



When Servers Meet Species: A Fab-to-Grave Lens on Computing's Biodiversity Impact

TIANYAO SHI, Purdue University, USA

RITBIK KUMAR, Purdue University, USA

INEZ HUA, Purdue University, USA

YI DING, Purdue University, USA

Biodiversity loss is a critical planetary boundary, yet its connection to computing remains largely unexamined. Prior sustainability efforts in computing have focused on carbon and water, overlooking biodiversity due to the lack of appropriate metrics and modeling frameworks. This paper presents the first end-to-end analysis of biodiversity impact from computing systems. We introduce two new metrics—Embodied Biodiversity Index (EBI) and Operational Biodiversity Index (OBI)—to quantify biodiversity impact across the lifecycle, and present FABRIC, a modeling framework that links computing workloads to biodiversity impacts. Our evaluation highlights the need to consider biodiversity alongside carbon and water in sustainable computing design and optimization. The code is available at <https://github.com/TianyaoShi/FABRIC>.

CCS Concepts: • **Social and professional topics** → **Sustainability**; • **Hardware** → **Impact on the environment**.

Additional Key Words and Phrases: Biodiversity, Sustainability, Life Cycle Assessment, HPC

1 Introduction

The environmental impact of computing has grown rapidly with the rise of AI, hyperscale datacenters, and semiconductor fabs [45]. Prior work has focused on carbon [1, 15, 26, 45] and water [11, 14, 18, 22] dimensions of sustainability in computing. In contrast, biodiversity impact, a key endpoint in life cycle assessment (LCA) frameworks [8], remains largely unexamined in computing. From pollutant emissions during chip fabrication to toxic waste at end-of-life, computing systems generate stressors that contribute directly to species loss and ecosystem degradation [4].

Despite its importance, biodiversity impact has remained absent from sustainable computing research. We identify two key barriers: (1) the lack of quantifiable metrics that can attribute biodiversity and ecosystem damage to specific computing activities; and (2) the lack of a practical modeling framework that links specific computing workload and system to concrete biodiversity impacts across their lifecycle—from manufacturing to operation to end-of-life.

To overcome these challenges, we introduce two new metrics: the Embodied Biodiversity Index (EBI) covering manufacturing, transportation, and end-of-life, and the *Operational Biodiversity Index (OBI)* covering electricity use. They together provide a quantifiable basis for attributing biodiversity impact to computing systems.

Building on these metrics, we present FABRIC (Fabrication-to-grave Biodiversity Impact Calculator), a new modeling framework that traces the fab-to-grave lifecycle of computing components (e.g.,

CPU, GPU, memory, storage) across manufacturing, transportation, use, and end-of-life stages. The key insight behind FABRIC is that biodiversity impact can be decomposed and attributed using a combination of LCA midpoint indicators (e.g., acidification, eutrophication, ecotoxicity). This enables us to directly connect real-world computing activities to measurable ecological consequences.

We evaluate FABRIC on 7 HPC workloads across 3 computing platforms with diverse devices. Key findings include:

- Manufacturing and acidification dominate embodied biodiversity loss. Newer devices have lower impact per unit performance.
- Operational electricity outweighs manufacturing in biodiversity lifecycle impact under conservative assumptions.
- Low carbon does not always mean low biodiversity impact.

Our contributions are summarized as follows.

- Conduct the first analysis of biodiversity impact in computing.
- Introduce two new metrics—EBI and OBI—to quantify biodiversity impact from manufacturing, transportation, end-of-life and use.
- Present FABRIC, the first modeling framework that connects computing workloads to lifecycle biodiversity impact.
- Evaluate biodiversity impact across devices, systems, workloads, and geographic deployment.

2 Background

In this section, we review the LCA framework for assessing biodiversity impacts of computing and define the scope of our study. We focus on two hierarchical levels: *mid-point* impacts, which capture specific pollutant-driven mechanisms, and *end-point* impact, which quantifies resulting damage to ecosystems and species.

2.1 Mid-point Impact Categories

In this paper, we focus on three midpoint impacts: acidification, eutrophication, and ecotoxicity; as sustainability reports from key vendors disclose emission data that directly map to these impacts. Other midpoints like ozone depletion and smog formation are excluded due to limited data [16] and lower contribution to endpoint damage [43].

Acidification (AP), measured in kilograms of SO₂ equivalents, reflects how much emitted gases—SO₂, NO_x, and NH₃—can lower environmental pH. These gases form acid rain, which leaches nutrients from soil and harms plants and aquatic life. In computing, AP largely results from SO₂ and NO_x emissions from fossil fuel combustion in chip manufacturing and datacenter electricity use.

Eutrophication (EP), reported in kilograms of phosphate equivalents (kg PO₄³⁻ eq), measures how excess nitrogen and phosphorus enrich ecosystems [34], often leading to oxygen-depleting algal

Authors' Contact Information: Tianyao Shi, Purdue University, West Lafayette, IN, USA, shi676@purdue.edu; Ritvik Kumar, Purdue University, West Lafayette, IN, USA, kumar684@purdue.edu; Inez Hua, Purdue University, West Lafayette, IN, USA, hua@purdue.edu; Yi Ding, Purdue University, West Lafayette, IN, USA, yiding@purdue.edu.

