Received: 20 November 2023 | Accepted: 13 December 2023

DOI: 10.1002/cpe.8003

EDITORIAL

WILEY

Special Issue on the pervasive nature of HPC (PN-HPC)

1 | INTRODUCTION

It is with great pleasure that we introduce this special issue of the journal *Concurrency and Computation: Practice and Experiences*, focusing on "The Pervasive Nature of High-Performance Computing (PN-HPC)." The idea for this collection of high-quality articles originated during the sixth Workshop on Models, Algorithms, and Methodologies for Hybrid Parallelism in New HPC Systems (MAMHYP-22), which took place in Gdansk, Poland, in September 2022, in conjunction with the 14th International Conference on Parallel Processing and Applied Mathematics (PPAM 2022).

Since its inception in 2011, the MAMHYP workshop has evolved into a biennial gathering for scientists and engineers to exchange their insights and experiences concerning the integration of various forms of parallelism in emerging HPC systems. MAMHYP-22 continues this tradition and expands its scope to encompass new areas of interest, including the Internet of Things and High-Performance Cloud Computing. This special issue serves as an extension of the most valuable contributions presented at MAMHYP-22, along with new original research papers, with the goal of consolidating the current state of knowledge in this field.

After several years, we find ourselves in an exciting era for HPC. When we launched this Special issue, the Top500 list had just crowned the first supercomputer capable of reaching 1 Exaflop, and the European Commission had recently funded a Exascale-Class Supercomputing Ecosystem in Europe through the European High-Performance Computing Joint Undertaking program. Even the Turing Award awarded to Jack Dongarra in 2022 by the Association for Computing Machinery (ACM) should be seen as the definitive consecration of the central role of HPC for the advancement of society.

However, HPC is no longer confined to large-scale infrastructures. Today, it is recognized as a pervasive technology that spans the entire Computing Continuum spectrum. Over the past decade, we have witnessed the emergence of Grid and Cloud Computing environments, implemented through advanced middleware acting as operating systems, orchestrating the efficient resource management. More recently, the Internet of Things and Edge Computing environments aim to make low-power and high-performance resources available in a transparent and user-friendly manner, found ubiquitously.

These diverse environments are characterized by their differences and require sophisticated programming models to achieve high performance, with an increasing emphasis on energy efficiency. This special issue focuses on architectures, methodologies, and algorithms for high-performance computing, paying particular attention to all forms of parallelism and their integration in emerging environments. Its aim is to provide a platform for discussion on these topics and to aggregate the current state of knowledge in the field.

The selected papers cover a wide range of research areas in High-Performance Computing in all its facets, showcasing the challenges inherent in modern platforms and, hopefully, offering insights to interested readers about potential directions for further research in this field.

2 | CONTENTS OF THE SPECIAL ISSUE

This special issue comprises seven articles, meticulously selected after undergoing at least two rounds of rigorous peer review, from a pool of 11 submissions originating from eight different countries, resulting in an acceptance ratio of approximately 63%. During the first review round, each submission received evaluations from two external, anonymous referees. Cases with uncertainty (i.e., papers at the borderline or with conflicting recommendations) received additional individual scrutiny from the guest editors of the special issue. Authors of papers selected in the first round were then invited to revise their manuscripts based on reviewer feedback and resubmit them for a second review round. The guest editors express their regret that they were unable to include more of the many high-quality papers submitted to this special issue.

Specifically, the articles accepted for publication in this special issue, along with their unique contributions, are as follows:

- In Reference 1, the author analyzes the principles behind two metrics for performance portability, comparing their strengths and weaknesses.
- In Reference 2, the authors propose a hybrid sparse matrix layout that combines the flexibility of some well-known sparse formats to offer a number of appealing properties.
- In Reference 3, the authors propose a parallel variation of the Needleman–Wunsch algorithm that enables scalable global sequence alignment with customizable scoring schemes.

