Artifacts Available V1.1
Artifacts Evaluated Functional V1.1
Results Reproduced V1.1

ThirstyFLOPS: Water Footprint Modeling and Analysis Toward Sustainable HPC Systems




Yankai Jiang

Raghavendra Kanakagiri

Rohan Basu Roy

Devesh Tiwari

Northeastern University

भारतीय प्रौद्योगिकी संस्थान तिरुपति
IIT TIRUPATI

THE UNIVERSITY OF UTAH

Supercomputing Centers: Energy, Carbon, and Water

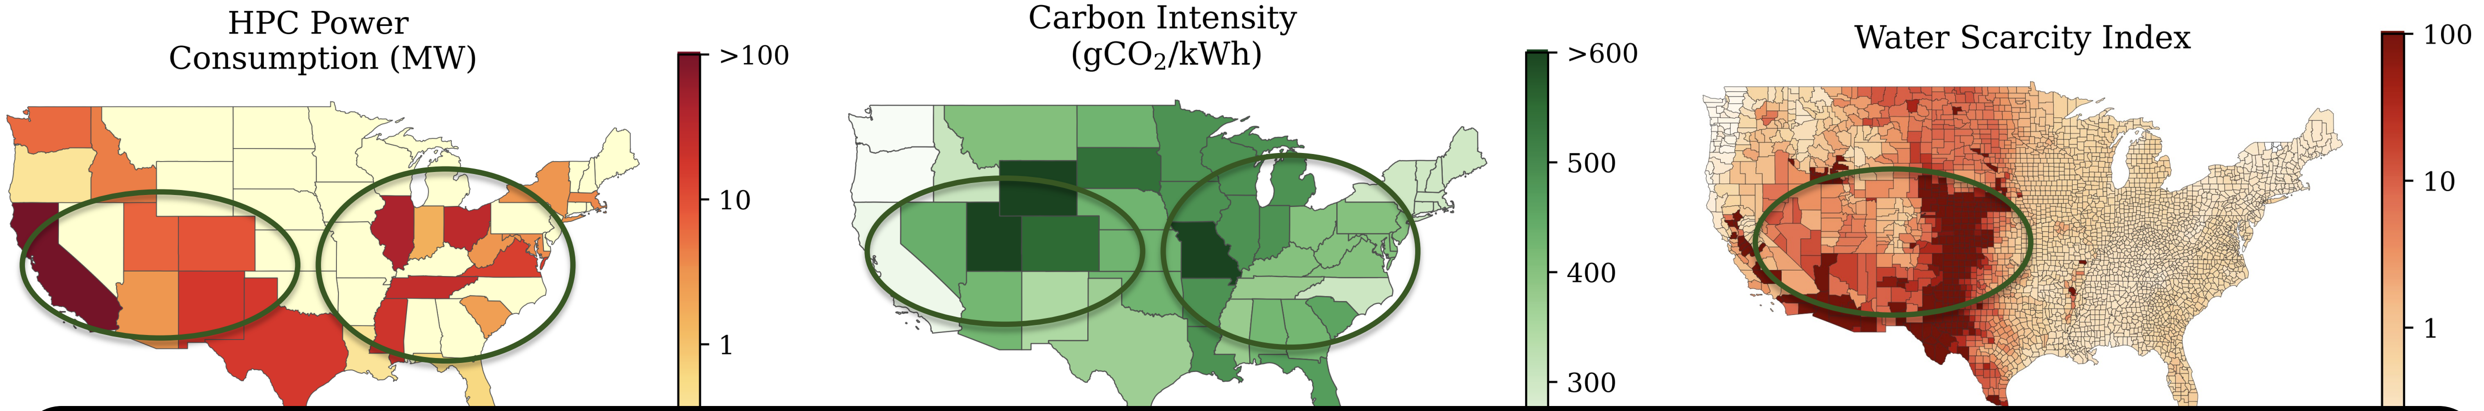

How do we even quantify and understand the water footprint of our HPC systems?

The current HPC centers may not always be located in the most carbon-friendly or water-rich places. In fact, some HPC datacenters may be located in relatively water-scarce regions.

ThirstyFLOPS: Executive Summary

ThirstyFLOPS
A tool for estimating the water footprint of HPC

✓ **A Novel Toolset for Estimating the Water Footprint:** quantification of relative importance and contribution of different factors (indirect vs. direct water consumption).

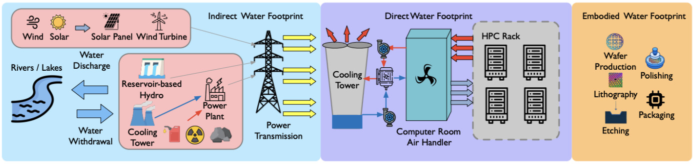

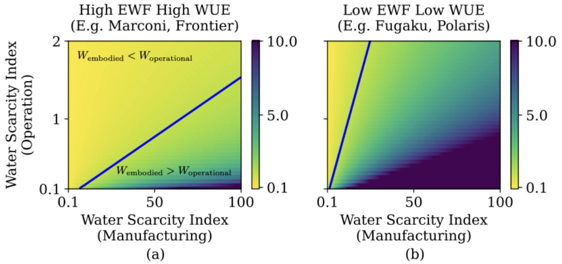

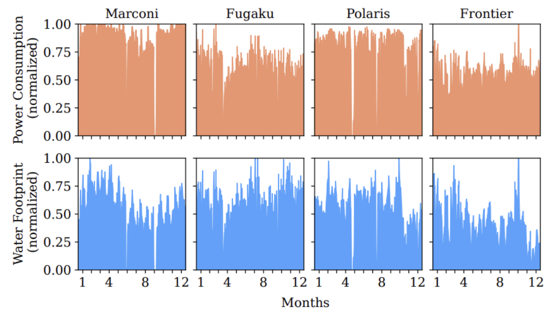

ThirstyFLOPS: Executive Summary

ThirstyFLOPS
A tool for estimating the water footprint of HPC

- ✓ **A Novel Toolset for Estimating the Water Footprint:** quantification of relative importance and contribution of different factors (indirect vs. direct water consumption).

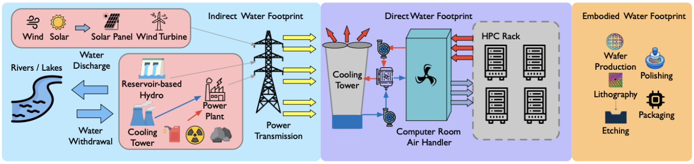

- ✓ **Impact of Regional Water Scarcity on Future HPC Data Center Site Selection:** The utility of the same amount of water volume varies across different geographical locations.

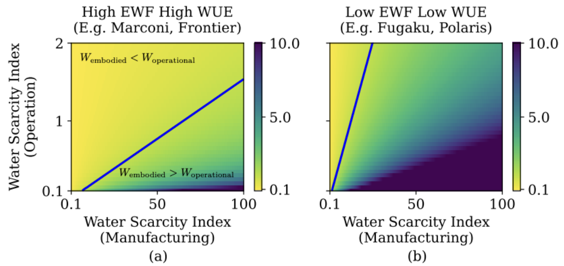

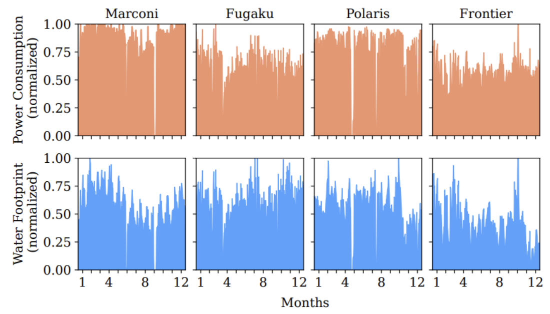

ThirstyFLOPS: Executive Summary

ThirstyFLOPS
A tool for estimating the water footprint of HPC

- ✓ **A Novel Toolset for Estimating the Water Footprint:** quantification of relative importance and contribution of different factors (indirect vs. direct water consumption).

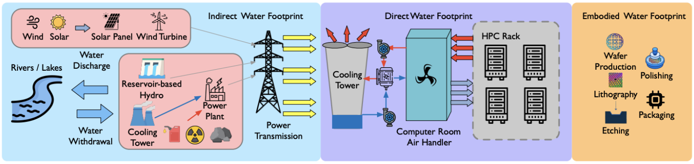

- ✓ **Impact of Regional Water Scarcity on Future HPC Data Center Site Selection:** The utility of the same amount of water volume varies across different geographical locations.

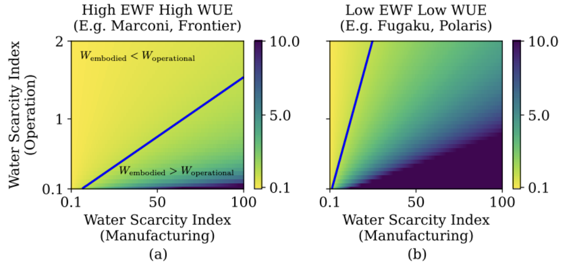

- ✓ **Competing Sustainability Goals (Carbon & Water):** Greener energy sources often come with higher water intensity, and hardware components with lower manufacturing carbon footprint can have higher water footprint.

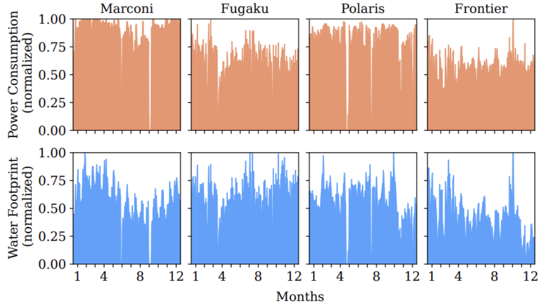

ThirstyFLOPS: Executive Summary

ThirstyFLOPS
A tool for estimating the water footprint of HPC

- ✓ **A Novel Toolset for Estimating the Water Footprint:** quantification of relative importance and contribution of different factors (indirect vs. direct water consumption).

