



HAL
open science

Green HPC: An analysis of the domain based on Top500

Abdessalam Benhari, Yves Denneulin, Frédéric Desprez, Fanny Dufossé, Denis Trystram

► **To cite this version:**

Abdessalam Benhari, Yves Denneulin, Frédéric Desprez, Fanny Dufossé, Denis Trystram. Green HPC: An analysis of the domain based on Top500. 2024. hal-04519645

HAL Id: hal-04519645

<https://hal.science/hal-04519645v1>

Preprint submitted on 25 Mar 2024

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.



Distributed under a Creative Commons Attribution 4.0 International License

Green HPC: An analysis of the domain based on Top500

Abdessalam Benhari^{1,2} Yves Denneulin¹
Frédéric Desprez³ Fanny Dufossé¹ Denis Trystram¹

¹Univ. Grenoble Alpes, Grenoble INP, Inria, CNRS, LIG, France

²Bull SAS (Eviden, Atos group), France

³Inria, Grenoble, France

March 25, 2024

Abstract

The demand in computing power has never stopped growing over the years. Today, the performance of the most powerful systems exceeds the exascale and the number of petascale systems continues to grow. Unfortunately, this growth also goes hand in hand with ever-increasing energy costs, which in turn means a significant carbon footprint. In view of the environmental crisis, this paper intends to look at the often hidden issue of energy consumption of HPC systems. As it is not easy to access the data of the constructors, we then consider the Top500 as the *tip of the iceberg* to identify the trends of the whole domain.

The objective of this work is to analyze Top500 and Green500 data from several perspectives in order to identify the dynamic of the domain regarding its environmental impact. The contributions are to take stock of the empirical laws governing the evolution of HPC computing systems both from the performance and energy perspectives, to analyze the most relevant data for developing the performance and energy efficiency of large-scale computing systems, to put these analyses into perspective with effects and impacts (lifespan of the HPC systems) and finally to derive a predictive model for the weight of HPC sector within the horizon 2030.

keywords: High Performance Computing, Energy efficiency , Top500, Green500

1 Introduction

With the climate crisis and the global warming, humanity faces an unprecedented challenge [12]. It is also well-established that the ICT sector (Information and Communication Technologies) plays a significant role in this crisis

through the huge energy consumption of computing, storage and interconnection devices. A recent survey of studies dealing with methodologies and tools for estimating the carbon footprint of the domain [9] gives a value between 2.1 et 3.9% of the global GHG emissions. In this wide variety of electronic and computing systems, data centers account for around one percent of the world's electricity consumption [18]. In France, the domain is estimated to reach 10% by 2030 [3].

It is very difficult to have a precise view of the whole HPC domain. Many researchers in the field and political leaders believe that HPC will provide solutions to the warming/environmental crisis. Large-scale climate simulations are the only way to understand global dynamics and gain a clearer view of further decisions. Optimizing components in the fields of transportation, intelligent buildings, energy or health requires considerable computing power [13]. On another hand, HPC is also part of the problem as it belongs to the digital world. In this paper, we develop an analysis that should contribute to understand the carbon footprint of the HPC domain. Our main objective is to identify and clarify the impacts of HPC from the perspective of energy consumption. The main contributions are the following:

- analyze the empirical laws governing the evolution of HPC computing systems both from the performance and energy perspectives,
- analyze the most relevant data to study the performance and energy efficiency of large-scale computing systems,
- put these analyses into perspective with effects and impacts (lifespan of HPC systems),
- derive an estimation of GHG emissions for the HPC domain within the horizon 2030.

2 Background

Top500 was introduced in 1993 to rank the most powerful systems in the world using a common set of applications and criteria. It highlights the evolution and trends in HPC over the years. Including varying parameters on the architecture, the location and the capacity of systems, it has evolved to include the power consumption and the architecture's heterogeneity. It has rapidly gained in popularity and is now considered as a major source of information on the domain. Thoroughly studied over the years, many studies have observed the continuous evolution of computing capacity and energy efficiency on these systems. It was at the beginning fairly exhaustive but in the past 10 years has become less so because many powerful systems are not included, for example most hyperscalers ones. Nevertheless, the large duration and stability of the systems ranked in the Top500 makes it the most comprehensive and reliable source to study the evolution of high end powerful computers in the past 30 years. Over the years,

new ranking lists have been released based on other benchmarks. Top500 introduced in 2017 a HPCG ranking based on a High Performance Conjugate Gradient benchmark, to bring more diversity to its evaluation. At the same epoch, Graph500 was released with more data-intensive benchmarks. We focus in this study on Top500 data and we draft a comparison of evolution trajectory with Graph500.

