Received: 23 December 2017 | Accepted: 23 December 2017
DOI: 10.1002/cpe.4427

WILEY

EDITORIAL

Special issue on high performance computing in modeling and simulation

In the pure technological era we are living, the need for appropriate tools, methods, and approaches that could boost and skyrocket real world various applications is of paramount importance even for daily life. Toward this direction, in the up-to-date literature, several computational tools are offered, new advanced nearly real-time performing techniques are introduced, almost every day, and powerful computing approaches are promising to tackle the issues of performance, energy efficiency, and computational burden, with many different fruitful ways. Nevertheless, most of these demands, trends, and perspectives would have never met the expected outcome without the help of modern high performance computing systems able to model and simulate computationally intensive scientific applications in the most efficient and appropriate way. Consequently, numerous and various high performance computing approaches like multi-/manycore systems, accelerators, compute clusters, and massively parallel machines, when combined with efficient numerical methods for differential equation systems and native computational paradigms, enable scientists and researchers worldwide to significantly advance the application of computing methodologies in research and industry applications, both in qualitative but mainly in quantitative way.

In this aspect, this Special Issue aimed to offer both scientists and engineers in academy and industry an opportunity to express and discuss their views on current trends, challenges, and state-of-the art solutions to various problems in High Performance Computing for Modeling and Simulation. Moreover, it was highly related to the corresponding Special Session on High Performance Computing in Modeling and Simulation (HPCMS), within the 23rd Euromicro International Conference on Parallel, Distributed and network-based Processing (PDP), held in Turku, Finland on March 4-6, 2015, and its relevant topics. Eventually, a major part of these topics is covered by the content of the fore-coming Special Issue of *Concurrency and Computation: Practice and Experience* through eight (8) finally selected papers, all thoroughly reviewed and revised properly as a detailed major extension of their conference papers earlier published in the PDP 2015 proceedings.

More specifically, in this Special Issue, both theoretical aspects of high performance computing systems, like libraries for the reduction of the programming burden of numerical models on heterogeneous parallel architectures, hybrid programming model MPI/OpenMP for tackling the communication load imbalance issues, and applications starting with parallel shared-memory version of the Space Saving algorithm for mining items, approximate and semi-asynchronous parallel model for supporting Parallel and Discrete Event Simulation, parallel execution pipeline of an existing description algorithm capable of characterizing both color and texture information of a given feature point for robotic visual place recognition, and parallel and hardware acceleration of detection of ambiguous objects for surveillance reasons, as well as optimization techniques to be parallelized such as Imperialist Competitive Algorithm, are fully considered in a fruitful and plausible way.

In more details, the article by Chakroun et al[1] described ExaShark,[2] an open source library with the aim of reducing the programming burden of numerical models on heterogeneous parallel architectures. The presented library offers a global-array-like interface, whereas its run-time can be configured to use shared memory threading techniques, inter-node distribution techniques, or combinations of both. ExaShark takes advantage of the latest HPC technologies, helping to scale to future generation systems. The article demonstrates the usefulness of the ExaShark library through several experiments, including stencil codes, solvers, and matrix factorization algorithms.

Utrera and co-authors[3] analyzed a significant problem in the field of HPC applied to modeling and simulation, that is, the communication load imbalance generated by irregular-data applications running in a multi-node cluster. The study targets, in particular, a hybrid programming model MPI+OpenMP, where several approaches to diminish communication load imbalance are adopted, like computation-communication overlap, issuing communications in parallel, and a new approach based on message fragmentation in order to take advantage of the eager-protocol. The article includes a number of interesting results of experiments, in which the performance of overlapped and non-overlapped approaches are quantified, including the impact due to network latency.