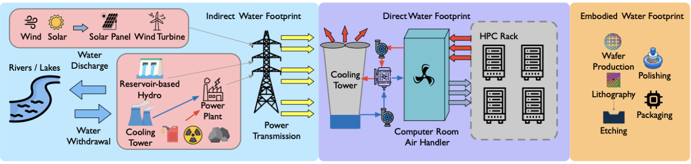

- ✓ **Impact of Regional Water Scarcity on Future HPC Data Center Site Selection:** The utility of the same amount of water volume varies across different geographical locations.

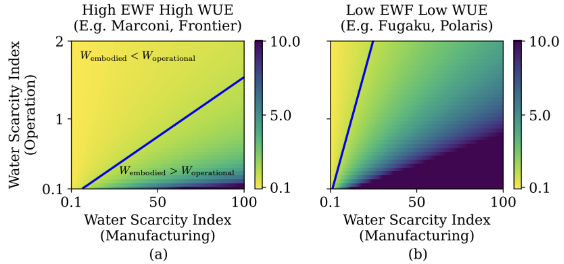

- ✓ **Competing Sustainability Goals (Carbon & Water):** Greener energy sources often come with higher water intensity, and hardware components with lower manufacturing carbon footprint can have higher water footprint.

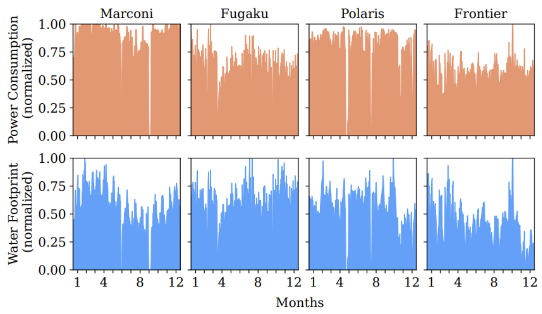

- ✓ **Nuclear-powered HPC and Water Footprint:** Small modular nuclear reactors (SMRs) are promising for exploding energy needs, but SMRs have significantly higher water footprint.

HPC Water Footprint Modeling

$$W = W_{\text{embodied}} + W_{\text{operational}} = W_{\text{embodied}} + W_{\text{direct}} + W_{\text{indirect}}$$

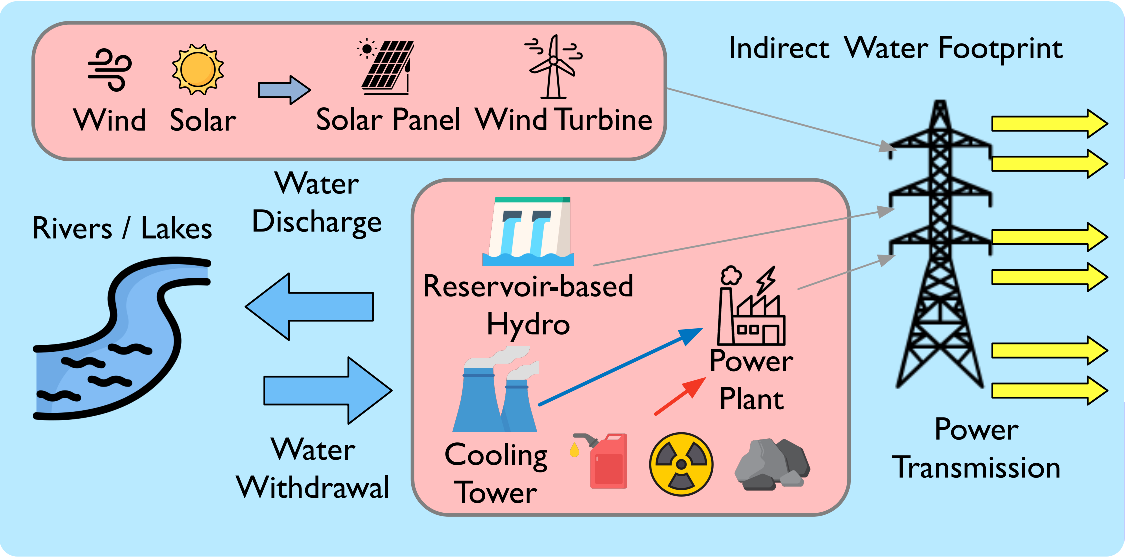

$$W_{\text{indirect}} = E \cdot \text{PUE} \cdot \text{EWF}$$

$$\text{EWF} = \frac{\text{energy generation water}}{\text{power generated}} = f\,(\text{mix\%}, \text{EWF}_{\text{energy}})$$

Indirect Water Footprint is the amount of water consumed during energy generation. The energy water factor (EWF) is the metric that determines the indirect water footprint.

HPC Water Footprint Modeling

$$W = W_{\text{embodied}} + W_{\text{operational}} = W_{\text{embodied}} + W_{\text{direct}} + W_{\text{indirect}}$$

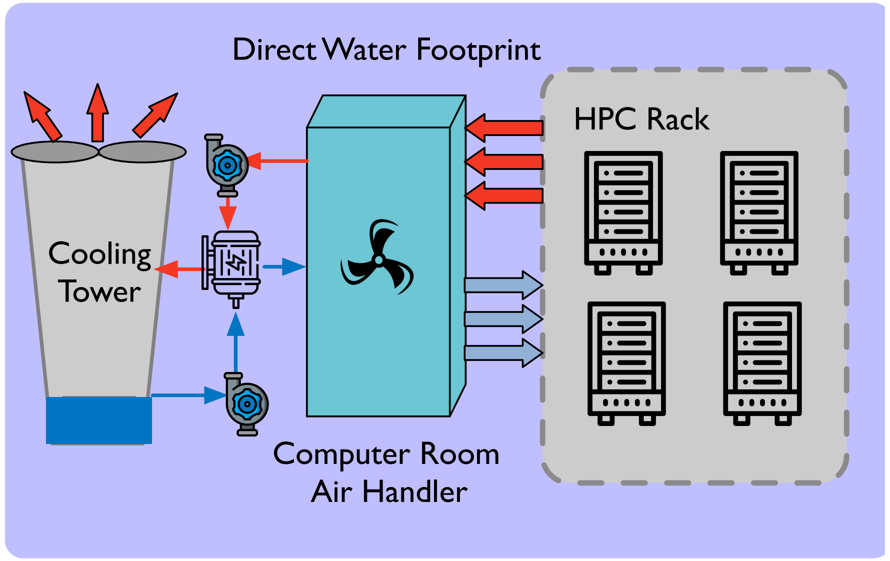

$$W_{\text{direct}} = E \cdot \text{WUE}$$

$$\text{WUE} = \frac{\text{cooling water}}{\text{power consumption}} = f\left(\text{Air}_{\text{temperature}}, \text{humidity}\right)$$

Direct Water Footprint is the water consumed during the cooling process in the HPC systems. Water Usage Effectiveness (WUE) quantifies the amount of water needed to cool one unit of energy (lower is better).

HPC Water Footprint Modeling

$$W = W_{\text{embodied}} + W_{\text{operational}} = W_{\text{embodied}} + W_{\text{direct}} + W_{\text{indirect}}$$

