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

Tianyao Shi, Ritbik Kumar, Inez Hua, Yi Ding

The 4th Workshop on Sustainable Computer Systems (HotCarbon'25)

PURDUE UNIVERSITY

We All Have Seen Biodiversity Loss...

But what you may not know is

Computing Can Be Part of the Reason

for biodiversity loss

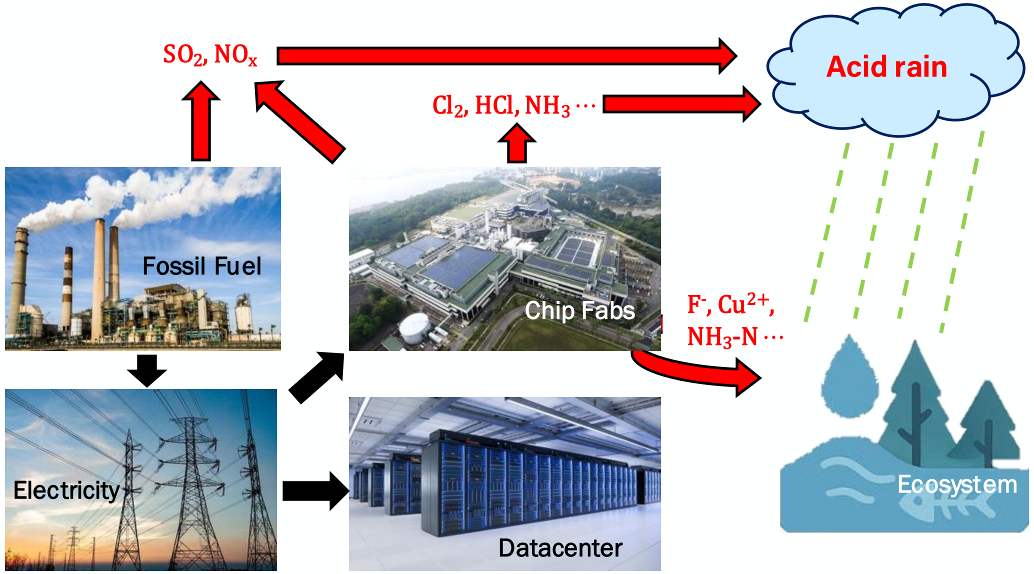

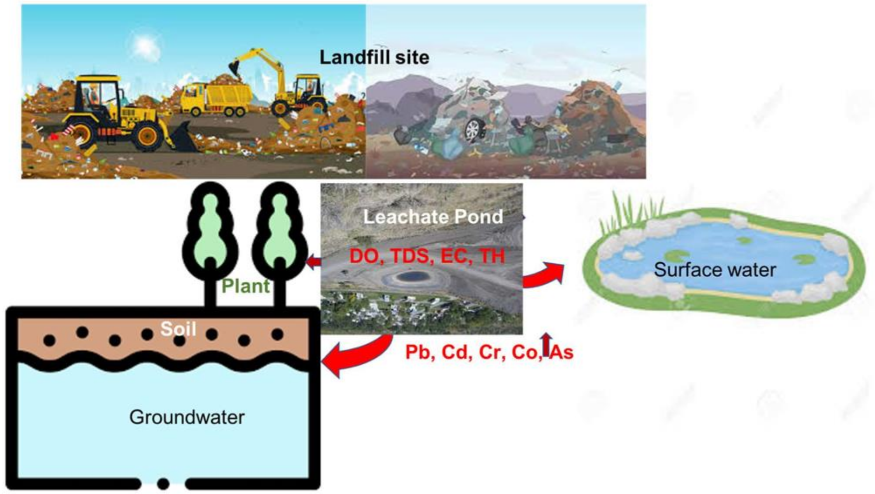

Rysul H. Hredoy et al., Impacts of Landfill Leachate on the Surrounding Environment: A Case Study on Amin Bazar Landfill, Dhaka (Bangladesh)

- **Computing's lifecycle generates direct environmental stressors.**
- Examples:
  - Chip Fabrication: Releases air/water pollutants.
  - Datacenter Power: Fossil fuel use drives acidification (acid rain).
  - End-of-Life: E-waste leaches heavy metals into ecosystems.

