High-Performance Computing and Software Sustainability: Toward Green Software Development

Date: 22 January 2025
Presented by: Dr. Antigoni Georgiadou (Oak Ridge National Laboratory)

(The slides are available via the link in the page's sidebar.)

Q: Have you had any experience with the GreenAlgorithms project? You can get estimates of your HPC carbon footprint and I wonder if just putting it in understandable terms (e.g. it reports car-kilometer equivalents) could motivate researchers to use compute resources more responsibly? Or is the main limitation that they lack the knowledge of how to do more efficient computing? (sorry, had to leave early so apologies if you addressed this in your talk!)

A: This seems interesting and I believe a discussion of the tool would be a good addition to the video series.

In our SC24 paper (https://doi.org/10.1109/SCW63240.2024.00225) we mention that "The transition to energy-efficient scientific computing presents significant challenges, particularly in how it affects current allocation models and user incentives within leadership-class computing facilities", "...face challenges in integrating with complex build and test environments and in presenting a user-friendly learning curve". The additional profiling tools and new policies will certainly seem novice for the leadership computing facility users but typically OLCF users are aware and capable of pursuing the potential optimizations needed for energy efficient scientific codes.

Q: Nuclear reactors are green. The top few HPC machines on the TOP500 list have 20-30MW power ratings. A nuclear powered aircraft carrier is run by A1B reactor of about 700MW in the scheme of things. A small nuclear reactor alongside an HPC machine is green. Why does the speaker not mention nuclear reactor as a green solution?

A: I list below a few artifacts and media including DOE media discussing the option of nuclear reactors as an alternative for energy usage at HPC and data centers for interested parties. This talk was centralized around the topics that a leadership computing facility will need to address now and within the upcoming years when the discussion of the status and nature of the next HPC system will arise. Within this

context of the facility users, their experience and contribution, hardware and software required to account for energy measurements, and the assessment on the science being produced, I used this opportunity to discuss the steps starting now and trying to extend the discussion for the post-exascale era for HPC.

- https://doi.org/10.1016/j.pnucene.2021.104080
- https://www.energy.gov/ne/articles/5-ways-us-nuclear-energy-industry-evolving-2024

Moderator's comment: It is not only about climate impacts. Electricity costs money. So if we can reduce the power consumption of an HPC system, it will cost less to operate.

Q: Should we be discouraging users from using inefficient codes written in Python or R? Or are those insignificant blips in the grand scheme of exascale?

A: LCF users are increasingly using python as a language for many processes and overall science runs. Julia programming language has taken steps to make a distributed and accelerated version to minimize the time to solution. Many applications have been using Julia including scientific domains, e.g. in climate. Thorough energy cost and optimization potential surveys across programming languages are required. I include here a list of articles and publications that are related to the topic.

- Caballar, R.D.: We need to decarbonize software: The way we write software has unappreciated environmental impacts. IEEE Spectrum 61(4), 26–31 (2024). DOI:10.1109/MSPEC.2024.10491388
- Couto, M., Pereira, R., Ribeiro, F., Rua, R., Saraiva, J.: Towards a green ranking for programming languages. In: Proceedings of the 21st Brazilian Symposium on Programming Languages. pp. 1–8 (2017). DOI:10.1145/3125374.3125382
- https://greenlab.di.uminho.pt/wp-content/uploads/2017/10/sleFinal.pdf
- https://dl.acm.org/doi/10.1145/3624062.3624278
- https://sc24.supercomputing.org/proceedings/poster/poster_files/post113s2-file3.pdf

Q:Can you elaborate on the tools you are using for energy efficiency calculations? (Which languages does it support? Architecture dependent/how are they calibrated?)

A: In our study we use hardware specific tools and application metrics to assess energy consumptions on Frontier. A short list of those are Omnistat, crayPat, crayMP, saact.

Q: In one of the slides you mentioned numbers for the consumed energy, coming from what looked like a SLURM job. What units are these numbers in please?

A: Joules

Q: Is it rather recommended to use full precision?

A: The usage of the type of precision used is domain specific and with respect to the policies imposed by a LCF, there are no restrictions or encouragements with respect to the precision used by users.

I include here some interesting papers that discuss the topic for specific domains but also the specificity of using the OLCF facility.

- https://doi.org/10.1016/j.jpdc.2023.104746
- https://arxiv.org/abs/2412.19322

Q: How much electricity do we save by writing efficient codes?
Powering a supercomputer is challenging and within this context producing cost effective science is the key to being able to utilize the computing resources and allowing additional projects and teams to run on supercomputing systems.

A: Frontier consumes around 21 megawatts (MW) (which is equivalent to the power needed for 15,000 single-family homes), compared to its predecessor Summit's 13 MW. In that sense, optimized software will indirectly assist to the cost of running on a supercomputer.

Q: I was really hoping to see some actual example code with before and (green) after. Where can I go to see this?

A: Suggestion: https://doi.org/10.1109/TPDS.2018.2872992.

Q: At the outset it was mentioned that electrical energy is used for the calculations, this turns into heat, and then energy is used to remove the heat. Are you aware of any projects that repurpose the heat, rather than remove it? For instance, in cool climates heat energy can be captured (say in a thermal battery in the ground) and used to heat buildings during the cold months.

A: There are benefits in repurposing the heat of a supercomputer and reusing it for other purposes. Many industry companies have been following this practice with their data centers. This is a great alternative which of course possesses its own challenges and one might think is not straightforward. I include here some highlights and publications that might be of interest to this topic:

- https://www.techtarget.com/searchdatacenter/tip/Data-center-heat-reuse-How-to-make-the-most-of-excess-heat
- https://doi.org/10.1016/j.enconman.2024.118408

Q: Dr. James Demmel (of LAPACK fame) conducts research related to communication-avoiding (thus, energy-efficient) algorithms at University of California, Berkeley. Would it be useful to contact his group regarding HPC sustainable software or have you already been in contact with them?

This sounds to be an interesting addition to the discussion. I will reach out to them.