



Carbon Topography Representation: Improving Impacts of Data Center Lifecycle

OLIVIER WEPPE, Univ. Rennes, INSA Rennes, CNRS, IETR - UMR 6164, France

DAVID BEKRI, HES-SO, Switzerland

THIBAUT MARTY, Univ. Rennes, INSA Rennes, CNRS, IETR - UMR 6164, France

LOÏC GUIBERT, HES-SO, Switzerland

LOUISE AUBET, Resilio, Switzerland

JEAN-CHRISTOPHE PRÉVOTET, Univ. Rennes, INSA Rennes, CNRS, IETR - UMR 6164, France

MAXIME PELCAT, Univ. Rennes, INSA Rennes, CNRS, IETR - UMR 6164, France

SÉBASTIEN RUMLEY, HES-SO, Switzerland

The globally growing data centers energy consumption and their carbon emissions pose significant environmental challenges. In this context, discerning server fabrication and operational energy contribution to data center carbon footprint is key to identifying effective mitigation strategies. This study partitions server carbon footprints into a) fabrication (embodied carbon), b) static operational power, and c) dynamic operational power, and proposes a novel 2D representation for analyzing data centers carbon impacts. This representation highlights the contributions of these three a-c factors, for any server load and any carbon intensity of electricity. To showcase our methodology and representation, we conducted experimental power measurements on four diverse servers under various load conditions, and combined them with Life Cycle Assessment (LCA) methods for their embodied carbon. Our results show that operational energy generally dominates the total footprint. Indeed, high static power consumption, due to poor energy proportionality in current hardware, is a major carbon emission factor, especially at low loads. We conclude that optimization efforts should follow this sequence: 1) improve server utilization, 2) prioritize low-carbon electricity, 3) maximize server lifetime. Hence, fabrication impact is primarily relevant only when servers are powered by low-carbon electricity. Our representation shows that reducing static power waste through future hardware with better energy proportionality is a priority to design and operate sustainable data centers.

CCS Concepts: • **Hardware** → **Interconnect power issues**; • **Information systems** → **Web services**; • **General and reference** → **Metrics**; **Experimentation**; **Reference works**; • **Computing methodologies** → **Artificial intelligence**.

Additional Key Words and Phrases: Energy, Throughput computing, ICT carbon footprint

1 INTRODUCTION

Data centers are now a central backbone of our societies, providing online services we daily use. As the demand for data, processing and storage continues to surge, these facilities require increasing hardware and energy. However, data centers operation consume vast amounts of energy [4] and contribute to high carbon emissions. Moreover, data center usage is fast growing: they consumed about 4.4 % of total US electricity in 2023 and their energy consumption is expected to reach between 7 % and 12 % of total US electricity by 2028 [27]. The massive adoption of Artificial Intelligence (AI) is one of the main reasons for the observed increase of energy requirements of data centers [6, 29]. Moreover, trends show that bigger AI models lead to increased energy needs [33].

To handle such demand increase, new data centers are built, pre-existing data centers operators renew their servers to offer more

computing performance. Building or expanding data centers requires new hardware acquisition, which also demands significant energy and raw materials for manufacturing, leading to large greenhouse gas emissions [36]. The environmental impact of hardware manufacturing can be estimated through standardized Life Cycle Assessment (LCA) methods, which evaluate the impacts associated with each stage of the lifetime of a hardware component, from raw material extraction and manufacturing, to operation, maintenance, and device disposal or recycling [1].

A concerning aspect of data centers is underused hardware, such as idle servers that remain powered on while awaiting requests, resulting in resource waste, as energy and computational power is consumed without processing any data [2, 3]. These setups not only increase operational costs for data centers through static power but also exacerbate their environmental impact.

This paper aims at better discerning the impact of the fabrication phase of servers from their dynamic and static energy impacts. Analyzing these components is complex, as their relative importance shifts dramatically depending on server utilization and the grid Carbon Intensity of Electricity (CIE). To exhibit the combined influence of these two factors, we propose and utilize a novel 2D representation, called Carbon Topography Representation (source code available [34]). We then feed this Representation with eight series of power measurements, corresponding to four servers and two workloads. The Representation reveals the influence of different operational scenarios on the overall carbon impact, enabling the identification of the dominant factors.