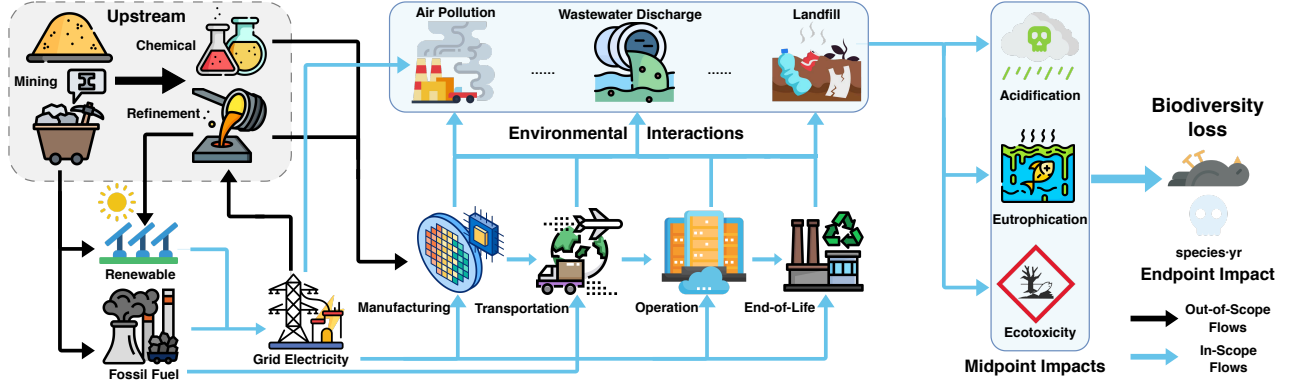


Fig. 1. The Fabrication-to-grave Biodiversity Impact Calculator (FABRIC) modeling framework.

blooms. In computing, EP stems from wastewater in chip production, chemical slurries used in processing, and nitrate/ammonia emissions from power generation.

Freshwater ecotoxicity potential (FETP), expressed in comparative toxic units for ecosystems (CTU_e), captures the toxic impact of chemicals on aquatic ecosystems [27]. In computing, key sources include heavy metals (e.g., Cu, Ni), photoresist residues, fluorinated surfactants, and waste from both fabs and material extraction.

2.2 End-point Impact Metric

End-point biodiversity impact in **species-yr**, which quantifies the expected fraction of species locally lost per year due to cumulative environmental stressors. For example, a value of 1×10^{-4} species-yr means one ten-thousandth of a species is statistically lost in a region over a year. To convert acidification, eutrophication, and ecotoxicity impacts into species-yr, we use the ReCiPe 2016 ecosystem damage model [17] — chosen over ILCD 2011 [7] and IMPACT World+ [5] for its comprehensive treatment of ecological impacts. This unified metric enables us to compare different environmental stressors and assess the biodiversity risks embedded in computing infrastructure.

2.3 Scope Discussion

In this paper, we use a *fab-to-grave* system boundary that includes the manufacturing, transportation, use, and end-of-life stages (Figure 1). We exclude upstream mining and material refinement despite their significance, due to fragmented and opaque global supply chains that limit reliable impact tracing. Following the Dell R740 LCA report [6], we assume packaging uses fully recycled materials with negligible biodiversity impact. We focus on five key computing components—CPU, GPU, DRAM, SSD, and HDD—which account for most semiconductor wafer throughput, operational energy use, and pollutant emissions driving AP, EP, and FETP. As such, our reported results reflect a lower bound on biodiversity impact, and will become more complete as public data expands.

3 The FABRIC Modeling Framework

Figure 1 shows the workflow of FABRIC. We first introduce the two newly defined metrics for quantifying biodiversity impact, and then explain their application across different lifecycle stages.

3.1 Definitions of EBI and OBI

The total biodiversity impact is the summation of EBI and OBI. Next, we introduce them separately.

Embodied Biodiversity Index (EBI) refers to the biodiversity damage incurred once over the physical lifecycle of a device that covers the manufacturing (Mfg), transportation (Trans), and end-of-life (EoL) stages. The device harms the environment in two key ways: (1) directly releasing toxic substances through leaks, discharge, or e-waste; and (2) indirectly causing emissions by driving fossil fuel use for transport and electricity generation. The emissions in each stage $l \in \{\text{Mfg}, \text{Trans}, \text{EoL}\}$ of a device d lead to multifaceted midpoint impacts $M_{c,l}(d)$, where $c \in C = \{\text{AP}, \text{EP}, \text{FETP}\}$ is the midpoint impact category. The EBI of a device is calculated by summing impacts across stages and impact categories:

$$B_{\text{emb}}(d) = \sum_{c,l} M_{c,l}(d) \Phi_c, \quad (1)$$

where Φ_c is the midpoint-to-endpoint conversion factor for c .

Similar with the definition of embodied carbon, EBI of a specific workload is proportional to the execution time amortized over device lifetime. Let LT_d be the lifetime of device d , the EBI of a specific workload w with execution time t_w is given by

$$B_{\text{emb}}(w) = \sum_{d \in \mathcal{D}} \frac{t_w}{LT_d} B_{\text{emb}}(d), \quad (2)$$

where \mathcal{D} is the set of all devices involved in the execution.