The article by Majd et al[4] concerned another relevant aspect often involved in modeling and simulation, that is, optimization. In particular, the authors focus their work on parallelizing a relatively new evolutionary optimization approach, namely, the Imperialist Competitive Algorithm.[5] The proposed parallelizations include a master-slave version and a more sophisticated multi-population strategy, both exploiting the well-known Message Passing Interface. The article describes a variety of experiments and comparisons, based on two different computing platforms, and proved that the developed parallel algorithms can achieve significant performances in both optimization and speed of execution.

Rousset et al[6] presented nested graphs as an approach to model parallel and distributed multi-agent simulations aiming at facilitating the dynamic distribution of computations among parallel machines. The aforementioned task is successfully achieved due to finer granularity on multiple levels of abstraction. In the proposed approach, a common and generic framework, which represents the agent models, as well as their distribution, is efficiently presented. In addition, the proposed PDMAS framework includes a more graphical method to model parallel and distributed multi-agent simulations in a rather elucidated manner. Finally, in order to demonstrate and validate the advantages of using Nested Graph data structures, several models have been implemented and their performance assessed on HPC resources.

In their article, Cafaro and co-authors[7] proposed a parallel shared-memory version of the Space Saving algorithm[8] for mining frequent items. The proposed approach is investigated with regard to accuracy and performance on many and multi-core processors, including the Intel Phi accelerator. Moreover, the article investigates a hybrid MPI/OpenMP version and proves that it significantly enhances the performance of the earlier pure MPI version of the same algorithm. Among the other interesting results, the authors found that the Intel Phi accelerator is not suitable for the proposed algorithm due to both the highly limited data locality and the non-contiguous memory access that limit the exploitation of the cache hierarchy.

Inostrosa-Psijas and co-authors[9] proposed an approximate and semi-asynchronous parallel strategy for PDEVS, an extension of DEVS[10] for supporting Parallel and Discrete Event Simulation. In particular, the proposed strategy aims at simulating Web search engines operating on large clusters of processors. In brief, the authors used a dynamic algorithm to determine the elapsed time of two consecutive window barriers, in order to reduce the events simulated in non-chronological order (straggler events). This allows improving the trade-off between the running time of the simulation and the accuracy of obtained results. In the proposed approach, a specifically devised semi-asynchronous algorithm aims at reducing the number of straggler events and the number of global synchronization barriers. In addition, a load balance algorithm automatically allocates among processors the logical processes composing the discrete event simulation. Overall, the described strategy proved able to effectively reduce both execution times and memory usage, at the cost of acceptable errors.

The article proposed by Bampis et al[11] studied a significant problem related to Visual Place Recognition, where a robot has to recognize places that it has encountered in the past, using only visual sensors. In this context, the proposed article addresses the so-called loop closure detection task, which is particularly relevant to identify revisited regions of a robot's trajectory. More, in detail, the authors proposed a parallel approach based on an adaptation of an existing description algorithm capable of characterizing both color and texture information of a given feature point, ie, the “Color and Edge Directivity Descriptor.”[12] To enable real-time applications, the authors used an effective parallel execution pipeline based on GPGPU computing. According to the presented experiments, the proposed approach produces fast and highly accurate results and proved capable of running on the Google's project Tango device in real time.

Amanatiadis et al[13] addressed an issue of great importance for surveillance systems based on Unmanned Aerial Vehicles, that is, the detection of ambiguous objects. The proposed approach is based on a moment-based blob detection method exploiting the thermal footprint obtained from single infrared images, and it is able to distinguish human or fire sized and shaped figures. The authors investigate the implementation using hardware acceleration devices, such as GPGPUs and FPGAs, in order to achieve real-time performances and energy efficiency. Using the OSU thermal image data set,[14] the proposed parallel approaches proved effective and suitable for real life missions.

Overall, we think that the contributions summarized above provide interesting and useful insight into various aspects of HPCMS research. We would like to thank the authors for their interesting contributions and all the reviewers for helping us to assemble such a high-quality selection of articles. Moreover, we would also like to thank the editorial office and especially the Journal's editor in chief Professor Geoffrey C. Fox for their valuable support in realizing this Special Issue.