2.1 Top500

Top500 is a project [2] whose objective is to publish twice a year a ranking of the 500 most powerful HPC systems. Throughout this time, it published more than fifty lists, in which we count over 10,000 supercomputers from more than 2,800 institutions around the world. To rank supercomputers, the *Top500 project* uses the well-known High-Performance *Linpack* [6] benchmark. A program supplied by Top500 has to be run on the HPC system. It solves a dense system of linear equations and provides the number of floating point operations per second (FLOPS) during its execution.

Top500 provides attributes about the identification and location of each system, its architecture and its performance. An HPC system is identified by its name, first and previous ranking, release year, and site. Two attributes characterize each supercomputer’s performance: R_{max} , the maximal Linpack performance achieved and R_{peak} , the theoretical peak performance calculated using the advertised clock rate of the CPU. Concerning the architecture, it includes the number of cores, memory capacities, CPU/GPU types, constructors, etc. Table 1 shows an example of Top500 main attributes for a recent HPC system.

2.2 Green500

In 2008 energy efficiency became a concern in HPC. The monetary cost of running HPC systems exceeded the cost of purchasing and maintaining them [4]. This highlighted the importance of improving the energy efficiency of the upcoming supercomputers.

To raise awareness on that matter, the *Green500* was brought to light in 2009 [21]. This project started with three exploratory lists (Little Green500, HPCC Green500, and Open Green500), then became what we officially know as Green500. The official ranking published in the Top500 website starts from 2013 and is updated at the same pace as Top500. It uses the same Linpack benchmark as the Top500 project, but gives more attention to the energy efficiency metric to compare the different supercomputers displayed in the list. Other than the energy efficiency and the "green" ranking, Green500 contains the same attributes as Top500.

Attribute Name	Values
Name	Frontier
Country	United States
Rank	1 (year 2023)
Installation site	DOE/SC/Oak Ridge National Laboratory
Total cores	8699904
Accelerator cores	8138240
Processor	AMD Optimized 3rd Generation EPYC 64C 2GHz
R_{peak}	1679818.75 [TFlops]
R_{max}	1194000 [TFlops]
Power	22703 [kW]

Table 1: **Example of Top500 (and Green500) main attributes.**

2.3 Graph500

Graph500 [1] is another ranking list of HPC systems based on more irregular computations related to graphs. It appeared in 2010 on BFS benchmark (for Breadth-First Search) including a GreenGraph500 variant. An extension list, namely SSSP (Single-Source Shortest Paths), was released in 2017. These benchmarks are data-intensive and not well suited for GPU execution. They are thus a counterpoint to the Linpack benchmark. The objective was to focus on real life applications with large-scale data-set from domains as cyber-security, medical data-base or symbolic network problems [20]. The metric of Graph500 is the Traversed Edges Per Second (TEPS). Graph500 is however not as popular as Top500, and counts 231 systems in November 2023. The focus of this paper is limited to the Top500. Graph500 will be only used for comparing the respective growth rates on different performance metrics.

2.4 Restrictions

The main advantage of the Top500 is that it provides a stable view of the same features collected over many years and is so a valuable source to measure the evolution of HPC architectures. Beside the fact that it only concerns a rather small part of the whole computing systems, it has two drawbacks:

its declarative character the Top500 systems included in the list are submitted on a voluntary (not contractual) basis. The performances are self-validated and self-reported with almost no independent checks. Many HPC systems are not included mainly for geopolitical or economical reasons. Top500 does not contain all the most powerful systems, in particular, some large companies do not have interest and thus, do not publish in the list.

its orientation towards large scale HPC systems this is especially true for the green counterpart of Top500 where highly innovative but small scale systems have no chance to appear.