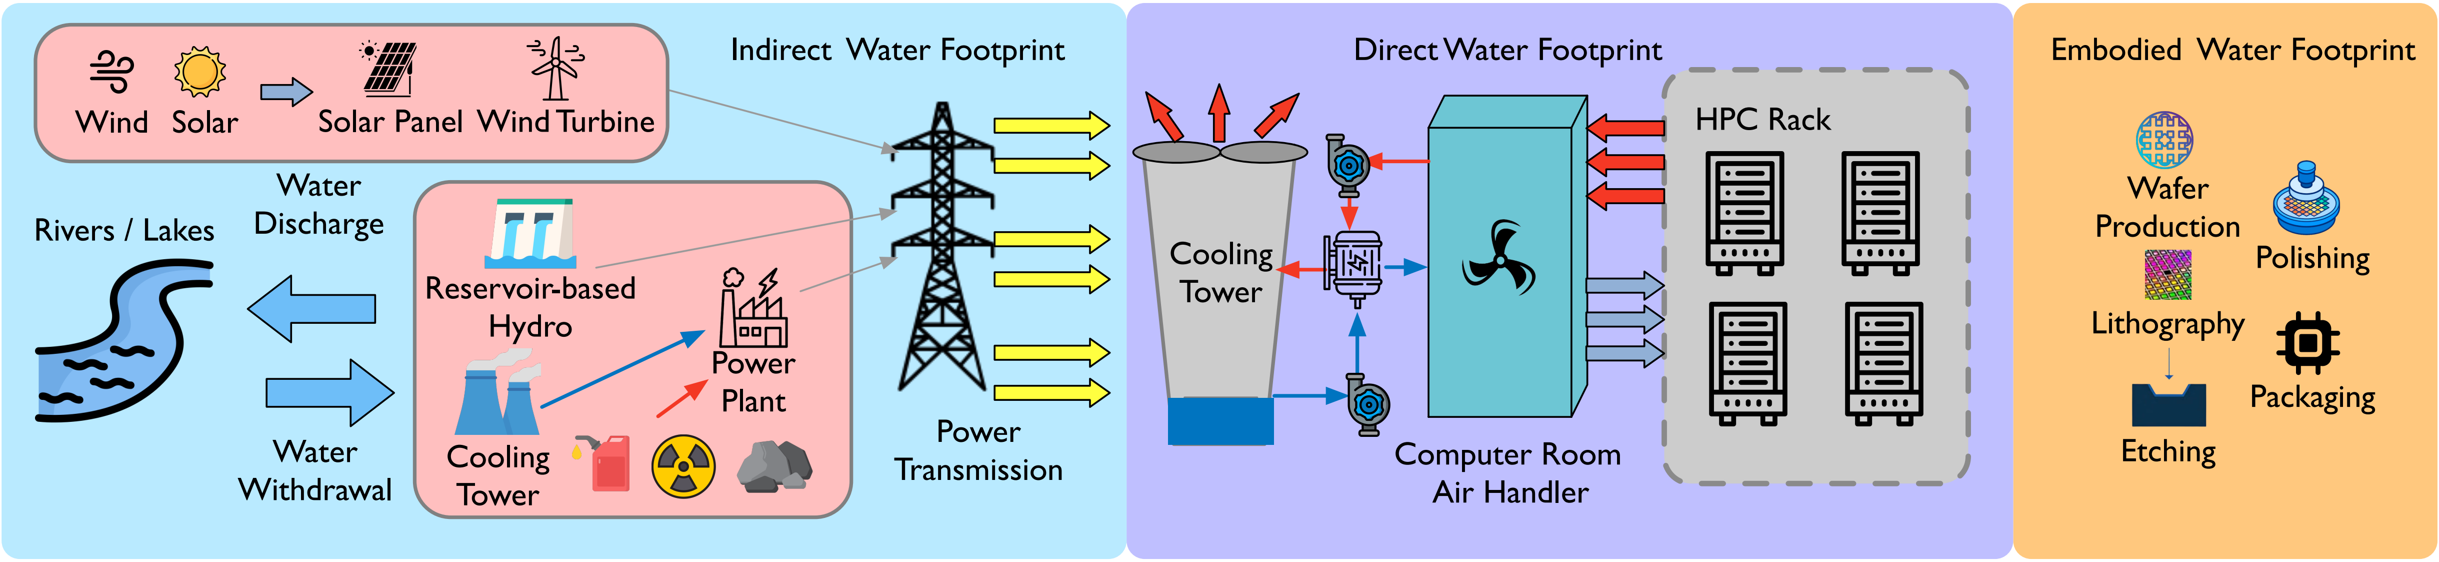

Embodied Water Footprint: water consumed during hardware packaging and manufacturing.

Refer to the papers for detailed embodied water footprint modeling.

HPC Water Footprint Modeling

$$W = W_{\text{embodied}} + W_{\text{operational}} = W_{\text{embodied}} + W_{\text{direct}} + W_{\text{indirect}}$$

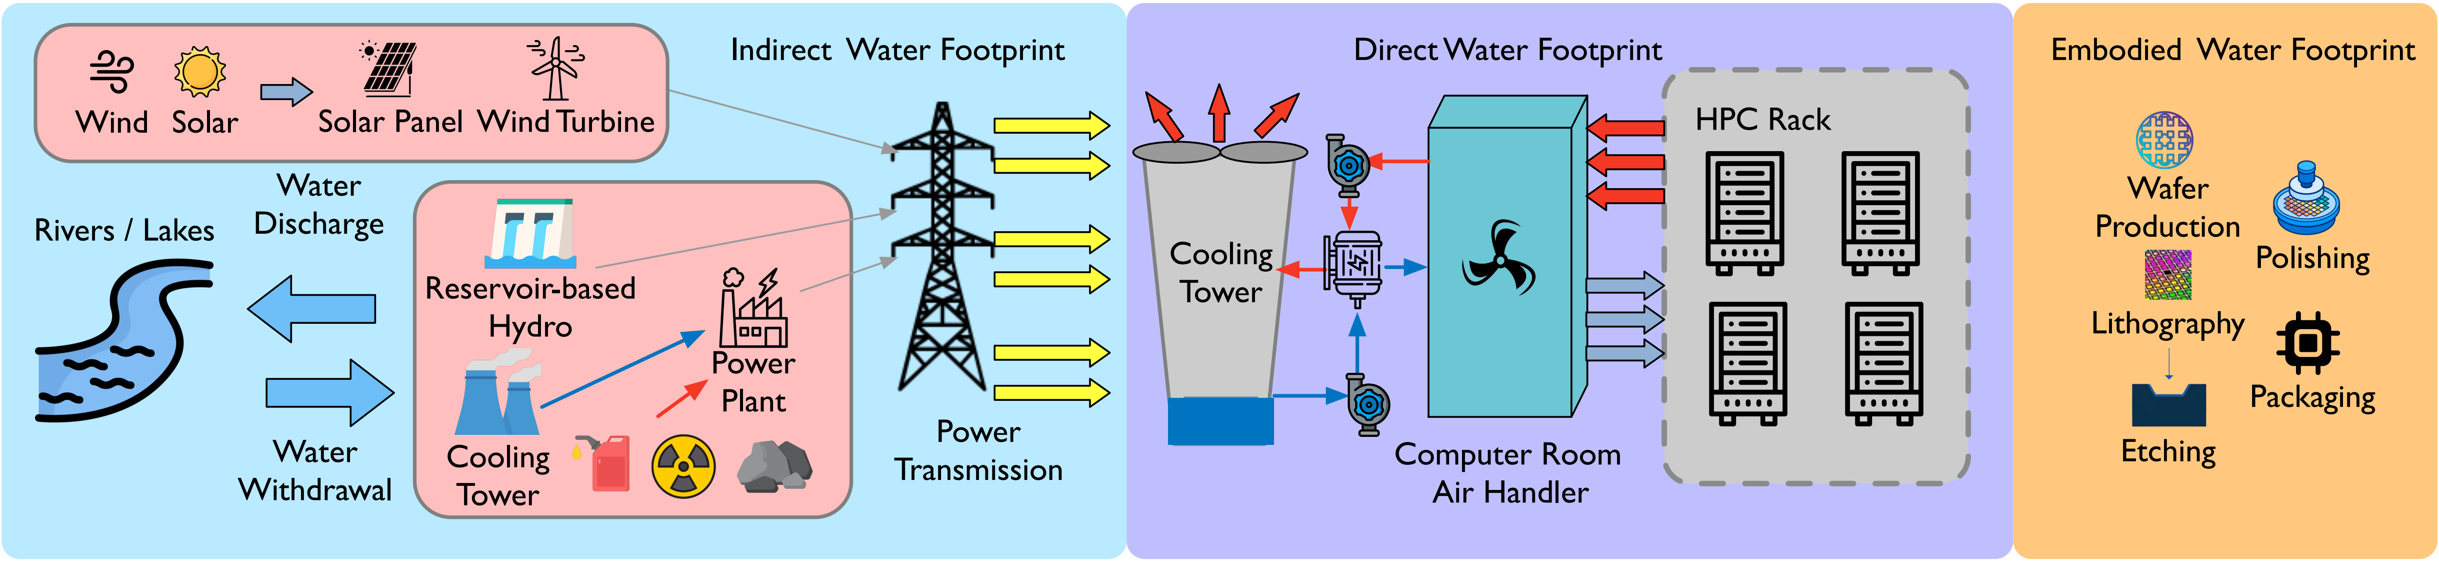

$$\begin{aligned} W_{\text{operational}} &= W_{\text{direct}} + W_{\text{indirect}} \\ &= E \cdot \mathrm{WUE} + E \cdot \mathrm{PUE} \cdot \mathrm{EWF} \\ &= E \cdot (\mathrm{WUE} + \mathrm{PUE} \cdot \mathrm{EWF}) = E \cdot \mathrm{WI} \end{aligned}$$

$$\text{Water Intensity (WI)} = \mathrm{WUE} + \mathrm{PUE} \cdot \mathrm{EWF}$$

Water Intensity: it can simplify the calculation of operational water footprint.

HPC Water Footprint Modeling

$$W = W_{\text{embodied}} + W_{\text{operational}} = W_{\text{embodied}} + W_{\text{direct}} + W_{\text{indirect}}$$