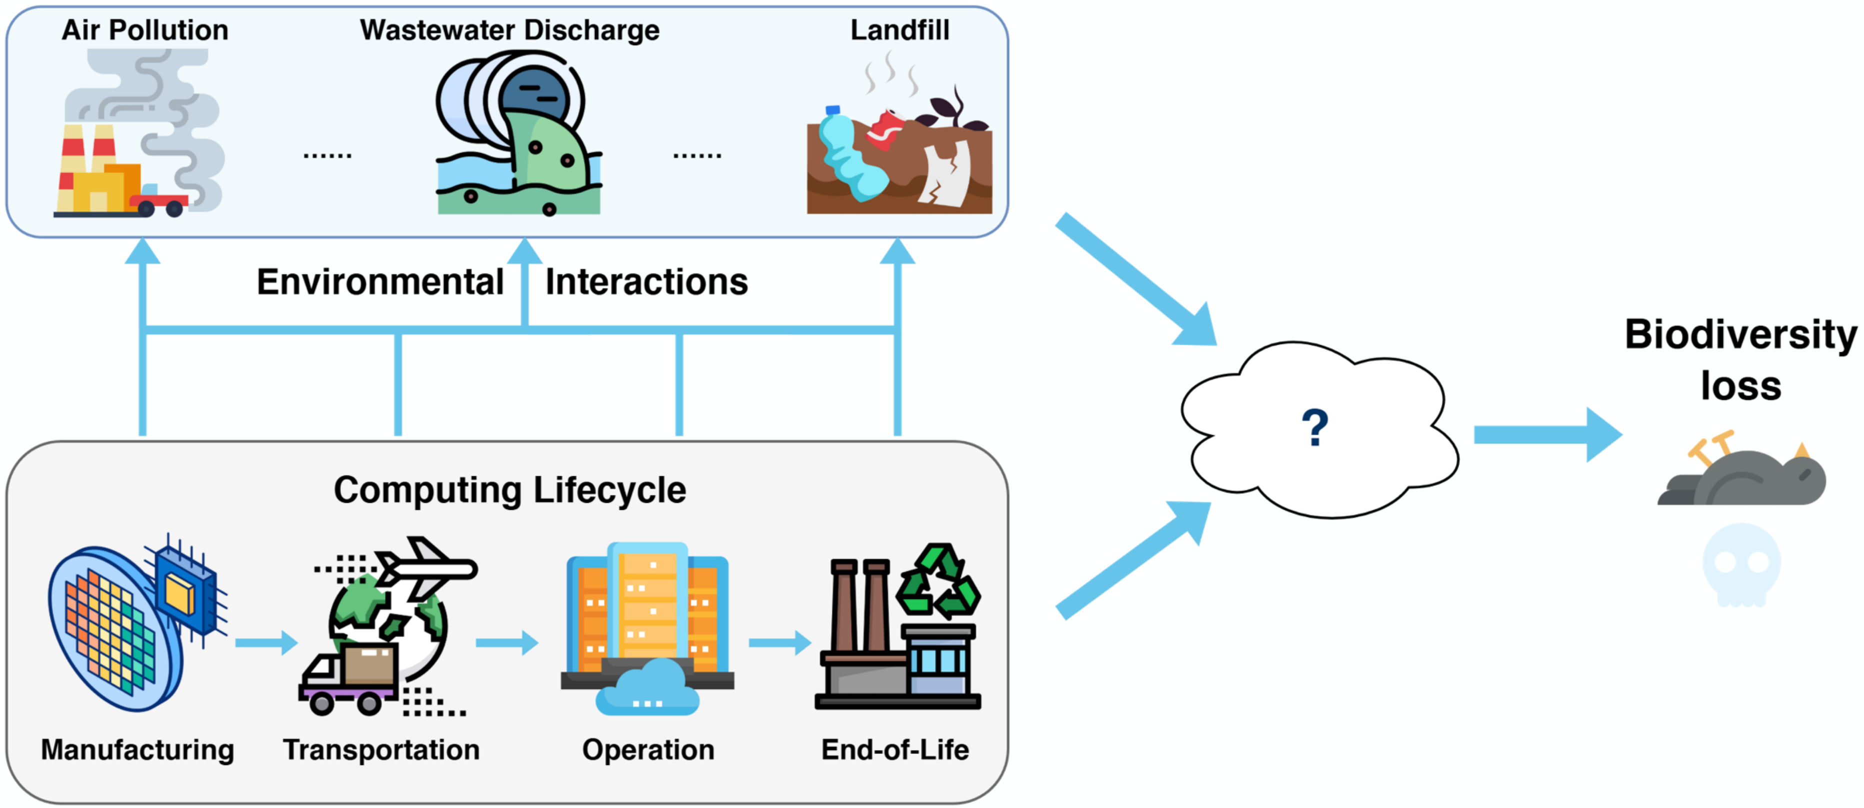

Why has no one ever explored it?

Research Challenges

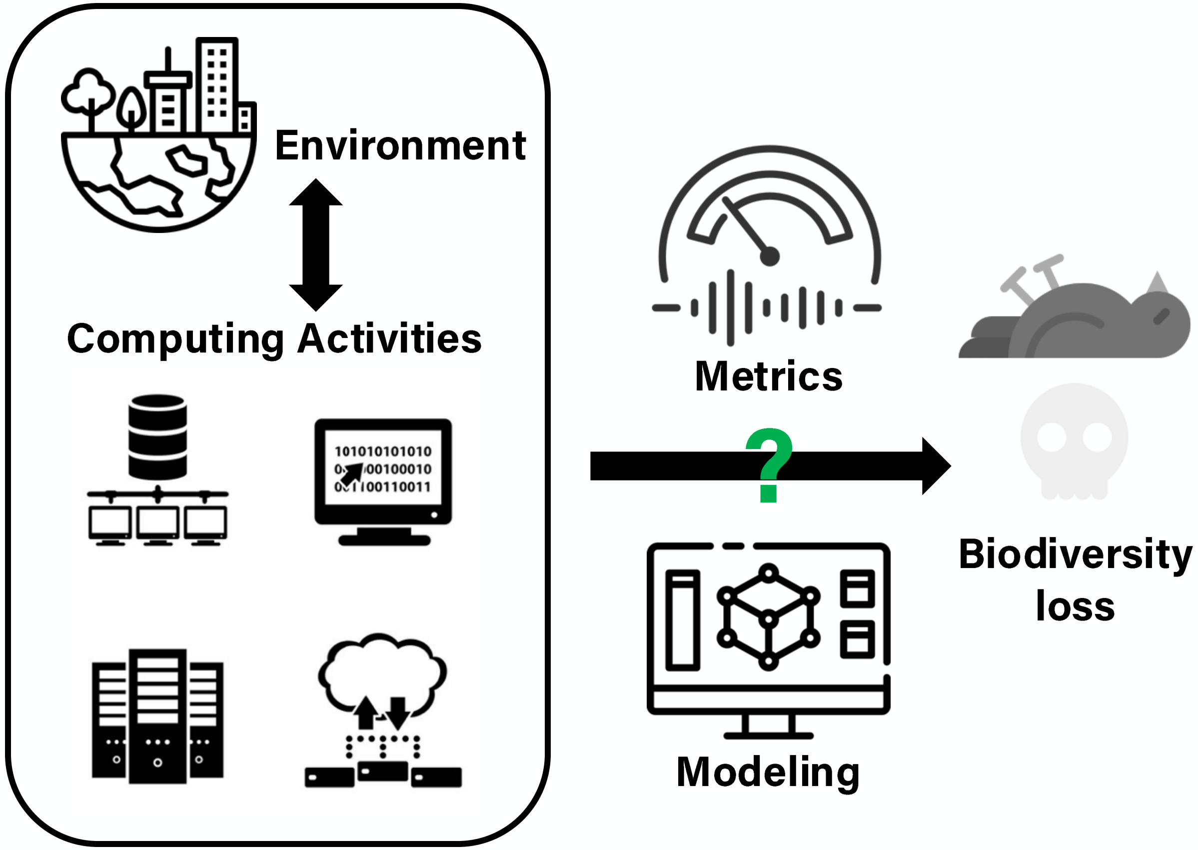

1. **The Metrics Barrier**: No quantifiable way to attribute biodiversity loss to a specific computing activities.

2. **The Modeling Barrier**: No modeling framework to connect specific computing workload and system to concrete biodiversity impacts across their lifecycle—from manufacturing to operation to end-of-life.

