EoCoE Performance Benchmarking Methodology for Renewable Energy Applications

Paul Gibbon[1], Matthieu Haefele[2], Daniel Rohe[1], Sebastian Lührs[1], Yacine Ould-Rouis[2], Guillaume Latu[2], Thomas Breuer[1], Rene Halver[1], Abel Marin-Laflèche[2], Mathieu Lobet[2], Wendy Sharples[1], Nathalie Girard[2], Edouard Audit[2]
[1]Jülich Supercomputing Centre, Institute for Advanced Simulation, Forschungszentrum Jülich D-52425 Jülich, Germany
{p.gibbon,d.rohe,s.luehrs,t.breuer,r.halver,w.sharples}@juelich.de
[2]Maison de la Simulation, CEA, CNRS, Univ. Paris-Sud, UVSQ, Université Paris-Saclay, F-91191 Gif-sur-Yvette, France
{matthieu.haefele,yacine.ould-rouis,abel.marin-lafleche,mathieu.lobet,nathalie.girard,edouard.audit}@cea.fr

ACM Reference format:
Paul Gibbon[1], Matthieu Haefele[2], Daniel Rohe[1], Sebastian Lührs[1], Yacine Ould-Rouis[2], Guillaume Latu[2], Thomas Breuer[1], Rene Halver[1], Abel Marin-Laflèche[2], Mathieu Lobet[2], Wendy Sharples[1], Nathalie Girard[2], Edouard Audit[2]. 2017. EoCoE Performance Benchmarking Methodology for Renewable Energy Applications. In *Proceedings of ACM Conference for High Performance Computing, Networking, Storage and Analysis , Denver, Colorado USA, November 2017 (SC'17),* 2 pages.
https://doi.org/10.475/123_4

1 INTRODUCTION

The global transition away from fossil fuels towards a sustainable, decarbonized energy ecosystem will rely heavily on digitization to drive necessary innovations in production and storage technologies, mitigate power source intermittency and manage its distribution via a complex grid hierarchy. At the same time, supercomputing is also experiencing a major paradigm shift: future exascale technologies will open up unprecedented opportunities to tackle complex physical problems – such as the design of wind farms or smart materials for photovoltaics and batteries – but will demand major restructuring of application software, numerical algorithms and programming models. These challenges motivated the creation of the Energy Oriented Centre of Excellence (EoCoE) two years ago; an EU-funded consortium twenty-one partners across eight countries with strong engagements in both the HPC and energy fields.

This poster presents an optimisation strategy developed by the Energy Oriented Centre of Excellence (EoCoE) for computational models used in a variety of renewable energy domains. It is found that typical applications in this comparatively new sector cover the widest possible range of HPC maturity, from simple parallelization needs to near-exascale readiness. A key part of this process has therefore been the quantitative, reproducible performance assessment of applications consolidated by follow-up actions by code-teams comprising members of both developer groups and HPC centres involved with the EoCoE consortium. Examples of early successes achieved with this practice are given, together with an outlook on challenges faced for energy applications with next-generation, pre-exascale architectures.

SC'17, November 2017, Denver, Colorado USA

ACM ISBN 123-4567-24-567/08/06…$15.00
https://doi.org/10.475/123_4

2 EOCOE METHODOLOGY

During its first 18 months of operation, EoCoE has established a new network of expertise from four strategically important carbon-free energy domains in meteorology, materials, hydrology and fusion, linking them all together through a transversal, multi-layered HPC platform. A central operational feature of EoCoE to initiate and maintain interaction between the HPC teams and the energy domain developer groups has been the introduction of a series of hands-on workshops targeted at *real-world* energy applications, with the immediate aim of enabling faster, more complex simulations to be carried out. At these workshops, jointly run with the Performance Optimisation and Productivity (POP) CoE [1], a standardized performance evaluation of the applications is a key instrument, permitting us to: i) establish the precise status of an application when it is first examined by EoCoE personnel; ii) verify each code modification during the optimization process; and iii) quantitatively assess the impact of such support activity in its respective energy domain.

The performance status of an application is important to define objectively, whether measured by domain scientists and computer scientists, or for comparison to other applications in the same scientific area. For this reason the first step in setting up the EoCoE performance evaluation process was to establish a comprehensive set of 28 performance metrics obtained with performance tools such as the scalable profiler Scalasca, the powerful visual tracing toolset Paraver/Extrae and the I/O characterization tool Darshan: a complete description of the tools and associated metrics can be found in Ref. [2]. To manage this process the benchmarking environment JUBE [3] was extended in order to compile, run and *automatically extract* metrics for any application to ensure reproducible results at different points in time and location with the same or improved code versions.