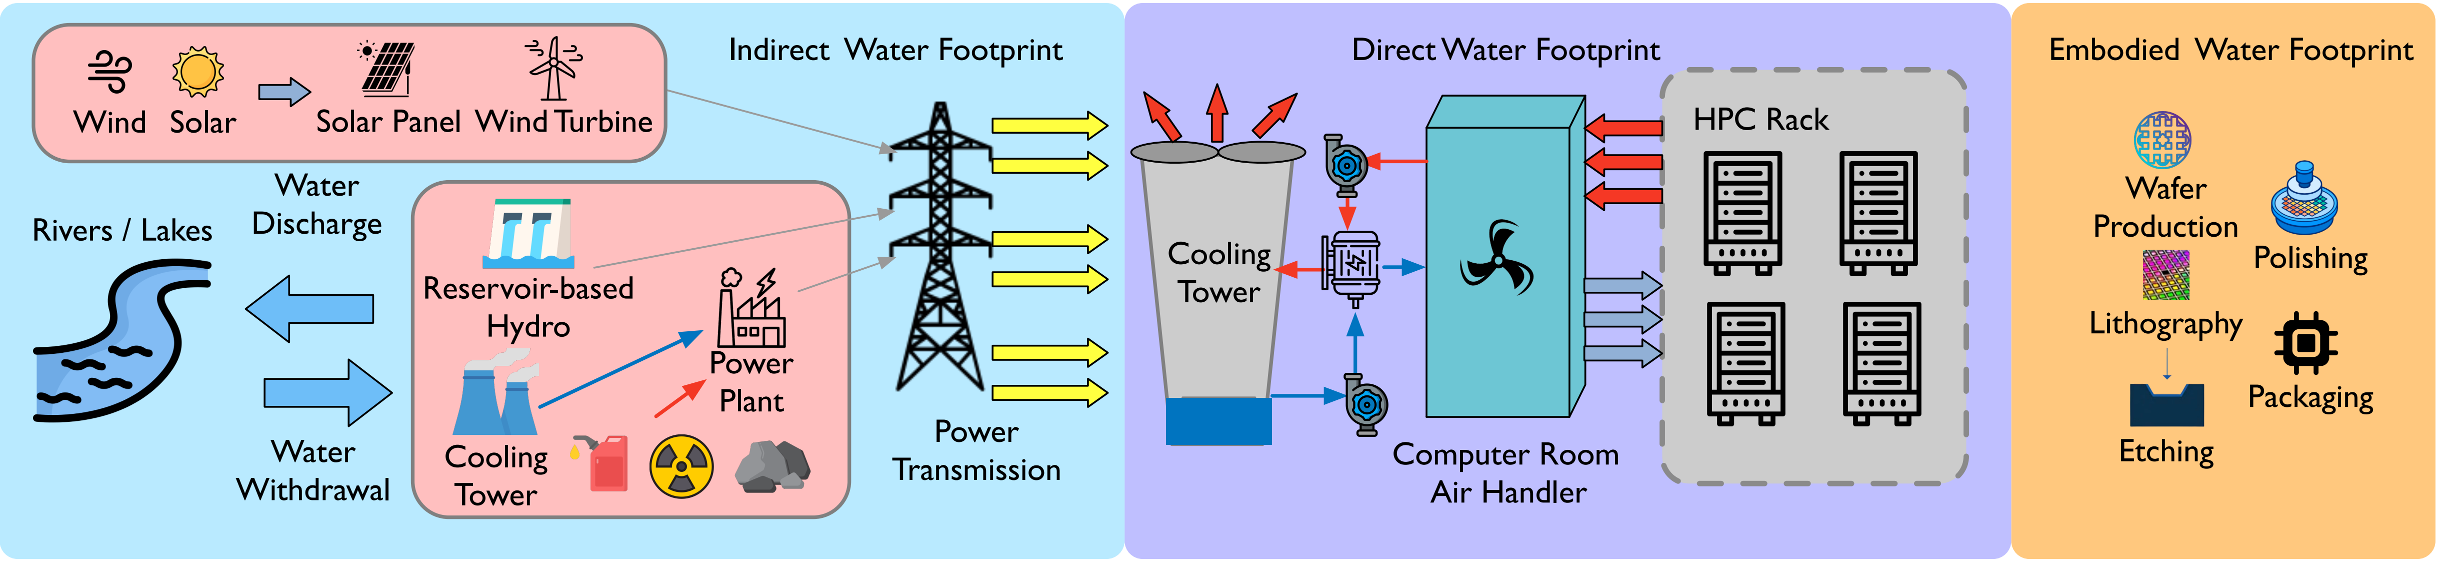

$$\begin{aligned} W_{\text{operational}} &= W_{\text{direct}} + W_{\text{indirect}} \\ &= E \cdot \text{WUE} + E \cdot \text{PUE} \cdot \text{EWF} \\ &= E \cdot (\text{WUE} + \text{PUE} \cdot \text{EWF}) = E \cdot \text{WI} \end{aligned}$$

$$\text{Water Intensity (WI)} = \text{WUE} + \text{PUE} \cdot \text{EWF}$$

$$\text{WI}^{\text{WSI}} = \text{WI} \cdot \text{WSI}$$

Water Intensity: it can simplify the calculation of operational water footprint.

Water Scarcity Index (WSI): accounts for geographic variation by applying weighting factors to volumetric water use.

ThirstyFLOPS: Tool Parameters and Case Study

Marconi

Fugaku

Polaris

Frontier

	Parameter	Parameter Description	Input ❍/Derive ▲	Data Range	Data Source	Unit	Reference
W_{embodied}	N_{IC}	Number of ICs (CPU/GPU/memory/storage)	❍	9-26 (Vary across hardware)	From hardware design	None	[31, 49]
	W_{IC}	Packaging water overhead	▲	0.6	From manufacturer	L	[69, 83]
	A_{die}	Die size of processors (CPU/GPU)	❍	Vary across hardware	From CPU/GPU design	mm^2	[53, 57, 78]
	Yield	Fab yield rate of hardware manufacturing	❍	0-1 (0.875 as default)	From manufacturer	None	[31]
	Location	Manufacturing location of hardware	❍	TSMC or GlobalFoundries	From manufacturer	None	[90]
	Process Node	Semiconductor manufacturing process of CPU/GPU	❍	3-28 (Vary across hardware)	From CPU/GPU design	nm	[10]
	UPW	Ultrapure water usage during manufacturing	▲	5.9-14.2 (Vary across process node)	From manufacturer	L	[10]
	PCW	Process cooling water during manufacturing	▲	Vary across locations and process node	From manufacturer	L	[10]
	WPA	Water for power generation during manufacturing	▲	Vary across locations and process node	From manufacturer	L	[10]
	WPC	Water footprint per capacity of DRAM, HDD, SSD	▲	0.8 (DRAM), 0.033 (HDD), 0.022(SSD)	From manufacturer	L/GB	[64, 65, 70]
	Capacity	Capacity of DRAM, HDD, SSD	❍	Vary across hardware	From manufacturer	GB	[21, 28, 44, 46]
$W_{\text{operational}}$	E	Energy consumption	❍	Vary across applications/hardware	From hardware profiling	kWh	[71, 76, 85]
	Wet bulb temperature	Site-related wet bulb temperature	❍	Vary across HPC locations	From weather report	°C	[54]
	WUE	Water usage effectiveness	▲	>0.05	From wet bulb temperature	L/kWh	[30]
	PUE	Power Usage Effectiveness	❍	≥1 (Marconi: 1.25, Fugaku:1.4, Polaris:1.65, Frontier: 1.05)	From HPC report	None	[20, 45, 76, 79]
	mix%	Percentage energy mix usage	❍	0-100	From power grid	%	[52]
	$\text{EWF}_{\text{energy}}$	energy water factor of energy sources	▲	1-17	From environment report	L/kWh	[51, 61]
	EWF	energy water factor of HPC system	▲	Vary across locations	From mix% and $\text{EWF}_{\text{energy}}$	L/kWh	[52]
	$\text{WSI}^{\text{direct}}$	Direct water scarcity index	❍	0.1-100	From WSI report	None	[14, 37, 48]
	$\text{WSI}^{\text{indirect}}$	Indirect water scarcity index	❍	0.1-100	From WSI report and power plant locations	None	[14, 37, 48]

ThirstyFLOPS: Toward Water-aware HPC Systems

Water Footprint Analysis

Energy, Carbon, and Water

Nuclear Reactor Powered HPC

ThirstyFLOPS: Toward Water-aware HPC Systems

Water Footprint Analysis

Energy, Carbon, and Water

Nuclear Reactor Powered HPC

Embodied Water Footprint of HPC Systems (Built-in Water Tax)

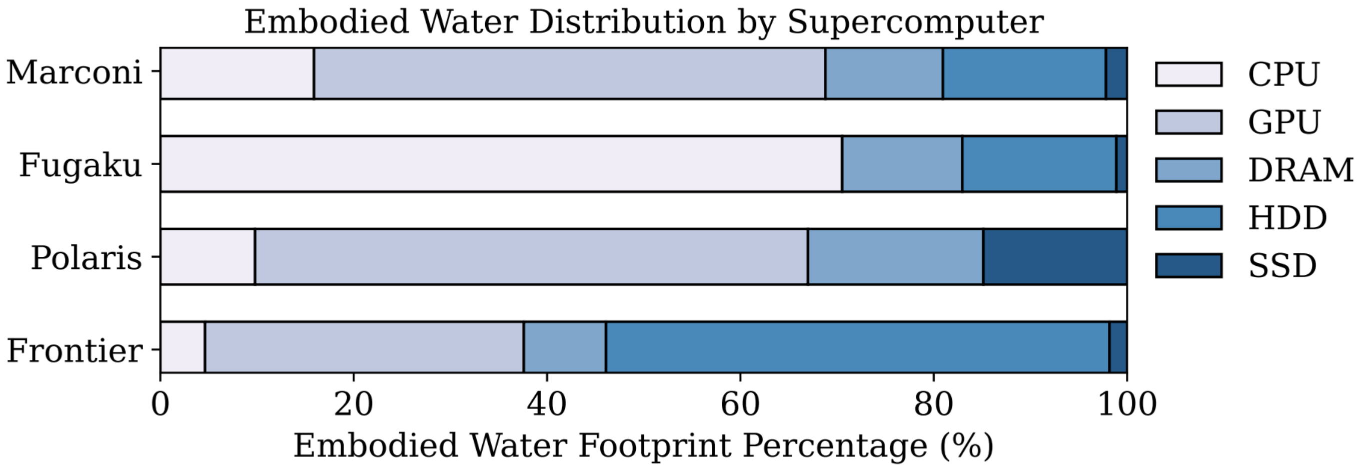

Name	Location	Processor (CPU/GPU)	Start Year
Marconi [21]	Bologna, Italy CINECA	IBM Power9 AC922 NVIDIA V100 SXM2	2019
Fugaku [28]	Kobe, Japan Riken CCS	Fujitsu A64FX 48C No GPU	2020
Polaris [44]	Lemont, IL, US Argonne National Lab	AMD EPYC 7532 NVIDIA A100 PCIe	2021
Frontier [46]	Oak Ridge, TN, US Oak Ridge National Laboratory	AMD EPYC 7A53 AMD Instinct MI250X	2021

GPUs contribute significantly to the overall embodied water.

The HPC systems with a large storage capacity, backed by traditional hard disk drives, have a high embodied water coming from hard drives.

Hard drives have lower carbon footprint, but higher water footprint.

Embodied Water Footprint of HPC Systems (Built-in Water Tax)

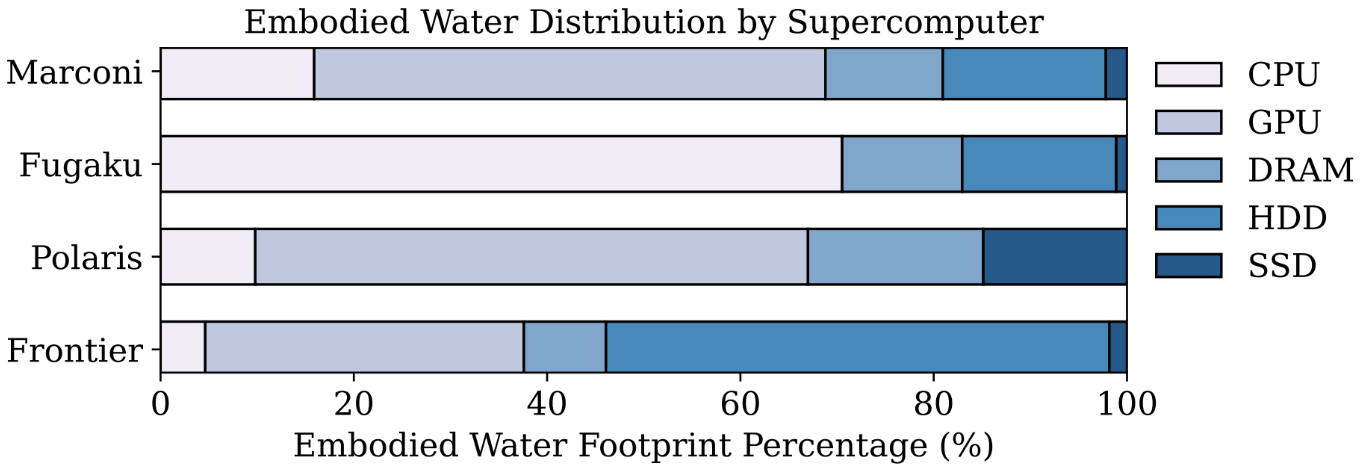

Name	Location	Processor (CPU/GPU)	Start Year
Marconi [21]	Bologna, Italy CINECA	IBM Power9 AC922 NVIDIA V100 SXM2	2019
Fugaku [28]	Kobe, Japan Riken CCS	Fujitsu A64FX 48C No GPU	2020
Polaris [44]	Lemont, IL, US Argonne National Lab	AMD EPYC 7532 NVIDIA A100 PCIe	2021
Frontier [46]	Oak Ridge, TN, US Oak Ridge National Laboratory	AMD EPYC 7A53 AMD Instinct MI250X	2021

Achieving practical environmental sustainability of an HPC system is challenging for facility designers – different HPC components rank differently on different sustainability metrics (e.g., carbon vs water).

