Water Footprint Modeling, Characterization, and Analysis Toward Water-aware HPC System Design and Operations

Yankai Jiang[1], *Raghavendra Kanakagiri*[2], *Rohan Basu Roy*[3], *Devesh Tiwari*[1], Northeastern University[1], IIT Tirupati[2], University of Utah[3]

High-performance computing (HPC) systems are becoming increasingly water-intensive due to their reliance on water-based cooling and the energy used in power generation. However, the water footprint of HPC remains relatively underexplored—especially in contrast to the growing focus on carbon emissions. In this paper, we present a comprehensive water footprint analysis for HPC systems. Our approach incorporates region-specific metrics, including Water Usage Effectiveness (WUE), Power Usage Effectiveness (PUE), and Energy Water Factor (EWF), to quantify water consumption using real-world data. By evaluating four representative HPC systems, Marconi, Fugaku, Polaris, and Frontier, we provide implications for HPC system planning and management. Furthermore, we explore the impact of regional water scarcity and nuclear-based energy strategies on HPC sustainability. Our findings aim to advance the development of water-aware, environmentally responsible computing infrastructures.

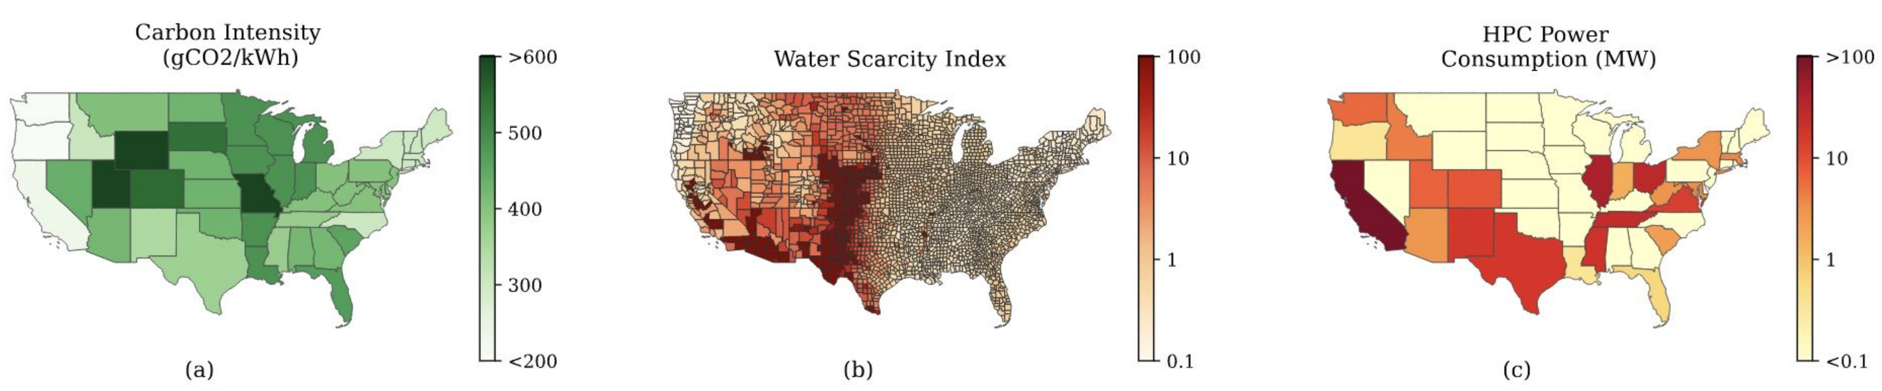

Motivation. *Some HPC centers are located in relatively water-scarce regions. Currently, the carbon footprint of energy sources as main influencing factors when determining where HPC centers should be located and how they should be operated.*