3 RESULTS: ACCELERATING RENEWABLE ENERGY MODELLING

Equipped with this powerful combination of tools and methodology, 21 energy-oriented applications, including 7 from clients external to the EoCoE consortium have been welcomed to three performance

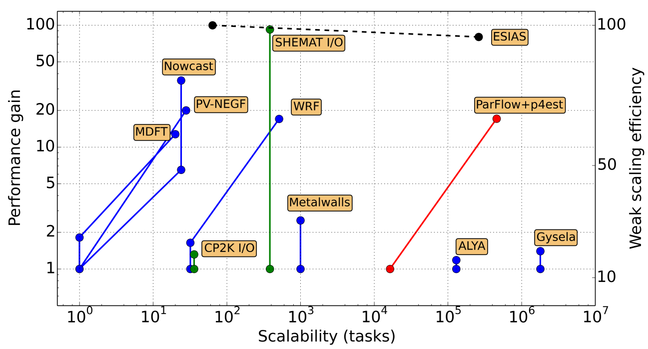

Figure 1: Summary of improvements made to selected applications during the first 18 months of the EoCoE project: optimisation and parallelisation (blue), parallel I/O library (green), kernel refactoring (red). The performance gain indicates the time-to-solution improvement since the first benchmarking workshop.

evaluation workshops in the first half of the project. The objective of these workshops is two-fold. First, this rare opportunity for face-to-face interaction between HPC experts and code owners is made the most of to establish a mutually agreed baseline for a reproducible performance evaluation process. Second, POP experts train the present code developers in the usage of state-of-the-art performance tools and help developers find potential bottlenecks in their code. The final outcome of such a workshop is a performance report that encapsulates these two parts, i) a set of quantitative, reproducible metrics describing the status of an application and its evolution in time and ii) an optimisation roadmap deduced from the analysis of the results generated by the performance tools.

The user engagement promoted by the EoCoE consortium does not end with the conclusion of a workshop, however. Bilateral *code-teams* comprising members of the developer group and the EoCoE HPC staff continue to monitor the code progress over subsequent months as optimisation efforts are undertaken. The latter may be carried out within the developer team, but also on request to other members of the consortium equipped with the appropriate expertise. In practice this can include simple parallelization work, replacement of I/O libraries, or refactoring of code kernels to deploy more efficient linear algebra packages. During the first half of the project around a dozen application support activities have taken place originating from all energy domains. These presented a wide variety of HPC maturity ranging from serial applications (in photovoltaics, battery research and near-term cloud cover prediction) to potential exascale candidates (in ensemble weather prediction, wind farm modelling and fusion plasma turbulence) already able to run efficiently on the full BlueGene/Q system at JSC [4], and are summarized in the progress chart of performance gain vs scalability also shown in the top right of the poster.

These enhancements fall essentially into two categories: first, a number of serial applications (MDFT[5], PV-Negf[6], Nowcast) could be successfully parallelised with OpenMP and are now able to efficiently leverage the computing power of a single multi-core node, typically reducing by almost an order of magnitude the application's time to solution with an immediate impact on scientific productivity. The exposure to high-level HPC expertise and application developer groups further along the scalability path has also inspired some consortium members to explore more radical overhauls of their applications in order to exploit distributed memory and even next generation architectures. A good example of this is PV-Negf , where a 1D distributed-memory version has already been developed as part of a separately funded PhD thesis, which scales to the complete JUQUEEN partition[4]. At the other side of the spectrum, the performance improvements of highly scalable applications (ESIAS, ALYA, Gysela and Parflow[7]) were more modest, but led to significant aggregate core-hour savings and serve as preparation for testing on next-generation architectures. All of these gains have largely been achieved with minimum impact on the source code, ensuring that the improved version was accepted back by the code developer and immediately put into production. This approach has been rewarded with a series of successes reported by the project partners in each energy domain [8].

4 OUTLOOK: EOCOE EXASCALE PATHWAY

Needless to say, future work with these applications aimed at exploiting forthcoming exascale architectures will demand a deeper refactoring of data management and kernel algorithms so that they are able to exploit multiple levels of memory access and parallelism. This will entail a concentration of resources working on a single code in close collaboration with the developers, but can still call on the advanced software engineering tools and methodologies available within the EoCoE network. A major outcome of this proof-of-concept phase has been the conception of a user-driven, structured support pathway for enhancing the HPC capability of energy-oriented numerical models, from entry-level parallelism to fully-fledged exascale readiness, ultimately establishing a portfolio of HPC services which may be of interest to SME stakeholders outside the consortium [9].

ACKNOWLEDGMENTS

The authors would like to thank the POP team Judit Gimenez, Brian Wylie, Ilya Zukhov, Michael Knobloch and Gemàn Llort for their invaluable contributions to the EoCoE-POP application benchmarking workshops and for helping to devise a consistent set of performance metrics. This work was supported by the Energy oriented Centre of Excellence (EoCoE), grant agreement number 676629, funded within the Horizon2020 framework of the European Union.

REFERENCES

[1] https://pop-coe.eu
[2] http://www.eocoe.eu/d115-application-performance-evaluation
[3] S. Lührs, D. Rohe, A. Schnurpfeil, K. Thust, W. Frings, in: Advances in parallel computing 27, (2016). doi:10.3233/978-1-61499-621-7-431
[4] http://www.fz-juelich.de/ias/jsc/EN/Expertise/High-Q-Club/_node.html
[5] M. Burbano, D. Carlier, F. Boucher, B. J. Morgan and M. Salanne, Phys. Rev. Lett., **116**, 135901 (2016). doi:10.1103/PhysRevLett.116.135901
[6] U. Aeberhard, J. Comp. Electron. **15**, 1219 (2016). doi:10.1007/s10825-016-0877-3
[7] C. Burstedde, J. A. Fonseca and S. Kollet, https://arxiv.org/abs/1702.06898
[8] http://www.eocoe.eu/success-stories
[9] http://www.eocoe.eu/services