Water Footprint: Hardware Manufacturing vs Operations

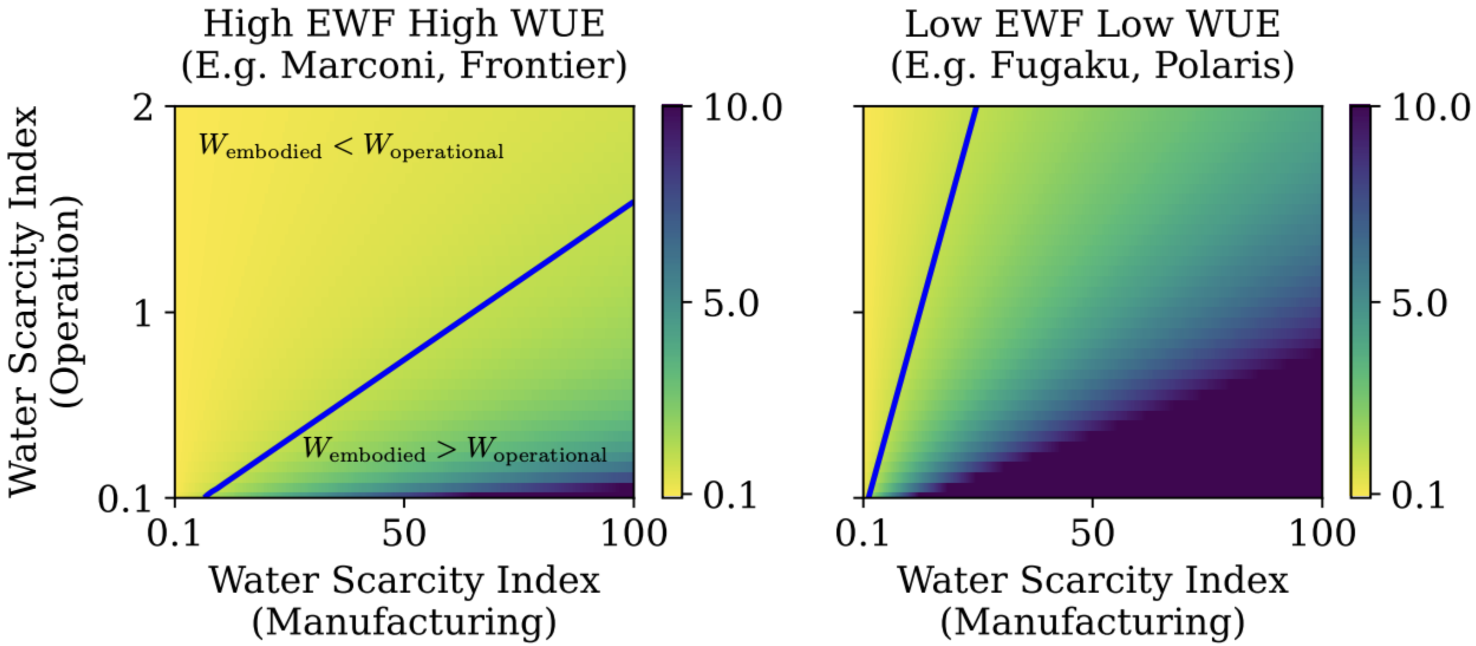

Both the hardware manufacturer's location and the physical location of the HPC data center influence the overall water footprint of the HPC systems.

Building fabrication facilities in water-scarce regions can lead to high embodied water footprint. HPC site selection must account for water-intensity.

Secret of Water Consumption During Electricity Generation

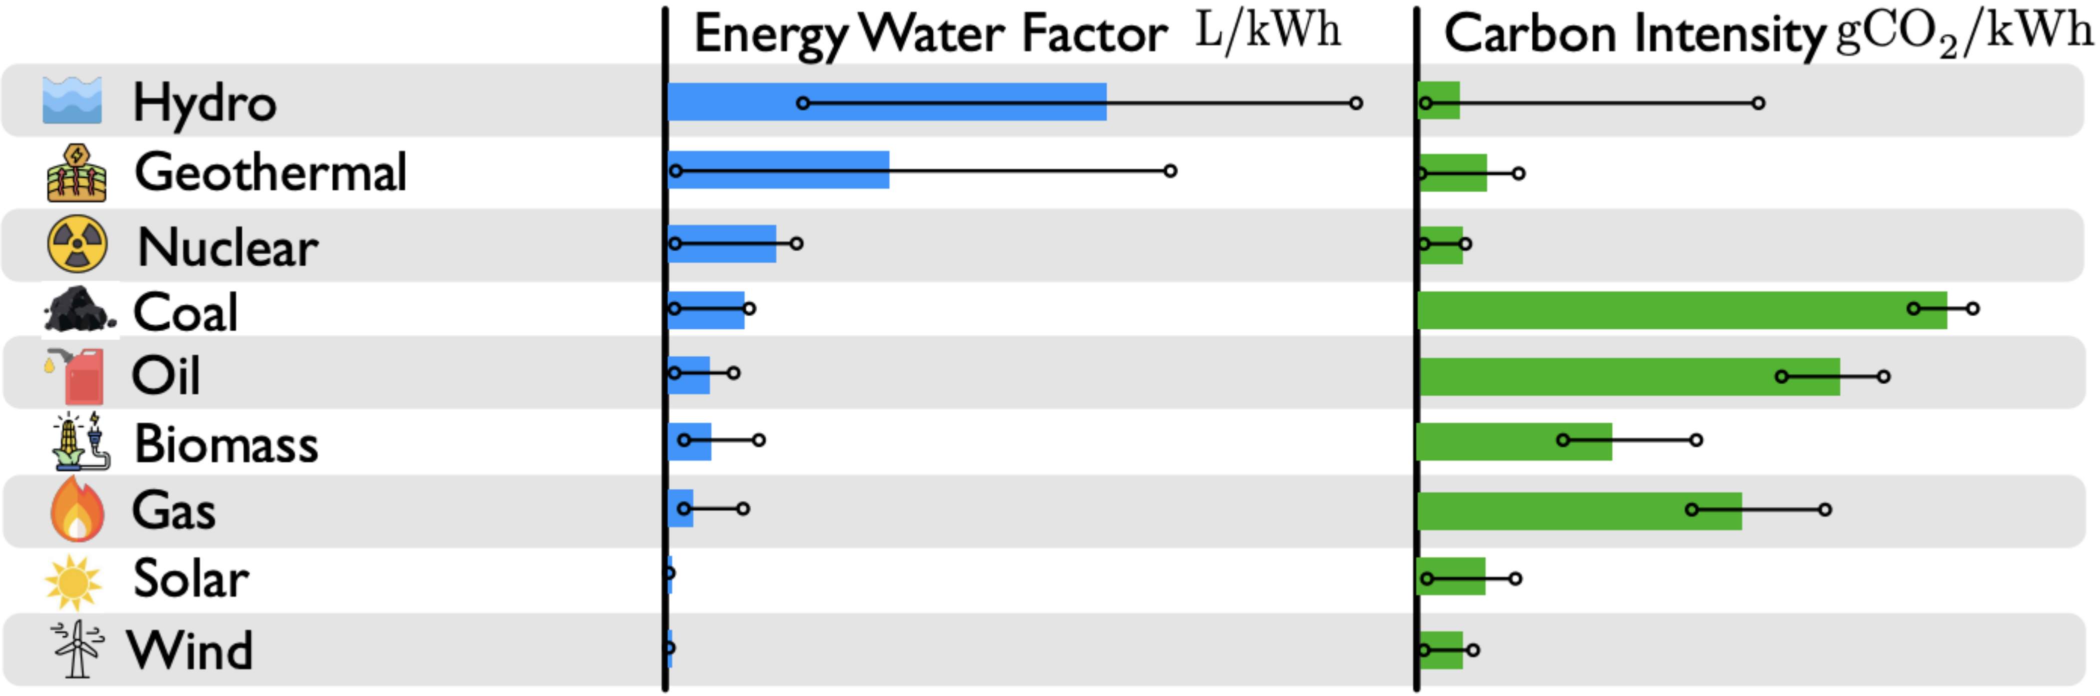

Water consumption during electricity generation can be significant, and unfortunately, energy sources that are carbon-friendly are not necessarily water-friendly.

