

Guest Editorial for Hybrid Parallelism in New HPC Systems

**Giuliano Laccetti¹ · Ian Foster² ·
Marco Lapegna¹ · Paul Messina³ ·
Raffaele Montella⁴ · Almerico Murli⁵**

Published online: 5 October 2016
© Springer Science+Business Media New York 2016

It is a great pleasure to introduce this special issue of the International Journal of Parallel Programming on Hybrid Parallelism in new HPC Systems. The idea of this publication is born during the “organization” of the Third Workshop on Models, Algorithms and Methodologies for Hybrid Parallelism in New HPC Systems MAMHYP-15, held in Krakow (Poland) in September 2015, jointly with the 11th International Conference on Parallel Processing and Applied Mathematics (PPAM 2015).

Since 2011, the MAMHYP workshop has become a biennial meeting for scientists and engineers to share their ideas and experiences about the combination of all forms of

✉ Giuliano Laccetti
giuliano.laccetti@unina.it

Ian Foster
foster@cs.uchicago.edu

Marco Lapegna
marco.lapegna@unina.it

Paul Messina
messina@alcf.anl.gov

Raffaele Montella
raffaele.montella@uniparthenope.it

Almerico Murli
almerico.murli@unina.it

- ¹ University of Naples Federico II, Naples, Italy
- ² University of Chicago, Chicago, USA
- ³ Argonne National Laboratory, Argonne, USA
- ⁴ University of Naples Parthenope, Naples, Italy
- ⁵ SPACI, Naples, Italy

parallelism in the emerging HPC systems. The MAMHYP-15 workshop continues this tradition and extends its area of interest to the new topics of the Internet of Things and the High Performance Cloud Computing. This special issue then collects an extension of the most valuable works presented at MAMHYP-15 as well as new original papers, with the purpose of gathering the current state of knowledge in the field.

The relevance of these topics is confirmed by the observation that 94 systems out of 500 in the 47th edition of the TOP500 list (June 2016) need the integration of at least three forms of parallelism: among the nodes in the system, among the cores in a single node and among the computing units in a single accelerator device. Furthermore, the advent of new computing models, like the Internet of Things and the High Performance Cloud Computing paradigms, need the management of new and more exotic forms of parallelism related to the management of highly heterogeneous computing devices connected by large networks. For all these reasons the development of efficient algorithms and applications for hybrid HPC systems is still a challenging problem.

This special issue contains 11 articles selected, after two review rounds, from 25 papers submitted from 13 different countries, with an acceptance ratio of less than 50%. In the first review round, each submission received two reviews from external anonymous referees. To the doubtful cases (i.e. borderline papers or with diverging recommendations) additional individual attention by the guest editors of the special issue was given. The authors of papers selected in the first round were then asked to revise their papers based on the reviewers feedback and submit a revised manuscript for a second review round.

The guest editors regret that they could not include more of the many high quality papers submitted to the special issue.

The papers in this special issue cover a broad range of topics, so that they could be classified in three large classes:

- Methods and tools for hybrid HPC systems with GPU technologies or other computing accelerators;
- Environments and technologies for the new paradigms of the Internet of Things and High Performance Cloud Computing including virtualization;
- Solution of real world applications with new hybrid HPC systems.

We believe that all of these papers not only provide novel ideas and state-of-the-art techniques, but will also stimulate future research and technology improvement in the field.

Methodologies and tools for hybrid HPC systems with GPU technologies or other accelerators

In [1] the author presents analytical models based on scaled power metrics to analyze the impact of various architectural design choices, including hybrid systems, on scaled performance and power savings. The power consumption implications of different processing schemes and various chip configurations were also analyzed.

The main topic of paper [2] is the software package General, Hybrid and Optimized Sparse Toolkit (GHOST). This library can be classified as an approach towards a

highly scalable, and truly heterogeneous sparse linear algebra toolkit with a key target in the development process being optimal performance on all parts of heterogeneous systems.

The paper [3] is aimed to an Intel Xeon Phi implementation of the explicit fourth-order Runge-Kutta method (RK4) for very sparse matrices with very short rows, arising from a Markovian modelling of computer and telecommunication networks. The performance evaluation of an implementation based on Intel Math Kernel Library (Intel MKL) routines and authors' own implementation, were investigated.

The paper [4] deals with parallelization of computing similarity measures between large vectors. Rather than focusing on optimization of the algorithm itself, assuming specific measures, the paper assumes a general scheme for finding similarity measures for all pairs of vectors and investigates optimizations for scalability in a hybrid multicore CPU + Xeon Phi system.

In [5] the authors focus their attention on computing acceleration of algorithms, for content authentication used for digital images and video sequences, on hybrid multiprocessor environments. Aim of this work was to explore and take advantage of mixed CPUs and GPUs processing approach and to investigate possibilities to develop optimal authentication algorithm.

Environments and technologies for the new paradigms of the Internet of Things and High Performance Cloud Computing including virtualization

In [6] the authors propose a Data-Locality Aware Workflow Scheduling (D-LAWS) technique and a locality-aware resource management method for data-intensive scientific workflows on HPC cloud environments. D-LAWS applies data-locality and data transfer time based on network bandwidth to scientific workflow task scheduling and balances resource utilization and parallelism of tasks at the node-level.