2 DATA CENTER STUDY METHODOLOGY

The closest work to Carbon Topography Representation we are aware of is the *GreenChip* 2D representation that relates a *break-even time* for a device (in months or years) to its *sleep ratio* (share of time during which neither static nor dynamic power are drawn) and its *activity ratio* (share of non-sleep time where dynamic power is drawn, i.e. truly useful share of processing time [16]). Break-even time designates the minimum time to be waiting before replacing a device (e.g. System-on-Chip (SoC) or memory) for the gain in its energy efficiency to compensate its embodied energy i.e. the energy spent to produce the new device. Carbon topography representation is complementary to GreenChip representation in that

it concentrates on carbon footprint and incorporates the notion of grid CIE.

In order to build the carbon topography representation, we rely on direct measurement of power consumption for operational carbon emissions and LCA for embodied carbon emissions.

2.1 Server Energy Consumption Analysis

To evaluate a server operation carbon emissions, which stems from its power consumption, we propose the experimental setup illustrated in Figure 1 that comprises three key components: the Hardware Under Test (HUT), a Load Tester Device (LTD), and a Power Distribution Unit (PDU). The HUT hosts an Service Under Test (SUT) exposing a REST API endpoint that accepts parameters. The load imposed on the HUT can be tuned through the requests sent per unit of time.




Fig. 1. Data center energy efficiency testbed setup.

The HUT can take various forms, provided it meets three essential criteria. First, the HUT must offer a SUT over HTTP/IP as described. Second, it must provide a low-level monitoring tool to track hardware and OS performance metrics. Lastly, the HUT must come with a power cable pluggable in a PDU. Since our focus is on electricity utilization, the HUT power supply must be measurable in real-time. To achieve this, the PDU must integrate a network-accessible power monitoring system, providing automated data collection.

The LTD serves as an experiment controller and operates three core functions: an HTTP client, a HUT metrics collector, and a PDU metrics collector. The HTTP client component simulates user interactions with the SUT by periodically querying the HUT while concurrently gathering performance and power metrics. We use the well-established Grafana k6 tool [17] to precisely control the number of requests concurrently reaching the SUT. As metrics collector, we use Prometheus Node Exporter [26]. We wrote custom code to retrieve power metrics from the PDU, an RXN UDPU® [21].

It is important to note that this setup isolates server-level power measurements, deliberately excluding broader data center consumption to maintain focus on the power consumption of the sole HUT, coherently with the LCA that also captures the HUT exclusively.

We collected our measurements using four servers as HUT, listed by performance in descending order: an ARM-based QuantaGrid S74G-2U (QCT), an Intel-based ProLiant DL360 Gen9 (HPE), an

Intel-based PowerEdge T430 (Dell), and a Raspberry Pi Model 4B, used as a low power platform reference. We implemented two SUT: the *Count service* simply increments a counter until reaching a limit provided as parameter; the *AI service* performs a prediction of the most probable next token of an input sentence. Measurement results are available in full details in [12].

2.2 Embodied Carbon

As a concern for carbon impact, the energy consumed during the utilization of the considered computing device does not reflect the total environmental impact over the lifetime of the product. To account for a wider share of the carbon emissions of a given service, the embodied impact, caused by the fabrication of the device and its end-of-life process, should also be considered [16].

The LCA methodology generates a comprehensive estimate of the environmental impacts of a given product [11]. State-of-the-art methodologies rely on this principle, like the PAIA tool [24] used by manufacturers to provide the Product Carbon Footprint (PCF) associated with their products, or more recently the methodologies described by Malmudin et al. [22] and Li et al. [19].

Because of the obvious hardware scale difference, distinct methodologies are used for the servers and the Raspberry Pi.

2.2.1 Servers. Given the lack of detailed LCAs for every model, we approximate the server embodied carbon (G_{Fab}) by isolating the impact of major computational and storage components ($G_{ICs} + G_{HDD}$) from a baseline representing common elements (G_{Common}): $G_{Fab} = G_{Common} + G_{ICs} + G_{HDD}$

As shown in Table 1, G_{Common} is estimated at 204.9 kg CO₂e using data from the Dell R740 LCA [8] for parts like the chassis, PSUs, and cooling. We thus assume these common contributions are consistent across the servers studied. This approach, while introducing some uncertainty, enables a focused comparison based on the variable high-impact components like processors and memory.