Our Contributions

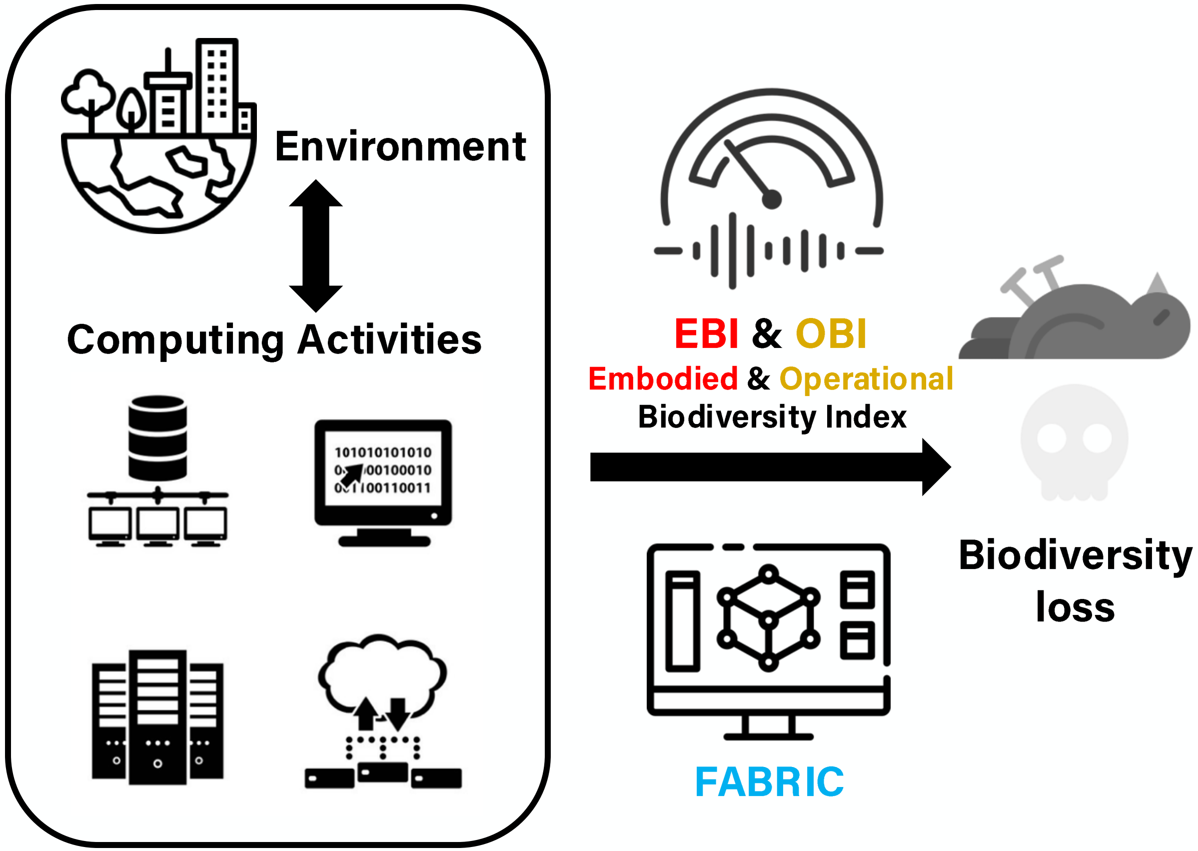

- Conduct the first analysis of biodiversity impact in computing.
- Introduce two new metrics-**EBI** and **OBI**-to quantify biodiversity impact from manufacturing, transportation, end-of-life and use.
- Present **FABRIC**, the first modeling framework that connects computing workloads to lifecycle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

Study Scope: Midpoints & Endpoints

- **Midpoint Impacts**: quantify direct environmental pressures induced by toxic substances.
  - **Acidification** (kg $SO_{2\,eq}$): how much emitted gases can lower environmental pH
  - **Eutrophication** (kg $PO_4^-{}_{eq}$): how excess nitrogen and phosphorus enrich ecosystems
  - **Freshwater Ecotoxicity** (CTUe): toxic impact of chemicals on freshwater ecosystems
- **Endpoint Impact**: use the ReCiPe 2016[1] model to convert midpoints into a single, unified biodiversity metric.

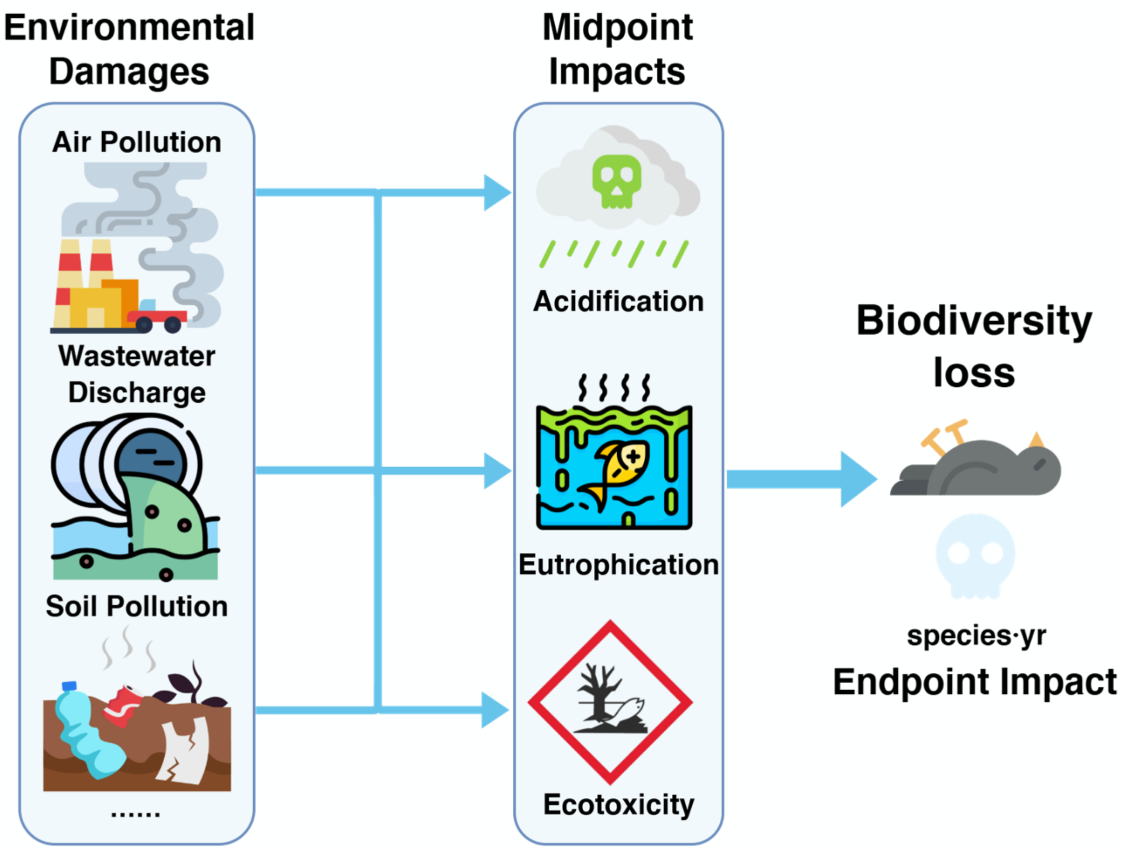

1: Mark A.J. Huijbregts et al., ReCiPe 2016: A Harmonized Life Cycle Impact Assessment Method at Midpoint and Endpoint Level. Report I: Characterization. RIVM National Institute for Public Health and the Environment, Bilthoven, The Netherlands.

