



HAL
open science

Foreword to the Special Issue of the Twenty Sixth International Heterogeneity in Computing Workshop (HCW) and to the Fifteenth International Workshop on Algorithms, Models and Tools for Parallel Computing on Heterogeneous Platforms (HeteroPar)

Jorge G. Barbosa, Emmanuel Jeannot

► To cite this version:

Jorge G. Barbosa, Emmanuel Jeannot. Foreword to the Special Issue of the Twenty Sixth International Heterogeneity in Computing Workshop (HCW) and to the Fifteenth International Workshop on Algorithms, Models and Tools for Parallel Computing on Heterogeneous Platforms (HeteroPar). *Concurrency and Computation: Practice and Experience*, 2018, pp.2. 10.1002/cpe.5007. hal-01903118

HAL Id: hal-01903118

<https://inria.hal.science/hal-01903118>

Submitted on 24 Oct 2018

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Foreword to the Special Issue of the Twenty Sixth International Heterogeneity in Computing Workshop (HCW) and to the Fifteenth International Workshop on Algorithms, Models and Tools for Parallel Computing on Heterogeneous Platforms (HeteroPar)

Jorge G. Barbosa¹ Emmanuel Jeannot² and

¹LIACC, Faculdade de Engenharia, Universidade do Porto, Portugal

²Inria, LaBRI, Univ. of Bordeaux, CNRS, Talence, France

Heterogeneity is emerging as one of the most profound and challenging characteristics of today's parallel environments. As most modern computing systems are heterogeneous, from the macro level, where networks of distributed computers, composed by diverse node architectures, are interconnected with potentially heterogeneous networks, to the micro level, where deeper memory hierarchies and various accelerator architectures are increasingly common, the impact of heterogeneity on all computing tasks is increasing rapidly.

Traditional parallel algorithms, programming environments and tools, designed for legacy homogeneous multiprocessors, will at best achieve a small fraction of the efficiency and the potential performance that we should expect from parallel computing in tomorrow's highly diversified and mixed environments. New ideas, innovative algorithms, and specialized programming environments and tools are needed to efficiently use these new and multifarious parallel architectures.

This special issue has the propose to collect extended version of contributions submitted to the Twenty Sixth International Heterogeneity in Computing Workshop (HCW) and to the Fifteenth International Workshop on Algorithms, Models and Tools for Parallel Computing on Heterogeneous Platforms (HeteroPar). After a thorough peer-reviewing process, three papers were selected for publication. The topics addressed in this issue are the performance of cloud function providers, a cross-architecture Kalman filter to accelerate the reconstruction computation of particles collision and, a methodology to model applications jointly with a fast and greedy algorithm to obtain realist mapping and scheduling solution on heterogeneous systems.

Cloud functions are becoming an increasingly popular method of running distributed applications, as they allow the developer for deploying their code in the form of a function to the cloud, which is then responsible for automatic resource provision and scaling. Kamil Figiela, Adam Gajek, Adam Zima, Beata Obrok and Maciej Malawski [Figiela, 2018] present a performance evaluation study of heterogeneous cloud functions where the major cloud function providers are evaluated, namely, AWS Lambda, Azure Functions, Google Cloud Functions and IBM Cloud Functions.

At the LHCb detector in the Large Hadron Collider, the reconstruction of particle collisions in high-energy physics detectors happens at an average rate of 30 million times per second, being the Kalman filter a fundamental element in this process. Due to iterative enhancements in the detector's technology, together with the projected removal of the hardware filter, the rate of particles that will need to

be processed in software in real-time is expected to increase in the coming years by a factor of 40. To tackle such a huge amount of data, Daniel Pérez and Omar Awile [Pérez, 2018], present a cross-architecture Kalman filter optimized for low-rank problems and modern SIMD architectures. Their implementation increases the throughput of the Kalman filter sequences, achieving a performance gain on existing hardware.

Nowadays, numerous services need the execution of complex applications, such as weather forecasting, search engines, big data medical analyzes or intelligent transportation systems. A large amount of energy is used to power them, and it is of primary importance to compute more efficiently to sustain the increasing demand of computing power while keeping energy consumption reasonable. As these applications are heterogeneous in terms of the computational paradigm, heterogeneous resources are seen as more adapted to efficiently execute parts of applications. Nevertheless, the exploitation of these platforms raises new challenges in terms of application management optimization. To tackle this problem, Lilia Zaourar, Massinissa Ait Aba, David Briand and Jean-Marc Philippe [Zaourar, 2018], propose a detailed modeling of the applications and the underlying hardware to be able to find realistic solutions. Two implementations are provided with state-of-the-art tools in order to compare the mapping and scheduling solutions to the fast greedy online resolution algorithm proposed in the paper.

These papers hence represent top of the shelf research results and we hope that you will find them informative, interesting and inspiring.

REFERENCES

[Figiela, 2018] Kamil Figiela, Adam Gajek, Adam Zima, Beata Obrok and Maciej Malawski, *Performance Evaluation of Heterogeneous Cloud Functions*, Concurrency and Computation: Practice and Experience [this issue].

[Pérez, 2018] Daniel Pérez and Omar Awile, *An Efficient Low-Rank Kalman Filter for Modern SIMD Architectures*, Concurrency and Computation: Practice and Experience [this issue].

[Zaourar, 2018] Lilia Zaourar, Massinissa Ait Aba, David Briand and Jean-Marc Philippe, *Task management on fully heterogeneous micro-server system: modeling and resolution strategies*, Concurrency and Computation: Practice and Experience [this issue].