The Relationship Between Sustainability and Parallelism

University of St Andrews

By Morgan Hunt

Background

There has been a rise in recent years in the need for more computational power and resources with demand through data centers and the consumer multidevice ecosystems. This increase in demand causes an increase in resource consumption like energy consumption which poses major sustainability concerns. Parallelism has the potential to help mediate this issues by distributing the workload of a program across multiple processor cores which can potentially decrease energy consumption compared to a sequential approach. However, the introduction of parallelism alone is not sufficient for sustainability as the choice of parallel model, granularity, and design methods significantly impact energy consumption.

Sequential versus Parallel

Energy Consumption Analysis of Matrix Multiplication Program

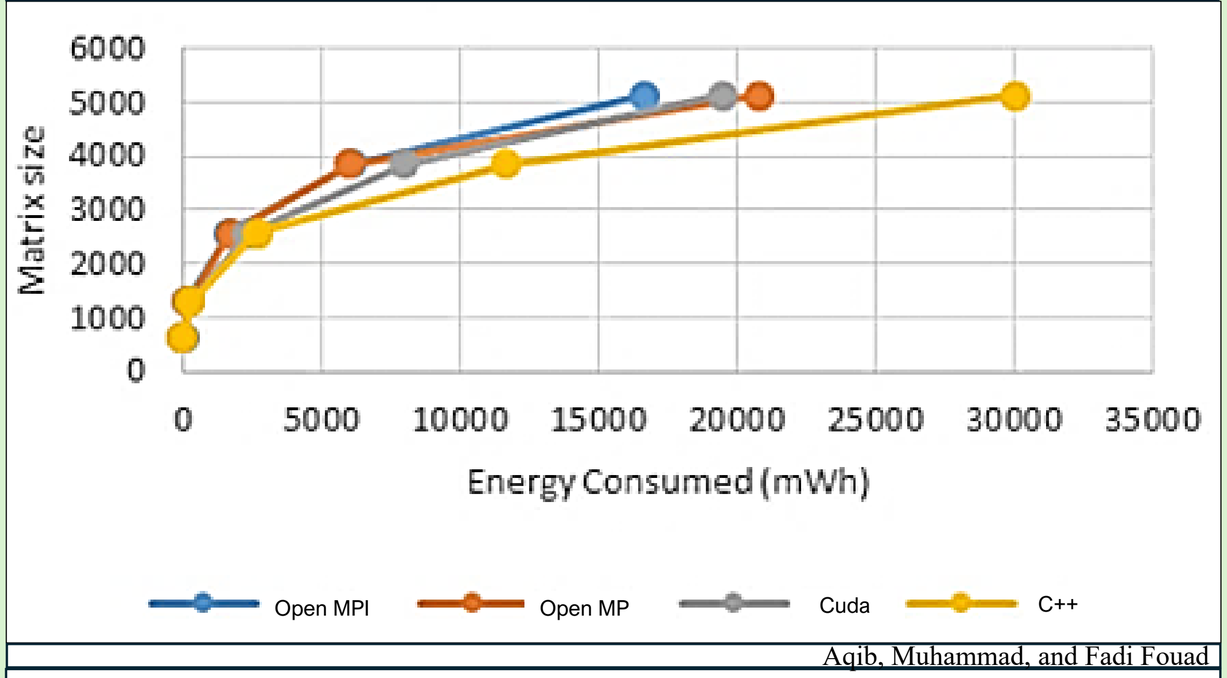

Aqib, Muhammad, and Fadi Fouad

The experiment by Muhammad Aqib and Fadi Fouad (above and below) was performed in two parts, matrix multiplication algorithm and a quick sort algorithm, using sequential C++ and a few parallel programming models, OpenMP, Open MPI, and CUDA. They recorded the execution time and the energy consumption for each program. The data here shows that parallelism works well to decrease energy consumption for tasks that can parallelize well and need to process a large amount of data. This demonstrates that the addition of parallelism is not an exclusive solution to decreasing energy consumption and must be done in contexts where different tasks can be done entirely independently from each other and the sample size is large. If the amount of work that needs to be done is small, the difference in energy consumption between a sequential and parallel approach is small.