Water Footprint: Energy Generation vs HPC System Cooling

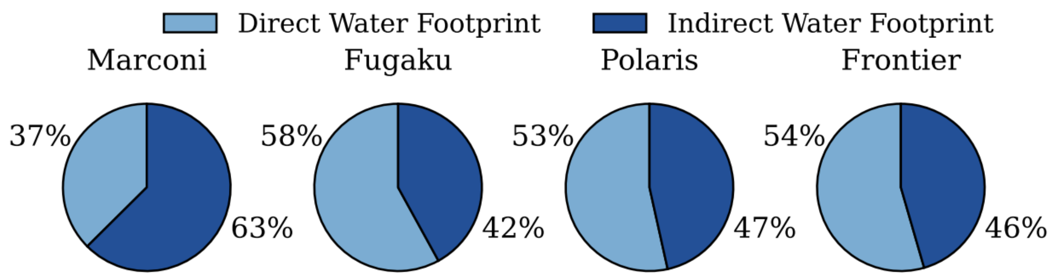

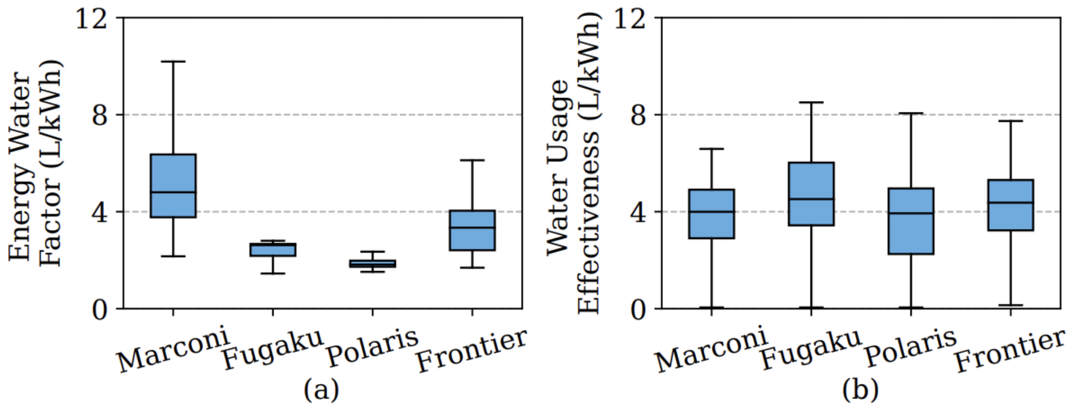

Indirect water footprint of HPC systems (water use during electricity generation) can be comparable to direct water footprint (water use during HPC system cooling), and both components vary across time and location.

Favorable climate condition for cooling the HPC datacenter does not necessarily lead to overall lower water consumption – since water consumed during energy generation may become dominant.

Water Footprint: Energy Generation vs HPC System Cooling

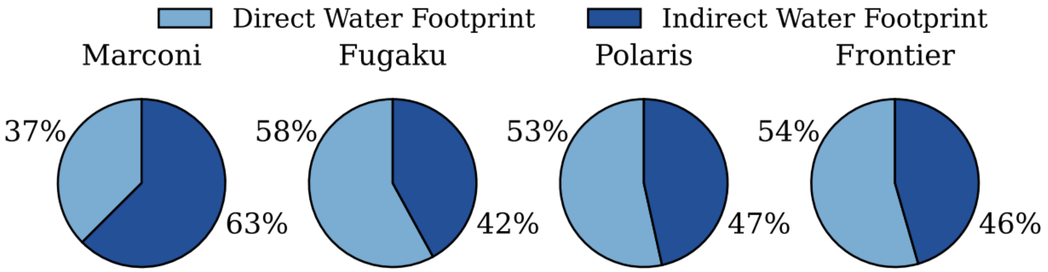

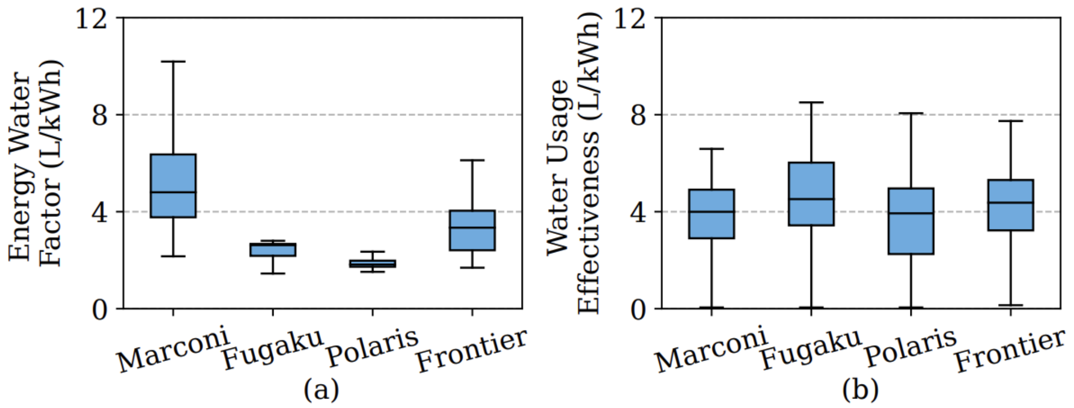

Implication: HPC facilities and city operators should dynamically determine what fraction of total water goes where ("water capping"), when water is a constrained resource – toward the cooling of the datacenter, or toward energy generation.

All Liters of Water Do Not Hold the Same Value!

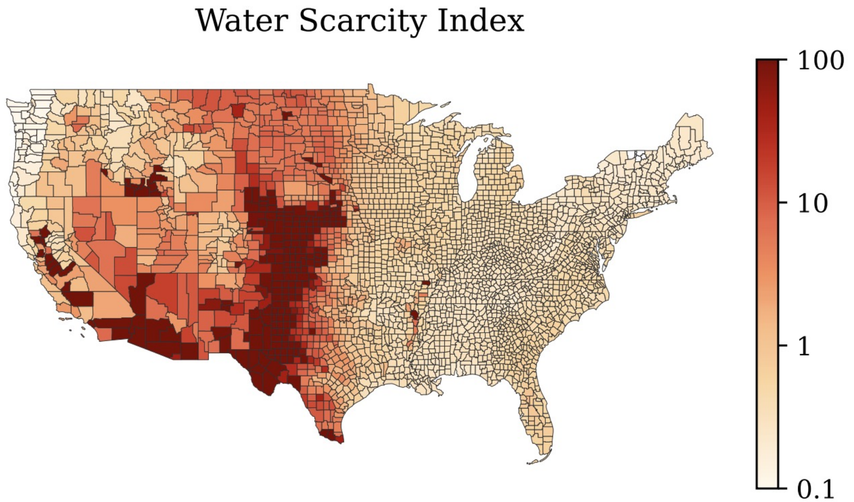

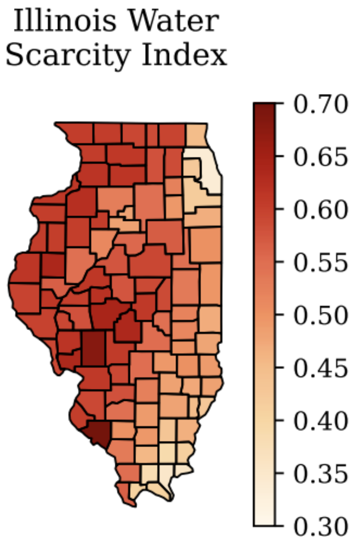

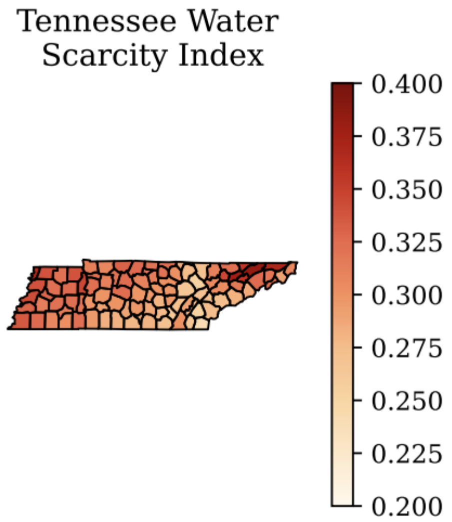

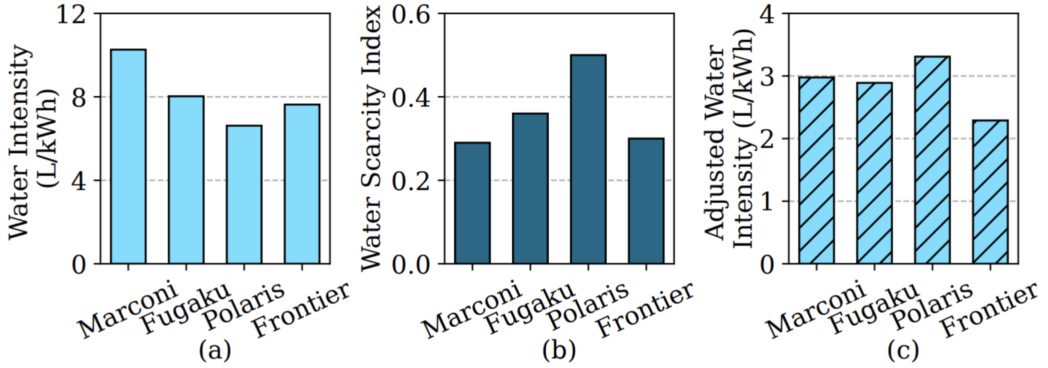

The geographical dependence of the water scarcity index affects the effective water footprint of different HPC systems.

All Liters of Water Do Not Hold the Same Value!