Water Consumption in HPC

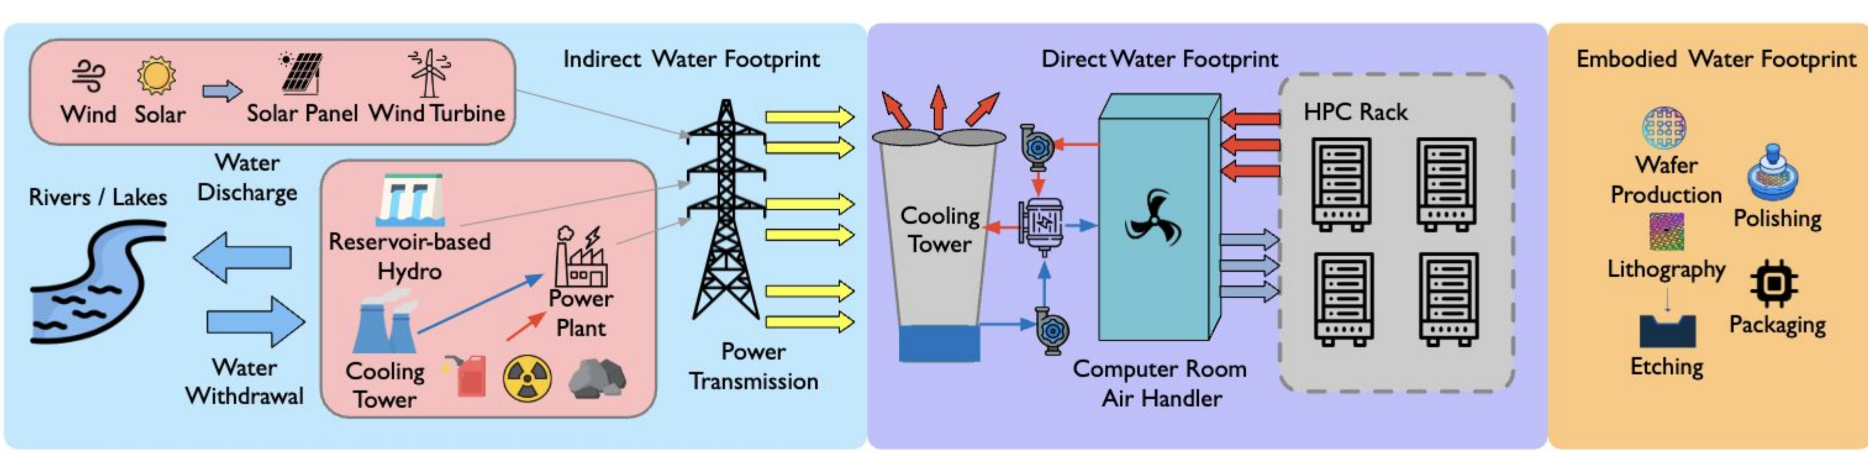

Indirect Water: Water used in power generation, *Direct Water*: Water used for system cooling
Embodied Water: Water used in hardware manufacturing

ThirstyFLOPS: Tool for Estimating Water Footprint

	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	[25, 40]
	W_{IC}	Packaging water overhead	▲	0.6	From manufacturer	L	[58, 69]
	A_{die}	Die size of processors (CPU/GPU)	❍	Vary across hardware	From CPU/GPU design	mm^2	[45, 47, 65]
	Yield	Fab yield rate of hardware manufaturing	❍	0-1 (0.875 as default)	From manufacturer	None	[25]
	Location	Manufacturing location of hardware	❍	TSMC or GlobalFoundries	From manufacturer	None	[74]
	Process Node	Semiconductor manufacturing process of CPU/GPU	❍	3-28 (Vary across hardware)	From CPU/GPU design	nm	[6]
	UPW	Ultrapure water usage during manufatruing	▲	5.9-14.2 (Vary across process node)	From manufacturer	L	[6]
	PCW	Process cooling water during manufatruing	▲	Vary across locations and process node	From manufacturer	L	[6]
	WPA	Water for power generation during manufatruing	▲	Vary across locations and process node	From manufacturer	L	[6]
	WPC	Water footprint per capacity of DRAM, HDD, SSD	▲	0.8 (DRAM), 0.033 (HDD), 0.022(SSD)	From manufacturer	L/GB	[53, 54, 59]
	Capacity	Capacity of DRAM, HDD, SSD	❍	Vary across hardware	From manufacturer	GB	[16, 22, 36, 38]
$W_{operational}$	E	Energy consumption	❍	Vary across applications/hardware	From hardware profiling	kWh	[60, 63, 70]
	Wet bulb temperature	Site-related wet bulb temperature	❍	Vary across HPC locations	From weather report	$°C$	[46]
	WUE	Water usage effectiveness	▲	>0.05	From wet bulb temperature	L/kWh	[24]
	PUE	Power Usage Effectiveness	❍	≥1 (Marconi: 1.25, Fugaku:1.4, Polaris:1.65, Frontier: 1.05)	From HPC report	None	[15, 37, 63, 66]
	mix%	Percentage energy mix usage	❍	0-100	From power grid	%	[44]
	EWF_{energy}	energy water factor of energy sources	▲	1-17	From environment report	L/kWh	[42, 43, 51]
	EWF	energy water factor of HPC system	▲	Vary across locations	From mix% and EWF_{energy}	L/kWh	[44]
	WSI^{direct}	Direct water scarcity index	❍	0.1-100	From WSI report	None	[9, 29, 39]
	$WSI^{indirect}$	Indirect water scarcity index	❍	0.1-100	From WSI report and power plant locations	None	[9, 29, 39]

Parameters of ThirstyFLOPS for estimating water.