Operational Biodiversity Index (OBI) is the biodiversity damage incurred from generating the electricity needed for the operational use of computing devices.

$$B_{\text{op}}(d) = \sum_c M_{c,\text{Use}}(d) \Phi_c. \quad (3)$$

Similar with operational carbon, the task-specific OBI of w is proportional to its energy consumption $E_{\text{el}}(w)$:

$$B_{\text{op}}(w) = \sum_c M_{c,\text{Use}}(w) \Phi_c, \quad (4)$$

$$M_{c,\text{Use}}(w) = \sum_k E_{\text{el}}(w) F_{\text{el},k,r} \Gamma_k^c. \quad (5)$$

Equation (5) reveals details on how to compute the midpoint impacts by summing grid-related emission loads: $k \in \{\text{SO}_2, \text{NO}_x, \text{NH}_3\}$ is the associated pollutant species, $F_{\text{el},k,r}$ is the emission factor for region r (in g/kWh), and Γ_k^c is the *characterization factor* to convert the pollutant k into the reference unit for impact category c . For example, 1 kg NH_3 is 0.93 kg SO_2 eq for AP [3].

Generally, the midpoint impact of a device d 's lifecycle stage l for impact category c is given by

$$M_{c,l}(d) = \sum_{j,k} E_j(d) F_{j,k,r} \Gamma_k^c, \quad (6)$$

where E_j denotes the measurable throughput of process j (e.g., kWh for electricity use, wafer-mask-layers for logic chips manufacturing, or ton-kilometers for transportation) and $F_{j,k,r}$ is the corresponding per-unit emission factor of pollutant species k in region r .

Summary. Since the biodiversity impact from electricity use falls under OBI and has been thoroughly covered above, we will focus next on modeling EBI—specifically from manufacturing, transportation, and end-of-life stages.

3.2 EBI from Manufacturing

Different devices have distinct manufacturing processes and have varying levels of data availability, so we calculate their EBIs separately. As implied by Equation (6), the key is to determine the values of process throughput E_j (i.e., how many production units it takes to manufacture a device) and the per-unit emission factors $F_{j,k,r}$ for each type of computing device.

3.2.1 Logic devices (CPU, GPU core). This is an example where both E_j and $F_{j,k,r}$ can be easily acquired:

$$E_{\text{Mfg}}(d) = \frac{\text{Die Size}}{\text{Wafer Area} \cdot \text{Yield}} \cdot \text{\#wafer-mask-layers},$$

where #wafer-mask-layers is 67, 78, 87, and 81 layers for 14, 10, 7, and 5nm node, respectively [28]. For $F_{j,k,r}$, the annual per-production-unit air pollutant emission, wastewater discharge, and electricity consumption is available in TSMC sustainability reports [37].

3.2.2 DRAM and SSD. This is an example where only fab-level annual emission data is available but the fab produces multiple types of devices. According to ISO 14044 [10], we use economic-value-based allocation to obtain per-unit emission factors. Let $p \in \{\text{DRAM, SSD}\}$ be the product class of d :

$$E_{\text{Mfg}}(d) = \frac{\text{Capacity}_d}{\text{Wafer Area} \cdot \text{Bit Density}_p / 8} \cdot \frac{1}{\text{Yield}},$$

$$F_{\text{Mfg},k,r}(d) = \frac{\text{Revenue Share}_p \cdot \text{Annual Total Emission}_k}{\text{Wafer Production Capacity}_p},$$

where $E_{\text{Mfg}}(d)$ means expected number of wafers to produce a single device d and $F_{\text{Mfg},k,r}(d)$ means per-wafer emission of pollutant species k . The emission from electricity use is derived similarly. For GPUs, the attached high-bandwidth memory (HBM) adopts the DRAM formula and the total EBI of GPU adds the two parts up.

3.2.3 HDD. Cradle-to-grave LCA studies [19, 29] already provide per-unit midpoint scores. We extrapolate these scores for the missing years in our study scope by assuming similar overall reductions as seen in carbon emissions [30]. The values are directed used as

$M_{c,l}(d_{\text{HDD}})$. Importantly, the broader system boundaries in LCAs only reinforce the conservativeness of our lower-bound estimates.

3.3 EBI from Transportation

The process throughput of transporting is

$$E_{\text{Trans},g}(d) = \sum_m \text{Shipping Mass}_d \cdot \text{Distance}_g,$$

where g is the transport mode (e.g., truck, ship). The mode-specific emission factors $F_{\text{Trans},g,k,r}$ are available in Ecoinvent database [44]. Note that emission factors vary over time due to stricter regulations on maritime fuel sulfur content and truck emissions [13]. For simplicity, we assume a fixed transport profile: 200km by truck (100km from the fab to the port and 100km from the port to the datacenter) and 14,000km by ship (East Asia to the U.S. west coast).

3.4 EBI from End-of-Life (EoL)

The end-of-life (EOL) stage requires a different modeling approach due to its complexity. Discarded servers and storage devices typically follow three disposal pathways: *recycling* (disassembly, metal shredding/smelting, and plastics recovery), *incineration*, and controlled *landfill*, with ratios denoted as R , I , L , respectively. Given the mass m_d , the midpoint impact is

$$M_{c,\text{EoL}}(d) = m_d [R F_{\text{rec}}^c + I(F_{\text{inc}}^c + p_{\text{ash}} F_{\text{ash}}^c) + L F_{\text{land}}^c], \quad (7)$$

where $R+I+L=1$, p_{ash} is the kilograms of bottom ash per kilogram of feed incinerated, and F_{\bullet}^c is the net impact for each pathway. However, direct F_{\bullet}^c values for mixed electronic waste are rarely reported, making detailed modeling intractable. Therefore, we use precise EoL impacts from the latest Fairphone 5 LCA study [35] as proxies for the recycling pathway Equation (6):

$$E_{\text{EOL}}(d) = m_d \cdot \sum_k F_{\text{EOL},k} C F_k^c = \frac{M_{c,\text{EOL}}(\text{Fairphone 5})}{m_{\text{Fairphone 5}}}.$$

We assume each category's impact scales linearly with device mass, as the bills of materials of computing devices distribute similar ratios to metal, plastic, etc. Then, we apply these estimates uniformly across all devices until more specific data become available.