Table 1. Estimated GWP for common server components

Component	GWP (kg CO ₂ e)	Note
Chassis	34.0	From R740 [8]
Mainboard	128.5	From R740 [8], where the share of CPU (26.6%) is removed
PSU	31.3	From R740 [8]
Cooling	11.1	From R740 [8]
Subtotal	204.9	

To account for the various characteristics of the main components (Processors, Memories, Storage) of the servers, a streamlined estimation technique has been used, as defined as $G_{Components} = \sum X_i \cdot K_i$, where X_i are the physical characteristics of components such as the die area or technology node for CPUs and GPUs, or the capacity in GB for memories and storage; each characteristic being associated with a K_i scaling factor that, depending on the type of component, converts this area into embodied carbon equivalent as detailed in [13, 22, 28]. The carbon estimates for the three types of servers is detailed in Table 2.

It should be noted that the SSD is a major contributor to the global GWP of the fabrication for the latest server (QCT) and may require special attention in terms of estimation accuracy. The methodology in [13], which uses the same proportional factor across all SSD technologies, may be too imprecise for this type of device. Improvements in flash technology, particularly 3D NAND, have led to a smaller carbon footprint per GB, as demonstrated in [35]. Using the more precise estimation described in this latter study, a value of 444.9 kg CO₂e is determined. Compared to the 633.6 kg CO₂e used in Table 2, the difference is significant but the total carbon footprint of the server remains in a comparable range.

Table 2. Embodied carbon estimates for sub-elements of computing devices in kg CO₂e

Item	QCT 2024	HPE 2015	Dell 2015	Quantity	K factor
Common	204.90	204.90	204.90	See Table 1	
HDD		96.00		4×1.2 TB	0.02 kg CO ₂ e/GB [28]
SSD	633.60			5.76 TB	0.11 kg CO ₂ e/GB [13]
			286.00	2.6 TB	0.11 kg CO ₂ e/GB [13]
GPU	11.88			8.14 cm ² , 4 nm	1.46 kg CO ₂ e/cm ² [5]
CPU	11.30			7.74 cm ² , 4 nm	1.46 kg CO ₂ e/cm ² [5]
		3.06		3.56 cm ² , 22 nm	0.86 kg CO ₂ e/cm ² [5]
			4.53	2×2.46 cm ² , 14 nm	0.92 kg CO ₂ e/cm ² [5]
RAM	139.20			480 GB, LPDDR5	0.29 kg CO ₂ e/GB [13]
	23.04			96 GB, HBM3	0.24 kg CO ₂ e/GB [13]
		37.12		128 GB, DDR4	0.29 kg CO ₂ e/GB [13]
			18.56	64 GB, DDR4	0.29 kg CO ₂ e/GB [13]
Total	1023.92	345.22	513.99		

As a point of comparison, the study from Loubet et al. [20] includes the LCA of a Dell low-end Server (Dell 3620) and high-end server (Dell 7920) with respective embodied carbon of 767.1 kg CO₂e and 280.8 kg CO₂e, which are similar to our estimates.

2.2.2 Raspberry Pi. The form factor, size and computing power of the single board computer (SBC) being different, the same methodology is not suitable to estimate embodied emissions. State-of-the-Art provides an assessment of the carbon footprint of the Raspberry Pi 4B [20], with an estimation of 14.3 kg CO₂e.

2.3 Carbon Topography Representation

Depending on the load of the server and on the Carbon Intensity of Electricity (CIE) where task execution takes place, the total server carbon footprint varies significantly. To address this disparity, we propose a 2D representation of our results, called Carbon Topography Representation. The x axis abstracts the load imposed on the server, and the y axis represents the quantity of carbon emitted per kWh. The absolute total server carbon footprint in kgCO₂e is displayed in black contour lines, but another aspect is how this carbon is attributable among the three types: Static consumption, Dynamic consumption and Fabrication. Using a color gradient, we superimpose this information in the same figure as explained in Figure 2. We further superimpose white contour lines showing where each type of consumption contributes more than 50 and 75 % in the total Global Warming Potential (GWP) of the HUT.

Figure 3 shows total carbon footprints for the count service (top) and the AI service (bottom) of the four platforms.

When applied to our measurements, the x axis becomes the number of requests per second. For each machine, the service load ranges from 0 (Idle) to the maximum achievable load, beyond which not all requests are being served.

On the y axis is the CIE which ranges from 0 (carbon neutral electricity) to 800 kg CO₂e/kWh (most emitting electricity mixes). Low carbon countries like Switzerland or France are in the 30–60 kg CO₂e/kWh range, while China is around 500 and the US, where most of the largest data centers are located¹, is around 400 with a large variation depending on the state (from 30 to 860)².