Our Approach: The FABRIC Framework

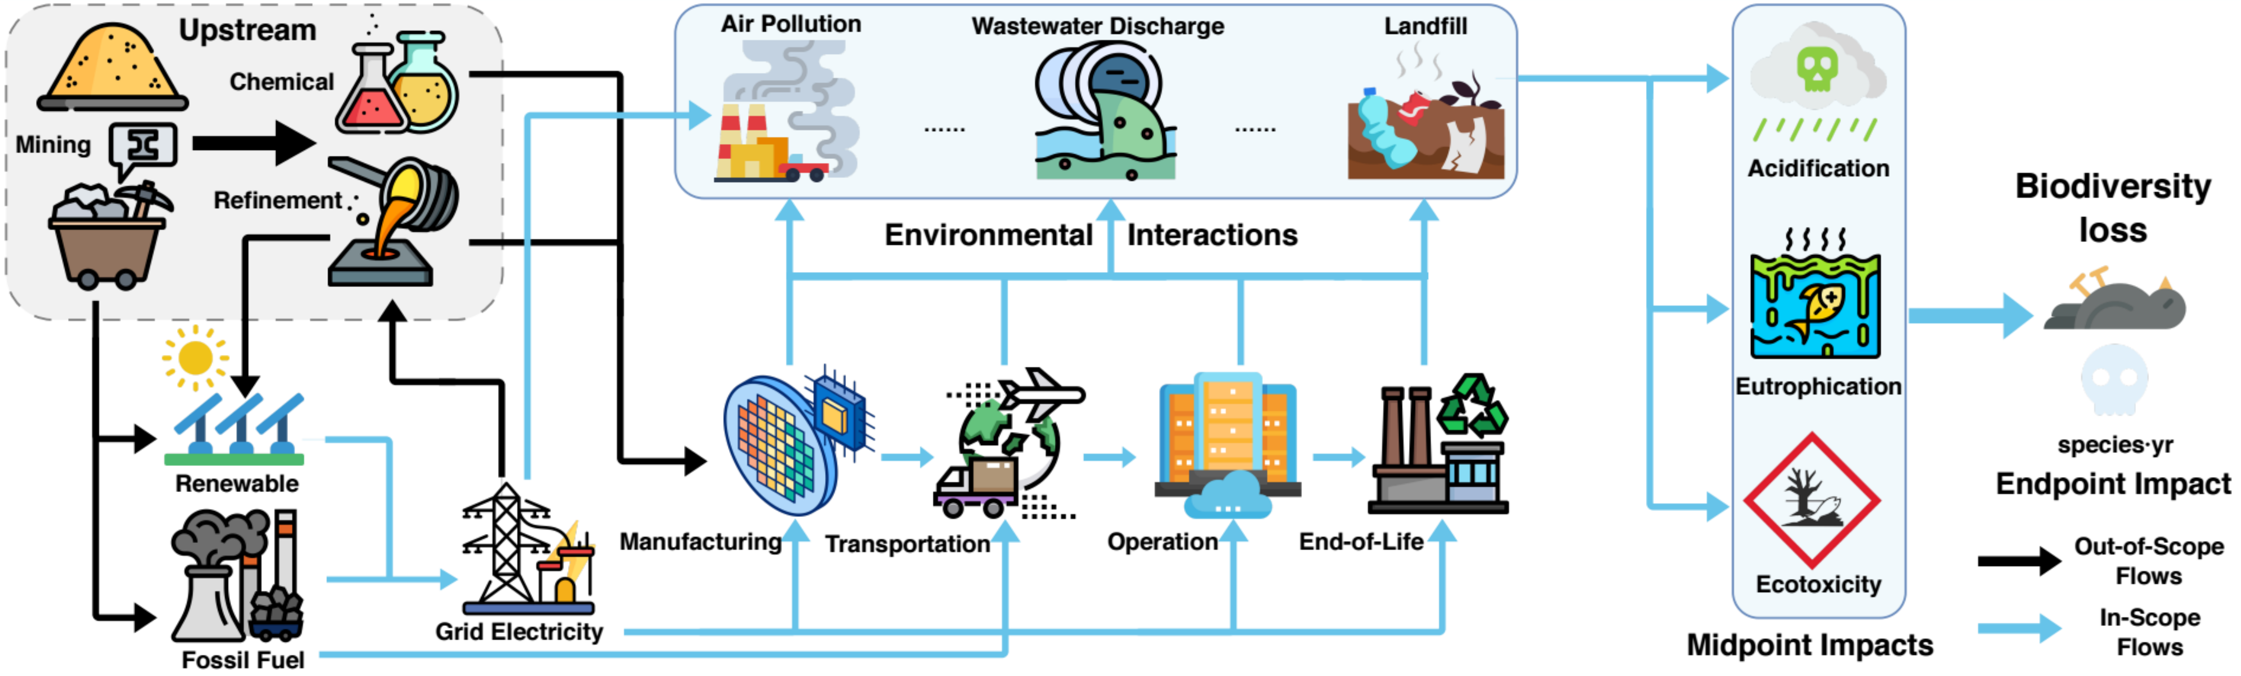

- **FABRIC: Fa**brication-to-grave **B**iodive**R**sity **I**mpact **C**alculator.

Our Approach: The FABRIC Framework

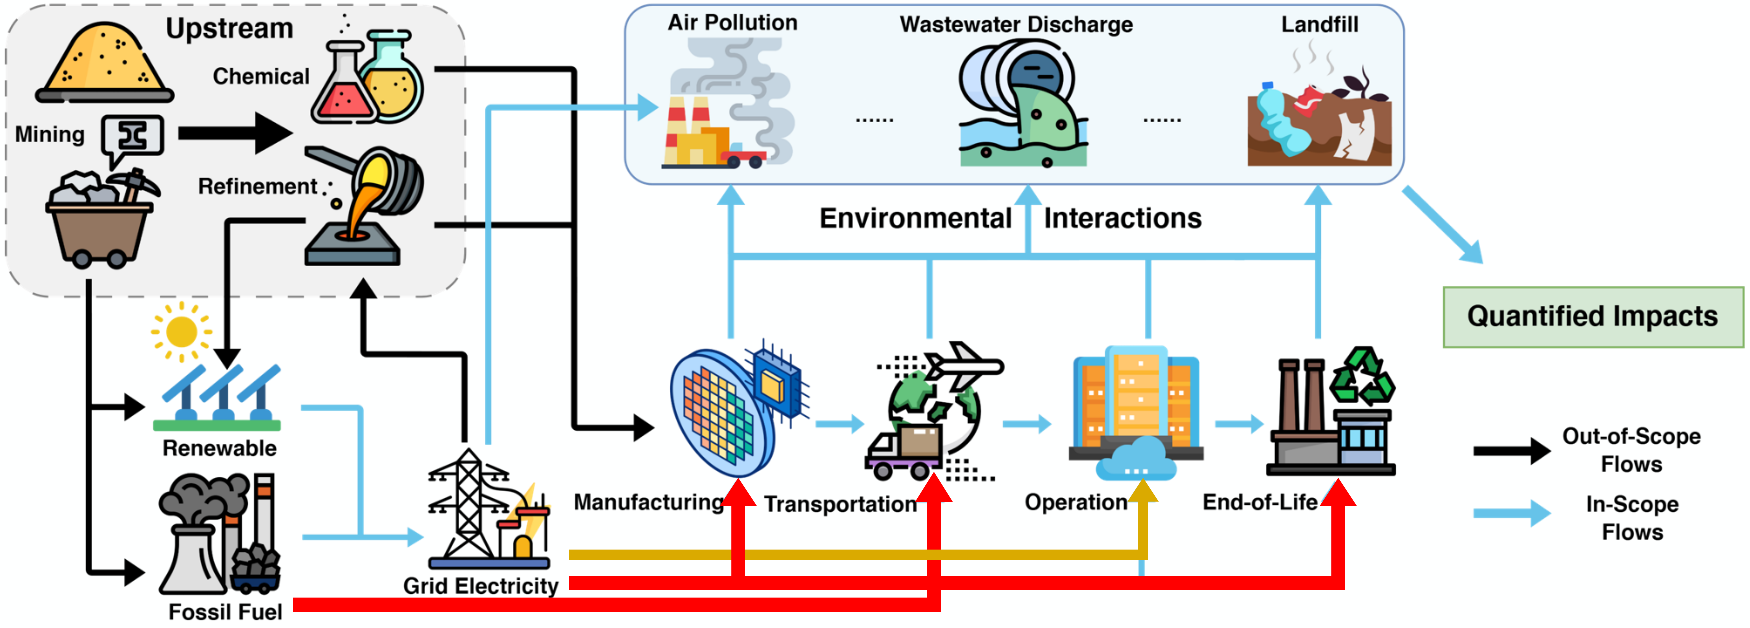

- **FABRIC: Fa**brication-to-grave **B**iodive**R**sity **I**mpact **C**alculator.
- Our Metrics to Overcome the Metrics Barrier:
  - **EBI** (Embodied Biodiversity Index): One-time impact of hardware's physical lifecycle, covers manufacturing, transportation, and end-of-life stages.
  - **OBI** (Operational Biodiversity Index): Ongoing impact of electricity use in the operational stage.