4 Evaluation

We examine the following research questions (RQs):

- RQ1:** Which lifecycle stage contributes most to computing's biodiversity impact? (§4.2)
- RQ2:** How do biodiversity impacts vary across HPC workloads and system platforms? (§4.3)
- RQ3:** How does deployment location affect computing's biodiversity impacts? (§4.4)

4.1 Evaluation Methodology

Device and system scopes. To analyze device-level trends over time, Table 1 lists the specifications of individual devices across a range of generations and manufacturing technologies used in our FABRIC analysis. To examine how biodiversity impacts scale from a single node to large-scale systems, we evaluate three computing setups shown in Table 2: a *local testbed server* representing a typical

Table 1. Individual devices studied in this paper.

CPU	Year	Node	Silicon Area	Cores
AMD EPYC 7B12	2019	7/14 nm	1 008 mm ²	64
AMD EPYC 7443	2021	7/12 nm	740 mm ²	24
AMD EPYC 7B13	2021	7/12 nm	1 064 mm ²	64
AMD EPYC 9B14	2023	5/6 nm	1 261 mm ²	96
GPU	Year	Node	Silicon Area	VRAM
NVIDIA T4	2018	12 nm	545 mm ²	16 GB
NVIDIA V100	2017	12 nm	815 mm ²	16 GB
NVIDIA L40	2022	5 nm	609 mm ²	48 GB
NVIDIA A100	2020	7 nm	826 mm ²	40 GB
NVIDIA H100	2023	5 nm	814 mm ²	80 GB
Device	Vendor	Type	Year	Capacity
DDR4 RDIMM	SK Hynix	DRAM	2020	64 GB
PE8111	SK Hynix	SSD	2022	15.36 TB
Exos X20	Seagate	HDD	2023	20 TB

edge or lab workstation, *Gautschi** representing a 2023-era community cluster in academia [25]¹, and *Perlmutter* representing a petascale supercomputer with mixed CPU/GPU architecture.

Workloads. We evaluate seven HPC workloads from the Phoronix Test Suite [24], spanning lossless compression/decompression, scientific/analytic kernels, media encoding, and software compilation (see x-axis in Figure 7). These workloads represent a broad range of compute, memory, and storage demands, capturing both legacy and emerging biodiversity stressors.

Data sources. We make our best effort to ensure the credibility of publicly accessible sources used in this study by selecting leading vendors’ sustainability reports, authoritative agency documents, and highly-recognized peer-review studies. The values for midpoint characterization factors Γ_k^c are determined based on TRACI documentation [2, 41] from the United States Environmental Protection Agency (EPA). The midpoint-to-endpoint conversion factors Φ_c are based on ReCiPe 2016 Hierarchist framework [17]. Logic-device data derive from TSMC’s sustainability reports [37], and memory data from SK Hynix [33]; revenue shares come from audited financials [32], wafer capacity from TrendForce [40], and bit densities from IEDM paper [20] and TechInsights [31, 38, 39]. All datasets are limited to the 2016–2023 window, aligning with the temporal scope of our study. The yield of devices is assumed to be a constant 0.875 as in prior work [15]. Electricity consumption is characterized with $\text{SO}_x/\text{NO}_x/\text{NH}_3$ factors from EPA eGRID [42], EDGAR [9], and East-Asian utilities’ (Taipower, KEPCO) sustainability reports [21, 36]. Although emission factors can differ across vendors due to varied manufacturing processes, TSMC and SK Hynix together dominate advanced logic and memory production, providing a reasonable reference for the sector. Any cross-vendor variance is therefore unlikely to overturn our trend-level conclusions.




Fig. 2. Midpoint EBI breakdown by lifecycle stages for CPUs and GPUs modeled.




Fig. 3. Endpoint EBI analyses for CPUs and GPUs modeled. Left: endpoint EBI per FP64 Teraflops. Middle: EBI contribution from different lifecycle stages. Right: EBI contributions from different midpoint impact categories.

4.2 Biodiversity Impact Breakdown

Mid-/End-point EBI analysis. Figure 2 shows that *manufacturing* overwhelmingly dominates device-level acidification (AP), eutrophication (EP), and freshwater-ecotoxicity (FETP), contributing 78–92% of total mid-point impacts. In contrast, transportation and end-of-life (EOL) together rarely exceed 20%. GPUs generally have higher absolute impacts than CPUs—e.g. the 12nm V100 emits nearly 0.020 kg SO₂ eq of AP, more than twice that of the 7nm EPYC 7B12—yet node scaling alone cannot offset larger package masses, as the 5nm, 700W H100 still tops the AP chart. When normalized by peak FP64 throughput (Figure 3, left), this trend reverses: the legacy double-precision GPUs (V100) incur $\sim 1.1 \times 10^{-8}$ species-yr TFLOPS⁻¹, roughly 15× higher than the tensor-dense H100 or modern CPUs—highlighting how architectural improvements and silicon scaling reduce embodied biodiversity damage. Figure 3 shows that manufacturing remains the largest contributor to endpoint biodiversity impact (55–75%), with acidification as the dominant pathway (60–85%), followed by EP and FETP. A similar trend holds for memory and storage (Figure 4): manufacturing dominates due to their minimal shipping mass per GB (< 1 gram), and capacity-normalized impacts drop nearly tenfold from 2016 to 2023.