Fig. 2. (Left) Ternary color plot used in our Carbon Topography Representation, showing the shares of fabrication (blue), static power (red), and dynamic power (green) in a cloud service's total carbon footprint, across request rates up to the server's maximum capacity. (Right) Electricity carbon intensity (g CO₂e/kWh) in countries with high (e.g., China) and low (e.g., France, Switzerland) carbon electricity.

In this ternary representation form, each point in the 2D graph has a RGB color indicating the relative contribution to the carbon footprint associated with Fabrication (blue component), Static power consumption (red component) and Dynamic (i.e. load-related) consumption (green component). Note that RGB components sums up such that $R + G + B = 100\%$. As a result, on the upper left corner of each graph, the red patch shows the role of the static power usage of the HUT, mostly visible at low load levels; as we move to the upper left to the upper right corner, the color *might* turn to green as the dynamic power increases with load. Finally, blue is mostly visible at the bottom : low CIE makes fabrication dominant in the carbon impact.

The rationale behind this representation is that analysing the complex interplay between server utilisation, grid carbon intensity and the resulting carbon footprint requires a holistic visualisation. While standard charts such as bar graphs can provide a precise breakdown of a single operating point, they cannot easily reveal overall trends and transition boundaries across continuous two-dimensional parameter spaces.

Carbon Topography Representation addresses this issue. Its primary purpose is to function as a 'map' of a server's carbon profile, allowing for the immediate identification of operating 'regimes' where fabrication, static power, or dynamic power is the dominant contributor. The boundaries between these regimes are quantitatively defined by contour lines. Although a ternary color gradient is used as a qualitative guide to identify these regimes, we recognize

¹According to information from <https://www.webopedia.com/technology/10-biggest-data-centers-in-the-world/> 2025-04-01

²According to data available from <https://app.electricitymaps.com/map/12mo/monthly> 2025-04-01




Fig. 3. Carbon Topography Representation of server carbon footprint for the count (top) and AI (bottom) services. Total carbon footprint is broken down into Static power (Red), Dynamic power (Green) and Fabrication (Blue). White contour lines provide break-even points (50%), and black contour lines gives the total carbon footprint of the service in absolute value over its lifespan.

its limitations for precise interpretation and for viewers with color vision deficiencies. Therefore, the key analytical features of this representation are the contour lines that delineate where each impact category contributes more than 50% and 75% of the total, providing a clear and accessible view of the trade-offs. This representation is designed to provide high-level strategic insights, which can then be complemented by detailed quantitative analysis of specific points of interest, as demonstrated by the bar charts in Section 3.

3 PARAMETER DEPENDENCIES OF TOTAL SERVER CARBON FOOTPRINTS

3.1 Insufficient Power Load Scaling

The representation underlines significant variations in carbon attribution profiles across different servers. Indeed, the share of dynamic power consumption is very different from one type of device to another, as shown by the distribution of the red and green areas. Even for HPE and Dell servers in Figure 3, which share the same x86-64 architecture and belong to a similar technological generation, the operational phase exhibits different energy efficiency patterns and can be caused by various reasons [32]. Notably, the red color dominance in the representation shows that load does not substantially alter power usage for QCT and Dell servers, especially for the count service. Similarly, the Raspberry Pi platform demonstrates only a moderate load-dependent effect on its carbon footprint. The HPE server, however, displays a pronounced dynamic power contribution, exceeding 50% under high load, especially for the AI service.

This shows that the 2007 observation by Barroso and Hölzle [2] that servers, unlike handheld devices whose power consumption

scales dynamically with computational demand, are primarily optimised for peak performance, still seems valid. While some applications operate continuously near saturation, other services experience significant workload fluctuations. In order to satisfy customer demand despite the cost and environmental impact, hardware infrastructures to deliver those services are often designed for the worst case, i.e. for the highest demand. Apart from these moments where the load is maximal, in our representations, the dominance of the red areas (static operational carbon) directly reflects the high static power consumption measured for these servers.

3.2 Carbon Optimization Strategies

Among available strategies to reduce the carbon footprint of a digital service is the prolongation of the computing device in order to spread the impact of fabrication over a larger period of time [7, 30]. Using our representation, as Figure 3 shows, apart from locations where CIE is low (below 100 g CO₂e/kWh), fabrication (blue areas) does not represent a major GWP contributor for the studied servers. This means that lifetime extension is not an efficient ecode-sign strategy when using highly carbonated electricity, where GWP is dominated by the impact of electricity usage (red and green areas).