Embodied Biodiversity Index (EBI)

- Two kinds of environmental damage in EBI:
  1. Directly releasing toxic substances
  2. Indirectly causing emissions by driving fossil fuel use
- The EBI of a device is calculated by summing impacts across stages and impact categories

$$B_{emb}(d) = \sum_{c,l} M_{c,l}(d)\Phi(c)$$

- EBI of a specific workload is proportional to the execution time amortized over device lifetime

$$B_{emb}(w) = \sum_{d} \frac{t_w}{LT_d} B_{emb}(d)$$

- More details can be found in the paper...

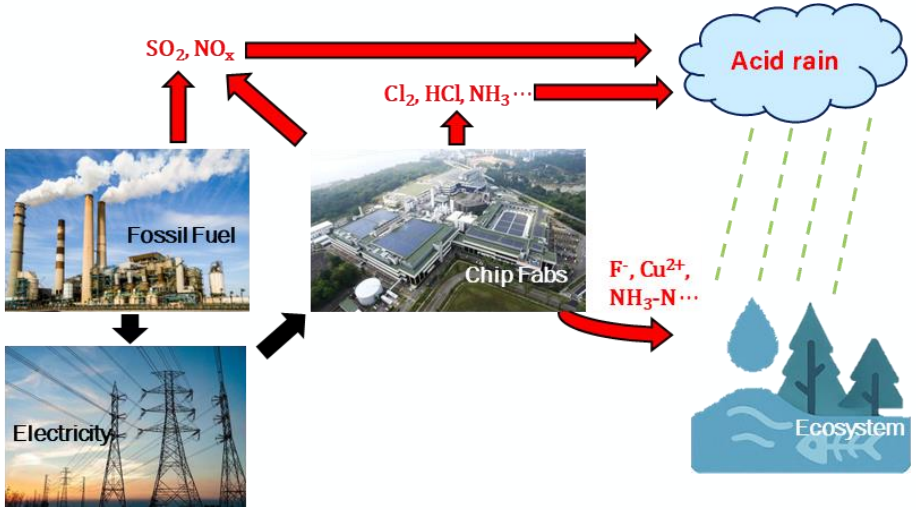

- EBI example:
  - *TSMC reports 39.5 t SO_2 emission in 2021*
  - $\Longrightarrow$ *0.057 g SO_2 per production unit*
  - $\Longrightarrow$ *0.029 g SO_2 per EPYC 7743 produced (based on die size, tech node, yield)*
  - $\Longrightarrow 2.9 \times 10^{-5}$ *kg $SO_{2\,eq}$ AP, 0 EP & FETP*
  - $\Longrightarrow 6.2 \times 10^{-12}$ *species·yr*
  - *...(Other substances)*

Operational Biodiversity Index (OBI)

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

$$B_{\mathrm{op}}(d) = \sum_{c} M_{c,\mathrm{Use}}(d)\Phi(c)$$

- Task-specific OBI of w is proportional to its energy consumption $E_{\mathrm{el}}\ (w)$

$$B_{\mathrm{op}}(w) = \sum_{c} M_{c,\mathrm{Use}}(w)\Phi(c)$$

characterization factor

$$M_{c,\mathrm{Use}}(w) = \sum_{k} E_{\mathrm{el}}\ (w) F_{\mathrm{el},k,r} \Gamma_k^c$$

regional grid emission factor (g/kWh)

- More details can be found in the paper...

Evaluation Methodology

- Devices

CPU	Year	Node	Silicon Area	Cores
AMD EPYC 7B12	2019	7/14 nm	1008 mm^2	64
AMD EPYC 7443	2021	7/12 nm	740 mm^2	24
AMD EPYC 7B13	2021	7/12 nm	1064 mm^2	64
AMD EPYC 9B14	2023	5/6 nm	1261 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

- Systems
  - *Local testbed server:* typical edge or lab workstation
  - *Gautschi**: 2023-era community cluster
  - *Perlmutter*: petascale supercomputer with mixed CPU/GPU architecture

- Workloads
  - Seven HPC workloads from Phoronix Test Suite
  - Scientific/Analytic kernel: `fft`, `npb`, `spark`
  - Encryption: `openssl`
  - Compilation: `build-Linux-kernel`
  - Compress: `compress-{pbzip2,lz4,zstd,gzip,lzma,xz}`
  - Video encoding: `x264`

- Research questions

1. Which lifecycle stage contributes most to computing's biodiversity impact?
2. How do biodiversity impacts vary across HPC workloads and system platforms?
3. How does deployment location affect computing's biodiversity impacts?

EBI Breakdown

RQ1: Which lifecycle stage contributes most to computing's biodiversity impact?

- Manufacturing overwhelmingly dominates embodied impact (55-75% of EBI).
- Acidification is the single largest midpoint impact (60-85%).
- Newer, more efficient devices have a much lower biodiversity impact per TFLOP.

