

# The case for accurate lifetime accounting in carbon metrics

Yujin Huang  
The Pennsylvania State  
University  
yqh5547@psu.edu

Timothy Zhu  
The Pennsylvania State  
University  
tuz68@psu.edu

Anshul Gandhi  
Stony Brook University  
anshul@cs.stonybrook.edu

## ABSTRACT

To represent the entire carbon footprint of computing devices, carbon metrics often include both an embodied cost (i.e., carbon cost to produce the device) and an operational cost (i.e., carbon cost to run the device). The embodied carbon cost is typically high, but it is amortized over the lifetime of the device. In this vision statement, we argue that for carbon metrics to be useful, we need (i) accurate metrics for lifetime, which are challenging for SSDs, and (ii) correct reasoning about carbon costs when using such metrics.

## 1. INTRODUCTION

Sustainable computing requires considering the environmental impact of the whole lifecycle of computing equipment. Research has shown that the amount of carbon generated to produce computers (embodied carbon) can be high compared to the operational carbon costs from running the computer during its lifetime, depending on how long it is used [6, 7, 9]. As we move to renewable energy, operational carbon costs will continue to shrink, so it becomes increasingly important to optimize for these embodied carbon costs.

Prior research has shown the importance of carefully considering the specific carbon metric and usage of the metric—e.g., for scheduling purposes, embodied carbon is a sunk cost that should not functionally impact scheduling decisions, but for procurement decisions, embodied carbon remains an important factor [1]. Typically, these carbon metrics take the form  $\frac{\text{embodied carbon}}{\text{expected lifetime}} + \text{operational carbon}$ , which incorporates how embodied carbon is amortized by the lifetime of the device [5, 4]. We argue that for any such carbon metric to be useful, we need to (i) have accurate metrics for lifetime, and (ii) exercise caution when applying such metrics.

**Context:** In this work, we focus our attention to lifetime/aging characteristics of SSDs, which are designed to wear out from usage; every write-erase cycle will damage the NAND flash cells, lowering their ability to retain data. As reported by Microsoft, Facebook, and Alibaba, a substantial percentage of SSDs (as much as 34%) fail in data-centers, thereby contributing to a significant percentage of server failures (about 5.6%) [13, 12, 16]. Coupled with the especially high embodied carbon costs of SSDs [15], there is a significant impact on carbon metrics and sustainability. Accurately tracking the impact in a SSD carbon metric depends on how we measure and account for their lifetime and

|                            | Switch after 2 years                                    | Switch after 7 years                                    |
|----------------------------|---|---|
| $X + Z$<br>total cost      | $35 + 2 \times 7 +$<br>$27 + 7 \times 4 = \mathbf{104}$ | $35 + 7 \times 7 +$<br>$27 + 7 \times 4 = \mathbf{139}$ |
| $Y + Z$<br>total cost      | $29 + 2 \times 8 +$<br>$27 + 7 \times 4 = \mathbf{100}$ | $29 + 7 \times 8 +$<br>$27 + 7 \times 4 = \mathbf{140}$ |
| $\frac{X+Z}{\text{years}}$ | $\frac{104}{9} \approx \mathbf{11.56}$                  | $\frac{139}{14} \approx \mathbf{9.93}$                  |
| $\frac{Y+Z}{\text{years}}$ | $\frac{100}{9} \approx \mathbf{11.11}$                  | $\frac{140}{14} = \mathbf{10}$                          |

**Table 1: Motivating example: to switch or not to switch.**

age. Lifetimes are impacted by a variety of factors such as workload intensity, SSD technology (e.g., SLC, MLC, TLC, QLC, PLC), and wear-leveling algorithms, so it is critical to develop appropriate SSD lifetime metrics that can be incorporated into a carbon metric.

**Motivating example:** Consider an SSD X with an embodied carbon cost of 35 and an operational carbon cost of 7/year, and an SSD Y with an embodied carbon cost of 29 and an operational carbon cost of 8/year; we omit units here for simplicity. If we have an expected lifetime of 5 years for both SSDs, then X has a combined carbon cost of 14/year whereas Y has a combined carbon cost of 13.8/year. Now suppose we were too conservative, and the actual lifetime is 7 years for both SSDs. Then X has a combined carbon cost of 12/year, whereas Y has a combined carbon cost of 12.14/year. The most carbon-efficient SSD to procure clearly depends on the value of the lifetime metric.