3.2.1 Lower the CIE. Instead, a first priority in carbon-intensive electricity locations is to reduce the CIE [10]. Figure 5a exhibits the breakdown of GWP for a given server (HPE) and service (AI inference) at various loads for different hardware lifetime values. The results have been attributed per request (in µg CO₂e) to reflect the environmental cost of the service usage.

The figure shows that even with moderately carbonated electricity (400 g CO₂e/kWh), electricity consumption is the main contributor to GWP of the service, whatever the load of the server. When the




Fig. 4. Focus on low carbon electricity scenarios (0 to 100 g CO₂e/kWh for QCT (left) and HPE (right) servers with hardware lifetime of 5 and 10 years. At low CIE, the service GWP is highly impacted by hardware fabrication (blue). Extending the lifetime to 10 years mitigates the contribution of manufacturing in the total carbon footprint.

server is close to idle state, static power drains a major share of GWP. On this same figure, the influence of the extension of the lifetime of the hardware device is negligible. This is due to the low contribution of fabrication in the total GWP.

3.2.2 Device Lifetime. The impact of low carbon electrical energy mix is illustrated in Figure 5b, showing the breakdown of GWP among fabrication, static and dynamic power the same HPE server running the same service with a lower CIE (50 g CO₂e/kWh). As shown by absolute values, switching to an 8× less carbonated electricity reduces the impact of each request by at least a factor 4. The share of fabrication in the total GWP is raised, and extending server lifetime becomes a relevant solution to decrease the carbon footprint of the service, which is visible in Figure 5b with a reduction between 17 and 27 % per request when extending the lifetime of the server from the original 5 years to 10 years.

This is also illustrated in Figure 4, which focuses on CIE below 100 g CO₂e/kWh, showing the dominance of fabrication contribution to the total carbon footprint of devices. Furthermore, it shows how device lifespan affects the environmental impact of the service, enabling a side-by-side comparison of a 5-year lifespan to 10-year scenario.

3.2.3 Device Efficiency. Relying on more energy-efficient platforms presents another option to diminish service carbon emissions. While newer servers might intuitively offer better efficiency, as suggested by Dennard scaling [9], this is not always true under all operating conditions. For instance, transitioning from older HPE (2015) to recent QCT (2024) hardware at the same request rate led to an increased GWP (Figure 5c). Although QCT servers provide higher peak operations per second and demonstrate superior efficiency at full load (Figure 5d), their significant static power consumption renders them inefficient under low utilization. Thus, realizing the benefits of such newer devices mandates operation near maximal capacity. Moreover, virtualizations techniques and Cloud capacities can allow better hardware usage factors, thus reducing energy consumption [14, 15]. The concept of a carbon break-even point [13, 16] helps navigate the trade-off between this load-dependent operational efficiency and the embodied carbon of new devices, enabling the determination of an optimal lifetime for maximal carbon efficiency.

4 FUTURE DIRECTIONS

The findings from this study need to be consolidated, particularly with regard to the hardware platforms used in our experiments. To cover a wide range of machines, we applied our methodology to recent, powerful hardware (QCT), legacy standard servers (Dell and HPE), and low-power platforms (Raspberry Pi). We also used different hardware architectures (ARM for QCT and Raspberry Pi, and x86_64 for Dell and HPE). However, it would be insightful to run the experiment on a larger variety of recent platforms to confirm that our findings are relevant to contemporary servers.

The concern for high idle power should be investigated further, particularly in relation to system configuration. Identifying the parameters that lead to better power optimisation could help to reduce the carbon footprint of servers. Particular care should be taken to consider the impact of power-saving tuning on responsiveness when switching from idle to processing mode when new requests are received, since deeper sleep states can increase the delay required to start processing a new task.

Our representation focuses only on the carbon footprint of a service. This is due to the current strong focus on global warming in the academic world and in the industry. Despite the concerning aspects of global warming, other environmental aspects must be studied. For example, an interest is growing about water usage [23]. Large quantities of water are used for electronic devices fabrication, but also for electricity production. The same type of study could be developed on the topic of water. However, contrary to carbon impact which has been well documented in the last decade, information about water usage is scarce, which makes the investigation challenging. This issue is even more pronounced regarding minerals usage, where almost no public information is disclosed.