The paper [7] shows the latest achievements and developments of Generic Virtualization Service (GVirtuS). Thanks to the new and improved remoting capabilities, GVirtuS now enables GPU sharing among physical and virtual machines based on x86 and ARM CPUs on local workstations, computing clusters and distributed cloud appliances.

The paper [8] introduces an Adaptive Scheduling Controller (ASC) for highly heterogeneous distributed environments, a complex system that, using an adaptive approach to the scheduling problem, is able to verify periodically if the job work flow has changed and if the actual system configuration (in terms of scheduler parameters) is good to maintain the same level of performances measured by means of classical efficiency and effectiveness metrics.

The paper [9] introduces a Hadoop Based Parallel Binary Bat Algorithm (HPBBA) method for efficient feature selection and classification to obtain optimized intrusion detection rate, by means of parallel computing model and nature inspired feature selection technique.

Solution of real world applications with new hybrid HPC systems

In [10] the authors deals with a Data Assimilation problem, applied to a oceanographic models. More precisely they introduce a decomposition of the Tikhonov Regularization (TR) functional which split this operator into several independent functionals,

that are then solved on a distributed memory multiprocessor (MIMD) with Graphic Processing Unit (GPU) accelerators.

Finally in [11] the authors investigate heterogeneous computing, as a “tool” for solving a computational cardiology problem. The simulation 3D tissue of the human cardiac ventricle is studied with a physiologically realistic cell model that has 10,000 calcium release units per cell and 100 ryanodine receptors per release unit.

All of the above papers address original research in the broad field of the hybrid computing on new HPC systems and environments. We collected them with the aim to prepare a volume that could serve as resource for education, information, and reference to professors, researchers, graduate students and the more general HPC community. The guest editors would like to express their gratitude to Prof. Alex Nicolau, the Editor-in-Chief of the International Journal on Parallel Programming, for inviting us to edit this Special Issue as Guest Editors, and to the his whole staff (academic and Springer). In addition, we are deeply indebted to the several anonymous referees for their dedication and expert work in reviewing these papers. Last but not least, we are grateful to all authors for their contributions: without them, the special issue, simply, would not have been produced! We hope that you will enjoy this special issue as much as we enjoyed preparing it.

The Guest Editors Committee:

Giuliano Laccetti

Ian Foster

Marco Lapegna

Paul Messina

Raffaale Montella

Almerico Murli

References

1. Marowka, A.: Energy-aware modeling of scaled heterogeneous systems. *Int. J. Parallel Progr.* (2017) (**this issue**)
2. Kreuzer, M., Thies, J., Röhrig-Zöllner, M., Pieper, A., Shahzad, F., Galgon, M., Basermann, A., Fehske, H., Hager, G., Wellein, G.: GHOST: building blocks for high performance sparse linear algebra on heterogeneous systems. *Int. J. Parallel Progr.* (2017) (**this issue**)
3. Bylina, B., Potiopa, J.: Explicit fourth-order Runge-Kutta method on Intel Xeon Phi coprocessor. *Int. J. Parallel Progr.* (2017) (**this issue**)
4. Czarnul, P.: Benchmarking performance of a hybrid Xeon/Xeon Phi system for parallel computation of similarity measures between large vectors. *Int. J. Parallel Progr.* (2017) (**this issue**)
5. Pietron, M., Glowacz, A.: Implementation of digital watermarking algorithms in parallel heterogeneous cluster environment. *Int. J. Parallel Progr.* (2017) (**this issue**)
6. Choi, J., Adufu, T., Kim, Y.: Data-locality aware scientific workflow scheduling methods in HPC cloud environments. *Int. J. Parallel Progr.* (2017) (**this issue**)
7. Montella, R., Giunta, G., Laccetti, G., Lapegna, M., Palmieri, C., Ferraro, C., Pelliccia, V., Hong, C., Spence, I., Nikolopoulos, D.S.: On the virtualization of CUDA based GPU remoting on ARM and X86 machines in the GVirtuS framework. *Int. J. Parallel Progr.* (2017) (**this issue**)
8. Carracciolo, L., Barone, G.B., Boccia, V., Bottalico, D., Campagna, R., Laccetti, G., Lapegna, M.: An approach to forecast queue time in adaptive scheduling: how to mediate system efficiency and users satisfaction. *Int. J. Parallel Progr.* (2017) (**this issue**)
9. Natesan Palanisamy, P., Rajalaxmi, R.R., Gowrison, G., Balasubramanie, P.: Hadoop based parallel binary bat algorithm for network intrusion detection. *Int. J. Parallel Progr.* (2017) (**this issue**)

10. D'amore, L., Arcucci, R., Carracciolo, L., Scotti, G., Laccetti, G.: A decomposition of the tikhonov regularization functional oriented to exploit hybrid multilevel parallelism. *Int. J. Parallel Progr.* (2017) (**this issue**)
11. Johannes, L., Lan, Q., Gaur, N., Cai, X.: Accelerating detailed tissue-scale 3D cardiac simulations using heterogeneous CPU-Xeon Phi computing. *Int. J. Parallel Progr.* (2017) (**this issue**)