Takeaway 1: Manufacturing and acidification drives embodied biodiversity impact for computing devices. While absolute damage increases with newer generations, species-yr per TFLOP decreases, highlighting that newer, more efficient devices are environmentally preferable on a per-performance basis.

¹We replaced 40 Intel CPUs with comparable AMD EPYC 9B14 processors to enable direct comparison.

Table 2. Three computing systems evaluated in this paper.

System	Main HW Year	CPUs (model × count)	GPUs (model × count)	Total DRAM	Storage Capacity
Local testbed server	2022	EPYC 7443 × 1	L40 × 4	520 GB	2 TB SSD + 20 TB HDD
Gautschi* [25]	2023	EPYC 9B14 × 442	H100 × 160 + L40 × 12	158 TB	1.5 PB SSD + 2.5 PB HDD
Perlmutter [23]	2021	EPYC 7B13 × 4 864	A100 80 GB × 1 024 + A100 40 GB × 6 144	1 984 TB	44 PB SSD




Fig. 4. Mid-/End-point EBI for memory and storage.




Fig. 5. EBI contributions from different components in three studied systems.

System-level EBI contributions. Figure 5 shows results for full systems. A single-socket *local server* carries an EBI of only 3.5×10^{-8} species-yr, whereas the GPU-heavy *Perlmutter* reaches 2.0×10^{-4} species-yr—five orders of magnitude higher. These differences reflect system composition. GPUs dominate impact when they dominate hardware: contributing 72% of *Perlmutter*'s total but only 17% in CPU-heavy *Gautschi**, where CPUs and HDDs (a by-product of the large file system) account for 40.3% and 24.3%, respectively. Memory contributes 10–31%, while SSDs remain below 8%.

Perlmutter highlights the long-term impact of operations. At 70% load and using U.S. grid average SO_2 and NO_x emission factors, its annual electricity use of 27.4GWh adds 2.51×10^{-3} species-yr—about 60× higher than its annualized EBI. This makes OBI dominate its lifecycle footprint. If 400 *Perlmutter*-class systems are active globally, the sector causes ≥ 1 species-yr of biodiversity impact per year, with a footprint that will grow proportionally with future scale.

OBI dominance. Figure 6 presents a full lifecycle view by stacking the embodied impacts (EBI) with annual operational impacts (OBI) across three representative duty cycles (30%, 50%, 70%). Assuming U.S. grid average emission factors, electricity use pushes all devices above their embodied impact for all cases. At 70% load, OBI outweighs EBI by nearly 100× for high-power CPUs and GPUs. Although the mass-heavy 20TB HDD shows the largest EBI across




Fig. 6. A lifecycle analysis for individual devices modeled.

all devices, its OBI is modest, occupying less than 10% total impact of logic devices. This suggests that improving energy efficiency and adopting cleaner electricity are the most effective ways to lower a system's total biodiversity impact.

Takeaway 2: When using conservative embodied estimates that exclude raw material extraction and refining, and assuming average U.S. grid emissions, operational electricity—not manufacturing—dominates a device's lifecycle biodiversity impact.

4.3 Workload-Level Impact Assessment

We test seven HPC workloads on three platforms: (1) *Google Cloud Platform* (GCP) n2d-standard-64 with 32 physical cores (half of EPYC 7B12) [12], (2) the local server in Table 2, and (3) GCP c3d-standard-60 with 30 physical cores (5/16 of EPYC 9B14), representing a generational progression of HPC systems every 2–3 years.

To ensure fair comparison, all platforms are assumed to have 256GB memory and 100GB HDD and located in the U.S. Midwest. Since all workloads are CPU-centric, we do not consider impacts from additional devices like GPUs on the local server. We normalize throughput, energy consumption per unit throughput, and total biodiversity impact per unit throughput using the oldest GCP n2d as the reference point. Results are shown in Figure 7. While performance increases—up to 55% with the newest VM—energy efficiency tells a clearer story: the local server consumes 1.3–3.2× more energy per unit work, likely due to less optimized power management. Consequently, biodiversity impact follows the same trend: the 5nm cloud VM reduces impact to 55–75% of baseline, while the local server increases it by 30–140%, nullifying performance gains.




Fig. 7. Normalized throughput, energy consumption, and total biodiversity impacts for HPC workloads.

