J Supercomput (2017) 73:1–3
DOI 10.1007/s11227-016-1922-5

High-performance computing: the essential tool and the essential challenge

Pedro Alonso[1] · **José Ranilla**[2] · **Jesús Vigo-Aguiar**[3]

Published online: 12 January 2017

High-performance computing (HPC) is nowadays an essential tool for the solution of many problems that arise in both scientific and engineering realms. HPC platforms are based on clusters of multicore nodes, and half of these facilities all around the world also include some type of accelerator device such as graphics processing units (GPUs) or the Intel Xeon Phi coprocessor. Many research interests are addressed to optimize applications that can get the most of these configurations.

At the same time, research on the HPC ecosystem (hardware, software tools, applications, etc.) is in the spotlight. In particular, *exascale* computing is receiving a major interest. The White House Office of Science and Technology Policy highlights the importance of exascale computing for the maintenance of US leadership over the coming decades, and it is for this reason that the United States is doing strategic investments in HPC to meet increasing computing demands and emerging technological challenges. Current and future research faces the natural problems that arise when concurrent resources become very large: huge electrical consumption, heat dissipation, and probability of failure, among others.

Many problems arise as long as we proceed in the way to developing exascale systems. One of them is the increase of failure rates. This special issue presents the

✉ Jesús Vigo-Aguiar
jvigo@usal.es

Pedro Alonso
palonso@upv.es

José Ranilla
ranilla@uniovi.es

1 Universitat Politècnica de València, Valencia, Spain

2 Universidad de Oviedo, Gijón, Spain

3 Universidad de Salamanca, Salamanca, Spain

developing of new fault tolerance techniques to reduce failures in the implementation of resilient MPI applications. Other source of problems relate to energy consumption and power dissipation of HPC applications. In this context we find several proposals to model the time and energy requirement of routines frequently used in HPC. For example, routines for Cholesky factorization, or other routines in which the adaptation of the throttling concurrence and the voltage-frequency scale reduce the demand of energy in the execution of this high efficiency linear algebra solvers. The concern of energy consumption and power dissipation also extends to integrated GPU–CPU systems in the context of the iterative computation. Other problems encountered in this way are derived from the exponential increase of data centers. Two of them have been addressed in this issue: the efficiency of broadcasting information inside a data center, and the management of the huge amount of data obtained from the real-time monitoring of thousand of resources. Writing accurate and reasonable parallel applications is a big concern too. A tool that automatically annotates a sequential code to produce a parallel one can represent a strong help in the right direction.

This special issue brings together a rich set of articles dealing with applications that require HPC capabilities. This set encompasses a varied range of topics, such as, e.g. the design of a heliostat field layout, the proposal of new parallel methods of information retrieval to high dimensionality representation spaces, the problem of finding a set of items/attributes characterizing a very large set of data, the design of efficient hydraulic solvers for the simulation of flows and pressures in water distribution systems, a hybrid distributed architecture proposed to improve scalability and performance of interactive 3D terrain visualization, or a system to calculate the wind field with the aim to mitigate forest fire propagation. Special attention has received the real-time video compressing in parallel, especially of the new standard HEVC, on both the shared and the distributed memory paradigms.

The use of heterogeneous systems continues hogging full attention in HPC applications. Hardware accelerators like NVIDIA GPUs are used for developing a fast variant of the Bitonic Merge Sort algorithm, and NVIDIA GPUs together with the Intel Xeon Phi are used in acoustic scattering problems, e.g. to predict noise problems. Also, the heterogeneous node as a whole is considered to design routines which automatically adapt to the conditions of the underlying computational system.

Low energy and mobile devices are attracting more and more attention in the HPC scene for applications that can benefit from its good performance, low energy consumption, and low price. Tablets, mobiles, or embedded systems usually contain processors with several cores implementing the ARM architecture and, sometimes, in turn, also include a GPU inside the chip die (SoCs). Some contributions contained in this issue present interesting applications for these devices, in particular, in the digital signal processing realm addressed to score following of musical sheets, or to building special sound environments.

This special issue gathers the most relevant contributions to the Minisymposium high-performance computing as a part of the 15th International Conference on Computational and Mathematical Methods in Science and Engineering (CMMSE 2015), which was held during 6–10 July 2015, in Cádiz. It was made possible thanks to the authors of the contributions and also to a significant number of reviewers who contributed the final result with their comments and suggestions for improvements. We

would like to conclude with our acknowledgments to all of them and we hope that the bottom line will be of your interest. We would also like to acknowledge to the "Ministerio de Educación y Ciencia" of Spain, for its support to the Spanish CAPAP-H5 network (HPC in Heterogeneous Systems, TIN2014-53522-REDT), and to the "Ministerio de Economía y Competitividad" from Spain/FEDER for supporting Grants TEC2015-67387-C4-1-R and TEC2015-67387-C4-3-R.

Finally, we would like to thank Prof. Hamid Arabnia, who with his professionalism has helped and encouragement us in this project.