- In Reference 4, the authors introduces a technique to load and retrieve the sliced multidimensional dataset on different cloud services such as Amazon Web Service (AWS), Google Cloud Platform, and Microsoft Azure.
- In Reference 5, the author proposes a framework that allows programming a parallel application for a multi-node system, with one or more graphical processing units (GPUs) per node, using an OpenMP+extended CUDA API.
- In Reference 6 the authors introduce a new benchmark for HPC heterogeneous resources inspired to the Linpack benchmark.
- In Reference 7 the authors explore new techniques to enhance the scalability of large codes for data assimilation models.

3 | CONCLUSIONS

All of the papers collected in this Special Issue address original investigations about the several forms of High-Performance Computing, ranging from high-end systems to small devices for the edge computing environments. We believe that they thoroughly represent the state-of-the-art in the field so that they can provide fresh ideas to scientists and practitioners for their studies. We put them together to prepare a volume that can be used in several contexts, from teaching to research and, more generally, as a reference for graduate students, researchers, and experts belonging to the HPC community.

On behalf of every reader of this Special Issue, the guest editors would like to thank all the authors who submitted their papers and worked hard to respond to reviewers' requests in due time, all the anonymous reviewers who participated in the review process providing helpful suggestions, as well as the Editor in Chief and the entire staff of Wiley's Concurrency and Computation: Practice and Experience who oversaw the whole process.

Marco Lapegna[1]
Valeria Mele[1]
Raffaele Montella[2]
Lukasz Szustak[3]

[1]*Department of Mathematics and Applications Renato Caccioppoli, University of Naples Federico II, Napoli, Italy*
[2]*Department of Sciences and Technologies, University of Naples Parthenope, Napoli, Italy*
[3]*Department of Computer Science, Czestochowa University of Technology, Czestochowa, Poland*

Correspondence
Marco Lapegna, Department of Mathematics and Applications Renato Caccioppoli, University of Naples Federico II, Napoli, Italy.
Email: marco.lapegna@unina.it

ORCID
Marco Lapegna https://orcid.org/0000-0001-9953-1319
Valeria Mele https://orcid.org/0000-0002-2643-3483
Lukasz Szustak https://orcid.org/0000-0001-7429-6981

REFERENCES

1. Marowka A. A comparison of two performance portability metrics. *Concurrency and Computation: Practice and Experiences*. Vol 35. Art; 2023:e7868. doi:10.1002/cpe.7868
2. Aliaga J'I, Anzt H, Quintana-Ortí ES, Tomás AE. Sparse matrix-vector and matrix-multivector products for the truncated SVD on graphics processors. *Concurr Comput Pract Exp*. 2023;35(28):e7871. doi:10.1002/cpe.7871
3. Sadiq MU, Yousaf MM. A scalable parallel algorithm for global sequence alignment with customizable scoring scheme. *Concurr Comput Pract Exp*. 2023;35:e7888. doi:10.1002/cpe.7888
4. Mellone G, De Vita CG, Sánchez-Gallegos DD, Laccetti G. A novel approach for large-scale environmental data partitioning on cloud, edge, and on-premises storage. *Concurrency and Computation*. 2023;35:e7893. doi:10.1002/cpe.7893
5. Czarnul P. A multithreaded CUDA and OpenMP based power-aware programming framework for multi-node GPU systems. *Concurr Comput Pract Exp*. 2023;35:e7897. doi:10.1002/cpe.7897
6. Carracciuolo L, Mele V, Sabella G. Toward a new Linpack-like benchmark for computing heterogeneous resources. *Concurr Comput Pract Exp*. 2023;35:e7962. doi:10.1002/cpe.7962
7. Cacciapuoti R, D'Amore L. Scalability analysis of space-time domain decomposition approaches for solving data assimilation models. *Concurr Comput Pract Exp*. 2023;35:e7937. doi:10.1002/cpe.7937