2.5 Evolution laws

What we refer to as a 'law' in this section is a trend observation done over many years that became widely adopted. Such laws are used as a primary tool to provide a rough prediction about the evolution of a certain metric. The most popular ones are:

Moore's law states that the number of transistors in an integrated circuit doubles about every two years. It is also often extended to track the evolution of the performance of whole HPC systems [14, 22, 23].

Koomey's law is similar to Moore's law, but it targets the energy efficiency trend. It observes that the number of computations per joule of dissipated energy roughly doubles every 18 months. [16]

3 Related works

Some papers analyzed the performance of supercomputers based on Top500 and Green500. Many keynote talks also referred to the performance evolution of HPC systems. However, none of them provide hints about the impact in term of carbon intensity, estimations and lifespan.

3.1 Trends / Performance analysis

Each Top500 list is released every year in May during ISC High Performance conference and in November at SuperComputing. It is the occasion to observe the evolution of trends in HPC. Regarding the technology, the authors of Top500 published a study in 2015 [22] using the evolution of R_{max} over time to evaluate the increase of processors capacity compared to Moore's law. It describes key break points in the evolution which highlight a slowdown starting from 2008 (which is linked to growth rate of the number of components). This study is restricted to homogeneous systems and it was published before the era of GPUs. Milojevic et al. [19] studied the evolution trend of HPC from the architectural point of view. They showed that the HPC domain evolved to more customisation. It also highlights the performance slowdown mentioned in [22] but attributes it to the end of Dennard scaling. Dennard scaling envisions a physical limit of power per die area that restrains the performance gain due to increasing density of transistors. None of the above works really tackle the energy efficiency and its impact of the overall evolution of the HPC systems. Khan et al. [14] conducted an analysis that takes into account the architectural trends of the Top500. It focuses on the comparison between homogeneous and heterogeneous systems in term of performance and power consumption. It was mostly focused on the period 2009-2019, that includes the diffusion of GPU in HPC. In [15], the authors made in 2011 performance and energy consumption projections up to 2024, claiming that only heterogeneous machines will be able to reach exascale but with an energy consumption up to 100 MW. It seems that Exascale was obtained with "only" 20 MW, but indeed with heterogeneous architectures.

3.2 Energy efficiency

Green500 was introduced to study the evolution of energetic efficiency of HPC systems, that is the ratio of R_{max} over the total power consumption. In November 2023, the head of this list is Henri system, with more than 65 GFLOPS/watt. In 2011, Scogland et al. [21] described the evolution of the Green500 metric over the three first years to identify the design aspects that contribute to more power consumption in the objective of exascale. This study was extended in 2013 in [23]. It introduces a new (holistic) metric to give a unified information about the energy efficiency/performance of Green500 called the EXASCALAR. This question of correlation between performance and energy efficiency of systems is crucial for various studies. Khan et al. [14] evaluated the evolution of the Pearson’s correlation coefficient of the two lists over the period 2009-2019.

Mair et al. [17] analysed the energy efficiency of heterogeneous and homogeneous systems based on Green500, and proposed a new energy efficiency metric to avoid the bias induced by system sizes. Hsu et al. mentioned the same metric in [11]. They provide a 10 years retrospective on the evolution of energy efficiency metrics and their exploitation to evaluate HPC systems (mainly the Power Usage Efficiency and Performance-Power Ratio). It also highlights the main issues that need to be addressed in term of evaluation metrics and measurements. Fraternali et al. [8] took a closer look at the impact of heterogeneity on performance and energy consumption variability by conducting directly their experiments on one of the top system in Green500. Although their study does not address trend analysis, they were able to propose some guidelines (hardware-wise) to build sustainable supercomputers. Gao et al. [10] relayed on Top500 and Green500 to study the influence of design and architectural parameters on the Linpack scores in term of performance and energy efficiency. Then, they identified some development trends in supercomputer design.

4 Methodology

We present in this section the experimental protocol we followed for the results presented in this paper.

Dataset description We collected data from three main existing lists (Top500, Green500 and Graph500). Then, they were filtered by removing any incomplete items or outliers (non-GPUs heterogeneous systems for instance). We also changed some of the resulting data to unify measurement units. Finally, the data were gathered into a single table to conduct the experiments where some new metrics were added (such as $GFLOPS/watt$, R_{max}/R_{peak} , etc.).