Giuseppe A. Trunfio[1]
William Spataro[2]
Georgios Ch. Sirakoulis [3]
[1]*Department of Architecture, Planning and Design, University of Sassari, Alghero, Italy*
[2]*Department of Mathematics and Computer Science, University of Calabria, Rende, Italy*
[3]*Department of Electrical and Computer Engineering, Democritus University of Thrace, Xanthi, Greece*

Correspondence
Giuseppe A. Trunfio, Department of Architecture, Planning and Design, University of Sassari, Alghero, Italy.
Email: trunfio@uniss.it

REFERENCES

1. Chakroun I, Vander AaT, De Fraine B, Haber T, Wuyts R. A high level library for multi-dimensional arrays programming in computational science. *Concurrency Computat Pract Exper*. 2018;e4376. https://doi.org/10.1002/cpe.4376
2. Scalable hybrid array kit. https://github.com/ExaScience/shark
3. Utrera G, Gil M, Martorell X. Analyzing the impact of communication imbalance in high-speed networks. *Concurrency Computat Pract Exper*. 2017;e4394. https://doi.org/10.1002/cpe.4394
4. Majd A, Sahebi G, Daneshtalab M, Plosila J, Lotfi S, Tenhunen H. Parallel imperialist competitive algorithms. *Concurrency Computat Pract Exper*. 2018;e4393. https://doi.org/10.1002/cpe.4393
5. Atashpaz-Gargari E, Lucas C. Imperialist competitive algorithm: an algorithm for optimization inspired by imperialistic competition. Paper presented at: IEEE Congress on Evolutionary Computation; 2007; Singapore.
6. Rousset A, Herrmann B, Lang C, Philippe L, Bride H. Nested graphs: a model to efficiently distribute multi-agent systems on HPC clusters. *Concurrency Computat Pract Exper*. 2018;e4407. https://doi.org/10.1002/cpe.4407

7. Cafaro M, Pulimeno M, Epicoco I, Aloisio G. Parallel space saving on multi and many-core processors. *Concurrency Computat Pract Exper*. 2017;e4160. https://doi.org/10.1002/cpe.4160
8. Metwally A, Agrawal D, Abbadi AE. An integrated efficient solution for computing frequent and top-k elements in data streams. *ACM Trans Database Syst*. 2006; 31(3):1095-1133.
9. Inostrosa-Psijas A, Gil-Costa V, Marin M, Wainer G. Semi-asynchronous approximate parallel DEVS simulation of web search engines. *Concurrency Computat Pract Exper*. 2017;e4149. https://doi.org/10.1002/cpe.4149
10. Zeigler BP, Kim TG, Praehofer H. *Theory of Modeling and Simulation*. 2nd ed. San Francisco, CA: Academic Press; 2000.
11. Bampis L, Chatzichristofis SA, Iakovidou C, Amanatiadis A, Boutalis YS, Gasteratos A. A LoCATe-based visual place recognition system for mobile robotics and GPGPUs. *Concurrency Computat Pract Exper*. 2017;e4146. https://doi.org/10.1002/cpe.4146
12. Chatzichristofis SA, Boutalis YS. CEDD: color and edge directivity descriptor: a compact descriptor for image indexing and retrieval. Paper presented at: 6th International Conference on Computer Vision Systems (ICVS 2008); 2008; Santorini, Greece.
13. Amanatiadis A, Bampis L, Karakasis EG, Gasteratos A, Sirakoulis G. Real-time surveillance detection system for medium altitude long-endurance unmanned aerial vehicles. *Concurrency Computat Pract Exper*. 2017;e4145. https://doi.org/10.1002/cpe.4145
14. Davis J, Keck M. A two-stage template approach to person detection in thermal imagery. Paper presented at: Workshop on Applications of Computer Vision; 2005; Breckenridge, CO.