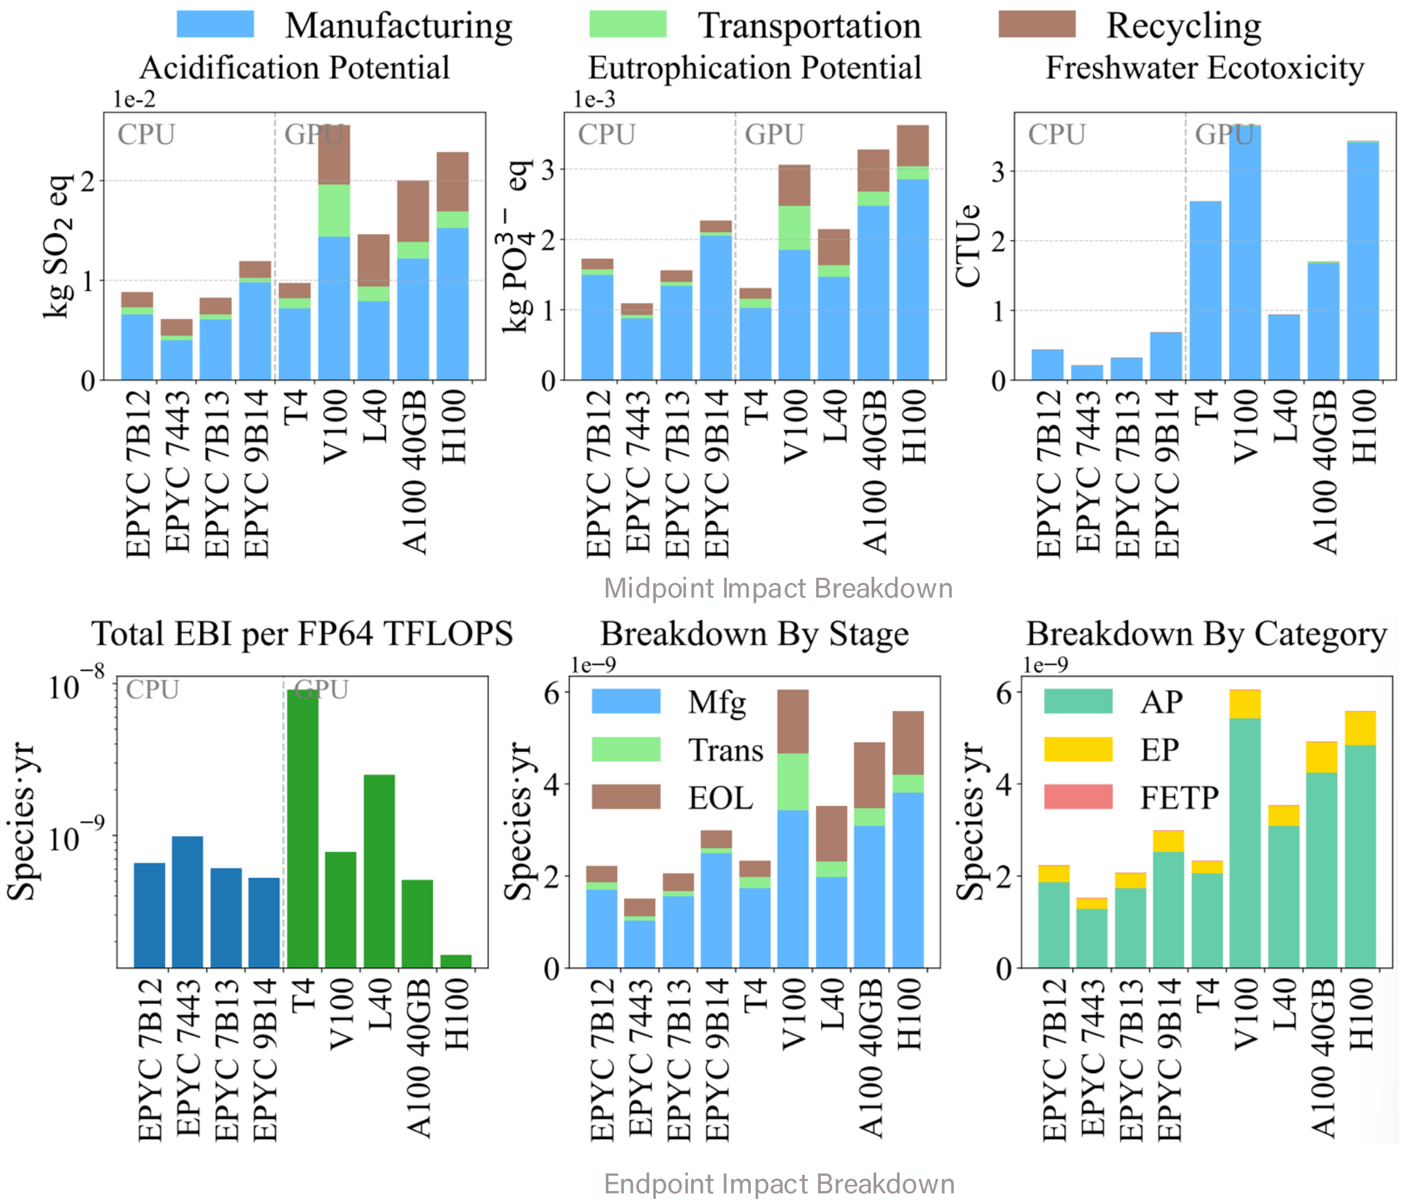

System-Level EBI Contributions

System	Main HW Year	CPUs (model x count)	GPUs (model x count)	Total DRAM	Storage Capacity
Local testbed server	2022	EPYC 7443 x 1	L40 x 4	520 GB	2 TB SSD + 20 TB HDD
Gautschi*	2023	EPYC 9B14 x 442	H100 x 160 + L40 x 12	158 TB	1.5 PB SSD + 2.5 PB HDD
Perlmutter	2021	EPYC 7B13 x 4864	A100 80GB x 1024 + A100 40GB x 6144	1984 TB	44 PB SSD

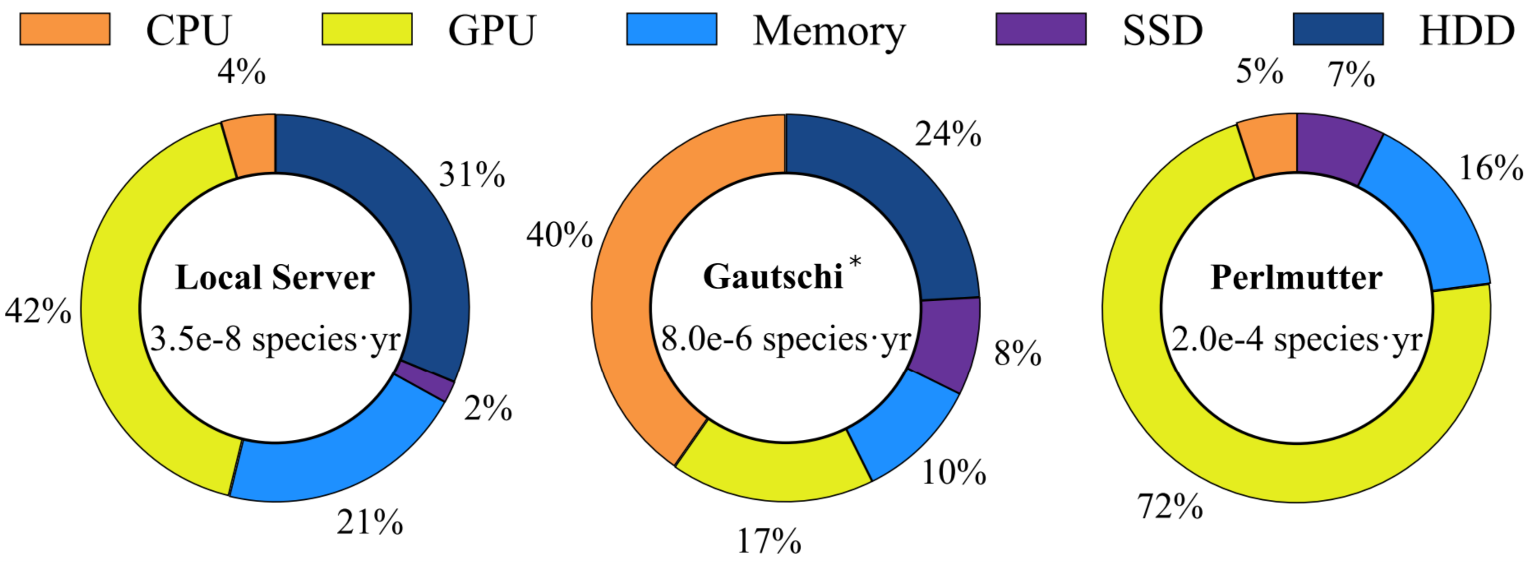

- System-level EBI depends both on the system scale and hardware composition.