Suppose after 2 years, a new SSD Z is available on the market with an embodied carbon cost of 27 and an operational carbon cost of 4/year. If Z has the same lifetime of 7 years, then its combined carbon cost of 7.86/year is much better than X and Y. In this case, a question now arises whether we should switch all our existing X and Y SSDs with Z. Table 1 shows the total carbon cost if switching immediately (once it is available after 2 years) or if switching after using X and Y over their lifetime (of 7 years). One might be tempted to normalize by the number of years and compare, which would indicate to continue using X and Y for their full lifespan. However, this is an incorrect analysis arising from a sunk cost fallacy [1]. The right way to reason about this metric is to consider the embodied carbon and the 2-year operational carbon of X and Y as sunk costs. We just consider their operational carbon and compare that to Z. But since Z has not been procured, the embodied carbon of Z is not a sunk cost and we need to factor it in along with its expected lifetime. Z’s combined carbon cost of 7.86/year is better than Y’s operational carbon of 8/year, so it makes

|  |
|--|
| Power On Hours   |
| Manufacturing Date                                     |
| Percentage Used  |
| Data Units Written (Total number of bytes written)     |
| Host Write Commands (Total number of write operations) |
| Available Spare  |
| Critical Warning                                       |
| Media and Data Integrity Errors                        |

Table 2: SSD age metrics.

sense to replace Y; but Z’s combined cost is not better than X’s operational carbon of 7/year, so it makes sense to continue using X. Thus, we argue that it is *important to accurately account for lifetime both from the perspective of the carbon metric and the SSD characteristics.*

**Challenges:** Defining an SSD lifetime/age metric is challenging because there is no single factor for SSD age. There are several metrics that represent the physical passing of time (e.g., power-on hours, manufacturing date), but they do not account for the workload. Product specifications often indicate a total write endurance (i.e., limit on total number of bytes written), but it is unclear whether many small random writes can cause as much damage as large sequential writes. A metric from SSD SMART data that seems most relevant is the “Percentage Used” metric [3], but our preliminary results show how this metric may not behave as expected.

Another challenge in defining an SSD lifetime/age metric is accounting for any end-of-life performance degradation. A common assumption in carbon metrics is that the performance of a device is constant over its lifetime, and if performance starts to degrade, it is unclear whether it makes sense to continue using the device. Thus, it is important to appropriately account for any age-related performance degradation in carbon metrics to ensure optimal decisions are made for deprecating hardware. For SSDs, prior work has suggested that SSDs that have experienced a lot of I/O activity are known to have worse performance compared to a pristine SSD [10]. We have not seen this effect in our preliminary results yet, so further research is warranted to understand if and why these behaviors occur. We propose to further study the impacts of various aging metrics on SSD lifetimes and performance.

## 2. PRELIMINARY RESULTS

To better understand factors that influence SSD lifetimes, our preliminary results study the effects of workloads on SSD aging. Unfortunately, prior work has not settled on a single metric for SSD age. Some works consider the number of months of operation [14, 11]. On the other hand, device manufacturer’s often quote lifetime guarantees in terms of total bytes written [2]. To accurately represent SSD age, we propose that a multi-dimensional metric be used.

Table 2 shows the metrics we have identified that are most relevant to the age of SSDs. Most of these metrics are easily accessible via SMART data [3]. While these metrics are correlated with each other, our preliminary results indicate that their relationship is not linear, and they change in unexpected ways based on the SSD workload and usage patterns.

In our preliminary results, we run a variety of workloads on a NVMe SSD (Crucial T705 1TB) and record age and

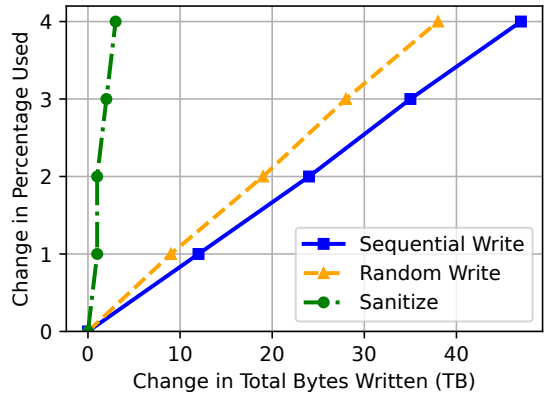


Figure 1: Number of bytes written to change the Percentage Used metric varies based on workload type.