Experimental setup All the experiments were produced using a jupyter notebook¹ with a Python3.8 kernel and take around 30 seconds to reproduce.

¹<https://jupyter.org/>

Reproducibility The code used to produce the different figures, along with all the steps to replicate these experiments is available on Github².

5 Experiments

The following section reports the multiple experiments performed on the previously mentioned dataset to investigate the performance and energy efficiency trends of large-scale computing systems.

5.1 Lifespan of systems in the Top500

We are interested here in the overall lifespan of HPC systems in Top500. We did an assessment of how long a system stays in Top500 after its first appearance by counting the number of apparitions of each system multiplied by 6, that corresponds to the interval between two successive publications of the list. Figure 1 – in which the *x-axis* represents the apparition date in Top500, and the *y-axis* represents the proportion of systems that stayed at (or more than) that time in Top500 – summarizes the observed lifespan of these systems in Top500.

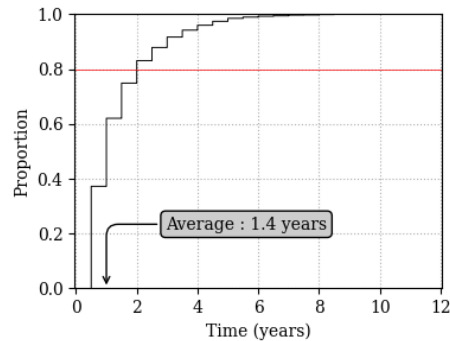


Figure 1: **lifespan of HPC systems in the Top500 (in years) from 1993. The red horizontal line means that 80% of the systems stay less than 2 years.**

Only a selected few of the highest supercomputers on the list can appear in it for up to 8 years while 80% of them do not exceed an appearance time of 2 years, which gives us an average of 1.4 years. Although these results validate the rapid evolution of large-scale HPC systems, they do not reflect the real lifespan since we lose track of most of the systems after they leave the Top500. However, it provides an insight about the duration, which is an important metric to assess their ecological impact. Complete life cycle analyses show that the usage cost of an HPC system represents only a part of the total impacts [5].

²<https://github.com/aBenhari/Green500-analysis.git>

Certain organizations choose to disclose information regarding their Supercomputer’s architecture and usage. For instance, the cluster *CURIE* displays its configuration with every upgrades or changes over time³. While the exact lifespan of most HPC supercomputers remains uncertain, it typically spans between 5 to 7 years [5]. Thus, once they no longer figure in the list, these systems are either upgraded to meet evolving computational demands (as is the case for *CURIE*), decommissioned and disassembled, or repurposed for other uses.

5.2 Performance efficiency

In this experiment, our aim was to observe closely the performance evolution of the Top500 systems. We focused on the R_{max} and R_{peak} metrics for all the HPC systems on one hand, and the first system from each list on the other hand. Figure 2 shows the evolution of both metrics in *GFLOPS* (y-axis) over time (x-axis).




Figure 2: **The evolution of the Top 1 [Purple lines] supercomputer’s performance metrics (R_{max}, R_{peak}) compared to the average evolution of all the Top500 supercomputer [Green lines] by date, based on Top500 along with Moore’s law projection.**

From this figure, a steady performance climb can be noticed through the years for the listed supercomputers, especially for those at the top. The R_{max} and R_{peak} have been increasing strongly since the beginning of the Top500 ranking. However, this exponential increase started slowing down since the second semester of 2013 where there is a clear breaking point. We compared this growth to Moore’s law projections. This comparison showed that overall, the performance evolution surpasses Moore’s law with an average doubling time of 1.87 years for the R_{max} value instead of 2 years. However, this performance doubling is getting very close to Moore’s projection after the previously mentioned break point.