Life-Cycle Analysis for Individual Devices

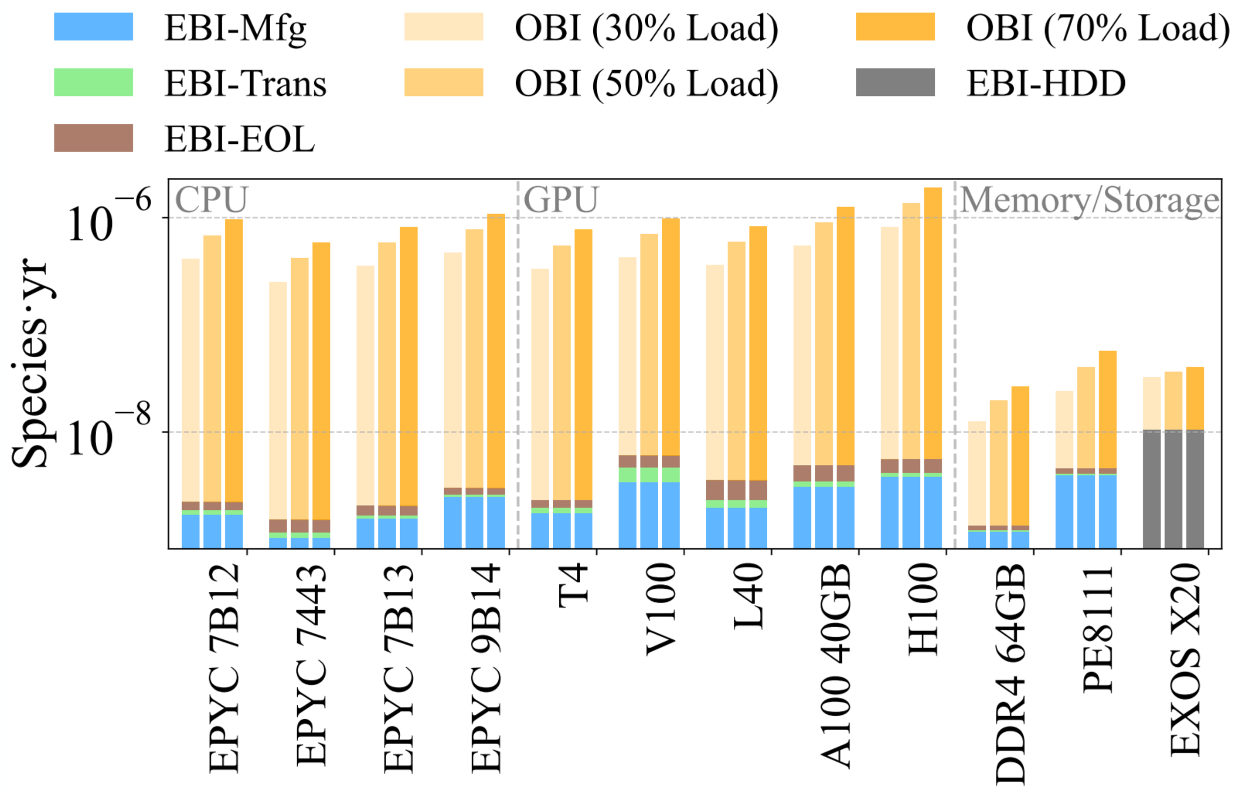

- At 70% load, OBI can be nearly **100x higher** than annualized EBI for CPUs/GPUs.

Takeaway:

When using conservative fab-to-grave estimate for EBI and U.S. grid average emission, operational electricity (OBI) dominates the lifecycle biodiversity impact.

System-Level Lifecycle Impact

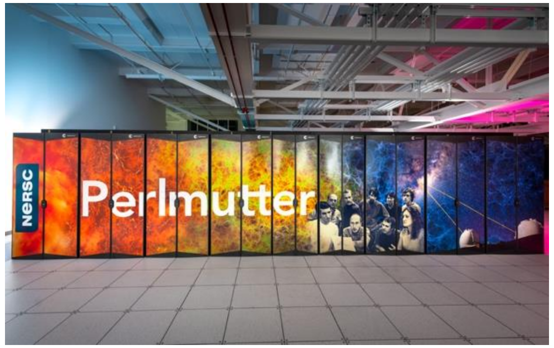

Ranks 25th in the Top 500 list

FABRIC

Annual EBI

$\mathbf{4 \times 10^{-5}}$ species·yr

60x

$\mathbf{2.5 \times 10^{-3}}$ species·yr

Annual OBI
(70% load, US grid)

TOP 500 The List.

$\gtrsim \mathbf{1}$ species·yr

- Assuming 70% system load and U.S. grid average emission factors, Perlmutter's annual OBI reaches 60x that of annual EBI.
- If 400 Perlmutter-class systems are active globally, the total biodiversity impact $\gtrsim$**1** species·yr is non-negligible.

Both How and Where You Compute Matter

RQ2: How do biodiversity impacts vary across HPC workloads and system platforms?

- GCP n2d-standard-64, 32 cores (1/2 EPYC 7B12, 2019)
- Local testbed server, 24 cores (1 EYPC 7443, 2021)
- GCP c3d-standard-60, 30 cores (5/16 EPYC 9B14, 2023)

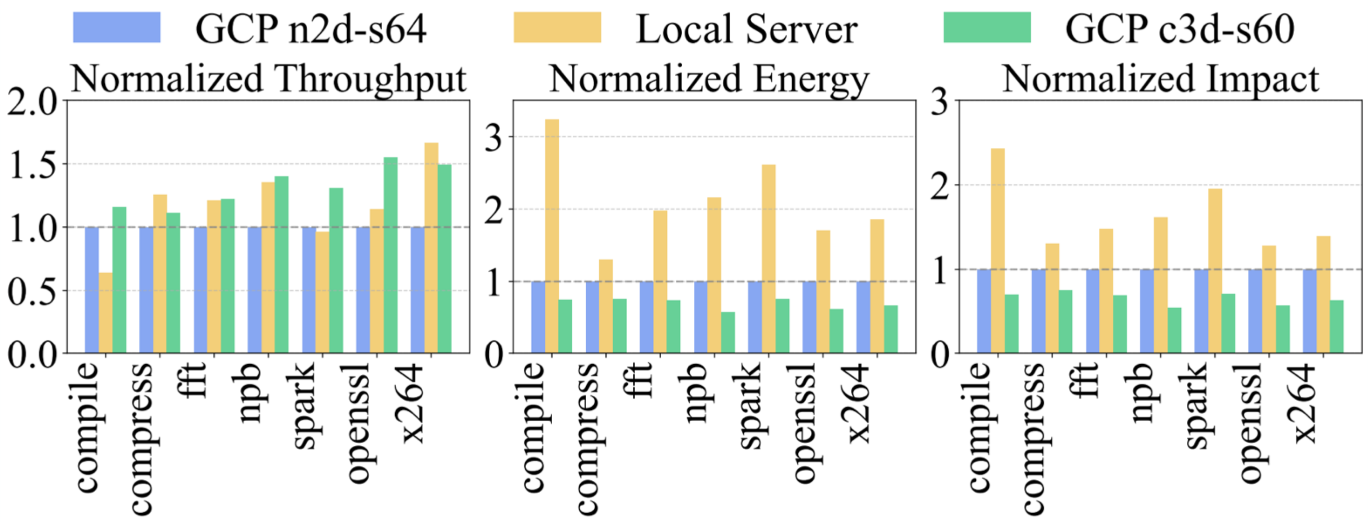

Takeaway:

- Modern cloud CPUs from GCP can **halve** the biodiversity impact per unit of work.
- An un-optimized local server can **double** it even the performance appears comparable.
- Datacenter-level power management is critical.

