

Guest editors note: Special issue on clusters, clouds, and data for scientific computing

The International Journal of High Performance Computing Applications
2023, Vol. 0(0) 1–2
© The Author(s) 2023
Article reuse guidelines:
sagepub.com/journals-permissions
DOI: 10.1177/10943420231180188
journals.sagepub.com/home/hpc



Jack Dongarra¹ and Bernard Tourancheau²

The research areas of cluster, cloud, and data analytics computing, which today provide fundamental infrastructure for all areas of advanced computational science, are being radically transformed by the convergence of at least two unprecedented trends. The first is the ongoing emergence of multicore and hybrid microprocessor designs, ushering in a new era of computing in which system designers must accept energy usage as a first-order constraint, and application designers must be able to exploit parallelism and data locality to an unprecedented degree. As the research community is rapidly becoming aware, the components of the traditional HPC software stack are poorly matched to the characteristics of systems based on these new architectures—hundreds of thousands of nodes, millions of cores, GPU accelerators, reduced bandwidth, and memory per core. The second trend is the dramatic escalation in the amount of data that leading edge scientific applications, and the communities that use them, are either generating or trying to analyze. A key problem in such data intensive science lies not only in the sheer volume of bits that must be processed and managed but also in the logistical problems associated with making the data of most current interest available to participants in large national and international collaborations, sitting in different administrative domains, spread across the wide area network, and wanting to use diverse resources—clusters, clouds, and data.

This special issue gathers selected papers of the Workshop on Clusters, Clouds and Data for Scientific Computing (CCDSC) that was held at *La Maison des Contes*, 69490 Dareize-France, on September 6–9, 2022. This workshop is a continuation of a series of workshops started in 1992 entitled Workshop on Environments and Tools for Parallel Scientific Computing, eventually Workshop on Clusters Computing and Grids for Scientific Computing. These workshops have alternate every 2 years between the U.S. and France, except for the 4 year gap caused by COVID in 2020.

The purpose of the workshop is to evaluate the state-of-the-art and future trends of scientific parallel computing, from multicore clusters to clouds, and their numerical and big data applications.

The following papers present research on performance evaluation, parallel workflows, quality of service ranking, task scheduling, and energy balancing, in the context of High Performance Computing.

The workshop was made possible, thanks to sponsorship from The National Science Foundation, AMD, Intel, Nvidia, Samsung, Science & Technology Facilities Council, UK, INRIA, and with the scientific support of the University of Tennessee in Knoxville, the Grenoble Alps University, and Vanderbilt University. We would like to express our warm thanks to them.

Author biographies

Jack Dongarra specializes in numerical algorithms in linear algebra, parallel computing, the use of advanced computer architectures, programming methodology, and tools for parallel computers. He holds appointments at the University of Manchester, Oak Ridge National Laboratory, and the University of Tennessee, where he founded the Innovative Computing Laboratory. In 2019, he received the ACM/SIAM Computational Science and Engineering Prize. In 2020, he received the IEEE-CS Computer Pioneer Award. He is a Fellow of the AAAS, ACM, IEEE, and SIAM; a foreign member of the British Royal Society; and a member of the U.S. National Academy of Science and the National Academy of Engineering. Most recently, he received the 2021 ACM A.M. Turing Award for his pioneering contributions to numerical algorithms and software that have driven decades of extraordinary progress in computing performance and applications.

¹Department of Electrical Engineering and Computer Science, The University of Tennessee Knoxville, TN, USA

²Grenoble Alps University and Grenoble Computer Science Laboratory, Grenoble, UK

Corresponding author:

Jack Dongarra, Department of Electrical Engineering and Computer Science, The University of Tennessee Knoxville, 1122 Volunteer Blvd, Knoxville, TN 37996-3450, USA.
Email: dongarra@icl.utk.edu

Bernard Tourancheau got MSc in Apply Maths from Grenoble University in 1986 and an MSc in Renewable Energy Science and Technology from Loughborough University in 2007. He was awarded best Computer Science PhD by Institut National Polytechnique of Grenoble in 1989 for his work on Parallel Computing for Distributed Memory Architectures. Working for the LIP laboratory, he was appointed as an assistant professor at Ecole Normale Supérieure de Lyon in 1989 before joining CNRS as a Senior Researcher. After initiating a CNRS-NSF collaboration, he worked two and half years on leave at the University of Tennessee on a senior researcher position with the US Center for Research in Parallel Computation at the ICL laboratory. He then took a professor position at University of Lyon in 1995 where he

created a research laboratory and the INRIA EPI RESO, specialized in High Speed Networking and HPC Clusters. In 2001, he joined SUN Microsystems Laboratories for 6 years sabbatical as a Principal Investigator in the DARPA HPCS project where he lead the backplane networking group. Back in academia, he oriented his research on sensor and actuator networks for building energy efficiency at LIP and CITI labs and associated INRIA EPI. He was appointed as a professor at University Joseph Fourier of Grenoble in 2012. Since then, he is developing research at the LIG laboratory Drakkar team about protocols and architectures for sensor networks and their applications to energy efficiency in buildings as well as GPGPU algorithms' optimization. He has authored more than a hundred peer-reviewed publications and owns 8 patents.