

Introduction to the Special Issue related to the Power-Aware Computing Workshop 2019—PACO 2019

The International Journal of High Performance Computing Applications 2021, Vol. 35(3) 209–210
© The Author(s) 2021



Article reuse guidelines:
sagepub.com/journals-permissions
DOI: 10.1177/10943420211008791
journals.sagepub.com/home/hpc



Peter Benner¹, Enrique Quintana-Orti² and Jens Saak¹

1. Background and history

Power-awareness in high-performance scientific computing has gained increased interest due to its non-negligible contributions to carbon-dioxide emissions and thus, to one of the main drivers of anthropogenic climate change. This is for instance recognized and popularized by the Green500 list,¹ which ranks the supercomputers from the TOP500 list in terms of energy efficiency by measuring performance per Watt.

In a joint project, funded 2015–2016 by the German Ministry of Education and Research (BMBF), the *Max Planck Institute for Dynamics of Complex Technical Systems* in Magdeburg (Germany) and the *Universidad de la República* in Montevideo (Uruguay) investigated numerical linear algebra algorithms for applications in systems and control theory with respect to power consumption and energy efficiency. As part of this effort, a first workshop on “Power-Aware Computing (PACO 2015)” was held in Magdeburg, July 6–7, 2015. The follow-up workshop PACO 2017 took place July 5–8, 2017, at Ringberg Castle in the south of Bavaria (Germany). PACO 2019, held November 5–6, 2019, again in Magdeburg, was the third instance in this series of workshops, and this special issue is dedicated to research results presented at this workshop.

The aims and scope of the PACO workshops comprise developments in power or energy savings in computational systems. The interests include, but are not strictly limited to:

- Exploitation of low-power hardware architectures;
- Development of algorithms that facilitate energy savings;
- Tools, libraries and environments for the development of power-aware software;
- Measurement and control of power consumption;
- High-performance power-aware applications;
- Communication-avoiding algorithms;
- High-performance implementations of tensor methods;
- Multi-precision algorithms in numerical linear algebra and related fields.

The focus of PACO 2019 was on strategies and algorithms that take into account minimizing the energy consumption, exemplified by the four plenary talks

- *Parallel solution of large sparse systems by direct and hybrid methods* by Iain S. Duff;
- *Iterative Refinement in Three Precisions* by Erin Carson;
- *Massively Parallel & Low Precision Accelerator Hardware as Trends in HPC—How to use it for large-scale simulations allowing high computational, numerical and energy efficiency with application to CFD* by Stefan Turek;
- *Parallel Algorithms for CP, Tucker, and Tensor Train Decompositions* by Grey Ballard.

In the following section, we briefly describe the contents of the papers submitted to this special issue.

2. Contents of this special issue

In their paper *Increased space-parallelism via time-simultaneous Newton-multigrid methods for nonstationary nonlinear PDE problems*, Jonas Dünnebacke, Stefan Turek, Christoph Lohmann, Andriy Sokolov, and Peter Zajac discuss how “parallel-in-space” and “simultaneous-in-time” Newton-multigrid methods can be designed such that the scaling behavior of the spatial parallelism is improved by reducing the latency costs. The idea is to compute many time steps at once and thereby solving fewer but larger systems. High parallel efficiency is then achieved by reordering the large-scale coefficient matrices and interpreting them as spatial discretizations. This allows to treat

¹Max Planck Institute for Dynamics of Complex Technical Systems, Magdeburg, Germany

²Universitat Politècnica de València, Spain

Corresponding author:

Peter Benner, Max Planck Institute for Dynamics of Complex Technical Systems, Sandtorstr. 1, Magdeburg 39106, Germany.
Email: benner@mpi-magdeburg.mpg.de

the linear systems using a multigrid algorithm with semi-coarsening in space and line smoothing in time direction.

The contribution *Energy efficiency of nonlinear domain decomposition methods* by Axel Klawonn, Martin Lanser, Oliver Rheinbach, Gerhard Wellein, and Markus Wittmann explores the energy efficient implementation of a nonlinear domain decomposition method featuring the solution of nonlinear problems by Newton's method on each subdomain embedded in asynchronous parallel iterations as opposed to the standard Newton-Krylov approach. The early termination of subproblems in the asynchronous iteration allows to put those nodes to sleep and thus save up to 77% of the energy consumption on some nodes compared to Newton-Krylov where all nodes are required to be active to the very end. Similar potential is reported for Additive Schwarz Preconditioned Inexact Newton.

Another contribution related to domain decomposition methods is the paper *Evaluating asynchronous Schwarz solvers on GPUs* by Pratik Nayak, Terry Cojean, and Hartwig Anzt. Here, the authors propose an asynchronous variant of the abstract Restricted Additive Schwarz method that introduces asynchronism between the sub-domains in order to eliminate the bulk synchronous execution pattern of the algorithm. For this purpose, the authors leverage the one-sided Remote Memory Access (RMA) in MPI in order to demonstrate the benefits of this asynchronous alternative, which renders appealing performance advantages when compared with the traditional synchronous schemes even for a well-balanced problem.

Although machine learning techniques are generally seen critical with respect to their energy footprint, due to the required expensive training phase, their clever combination with automatic tuning of sparse linear algebra tasks has excellent potential for energy savings. This is due to the vast number of applications of these operations in high-performance scientific computations in general. The two papers *Convolutional neural nets for estimating the run time and energy consumption of the sparse matrix-vector product* by Maria Barreda, Manuel F. Dolz and M. Asunción Castaño and *Selecting optimal SpMV realizations for GPUs via machine learning* by Ernesto Dufrechou, Pablo Ezzatti and Enrique S Quintana-Ortí discover ways to choose the right sparse matrix-vector product realization and determine energy consumption ahead of execution given the appropriate meta information about the sparse matrix structure and non-zero pattern.

Declaration of conflicting interests

The author(s) declared no potential conflicts of interest with respect to the research, authorship, and/or publication of this article.

Funding

The author(s) received no financial support for the research, authorship, and/or publication of this article.

Note

1. See <https://www.top500.org/lists/green500/>.