Another aspect we intended to highlight is the performance efficiency of the supercomputers listed in the Top500. For that, we calculated the ratio between

³<https://www-hpc.cea.fr/fr/Joliot-Curie.html>




Figure 3: The performance ratio average between the Linpack R_{max} and the theoretical R_{peak} over time for the first *Top500* system.

the maximum performance reached using the Linpack benchmark (R_{max}) and the maximum theoretical performance (R_{peak}). The results are plotted in Figure 3 to display the performance percentage (y-axis) evolution over time (x-axis). This graph shows a clear performance decrease in terms of percentage between the first Top500 supercomputer and the latter ones. Other than energy or budget limitations, this decrease in growth and performance efficiency is correlated to the rise of architectural complexity of high performance computing systems that evolved significantly over the years in response to emerging applications requirements and technological improvements.

5.3 Energy efficiency

Our purpose in this section is to highlight the progress made in terms of energy efficiency and to compare it with the performance. We focused on Koomey’s [16] law as a projection reference.




Figure 4: Maximum *Efficiency* of the *Green500* supercomputers by list date along with Koomey’s law projection starting at two different periods [2014 and 2019].

Figure4 illustrates a first experiment where we compare the energy efficiency evolution of HPC supercomputers in GFLOPs/Watt (*y-axis*) to Koomey’s law projections from two starting points (2014 and 2019) over time (*x-axis*).

We observe an increase in the energy efficiency, but at a lower rate than the performance. A huge progress have been made in that matter, with a maximum value increase from 4.5 *GFLOPS/watt* in the late 2013 to 65.39 *GFLOPS/watt* in 2023. However, this progress fails to follow Koomey’s law with a doubling time higher than 1.5 years, especially in the beginning of Green500. Fortunately, it seems to catch up in the last few years due to rising interest for energy efficiency in the HPC community and the emergence of heterogeneous supercomputers that appear to be more energy efficient.

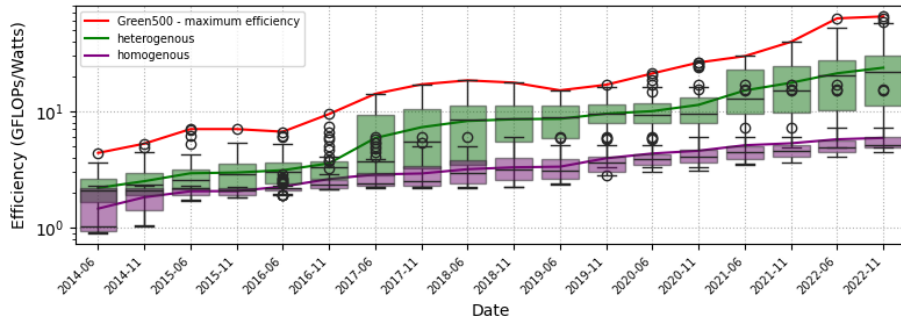


Figure 5: Maximum Energy efficiency growth *Green500* systems by date distinguished by architecture type (homogeneous vs heterogeneous)

To take a closer look on the impact of heterogeneity on the energy efficiency of Green500, we propose a second experiment where we compare the heterogeneous and homogeneous systems in that same list. Our results are displayed in Figure 5 where the y-axis represents the energy efficiency and the x-axis represents time. This figure definitely shows that homogeneous systems are less efficient than heterogeneous ones by an order of magnitude and that the gap between the two categories is increasing. A lot of architectural work has been done to develop energy efficient chips for vector computations. It confirms that using dedicated architectures for certain operations, typically vectorial ones, is a good way to increase efficiency. This is also a trend in lower end systems like smartphones with SoCs including components for data analysis and tensoral computations capabilities. Nonetheless heterogeneous systems are harder to program efficiently than homogeneous ones. Exploiting them to their full potential thus requires significant human expertise that all sites can not afford. We think that this explains the recent decrease of efficiency on average in the Top500 while the top system remains at a very high level as shown in Figure 3.

5.4 Does Top500 really represent real life?

The evolution of R_{max} performance on Top500 is mostly driven by technology improvements, as highlighted in previous figures. One can expect the same evolution on other benchmarks, including less regular ones. Table 2 presents

the improvement ratio over 12 years for various benchmarks. We consider the ranking lists corresponding to the two benchmarks of Top500: Linpack (R_{max}) and HPCG, and two benchmarks of Graph500: BFS and SSSP. For each list, we indicate the increase ratio over 3 periods of 4 years (11/2010-11/2014, 11/2014-11/2018 and 11/2018-11/2022) for the best system of the corresponding list (top 1), for the 5, 15 first systems. As HPCG list and SSSP list appeared in 2017, they are only evaluated over the second period. The list is extended to the 500 systems for only Top500 as Graph500 does not list 500 systems. BFS increase ratio is not given for the first period, as it only contained 10 systems in 2010, and the increase ratio was over 3,000 over the period for the first systems.