While this analysis focuses on the environmental impacts of computation within cloud services, these services depend on more than just CPU platforms. Other contributors to the overall environmental footprint of data centers include: cooling systems [18], which consume a substantial fraction of total energy and often utilize fluorinated gases [31] with high GWP; data storage, where the exponential growth in capacity demand, particularly for SSDs [35], increases the silicon manufacturing impact; and the network infrastructure required, which also adds to operational energy consumption [25].




Fig. 5. Optimizations strategies for AI service: the bars represent the absolute carbon footprint in $\mu\text{g CO}_2\text{e}$ broken down into Fabrication(Blue), Static(Red) and Dynamic(Green) power for server lifetime ranging from 5 to 10 years at 5, 50 and 100% loads. Starting from baseline (a), using lower carbon source of electricity (b) provides the most significant gains; using a more powerful and recent server (c) only provide improvement when used at full capacity (d).

As illustrated by our measurements, there is a need to mitigate energy consumption of data centers during periods of low utilization. One solution involves refining techniques for dynamically powering down idle servers or components, considering the potential impacts on service responsiveness. Furthermore, exploring the adoption of more energy-proportional hardware architectures should be considered. Processors inspired by mobile device designs, which often scale power consumption more effectively with computational load than traditional server CPUs, could offer significant energy savings if adapted successfully for data center workloads and performance requirements.

5 CONCLUSION

The carbon topography representation proposed in this paper provides a valuable tool for analyzing server carbon footprints, partitioning impacts from fabrication, static operational power, and dynamic load-based power using experiments and LCA across various loads and grid carbon intensities. Using this representation, we find that operational energy, particularly high static consumption due to limited energy proportionality in current hardware, dominates the footprint unless powered by low-carbon electricity. Consequently, accessing low-carbon grids is the primary mitigation strategy to be deployed on existing servers. The representation framework illustrates the conditions under which hardware lifetime extension becomes relevant: impacts and lifetime extension become significant factors only under such low-carbon conditions. However, the prevalence of static power waste, especially at low loads, highlights a need for future server hardware with improved energy proportionality and dynamic power scaling, potentially inspired by mobile architectures. Achieving sustainable digital infrastructure fundamentally requires advancements in hardware efficiency alongside decarbonized energy and optimized life-cycles.

ACKNOWLEDGMENTS

This study is funded by the France2030 ESOS Project - Electronics: Sustainable, Open and Sovereign, operated by French ANR (Agence Nationale de la Recherche).