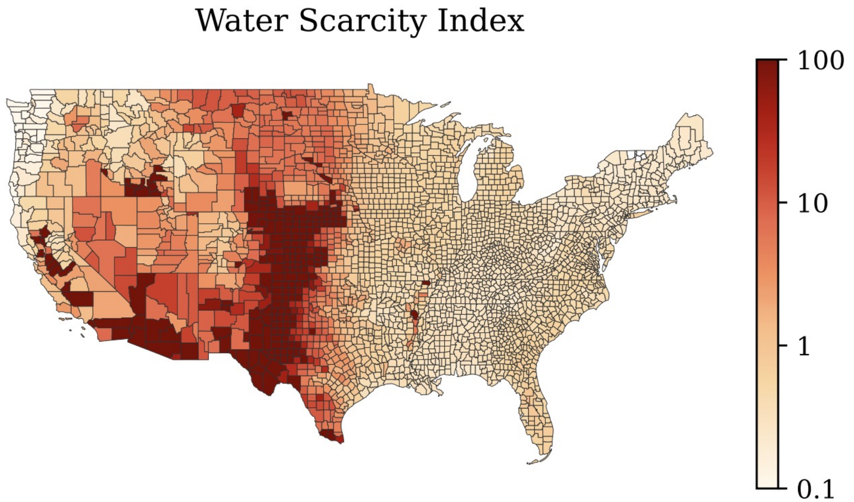

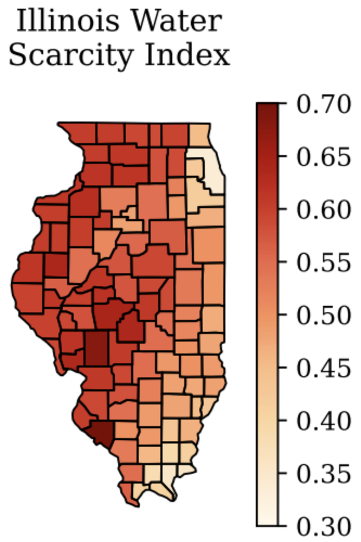

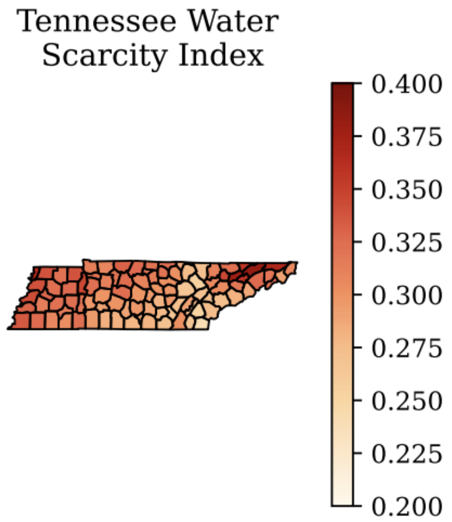

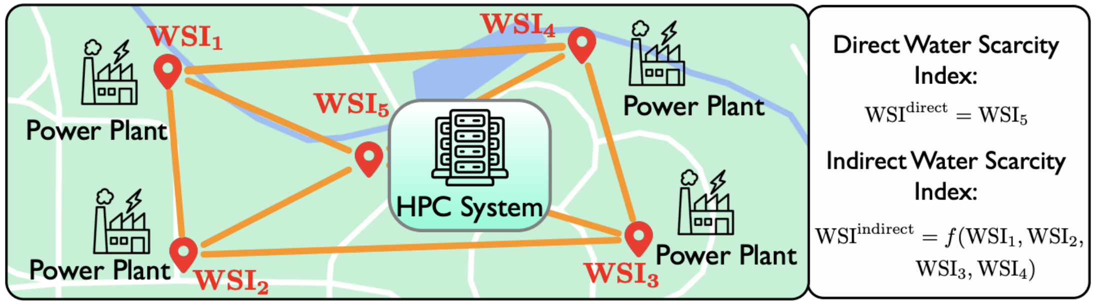

HPC center operations should consider accounting for the water scarcity index of all nearby power grids, besides electricity cost and renewability of the energy mix.

ThirstyFLOPS: Toward Water-aware HPC Systems

Water Footprint Analysis

Energy, Water and Carbon

Nuclear Reactor Powered HPC

Energy and Water Can Be at Odds!

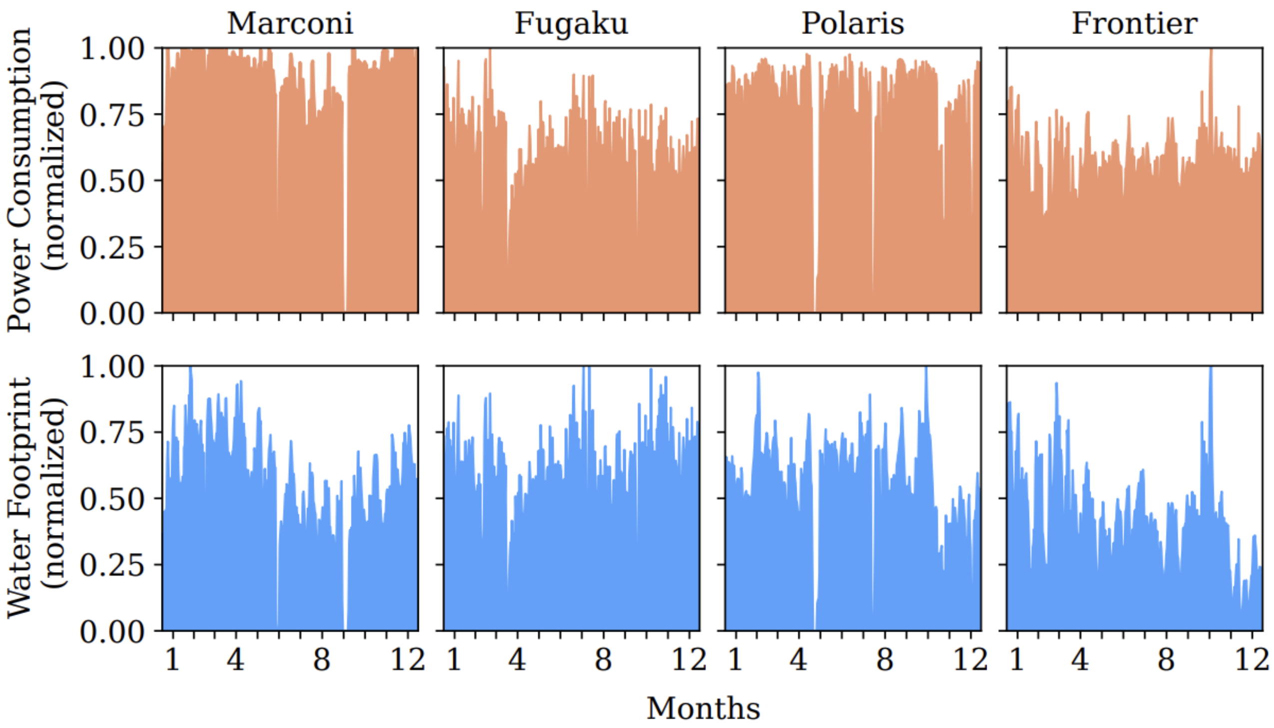

Energy-optimal operation does not necessarily mean water-optimal operation. Existing popular research strategies that attempt to minimize energy consumption may still lead to disproportionately high water use.

… unfortunately, so are Carbon and Water!

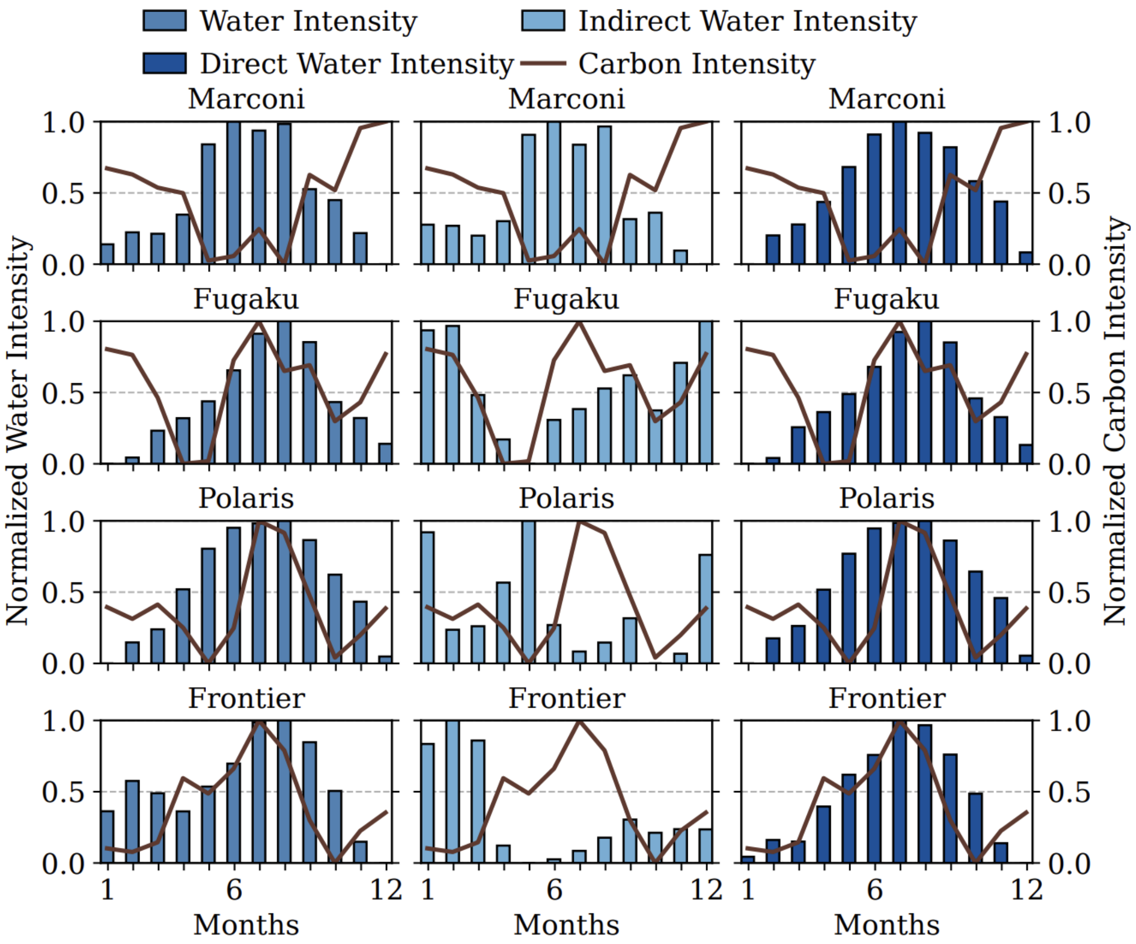

Decision-making for HPC and AI data centers should explicitly address the synergistic and competitive interactions between carbon and water impacts.