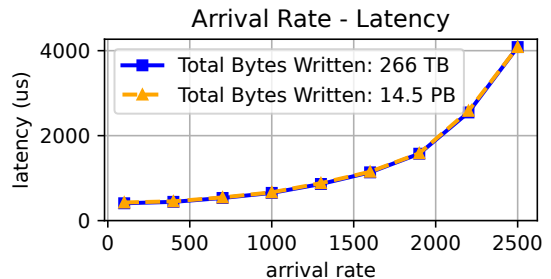


Figure 2: Write latency performance is similar when the Total Bytes Written metric is less than the Expected Endurance (600 TB) and more than 24 times the Expected Endurance.

performance metrics over time. Figure 1 shows the growth of the “Percentage Used” metric relative to the change in “Data Units Written” for a random write workload vs. a sequential write workload. We also include a workload where we periodically run a sanitize operation to wipe all the SSD contents. We see that, surprisingly, using the sanitize operation causes the “Percentage Used” metric to drastically increase. Even though the sanitize operation is quick (i.e., a few seconds), it supposedly wears out the device as fast as hours of continual sequential writes. It is unclear whether the rapid change in the Performance Used metric is an inaccuracy in how the SSD is internally tracking wear-out, or if the sanitize operation is actually degrading the SSD. Further research is necessary to understand the effects and implications of this behavior.

We also see that the growth rate of the “Percentage Used” metric depends on the workload. The random write workload increases this metric with fewer bytes written. Our results demonstrate that the *type* of workload has an impact on the various SSD aging metrics, suggesting that a single metric may not be sufficient to represent the SSD aging process.

Figure 2 shows the SSD performance when the “Total Bytes Written” metric is at 266 TB and 14.5 PB. Although

the manufacturer specifies an expected endurance of 600 TB Total Bytes Written, we see negligible difference between SSD's performance when it is within the expected endurance and after it has far exceeded the endurance. This indicates either "Total Bytes Written" is not the right metric to reflect the age of an SSD or SSDs may last much longer than expected. We also see that the "Available Spare" metric is 100% in both cases and the "Critical Warning" and "Media and Data Integrity Errors" metrics are both 0, indicating no signs of reliability degradation either. This motivates *potentially utilizing SSDs for longer lifespans than according to spec sheets, which would significantly impact any associated carbon metric.*

### 3. NEXT STEPS AND CONCLUSION

Based on our preliminary results, there are 3 immediate directions we plan to pursue as next steps.

1. We plan to perform a longitudinal study of SSD aging to study the effects of various workloads on SSD lifetimes. Specifically, we are interested in any performance degradation that may occur as SSDs start to wear out. Although our preliminary results do not show any performance degradation yet, we anticipate effects to appear with further device aging.
2. We plan to study how a multi-dimensional age metric could enable companies to utilize SSDs for longer times beyond what some metrics may indicate for an expected lifetime. Longer lifetimes will amortize the high embodied cost in carbon metrics, thereby potentially changing procurement decisions.
3. We plan to develop a framework for reasoning about lifetimes and carbon metrics. Given how SSD age/lifetimes are dependent on many factors such as workload, SSD technology, wear leveling algorithms, etc., we will show how expected lifetime should not be a constant number, but rather a function based on these factors. Furthermore, carbon metrics often incorporate a performance metric to represent the speed of the device. We believe that future work is needed to transform the performance metric from a constant number to a function based on age. Based on existing work on carbon depreciation models [8], we plan to develop a comprehensive carbon metric that accurately accounts for the age and lifetime characteristics of devices such as SSDs, which will help optimize sustainability decisions.

### ACKNOWLEDGMENTS

This work was supported by NSF grants CCF-2324858, CCF-2324859, CNS-2214980, and CNS-2106434.

