

Reducing the ecological impact of computing through education and Python compilers

To the Editor — We read with interest the Comment by Portegies Zwart on the ecological impact of high-performance computing in astrophysics¹. We fully agree with its take-home message: scientists should be mindful of their carbon footprint. One of the proposed solutions, however, is to avoid the Python programming language. We contend that this would be counterproductive and that scientific programs written in Python can be efficient and energy-friendly. We argue that human factors, such as education, and the advancement of compiler technology are much more important than choice of language.

To support his point, Portegies Zwart presents a benchmark of the N -body problem with an inefficient implementation in Python that runs 50 times slower than a C++ implementation. As Python users concerned about our ecological impact, we worked on similar benchmarks on the same problem. In contrast to Portegies Zwart, we (1) consider efficient implementations in Python and Julia and (2) measure the energy consumption with dedicated hardware equipped with wattmeters.

Python is one of the most used and loved languages for science and data analysis. It is a general-purpose, dynamic language oriented towards communication between humans and fast prototyping. The language was designed to boost productivity. Strong open-source communities use Python and have built a rich scientific ecosystem of efficient libraries.

Portegies Zwart characterizes languages as being ‘interpreted’ or ‘compiled’, but these categories make sense only for specific implementations. Some interpreters compile parts of the code during the execution (just-in-time compilation, JIT)² and dynamic languages can also be compiled ahead of time (AOT). However, the most standard way to execute Python code is to interpret it with the reference implementation called CPython. Due to the lack of a built-in JIT compiler, CPython is relatively slow. Note that this inefficiency of the interpreter has a weak effect on the overall performance of most programs since total elapsed time and energy consumption are often dominated by computations done in optimized libraries. For numerics, the NumPy library³ is used to describe algorithms with high-level code,

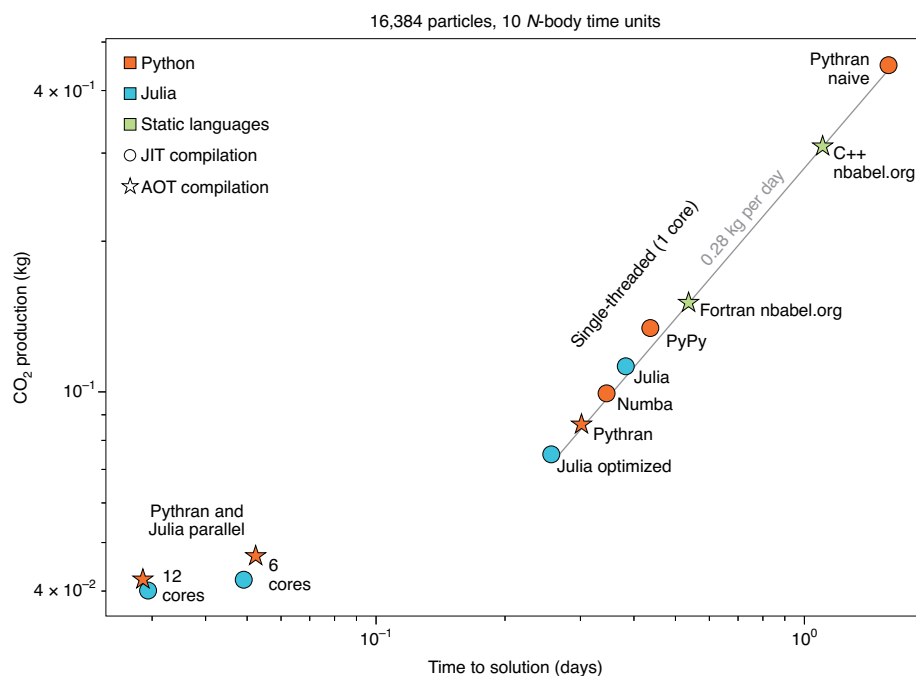


Fig. 1 | Efficiency in terms of CO₂ production and elapsed time for implementations in Python, Julia, C++ and Fortran. Energy consumption measurements were carried out on Grid’5000 clusters with 2.30 GHz Intel Xeon E5-2630 processors and converted from kWh to CO₂ using 283 g CO₂ kWh⁻¹. Optimizations were activated for all implementations with flags like `-Ofast`, `-march=native` and `-check-bounds=no`. We use gcc 8.3.0, Julia 1.5.3, Python 3.8.5, NumPy 0.19.5, Pythran 0.9.8, Transonic 0.4.7, Numba 0.52 and an unreleased version of PyPy including optimizations described in ref. ⁹.

which avoids too-frequent interactions with the interpreter.

In many cases, very few lines of code dominate the total computation. This is usually known as the 80/20 rule and provides support for two software development principles: (1) ‘premature optimization is the root of all evil’⁴ and (2) ‘measure, don’t guess’. These principles also apply for energy efficiency. For most Python programs, it would be counterproductive and expensive to manually rewrite them in C++, with a small gain-to-cost ratio.