Carbon Footprint vs Water Footprint

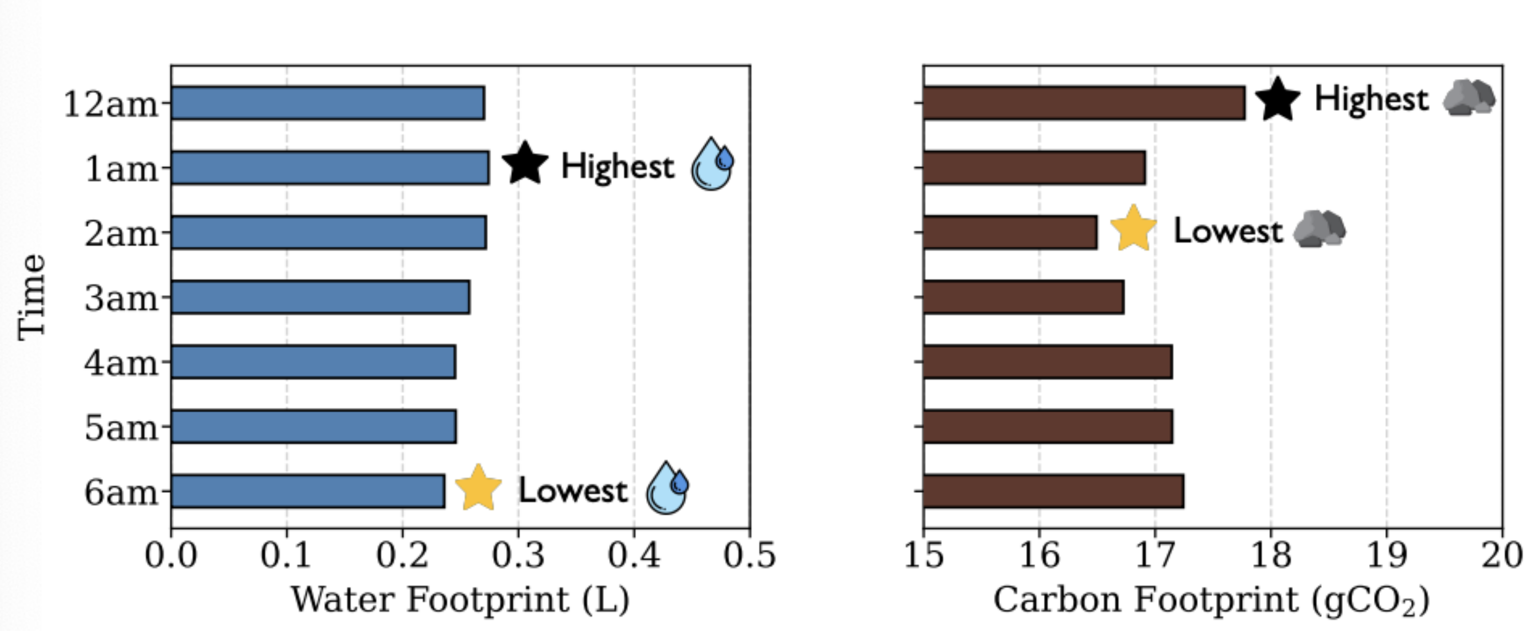

Fortunately, HPC programmers do not necessarily need to invent additional tools for optimizing water consumption. But new schedulers need to be developed if the HPC centers want to co-optimize for multiple sustainability metrics.

ThirstyFLOPS: Toward Water-aware HPC Systems

Water Footprint Analysis

Energy, Water and Carbon

Nuclear Reactor Powered HPC

Small Modular Nuclear Reactor for Powering HPC

Our first advanced nuclear reactor project with Kairos Power and Tennessee Valley Authority

Aug 18, 2025
2 min read

This public-private collaboration will help meet our data center electricity demand with advanced nuclear energy starting in 2030 and power the nuclear renaissance in Oak Ridge, Tennessee.

Amanda Peterson Corio
Global Head of Data Center Energy

Share

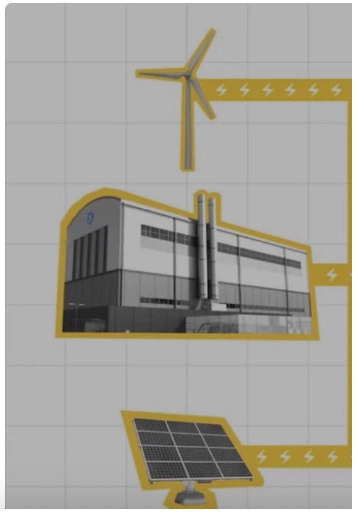

META

Meta and Constellation Partner on Clean Energy Project

June 3, 2025

LISTEN TO ARTICLE

OCTOBER 1, 2024

Data center owners turn to nuclear as potential electricity source

Nuclear power plants that have signed agreements to power data centers (as of Sep 2024)

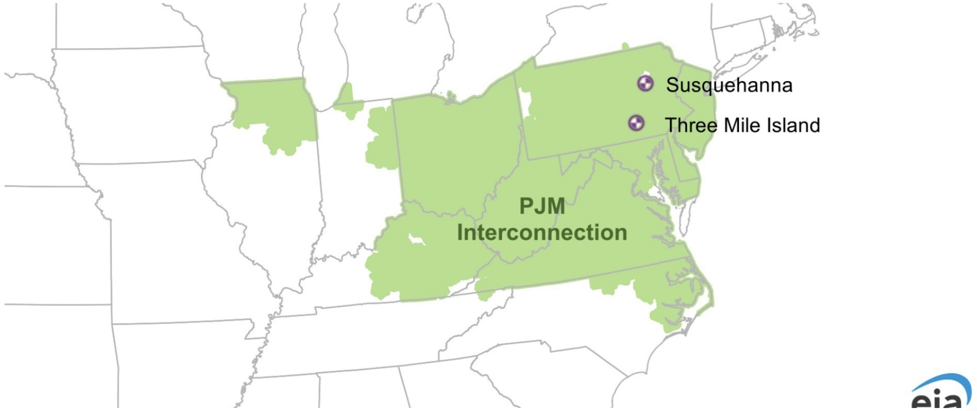

Hidden Water Footprint Of Nuclear Powered HPC Systems

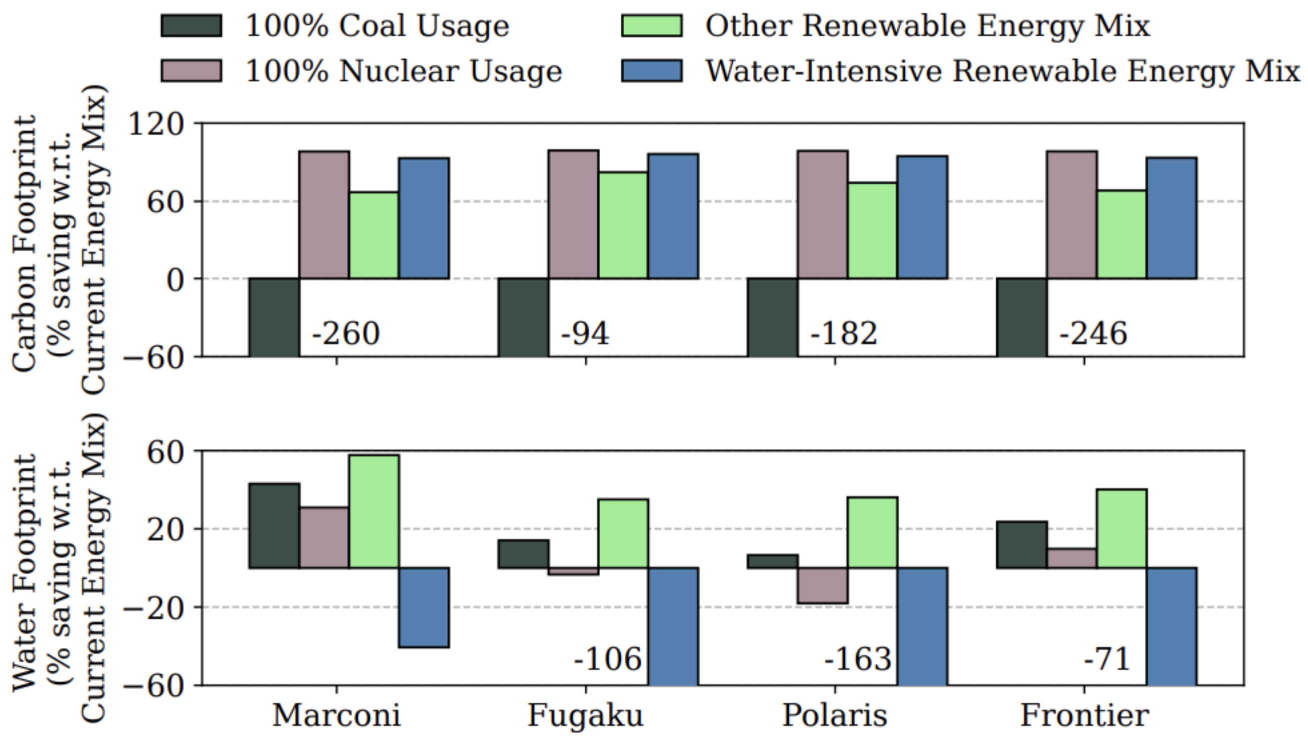

The water footprint of nuclear power reactors is significant – and, location-dependent. Naively employing nuclear reactors to power HPC centers, to mitigate energy and carbon footprint concerns, is a sub-optimal strategy.

ThirstyFLOPS: Summary Key Contributions

- ✓ ThirstyFLOPS is the first open-source tool for estimating the water footprint of HPC systems.
- ✓ ThirstyFLOPS incorporates regional factors and water scarcity to enable location-aware sustainability, inform HPC site selection decisions, and highlights the need for new water-carbon-aware HPC schedulers.
- ✓ ThirstyFLOPS demonstrates that "green" energy sources may be highly water-intensive, and small modular nuclear reactors can have significant water footprint.

ThirstyFLOPS
A tool for estimating the water footprint of HPC

Contact
Yankai Jiang
jiang.yank@northeastern.edu