### REFERENCES

- [1] N. Bashir, V. Gohil, A. B. Subramanya, M. Shahrad, D. Irwin, E. Olivetti, and C. Delimitrou. The sunk carbon fallacy: Rethinking carbon footprint metrics for effective carbon-aware scheduling. In *Proceedings of the 2024 ACM Symposium on Cloud Computing, SoCC '24*, page 542–551, New York, NY, USA, 2024. Association for Computing Machinery.
- [2] Crucial. Crucial T705 1TB PCIe Gen5 NVMe M.2 SSD with heatsink. <https://www.crucial.com/ssd/t705/ct1000t705ssd5>, 2025.
- [3] Crucial. SSDs and SMART Data. <https://www.crucial.com/support/articles-faq-ssd/ssds-and-smart-data>, 2025.
- [4] G. S. Foundation. Software Carbon Intensity (SCI) Specification — sci.greensoftware.foundation. <https://sci.greensoftware.foundation/>, 2025.
- [5] A. Gandhi, K. Ghose, K. Gopalan, S. Hussain, D. Lee, D. Liu, Z. Liu, P. McDaniel, S. Mu, and E. Zadok. Metrics for Sustainability in Data Centers. In *Proceedings of the 1st Workshop on Sustainable Computer Systems Design and Implementation, HotCarbon '22*, San Diego, CA, USA, 2022.
- [6] U. Gupta, M. Elgamal, G. Hills, G.-Y. Wei, H.-H. S. Lee, D. Brooks, and C.-J. Wu. Act: Designing sustainable computer systems with an architectural carbon modeling tool. In *Proceedings of the 49th Annual International Symposium on Computer Architecture, ISCA '22*, page 784–799, New York, NY, USA, 2022. Association for Computing Machinery.
- [7] U. Gupta, Y. Kim, S. Lee, J. Tse, H. S. Lee, G. Wei, D. Brooks, and C. Wu. Chasing carbon: The elusive environmental footprint of computing. In *2021 IEEE International Symposium on High-Performance Computer Architecture (HPCA)*, pages 854–867, Los Alamitos, CA, USA, mar 2021. IEEE Computer Society.
- [8] S. Ji, Z. Yang, A. K. Jones, and P. Zhou. Towards datacenter environmental sustainability using carbon depreciation models, 2025.
- [9] A. K. Jones, Y. Chen, W. O. Collinge, H. Xu, L. A. Schaefer, A. E. Landis, and M. M. Bilec. Considering fabrication in sustainable computing. In *2013 IEEE/ACM International Conference on Computer-Aided Design (ICCAD)*, pages 206–210, 2013.
- [10] M. Jung and M. Kandemir. Revisiting widely held ssd expectations and rethinking system-level implications. In *Proceedings of the ACM SIGMETRICS/International Conference on Measurement and Modeling of Computer Systems, SIGMETRICS '13*, page 203–216, Pittsburgh, PA, USA, 2013.
- [11] S. Maneas, K. Mahdavian, T. Emami, and B. Schroeder. Operational characteristics of SSDs in enterprise storage systems: A Large-Scale field study. In *20th USENIX Conference on File and Storage Technologies (FAST '22)*, pages 165–180, Santa Clara, CA, 2022.
- [12] J. Meza, Q. Wu, S. Kumar, and O. Mutlu. A large-scale study of flash memory failures in the field. In *Proceedings of the 2015 ACM SIGMETRICS International Conference on Measurement and Modeling of Computer Systems, SIGMETRICS '15*, page 177–190, Portland, OR, USA, 2015.
- [13] I. Narayanan, D. Wang, M. Jeon, B. Sharma, L. Caulfield, A. Sivasubramaniam, B. Cutler, J. Liu, B. Khessib, and K. Vaid. SSD Failures in Datacenters: What? When? And Why? In *Proceedings of the 9th ACM International Conference on Systems and Storage, SYSTOR '16*, Haifa, Israel, 2016.
- [14] B. Schroeder, R. Lagisetty, and A. Merchant. Flash reliability in production: The expected and the

- unexpected. In *Proceedings of the 14th USENIX Conference on File and Storage Technologies (FAST 16)*, pages 67–80, Santa Clara, CA, USA, 2016.
- [15] S. Tannu and P. Nair. The Dirty Secret of SSDs: Embodied Carbon. In *Proceedings of the 1st Workshop on Sustainable Computer Systems Design and Implementation, HotCarbon '22*, San Diego, CA, USA, 2022.
- [16] E. Xu, M. Zheng, F. Qin, Y. Xu, and J. Wu. Lessons and actions: What we learned from 10k SSD-Related storage system failures. In *2019 USENIX Annual Technical Conference (USENIX ATC 19)*, pages 961–976, Renton, WA, 2019.