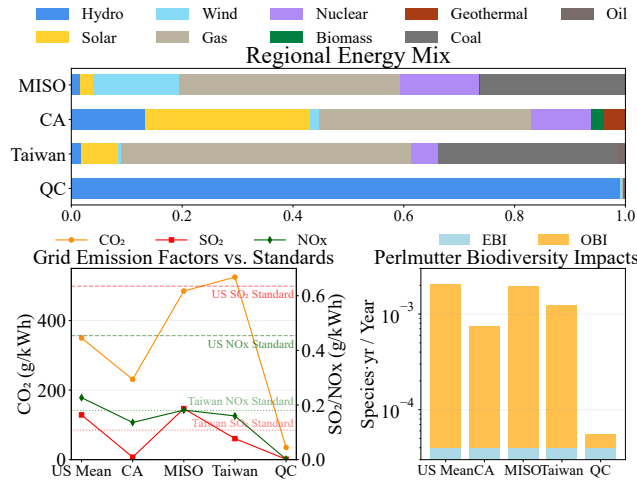


Fig. 8. Regional Grid Emissions vs. Biodiversity Impacts.

Takeaway 3: With the same memory and storage, newer energy-efficient CPUs in cloud environments reduce biodiversity impact of per unit of work by half, while a local server can double it—even if performance appears comparable. This indicates that datacenter-level power management may matter more than chip technology.

4.4 Geospatial Sensitivity

Because emission factors vary by location, we perform a geospatial sensitivity analysis by hypothetically placing the *Perlmutter* system in four distinct electricity markets: California (CA), the Midwest (MISO), Taiwan, and hydro-rich Québec (QC). Figure 8 shows the results. The findings highlight that low-carbon does not always mean low biodiversity impact. For example, CA’s solar-heavy grid and strict sulfur limits significantly reduce biodiversity harm. In contrast, MISO’s coal-heavy mix raises SO_2 emissions above the U.S. average despite only moderate CO_2 . Taiwan, dominated by gas, is more carbon-intensive but benefits from tight acid-gas controls, resulting in lower biodiversity impact than MISO. Relocating the system to QC’s nearly all-hydro grid cuts OBI to $\approx 1.5 \times 10^{-4}$ species-yr—comparable to EBI and about two orders lower than fossil-heavy grids.

Takeaway 4: Focusing solely on CO_2 is not enough—strict SO_2/NO_x limits and renewable-heavy grids can cut a supercomputer’s total biodiversity impact by an order of magnitude. This highlights the need for pollutant-specific policies.

5 Conclusion

This paper presents the first systematic study of computing’s biodiversity impact across the device lifecycle. By introducing EBI and OBI and developing the FABRIC framework, we provide a foundation for connecting workloads to biodiversity outcomes. As computing demand grows, biodiversity should become a first-class metric in sustainable system design and optimization.

Acknowledgments

We thank the anonymous reviewers for their valuable feedback. We acknowledge GCP for providing generous cloud research credits that supported the cloud CPUs used in this study. This work is supported by NSF CCF-2413870.

References

- [1] Bilge Acun, Benjamin Lee, Fiodar Kazhemiaka, Kiwan Maeng, Udit Gupta, Manoj Chakkaravarthy, David Brooks, and Carole-Jean Wu. 2023. Carbon explorer: A holistic framework for designing carbon aware datacenters. In *Proceedings of the 28th ACM International Conference on Architectural Support for Programming Languages and Operating Systems (ASPLOS), Volume 2*.
- [2] Jane Bare. 2011. TRACI 2.0: the tool for the reduction and assessment of chemical and other environmental impacts 2.0. *Clean Technologies and Environmental Policy* 13 (2011), 687–696.
- [3] BIO Intelligence Service. 2005. *A Study to Examine the Costs and Benefits of the ELV Directive – Final Report Annex 5: Environmental Impacts Analysed and Characterisation Factors*. Technical report. European Commission, Directorate-General for Environment, Brussels, Belgium. <https://ec.europa.eu/environment/pdf/waste/study/annex5.pdf> Accessed 19 May 2025.
- [4] Gordon B Bonan and Scott C Doney. 2018. Climate, ecosystems, and planetary futures: The challenge to predict life in Earth system models. *Science* 359, 6375 (2018), eaam8328.
- [5] Cécile Bulle, Manuele Margni, Laure Patouillard, Anne-Marie Boulay, Guillaume Bourgault, Vincent De Bruille, Viêt Cao, Michael Hauschild, Andrew Henderson, Sebastian Humbert, et al. 2019. IMPACT World+: a globally regionalized life cycle impact assessment method. *The International Journal of Life Cycle Assessment* 24 (2019), 1653–1674.
- [6] Andreas Busa, Malcolm Hegeman, Jeff Vickers, Natalia Duque-Ciceri, and Constantin Herrmann. 2019. *Life Cycle Assessment of Dell R740*. Technical Report. Dell Technologies. https://www.delltechnologies.com/asset/en-us/products/servers/technical-support/Full_LCA_Dell_R740.pdf Accessed 19 May 2025.
- [7] Kirana Chomkhamisri, Marc-Andree Wolf, and Rana Pant. 2011. International reference life cycle data system (ILCD) handbook: Review schemes for life cycle assessment. *Towards life cycle sustainability management* (2011), 107–117.
- [8] Michael Curran, Laura de Baan, An M De Schryver, Rosalie Van Zelm, Stefanie Hellweg, Thomas Koellner, Guido Sonnemann, and Mark AJ Huijbregts. 2011. Toward meaningful end points of biodiversity in life cycle assessment. *Environmental science & technology* 45, 1 (2011), 70–79.
- [9] European Commission Joint Research Centre (JRC) and Netherlands Environmental Assessment Agency. 2024. EDGAR 2024_GHG: Emissions Database for Global Atmospheric Research. https://edgar.jrc.ec.europa.eu/dataset_ghg2024. Accessed: 2025-05-19.
- [10] Matthias Finkbeiner, Atsushi Inaba, Reginald Tan, Kim Christiansen, and Hans-Jürgen Klüppel. 2006. The new international standards for life cycle assessment: ISO 14040 and ISO 14044. *The international journal of life cycle assessment* 11 (2006), 80–85.
- [11] Wedan Emmanuel Gnigba, Andrew A Chien, Anne Blavette, and Anne Cécile Orgerie. 2024. FlexCoolDC: datacenter cooling flexibility for harmonizing water, energy, carbon, and cost trade-offs. In *Proceedings of the 15th ACM International Conference on Future and Sustainable Energy Systems*. 108–122.
- [12] Google Cloud. 2025. General-purpose machine family for Compute Engine | Compute Engine Documentation. <https://cloud.google.com/compute/docs/general-purpose-machines>.