			2010-2014	2014-2018	2018-2022
Top500	R_{max} (FLOPS)	top 1	13.20	4.24	7.68
		top 5	11.19	4.72	5.26
		top 15	8.50	4.53	4.75
		top 500	7.07	4.58	3.44
	HPCG (FLOPS)	top 1	-	-	5.47
		top 5	-	-	6.11
		top 15	-	-	5.03
		top 500	-	-	4.99
Graph500	BFS (TEPS)	top 1	-	1.32	3.29
		top 5	-	1.50	1.79
		top 15	-	1.57	1.76
	SSSP (TEPS)	top 1	-	-	11.81
		top 5	-	-	54.18
		top 15	-	-	186.98

Table 2: **Ratio of increase over 4 years periods per list between 2010 and 2022 in Top500 and Graph500.**

A large difference can be observed between the different benchmarks. On the period 2018-2022, the increase ratio for vary between 1.76 for the top 15 of BFS benchmark and 186.98 for the top 15 of SSSP benchmark. It could be explained by the novelty of SSSP in 2018 and a wrong parametrization in the first years. However, BFS and R_{max} also behave differently. Top 1 R_{max} performance increased around 7.5 times during the same period, while BFS increased around 3.3 times only. BFS benchmark is more data-intensive than Linpack. The slowdown of this list compared to R_{max} could be the result of this point. On the contrary, for each list and period, we observe quite similar ratio for top 1, top 5, and top 15. If the best performance globally grows faster than the mean performance as observed in Figure 2 it is not the case for period 2014-2018 and for HPCG. Finally, a slowdown was observed after 2014. This focus on period 2014-2022 shows an acceleration on top systems after 2018. The comparison with Figure 4 highlights the impact of accelerators arrival in HPC.

6 Projection model

6.1 Footprint efficiency

The carbon footprint of HPC electricity consumption was the main objective behind Green500 release. It is however mostly studied for energy efficiency. We consider in this section the carbon footprint corresponding to Green500 power values. We compute the footprint of an HPC system by multiplying its power consumption by the electricity footprint of its country at the corresponding year⁴.

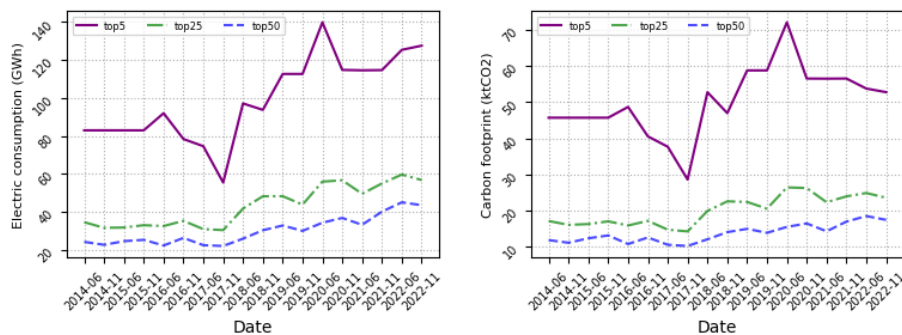


Figure 6: **Electric consumption & carbon footprint for the top 5, 25, 50 in Top500 over time.**

Figure 6 shows the evolution of mean electricity consumption of the first Top500 systems since 2014 and the corresponding mean footprint. It highlights the big difference of the top 5 systems and the following ones. The consumption was almost constant before 2018 and was multiplied by 2 until 2022. The increase is restrained for the footprint due to a reduction of electricity footprint in many countries at the same period. Both curves are very close except for the very recent period due to the energy mix.

Figure 7 shows the evolution of carbon efficiency in Top500, that is the number of FLOP per kg of CO₂. We observe a regular increase of the mean efficiency since 2014 with doubling period of 2.83 years on average. The first Top500 systems are close to the mean while the efficiency of the Green500 increases faster.

