the Maturity of Parallel Applications for Asymmetric Multi-Core Processors* [2] an extensive evaluation of the portability of HPC applications for asymmetric multicore architectures, with a focus on runtime level scheduler vs OS scheduling solutions. The experimental results in the paper show that scheduling is more effective when it takes place in the runtime supporting the parallel execution model as it improves the baseline by 23%, while the heterogeneous-aware OS

scheduling solution improves the baseline by 10%.

The similarity search in high-dimensional spaces is a core operation found in several online multimedia retrieval applications. With the popularity of these applications, they are required to handle very large and increasing datasets, while keeping the response time low. In [1], Andrade, Fernandes, Gomes, Ferreira and Teodoro propose a *Large-Scale Parallel Similarity Search with Product Quantization for Online Multimedia Services* that address these challenges with an efficient parallelization of the Product Quantization Approximate Nearest Neighbor Search (PQANNS) indexing. This method is capable of answering queries with a reduced memory demand; when coupled with the proposed distributed memory parallelization it can efficiently handle very large datasets. The paper also proposes mechanisms to minimize the query response times in online scenarios in which the query rates vary at run-time. For this sake, the strategies proposed in the paper tune the parallelism configurations and task granularity during the execution. The parallelism and granularity tuning approaches (ADAPT and ADAPT+G) have shown, for instance, to reduce the query response times by a factor of 6.4x in comparison with the best static configuration of parallelism and task granularity. Further, the distributed memory execution using 128 nodes/3584 CPU cores reached a parallel efficiency of 0.97 with a dataset of 256 billion SIFT vectors.

Zhang, Zhang, Snir, Yang and Hao propose in *Automatic Generation of Benchmarks for I/O-Intensive Parallel Applications* [11] a framework which automatically generates benchmarks for I/O intensive parallel applications, with trace merging and compression capabilities. The authors demonstrate the use of their framework on Taub and TianHe-2 supercomputers with four I/O-intensive parallel applications. The results show that the trace merging and compressing algorithms achieve better results than others, and that the generated benchmarks can accurately mimic the behaviors of original I/O-intensive parallel applications.

Zefreh, Lotfi, Lhanli and Karimpour propose in *Topology and Computational-Power Aware Tile Mapping of Perfectly Nested Loops with Dependencies on Distributed Systems* [10] a novel computation power and topology aware strategy for nested loops tile mapping on partially connected heterogeneous distributed systems, aiming to improve load balancing and reduce communication costs. The experimental results show that the proposed method improves the parallel execution time by up to 62% and 28% compared with the computational-power aware tile mapping and the topology aware tile mapping, respectively.

The paper *Minimizing Energy Consumption with Reliability Goal on Heterogeneous Embedded Systems* by Xu, Li, Li and Pan [9] proposes two energy-efficient scheduling algorithms for heterogeneous embedded systems that consider the reliability goal. Experiments with real parallel applications show that the proposed algorithms have significant improvements in energy efficiency compared with the state-of-the-art algorithms.

References

[1] Guilherme Andrade, André Fernandes, Jeremias M. Gomes, Renato Ferreira, and George Teodoro. Large-scale parallel similarity search with prod-

uct quantization for online multimedia services. *J. Parallel Distrib. Comput.*, 125:81–92, 2019.

[2] Kallia Chronaki, Marc Casas, Alejandro Rico, Mateo Valero, Miquel Moreto, Rosa Badia, and Eduard Ayguade. On the maturity of parallel applications for asymmetric multi-core processors. *J. Parallel Distrib. Comput.*, 127:105–115, 2019.

[3] João Paulo de Araujo, Luciana Arantes, Elias Procópio Duarte Jr., Luiz A. Rodrigues, and Pierre Sens. Vcube-ps: A causal broadcast topic-based publish/subscribe system. *J. Parallel Distrib. Comput.*, 125:18–30, 2019.

[4] Gurleen Kaur, Anju Bala, and Inderveer Chana. An intelligent regressive ensemble approach for predicting resource usage in cloud computing. *J. Parallel Distrib. Comput.*, 123:1–12, 2019.

[5] Thouraya Louati, Heithem Abbes, Christophe Cérin, and Mohamed Jemni. Lxcloud-cr: Towards linux containers distributed hash table based checkpoint-restart. *J. Parallel Distrib. Comput.*, 111:187–205, 2018.

[6] Gabriel Mateescu, Wolfgang Gentzsch, and Calvin J. Ribbens. Hybrid computing - where hpc meets grid and cloud computing. *Future Generation Computer Systems*, 27(5):440 – 453, 2011.

[7] Borja Perez, Esteban Stafford, Ramon Beivide, Sergi Mateo, Esteban Stafford, Xavier Turel, Eduard Ayguade, and Xavier Martorell. Autotuned opencl kernel co-execution in ompss for heterogeneous systems. *J. Parallel Distrib. Comput.*, 125:45–57, 2019.

[8] Rafael Sousa, Marcio Pereira, Fernando Pereira, and Guido Araujo. Dataflow analysis and optimization for data coherence in heterogeneous architectures. *J. Parallel Distrib. Comput.*, 130:126–139, 2019.

[9] Hongzhi Xu, Renfa Li, Keqin Li, and Chen Pan. Minimizing energy consumption with reliability goal on heterogeneous embedded systems. *J. Parallel Distrib. Comput.*, 129:14–35, 2019.

[10] Ebrahim Zarei Zefreh, Shariar Lotfi, Leyli Mohamad Lhanli, and Jaber Karimpour. Topology and computational-power aware tile mapping of perfectly nested loops with dependencies on distributed systems. *J. Parallel Distrib. Comput.*, 129:14–35, 2019.

[11] Weizhe Zhang, You Zhang, Marc Snir, Laurence Yang, and Meng Hao. Automatic generation of benchmarks for i/o-intensive parallel applications. *J. Parallel Distrib. Comput.*, 124:1–13, 2019.