- [13] Theodoros Grigoratos, Georgios Fontaras, Barouch Giechaskiel, and Nikiforos Zacharof. 2019. Real world emissions performance of heavy-duty Euro VI diesel vehicles. *Atmospheric environment* 201 (2019), 348–359.
- [14] Pranjol Sen Gupta, Md Rajib Hossen, Pengfei Li, Shaolei Ren, and Mohammad A Islam. 2024. A dataset for research on water sustainability. In *Proceedings of the 15th ACM International Conference on Future and Sustainable Energy Systems*. 442–446.
- [15] Udit Gupta, Mariam Elgamal, Gage Hills, Gu-Yeon Wei, Hsien-Hsin S. Lee, David Brooks, and Carole-Jean Wu. 2022. ACT: Designing sustainable computer systems with an architectural carbon modeling tool. In *ISCA*.
- [16] Alexandra C Hetherington, Aiduan Li Borrior, Owen Glyn Griffiths, and Marcelle C McManus. 2014. Use of LCA as a development tool within early research: challenges and issues across different sectors. *The International Journal of Life Cycle Assessment* 19 (2014), 130–143.
- [17] Mark A.J. Huijbregts, Zoran J.N. Steinmann, Pieter M.F. Elshout, Geert Stam, Francesca Veronesi, Marisa Vieira, Anne Hollander, Michiel Zijp, and Rosalie van Zelm. 2016. *ReCiPe 2016: A Harmonized Life Cycle Impact Assessment Method at Midpoint and Endpoint Level. Report I: Characterization*. Technical Report. RIVM National Institute for Public Health and the Environment, Bilthoven, The Netherlands. <https://www.rivm.nl/publicaties/recipe-2016-harmonized-life-cycle-impact-assessment-method-at-midpoint-and-endpoint-level>
- [18] Yankai Jiang, Rohan Basu Roy, Raghavendra Kanakagiri, and Devesh Tiwari. 2025. WaterWise: Co-optimizing Carbon- and Water-Footprint Toward Environmentally Sustainable Cloud Computing. In *Proceedings of the 30th ACM SIGPLAN Annual Symposium on Principles and Practice of Parallel Programming* (Las Vegas, NV, USA) (PPoPP '25). Association for Computing Machinery, New York, NY, USA, 297–311. doi:10.1145/3710848.3710891
- [19] Hongyue Jin, Kali Frost, Ines Sousa, Hamid Ghaderi, Alex Bevan, Miha Zakotnik, and Carol Handwerker. 2020. Life cycle assessment of emerging technologies on value recovery from hard disk drives. *Resources, Conservation and Recycling* 157 (2020), 104781.
- [20] Scotten W Jones. 2023. Modeling 300mm wafer fab carbon emissions. In *2023 International Electron Devices Meeting (IEDM)*. IEEE, 1–4.
- [21] Korea Electric Power Corporation (KEPCO). 2025. Sustainability Management Report. https://home.kepco.co.kr/kepco/EN/D/C/KEDCPP004.do?boardCd=BRD_000014&menuCd=EN040106. Accessed: 2025-05-19.
- [22] Pengfei Li, Jianyi Yang, Mohammad A. Islam, and Shaolei Ren. 2024. Making AI Less "Thirsty": Uncovering and Addressing the Secret Water Footprint of AI Models. *Commun. ACM* (2024).
- [23] NERSC. 2023. Perlmutter: High Performance Computing Optimized for Science. <https://perlmutter.carrd.co/>.
- [24] Phoronix Media. 2025. Phoronix Test Suite. <https://www.phoronix-test-suite.com/>. Accessed: 2025-05-18.
- [25] Purdue University Rosen Center for Advanced Computing. 2025. RCAC - Compute: Gautschi. <https://www.rcac.purdue.edu/compute/gautschi>.
- [26] Ana Radovanović, Ross Koningstein, Ian Schneider, Bokan Chen, Alexandre Duarte, Binz Roy, Diyu Xiao, Maya Haridasan, Patrick Hung, Nick Care, et al. 2022. Carbon-aware computing for datacenters. *IEEE Transactions on Power Systems* 38, 2 (2022), 1270–1280.
- [27] Ralph K Rosenbaum, Till M Bachmann, Lois Swirsky Gold, Mark AJ Huijbregts, Olivier Jolliet, Ronnie Juraske, Annette Koehler, Henrik F Larsen, Matthew MacLeod, Manuele Margni, et al. 2008. USEtox—the UNEP-SETAC toxicity model: recommended characterisation factors for human toxicity and freshwater ecotoxicity in life cycle impact assessment. *The International Journal of Life Cycle Assessment* 13 (2008), 532–546.
- [28] David Schor. 2020. TSMC Details 5 nm. *WikiChip Fuse* article. <https://fuse.wikichip.org/news/3398/tsmc-details-5-nm/> Last updated May 25, 2021.