However, some algorithms require low-level code and explicit loops. For example, for the N -body problem, the computation of the acceleration of each particle involves a loop over all other particles. A few lines of code are repeated $N^2/2$ times per time step. Portegies Zwart considered 10,000 time steps

and $N = 16,384$, so the program is dominated by 1,342,177,280,000 executions of a simple and inexpensive computation. Using CPython for this hot loop makes the whole program inefficient. Good news for Python: it is straightforward to use efficient alternatives. For our benchmark, we use three tools: (1) Pythran⁵, a Python-NumPy AOT compiler transpiling to C++, (2) Numba⁶, a Python-NumPy JIT compiler based on LLVM (the same compilation target as Julia) and (3) PyPy⁷, an alternative Python interpreter with a JIT compiler.

Figure 1 is similar in form to fig. 3 in ref. ¹: the CO₂ production is plotted as a function of the elapsed time for ten implementations. The C++ and Fortran implementations (green stars) are taken from www.nbabel.org and were used by Portegies Zwart. Note that these implementations could have been further optimized, but are representative of code

written by many scientists. We consider five implementations in Python (red markers). We would like to emphasize that (1) these implementations are fully written in Python. The implementations using Pythran and Numba are written in Python-NumPy but NumPy is only used for its arrays as a data structure and not for high-level functions. (2) Four implementations in Python are faster than the C++ and Fortran implementations. The implementation labelled 'Pythran naive' (simple NumPy code accelerated only by decorating one function with `@transonic.jit`⁸) is only three times slower than the Fortran implementation. (3) All Python implementations are simpler to understand, read and write than the C++ and Fortran implementations.

For comparison, we also consider three implementations in Julia (blue circles): the implementation labelled 'Julia' is comparable to the 'Pythran' and 'Numba' implementations and could have been written by scientists with similar skills. We did not include a Julia implementation similar to 'Pythran naive' because it is inefficient. 'Julia optimized' and 'Julia parallel' have been proposed by Julia users after a long discussion on a [Julia forum](#).

The four points close to the bottom-left corner correspond to two parallel implementations using Pythran+OpenMP and Julia executed using 6 and 12 CPU cores. We consider in Fig. 1 the energy consumption of the cores used (6 or 12 for these runs and 1 for the sequential jobs),

which makes sense on shared clusters wherein one can reserve only the needed cores. We see that parallelism with threads has only a moderate impact on energy consumption since the increase in power consumption partly counterbalances the decrease of elapsed time. For example, the 12-threaded Pythran version is 10 times faster than the single-threaded one but produces only two times less CO₂.

Our work shows that the performance of scientific programs depends less on languages than on the time spent on optimization and the developer skills to correctly use the right tools. These benchmarks demonstrate that dynamic languages like Python can actually be good solutions to easily obtain good performance while retaining simplicity and readability. We think that minimizing the ecological impact of scientific computing is mainly limited by human factors: time, work, knowledge and skills. For example, scientists have to be able to use job schedulers and shared clusters optimized in terms of energy consumption. They should know how to profile their code to discover which parts can potentially be optimized. Therefore, time and money should be invested in education and tooling to minimize the overall ecological impact of computing, irrespective of the underlying language. □

Pierre Augier ¹,
Carl Friedrich Bolz-Tereick ²,
Serge Guelton ³ and

Ashwin Vishnu Mohanan ⁴

¹University of Grenoble Alpes, CNRS, Grenoble INP, LEGI, Grenoble, France. ²Heinrich-Heine-Universität, Düsseldorf, Germany. ³Namek, Paris, France.

⁴Department of Meteorology, Stockholm University, Stockholm, Sweden.

✉ e-mail: pierre.augier@univ-grenoble-alpes.fr

Published online: 16 April 2021

<https://doi.org/10.1038/s41550-021-01342-y>

References

- Portegies Zwart, S. *Nat. Astron.* **4**, 819–822 (2020).
- Aycock, J. *ACM Comput. Surv.* **35**, 97–113 (2003).
- Harris, C. R. et al. *Nature* **585**, 357–362 (2020).
- Knuth, D. E. *ACM Comput. Surv.* **6**, 261–301 (1974).
- Guelton, S. et al. *Comput. Sci. Discov.* **8**, 014001 (2015).
- Lam, S. K., Pitrou, A. & Seibert, S. In *Proc. Second Workshop on the LLVM Compiler Infrastructure in HPC 7* (ACM, 2015).
- Bolz, C. F., Cuni, A., Fijalkowski, M. & Rigo, A. In *Proc. 4th Workshop on the Implementation, Compilation, Optimization of Object-Oriented Languages and Programming Systems* 18–25 (ACM, 2009).
- Augier, P., Mohanan, A. V. & Bonamy, C. *J. Open Res. Softw.* **7**, 9 (2019).
- Cheng, L., Ilbeyi, B., Bolz-Tereick, C. F. & Batten, C. In *Proc. 18th ACM/IEEE International Symposium on Code Generation and Optimization* 16–29 (ACM, 2020).

Acknowledgements

We thank D. Barreto, L. Martínez and A. Oskin for working on the optimized Julia implementations. We also thank Camille Nous (Cogitamus Laboratory) and the open-source communities who contributed to the tools used for this study. Part of this work was performed using resources provided by Grid'5000 (supported by a scientific interest group hosted by Inria and including CNRS, RENATER and several universities) and CINES under GENCI allocation number A0080107567.

Competing interests

The authors declare no competing interests.