REFERENCES

- [1] Yevgeniya Arushanyan, Elisabeth Ekener-Petersen, and Göran Finnveden. 2014. Lessons learned – Review of LCAs for ICT products and services. *Computers in Industry* 65, 2 (Feb. 2014), 211–234. <https://doi.org/10.1016/j.compind.2013.10.003>
- [2] Luiz André Barroso and Urs Hölzle. 2007. The Case for Energy-Proportional Computing. *Computer* 40, 12 (Dec. 2007), 33–37. <https://doi.org/10.1109/MC.2007.443>
- [3] Amlan Deep Borah, Deboraj Muchahary, Sandeep Kumar Singh, and Janmoni Borah. 2015. Power saving strategies in green cloud computing systems. *International Journal of Grid Distribution Computing* 8, 1 (2015), 299–306.
- [4] Rajkumar Buyya, Shashikant Ilager, and Patricia Arroba. 2024. Energy-efficiency and sustainability in new generation cloud computing: A vision and directions for integrated management of data centre resources and workloads. *Software: Practice and Experience* 54, 1 (2024), 24–38. <https://doi.org/10.1002/spe.3248> <https://onlinelibrary.wiley.com/doi/pdf/10.1002/spe.3248>
- [5] Interuniversity Microelectronics Centre. [n. d.]. Methodology - imec.netzero. <https://netzero.imec-int.com/methodology>
- [6] Sophia Chen. 2025. How much energy will AI really consume? The good, the bad and the unknown. *Nature* 639, 8053 (March 2025), 22–24. <https://doi.org/10.1038/d41586-025-00616-z> Bandiera_abtest: a Cg_type: News Feature Publisher: Nature Publishing Group Subject_term: Machine learning, Computer science, Sustainability, Energy.
- [7] Tingyu Cheng, Gregory D. Abowd, Hyunjooh Oh, and Josiah Hester. 2023. Transient Internet of Things: Redesigning the Lifetime of Electronics for a More Sustainable Networked Environment. In *Proceedings of the 2nd Workshop on Sustainable Computer Systems (HotCarbon '23)*. Association for Computing Machinery, New York, NY, USA, 1–8. <https://doi.org/10.1145/3604930.3605723>
- [8] DELL. 2019. Life cycle assessment of Dell R740.
- [9] Robert H. Dennard, Jin Cai, and Arvind Kumar. 2018. A Perspective on Today's Scaling Challenges and Possible Future Directions*. In *Handbook of Thin Film Deposition (Fourth Edition)*, Krishna Seshan and Dominic Schepis (Eds.). William Andrew Publishing, 3–18. <https://doi.org/10.1016/B978-0-12-812311-9.00001-3>
- [10] Jesse Dodge, Taylor Prewitt, Remi Tachet des Combes, Erika Odmark, Roy Schwartz, Emma Strubell, Alexandra Sasha Luccioni, Noah A. Smith, Nicole DeCaro, and Will Buchanan. 2022. Measuring the Carbon Intensity of AI in Cloud Instances. In *Proceedings of the 2022 ACM Conference on Fairness, Accountability, and Transparency (FAccT '22)*. Association for Computing Machinery, New York, NY, USA, 1877–1894. <https://doi.org/10.1145/3531146.3533234>
- [11] International Organization for Standardization. [n. d.]. ISO 14040:2006. <https://www.iso.org/standard/37456.html>
- [12] Loïc Guibert, Sébastien Reynaud, Olivier Weppe, Ciko Kristjon, Michael Welzl, and Sébastien Rumley. 2025. TRAFFIC: Testbed for Assessing energy efficiency in throughput Computing. In *Proceedings of the 21st ACM International Conference on Computing Frontiers (Cagliari, Italy) (CF '25)*. Association for Computing Machinery, New York, NY, USA. <https://doi.org/10.1145/3706594.3727961>
- [13] Udit Gupta, Young Geun Kim, Sylvia Lee, Jordan Tse, Hsien-Hsin S. Lee, Gu-Yeon Wei, David Brooks, and Carole-Jean Wu. 2020. Chasing Carbon: The Elusive Environmental Footprint of Computing. <http://arxiv.org/abs/2011.02839> arXiv:2011.02839 [cs].