Energy Consumption Analysis of Quick Sort Program

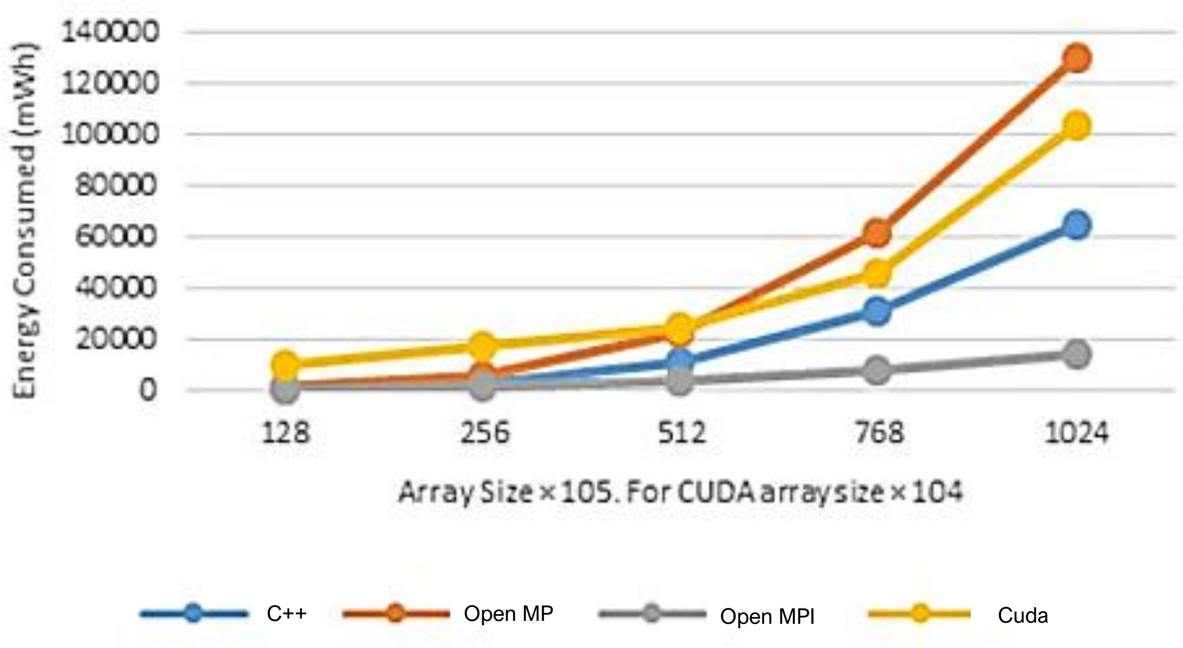

Aqib, Muhammad, and Fadi Fouad

Design of Parallel Programs

Asynchronous vs Synchronous Consumption

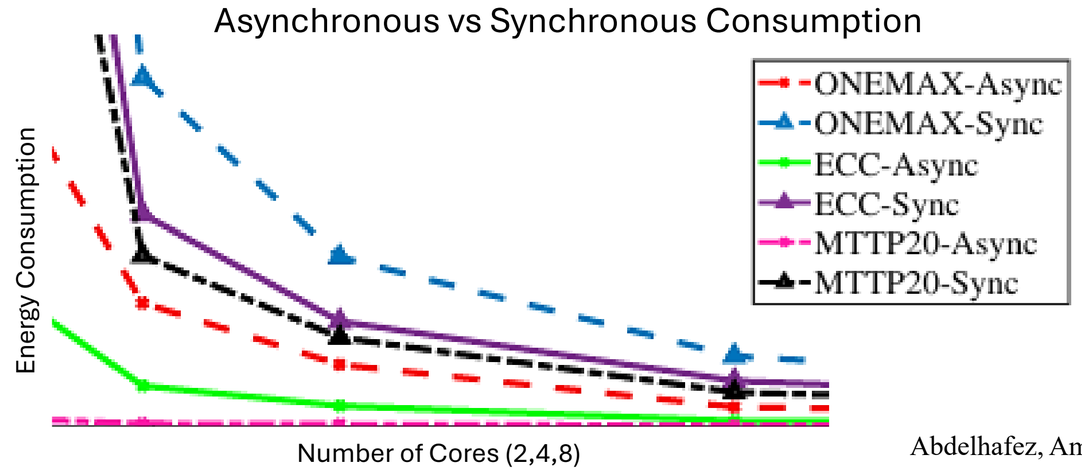

Abdelhafez, Amr, et al

Reducing energy consumption is also not exclusively about the context in which the developer chooses to parallelize. The design of the parallelization significantly impacts the energy consumption even in contexts where there are numerous independent tasks. The above experiment demonstrates this by comparing the same algorithms synchronously and asynchronously across varying numbers of processor cores. Overall, the synchronous algorithms consumed more energy than their asynchronous counterparts, but this difference decreased as the number of cores utilized increased. This shows that the method the developer decides to implement should depend on the resources of the system.

Additionally, the granularity of the parallelism is also important in controlling energy consumption. In the experiment performed by Fonseca, Alcides, and Bruno Cabral, they compared the energy consumption of tasks with different cutoff strategies and found that the choice of strategy significantly impacts the energy consumption of the program. The benefits of parallelism are not felt as much when the data is much smaller (at the end of the parallelization), so it makes sense to run the ends sequentially (the cutoff). Of course, when the cutoff happens must be balanced with this tradeoff.

Impact

Parallelism has the potential to reduce energy consumption in computing if in the correct context and when designed properly, and this impact has the most potential in the data centers necessary to support generative artificial intelligence (AI). These data centers use a large amount of energy on the global scale, and this energy usage is compounded by the growing popularity of generative AI. The choice of parallel framework, task granularity, and synchronous behavior decide whether these center use energy efficiently or wastefully.

Beyond generative AI, the modern world is full of various forms of technology in the consumer market from phones to personal computers. The appropriate use of parallelism on these machines means less battery usage which translates to less energy consumption overall. Even small improvements in the energy efficiency can yield a significant reduction in energy usage and carbon emission when applied at a global scale of the computing industry..

References

- Abdelhafez, Amr, et al. “A Component-Based Study of Energy Consumption for Sequential and Parallel Genetic Algorithms.” *The Journal of Supercomputing*, vol. 75, no. 10, 4 Apr. 2019, pp. 6194–6219, https://doi.org/10.1007/s11227-019-02843-4.
- Aqib, Muhammad, and Fadi Fouad. “The Effect of Parallel Programming Languages on the Performance and Energy Consumption of HPC Applications.” *International Journal of Advanced Computer Science and Applications*, vol. 7, no. 2, 2016, https://doi.org/10.14569/ijacsa.2016.070223. Accessed 26 Mar. 2025.
- Fonseca, Alcides, and Bruno Cabral. “Understanding the Impact of Task Granularity in the Energy Consumption of Parallel Programs.” *Sustainable Computing: Informatics and Systems*, vol. 17, Mar. 2018, pp. 69–80, https://doi.org/10.1016/j.suscom.2017.10.014. Accessed 9 Nov. 2021.