Location, Location, Location

RQ3: How does deployment location affect computing's biodiversity impacts?

- Low-carbon grids are not always low-impact for biodiversity.
- A hydro-based grid (**Québec**) cuts the operational impact by two orders of magnitude.

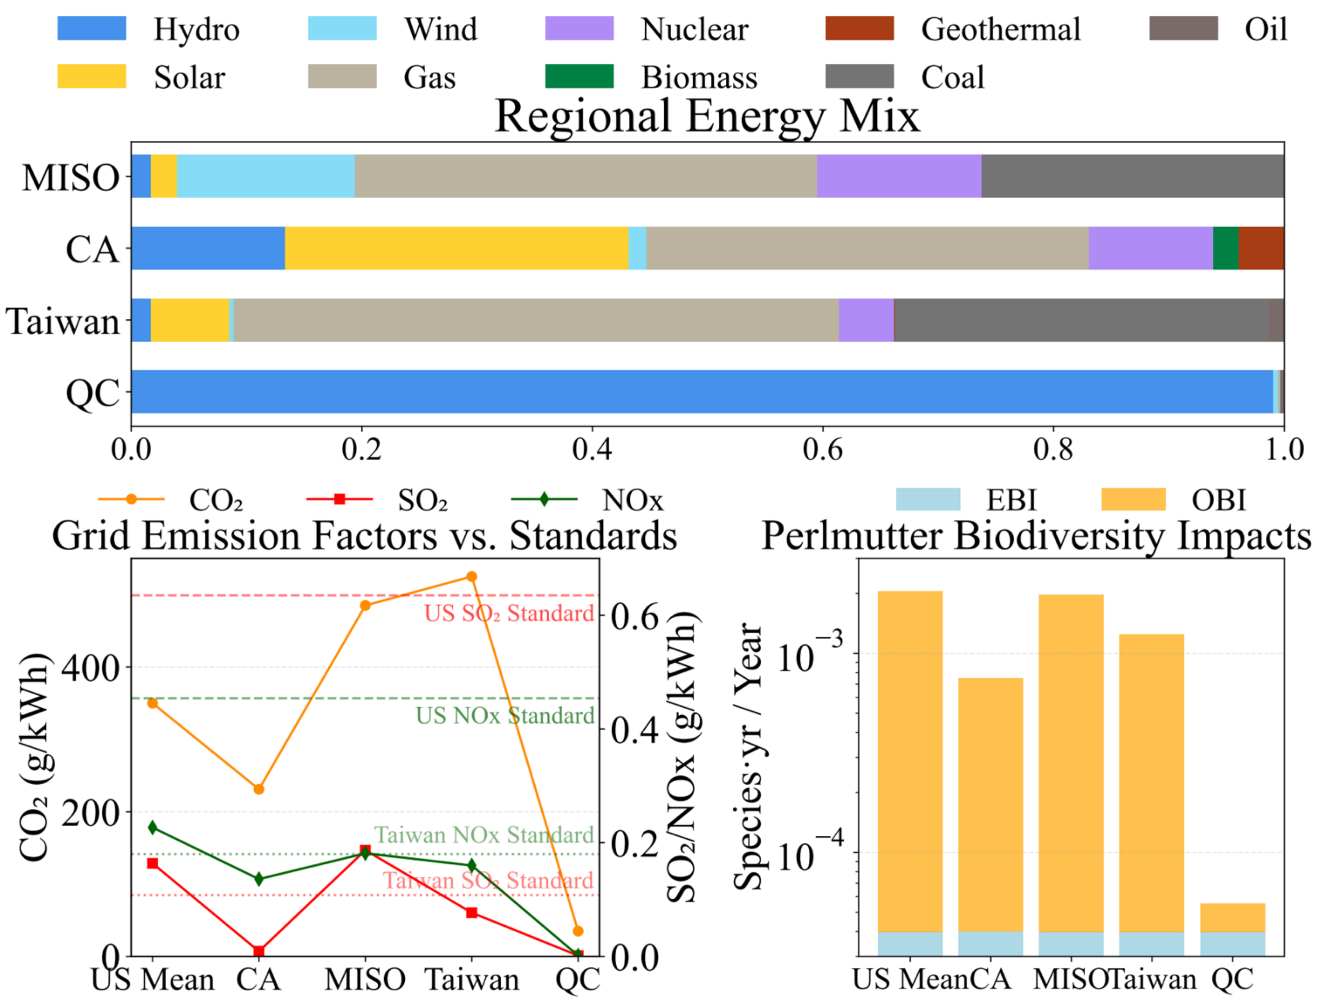

Takeaway:

- Focusing solely on CO_2 is not enough—strict SO_2/NO_x limits and renewable-heavy grids can also greatly cut the total biodiversity impact. This highlights the need for pollutant-specific policies.

Take-home Message

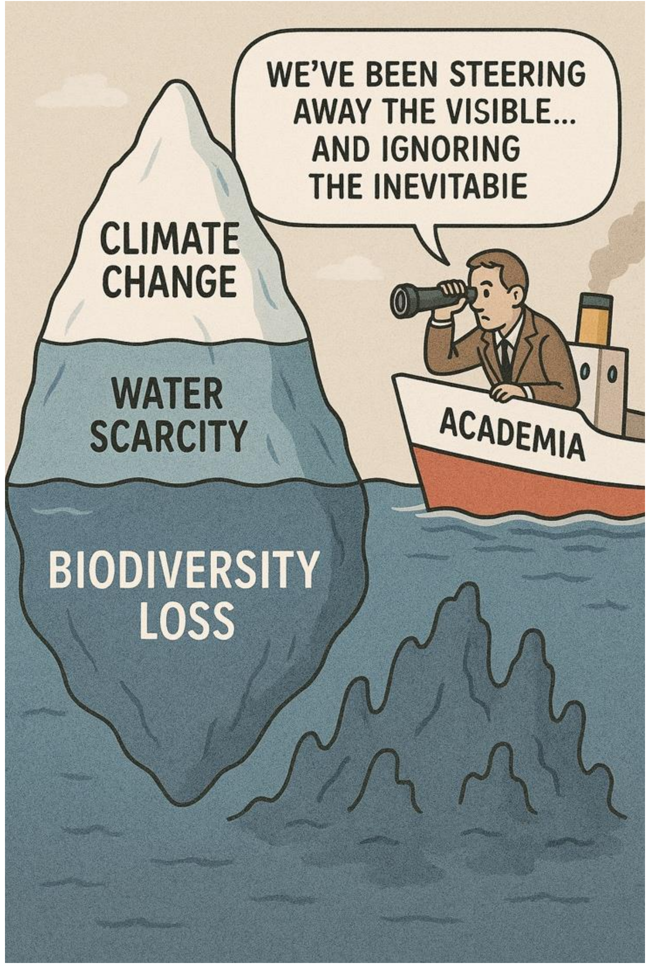

- By lifecycle stages, **manufacturing** dominates EBI.
- By impact categories, **acidification** dominates EBI.
- When using conservative fab-to-grave estimate for EBI and U.S. grid average emission, operational electricity (**OBI**) **dominates** the lifecycle biodiversity impact.
- **Strict SO_2/NO_x limits** and **renewable-heavy** grids can cut a supercomputer's total biodiversity impact by an order of magnitude.

THANK YOU!

Questions?