6.2 Scenarios

Based on the previous analysis, we provide an estimation of the electricity consumption of the domain in order to compare with the roadmap for decarbonization. Various scenarios exist, for instance those provided by IPCC[12]. The objective of the European Union Green Deal [7] is to reduce the GHG emissions

⁴<https://ourworldindata.org/grapher/carbon-intensity-electricity>

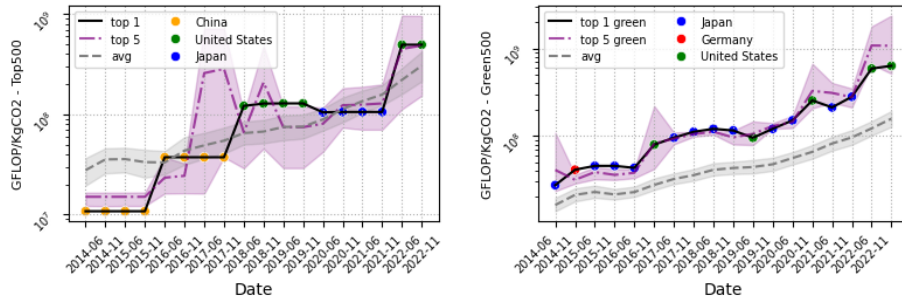


Figure 7: Evolution of the maximum performance in GFLOP per kgCO₂ for the top 1 & 5 systems in the Top500 [Left] and Green500 [right].

by 55% by 2030. We target the 2030 horizon since it is unreasonable to look any further. By extrapolating from the carbon efficiency data displayed in Fig. 7, we estimate the increase of the number of GFLOP per kgCO₂ to reach 1.64×10^9 in 2030 compared to 3.05×10^8 in 2022, representing an increase of 537% between 2022 and 2030 with a slope of 11.87% per issue date, which corresponds to 24.99% per year of the value recorded in 2022. Of course, this is a rough estimate, but our results have shown great stability of this improvement over a long period of time. The gap between this increase and the expected reduction is huge.

7 Conclusion

In this paper, we have studied the evolution of the HPC domain through data from the Top500 and its variants which we consider to be representative of the field. In-depth analysis of both performance and energy efficiency has enabled us to better understand the evolution and highlights several important results: (i) the lifespan of HPC systems in Top500 is in average lower than 2 years and has not evolved in the past decades, (ii) the performance and energy efficiency increase has diminished between 2014 and 2018, but keep increasing and (iii) heterogeneous systems are potentially more energy efficient but require more human expertise to be fully exploited.

In addition, we have proposed a prospective study up to the 2030. Consumption of HPC systems is set to continue rising, as it has been doing for a long time, at a time when GHG emissions need to be drastically reduced and the gap is huge. However, the calculations carried out on these large-scale systems may also help us to propose solutions to mitigate the crisis. Nowadays IPCC experts consider that it is still time to act against climate change. HPC domain needs to participate to the effort by improving its footprint. We hope this study will help estimating the necessary efforts to meet IPCC roadmaps.

8 Acknowledgments

Abdessalam Benhari is supported by the ANRT CIFRE convention 2022/0176. This work was supported by the research program on Edge Intelligence of the Multi-disciplinary Institute on Artificial Intelligence MIAI at Grenoble Alpes (ANR-19-P3IA-0003) and by the REGALE (H2020-JTI-EuroHPC-2019-1 agreement n. 956560) european project.