- [14] Yichao Jin, Yonggang Wen, and Qinghua Chen. 2012. Energy efficiency and server virtualization in data centers: An empirical investigation. In *2012 Proceedings IEEE INFOCOM Workshops*. 133–138. <https://doi.org/10.1109/INFOCOMW.2012.6193474>
- [15] Avita Katal, Susheela Dahiya, and Tanupriya Choudhury. 2022. Energy efficiency in cloud computing data centers: a survey on software technologies. *Cluster Computing* 26, 3 (8 2022), 1845–1875. <https://doi.org/10.1007/s10586-022-03713-0>
- [16] Donald Kline, Nikolas Parshook, Xiaoyu Ge, Erik Brunvand, Rami Melhem, Panos K. Chrysanthis, and Alex K. Jones. 2019. GreenChip: A tool for evaluating holistic sustainability of modern computing systems. *Sustainable Computing: Informatics and Systems* 22 (June 2019), 322–332. <https://doi.org/10.1016/j.suscom.2017.10.001>
- [17] Grafana Labs. [n. d.]. Grafana Labs k6 load tester. <https://www.k6.io>. Accessed: 2025-05-19.
- [18] Imran Latif, Muhammad Mubashar Ashraf, Umaima Haider, Gemma Reeves, Alexandrina Untaroiu, Fábio Coelho, and Denis Browne. 2025. Advancing Sustainability in Data Centers: Evaluation of Hybrid Air/Liquid Cooling Schemes for IT Payload Using Sea Water. *IEEE Transactions on Cloud Computing* 13, 1 (Jan. 2025), 184–197. <https://doi.org/10.1109/TCC.2024.3521666>
- [19] Yueying Li and Udit Gupta. 2024. Towards Carbon-efficient LLM Life Cycle. In *Workshop on Sustainable Computer Systems*. Cmbbridge, MA.
- [20] Philippe Loubet, Adrien Vincent, Annabelle Collin, Corinne Dejous, Anthony Ghiotto, and Christophe Jégo. 2023. Life cycle assessment of ICT in higher education: a comparison between desktop and single-board computers. *The International Journal of Life Cycle Assessment* 28, 3 (March 2023), 255–273. <https://doi.org/10.1007/s11367-022-02131-z>
- [21] Riedo Networks Ltd. [n. d.]. Riedo networks ltd Universal power distribution units. <https://www.rnx.ch/updu>. Accessed: 2025-05-19.
- [22] Jens Malmodin and Nina Lövehagen. 2024. A Methodology for Simplified LCAs of Electronic Products. In *2024 Electronics Goes Green 2024+ (EGG)*. 1–12. <https://doi.org/10.23919/EGG62010.2024.10631258>
- [23] David Mytton. 2021. Data centre water consumption. *npj Clean Water* 4, 1 (Feb. 2021), 1–6. <https://doi.org/10.1038/s41545-021-00101-w> Publisher: Nature Publishing Group.
- [24] Elsa Olivetti and Randolph Kirchain. 2012. A Product Attribute to Impact Algorithm to Streamline IT Carbon Footprinting. In *Design for Innovative Value Towards a Sustainable Society*, Mitsutaka Matsumoto, Yasushi Umeda, Keiji Masui, and Shinichi Fukushima (Eds.). Springer Netherlands, Dordrecht, 747–749.
- [25] Anne-Cécile Orgerie, Marcos Dias De Assuncao, and Laurent Lefevre. 2014. A survey on techniques for improving the energy efficiency of large-scale distributed systems. *Comput. Surveys* 46, 4 (April 2014), 1–31. <https://doi.org/10.1145/2532637>
- [26] Prometheus. [n. d.]. *GitHub - prometheus/node_exporter: Exporter for machine metrics*. https://github.com/prometheus/node_exporter Accessed: 2025-02-21.
- [27] Arman Shehabi, Sarah Josephine Smith, Alex Hubbard, Alexander Newkirk, Nuoa Lei, Md AbuBakar Siddik, Billie Holecek, Jonathan G Koomey, Eric R Masanet, and Dale A Sartor. 2024. *2024 United States Data Center Energy Usage Report*. Technical Report. Lawrence Berkeley National Laboratory. <https://doi.org/10.71468/P1WC7Q>
- [28] Swamit Tannu and Prashant J. Nair. 2022. The Dirty Secret of SSDs: Embodied Carbon. <https://doi.org/10.1145/3630614.3630616>
- [29] The International Energy Agency. 2025. Energy and AI. <https://www.iea.org/reports/energy-and-ai>
- [30] Amanda Tomlinson and George Porter. 2023. Something Old, Something New: Extending the Life of CPUs in Datacenters. *SIGENERGY Energy Inform. Rev.* 3, 3 (Oct. 2023), 59–63. <https://doi.org/10.1145/3630614.3630625>
- [31] Wen-Tien Tsai and Chi-Hung Tsai. 2023. A Survey on Fluorinated Greenhouse Gases in Taiwan: Emission Trends, Regulatory Strategies, and Abatement Technologies. *Environments* 10, 7 (July 2023), 113. <https://doi.org/10.3390/environments10070113> Number: 7 Publisher: Multidisciplinary Digital Publishing Institute.
- [32] Joakim von Kistowski, Hansfried Block, John Beckett, Cloyce Spradling, Klaus-Dieter Lange, and Samuel Kounev. 2016. Variations in CPU Power Consumption. In *Proceedings of the 7th ACM/SPEC on International Conference on Performance Engineering (Delft, The Netherlands) (ICPE '16)*. Association for Computing Machinery, New York, NY, USA, 147–158. <https://doi.org/10.1145/2851553.2851567>
- [33] Alex de Vries. 2023. The growing energy footprint of artificial intelligence. *Joule* 7, 10 (Oct. 2023), 2191–2194. <https://doi.org/10.1016/j.joule.2023.09.004> Publisher: Elsevier.
- [34] Olivier WEPPE. 2025. Python source code used to generate figures. <https://gitlab.insa-rennes.fr/esos/carbontopo.git>. Accessed: 2025-06-21.
- [35] Olivier Weppe, Thibaut Marty, Sébastien Toussaint, Nicolas Brusselmans, Jean-Christophe Prévotet, Jean-Pierre Raskin, and Maxime Pelcat. 2025. Embodied Carbon Footprint of 3D NAND Memories. In *Inter-national Conference on Computing Frontiers (CF'25)*. Cagliari (Sardinia), France. <https://doi.org/10.1145/3706594.3727962>
- [36] Eric Williams. 2011. Environmental effects of information and communications technologies. *Nature* 479, 7373 (Nov. 2011), 354–358. <https://doi.org/10.1038/>