- [29] Seagate Technology LLC. 2016. *Makara Enterprise HDD: Life Cycle Assessment Summary*. Technical Report. Seagate Technology LLC. <https://www.seagate.com/files/www-content/global-citizenship/en-us/docs/seagate-makara-enterprise-hdd-lca-summary-2016-07-29.pdf> Accessed 19 May 2025.
- [30] Seagate Technology LLC. 2025. Product Sustainability. <https://www.seagate.com/esg/planet/product-sustainability/>. Web page, accessed 19 May 2025.
- [31] Pete Singer and Jeongdong Choe. 2019. DRAM, NAND and Emerging Memory Technology Trends and Developments in 2019. <https://www.semiconductor-digest.com/dram-nand-and-emerging-memory-technology-trends-and-developments-in-2019/>. Accessed: 2025-05-19.
- [32] SK hynix Inc. 2025. Audit (Review) Report – Consolidated and Non-consolidated Financial Statements. https://www.skhynix.com/ir/UI-FR-IR12_T4. Investor-relations document, accessed 19 May 2025.
- [33] SK hynix Inc. 2025. Sustainability Reports & Policies. <https://www.skhynix.com/sustainability/UI-FR-SA1601/>. Web page hosting SK hynix annual sustainability reports; accessed 19 May 2025.
- [34] Val H Smith, G David Tilman, and Jeffery C Nekola. 1999. Eutrophication: impacts of excess nutrient inputs on freshwater, marine, and terrestrial ecosystems. *Environmental pollution* 100, 1-3 (1999), 179–196.
- [35] David Sánchez, Sarah Jane Baur, and Lara Eguren. 2024. *Life Cycle Assessment of the Fairphone 5*. Technical Report. Fraunhofer Institute for Reliability and Microintegration (IZM), Berlin, Germany. https://www.fairphone.com/wp-content/uploads/2024/09/Fairphone5_LCA_Report_2024.pdf Commissioned by Fairphone B.V..
- [36] Taiwan Power Company. 2025. Report – Taipower Sustainability Section. <https://service.taipower.com.tw/csr/en/report-download>. Accessed: 2025-05-19.
- [37] Taiwan Semiconductor Manufacturing Company Limited. 2025. ESG Resources and Documents. <https://esg.tsmc.com/en/resources/documents.html>. Web page, accessed 19 May 2025.
- [38] TechInsights. 2019. TechInsights memory technology update from IEDM18. <https://www.techinsights.com/blog/techinsights-memory-technology-update-iedm18>. Contributing authors: Dick James and Jeongdong Choe. Accessed: 2025-05-19.
- [39] TechInsights. 2024. *HBM (High Bandwidth Memory) Technology Trend: 2024 and Beyond*. Technical Report. TechInsights. Accessed: 2025-05-19.
- [40] TrendForce DataTrack. 2025. NAND Flash Makers Fab Capacity Breakdown by Brand. <https://datatrack.trendforce.com/Chart/groupContent/65/nand-flash-makers-fab-capacity-breakdown-by-brand>. Accessed: 2025-05-19.
- [41] United States Environmental Protection Agency. 2025. Tool for Reduction and Assessment of Chemicals and Other Environmental Impacts (TRACI) ver 2.2. https://www.epa.gov/system/files/documents/2024-01/traci_2_2.xlsx. Accessed: 2025-05-18.
- [42] United States Environmental Protection Agency (EPA). 2025. Emissions & Generation Resource Integrated Database (eGRID), eGRID2023rev1. <https://www.epa.gov/egrid>. Accessed: 2025-05-19.
- [43] Anne EM van den Oever, Stefano Puricelli, Daniele Costa, Nils Thonemann, Maeva Lavigne Philippot, and Maarten Messagie. 2024. Revisiting the challenges of ozone depletion in life cycle assessment. *Cleaner Environmental Systems* (2024), 100196.
- [44] Gregor Wernet, Christian Bauer, Bernhard Steubing, Jürgen Reinhard, Emilia Moreno-Ruiz, and Bo Weidema. 2016. The ecoinvent database version 3 (part I): overview and methodology. *The International Journal of Life Cycle Assessment* 21 (2016), 1218–1230.
- [45] Carole-Jean Wu, Ramya Raghavendra, Udit Gupta, Bilge Acun, Newsha Ardalani, Kiwan Maeng, Gloria Chang, Fiona Aga Behram, James Huang, Charles Bai, Michael K. Gschwind, Anurag Gupta, Myle Ott, Anastasia Melnikov, Salvatore Candido, David Brooks, Geeta Chauhan, Benjamin Lee, Hsien-Hsin S. Lee, Bugra Akayildiz, Maximilian Balandat, Joe Spisak, Ravi Kumar Jain, Michael G. Rabbat, and Kim M. Hazelwood. 2022. Sustainable ai: Environmental implications, challenges and opportunities. *Proceedings of Machine Learning and Systems* (2022).