References

- [1] Graph500 Website, <https://graph500.org/>
- [2] Top500 list from the top500 website. <https://top500.org/>
- [3] ADEME: Evaluation de l'impact environmental du numérique en france et analyse prospective (2023)
- [4] Belady, C.: In the Data Center, Power and Cooling Costs More than the IT Equipment It Supports. *Electron Cool* **13** (2007)
- [5] Benaben, D., Berthoud, F., Guennebaud, G., Ligozat, A.L., Valcke, S.: Estimation de l'empreinte carbone d'une heure de calcul sur un cœur CPU ou sur un GPU. Tech. rep. (2024), <https://labos1point5.org/les-rapports/estimation-empreinte-calcul>
- [6] Dongarra, J., Luszczek, P., Petitet, A.: The LINPACK Benchmark: Past, Present and Future. *Concurrency and Computation: Practice and Experience* **15**, 803–820 (2003)
- [7] European Commission: Delivering the European Green Deal (2021)
- [8] Fraternali, F., Bartolini, A., Cavazzoni, C., Benini, L.: Quantifying the Impact of Variability and Heterogeneity on the Energy Efficiency for a Next-Generation Ultra-Green Supercomputer. *IEEE Trans. on Parallel and Dist. Systems* **29**(7) (2018)
- [9] Freitag, C., Berners-Lee, M., Widdicks, K., Knowles, B., Blair, G.S., Friday, A.: The Real Climate and Transformative Impact of ICT: A Critique of Estimates, Trends, and Regulations. *Patterns* **2**(9), 100340 (2021)
- [10] Gao, Y., Zhang, P.: A Survey of Homogeneous and Heterogeneous System Architectures in High Performance Computing. In: *IEEE International Conf. on Smart Cloud*. pp. 170–175 (2016)
- [11] Hsu, C.H., Feng, W.C.: The Right Metric for Efficient Supercomputing: A Ten-Year Retrospective. In: *2016 IEEE International Parallel and Distributed Processing Symposium Workshops (IPDPSW)*. pp. 1090–1093 (2016)

- [12] Intergovernmental Panel on Climate Change (IPCC): Global Warming of 1.5°C, p. 1–24. Cambridge University Press (2022)
- [13] Kaack, L.H., Donti, P.L., Strubell, E., Kamiya, G., Creutzig, F., Rolnick, D.: Aligning Artificial Intelligence with Climate Change Mitigation. *Nature Climate Change* **12**(6), 518–527 (2022)
- [14] Khan, A., Sim, H., Vazhkudai, S.S., Butt, A.R., Kim, Y.: An Analysis of System Balance and Architectural Trends Based on Top500 Supercomputers. In: *The International Conference on HPC in Asia-Pacific Region. (HPC Asia)*, ACM (2021)
- [15] Kogge, P.M., Dysart, T.J.: Using the TOP500 to Trace and Project Technology and Architecture Trends. In: *Proceedings of 2011 International Conference for HPC, Networking, Storage and Analysis*. pp. 1–11 (2011)
- [16] Koomey, J., Schmidt, Z., Naffziger, S.: Supercomputing Performance and Efficiency: An Exploration of Recent History and Near-Term Projections. A report by Koomey Analytics, completed in collaboration with Advanced Micro Devices, Inc (2019)
- [17] Mair, J., Huang, Z., Evers, D., Chen, Y.: Quantifying the Energy Efficiency Challenges of Achieving Exascale Computing. In: *15th IEEE/ACM International Symposium on Cluster, Cloud and Grid Computing*. pp. 943–950 (2015)
- [18] Masanet, E., Shehabi, A., Lei, N., Smith, S., Koomey, J.: Recalibrating global data center energy-use estimates. *Science* **367**(6481), 984–986 (2020)
- [19] Milojevic, D., Faraboschi, P., Dube, N., Roweth, D.: Future of HPC: Diversifying Heterogeneity. In: *2021 Design, Automation & Test in Europe Conference & Exhibition (DATE)*. IEEE (2021)
- [20] Murphy, R.C., Wheeler, K.B., Barrett, B.W., Ang, J.A.: Introducing the Graph 500. *Cray Users Group (CUG)* **19**, 45–74 (2010)
- [21] Scogland, T., Subramaniam, B., Feng, W.c.: Emerging Trends on the Evolving Green500: Year Three. In: *IEEE International Symposium on Parallel and Distributed Processing Workshops and Phd Forum*. pp. 822–828 (2011)
- [22] Strohmaier, E., Meuer, H.W., Dongarra, J., Simon, H.D.: The TOP500 List and Progress in High-Performance Computing **48**(11), 42–49 (2015)
- [23] Subramaniam, B., Saunders, W., Scogland, T., Feng, W.c.: Trends in energy-efficient computing: A perspective from the Green500. In: *2013 International Green Computing Conference Proceedings*. pp. 1–8 (2013)