Analyze current HPC systems.

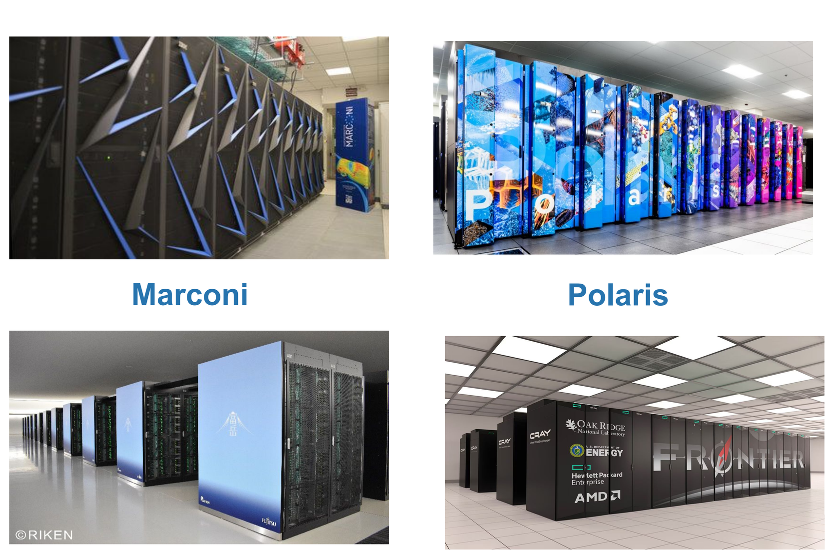

Marconi Polaris

Fugaku Frontier

Name	Location	Processor (CPU/GPU)	Start Year
Marconi [16]	Bologna, Italy CINECA	IBM Power9 AC922 NVIDIA V100 SXM2	2019
Fugaku [22]	Kobe, Japan Riken CCS	Fujitsu A64FX 48C No GPU	2020
Polaris [36]	Lemont, IL, US Argone National Lab	AMD EPYC 7532 NVIDIA A100 PCIe	2021
Frontier [38]	Oak Ridge, TN, US Oak Ridge National Laboratory	AMD EPYC 7A53 AMD Instinct MI250X	2021

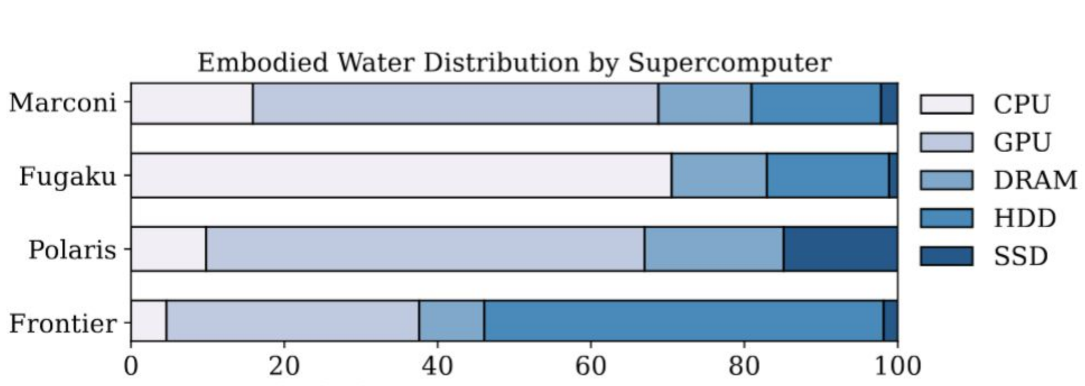

Observation 1. The HPC systems which have a large storage capacity backed by traditional hard disk drives, have a significant embodied water footprint coming from hard drives. SSDs, while more expensive, are favorable in terms of embodied water footprint compared to hard drives.

Insight: *Achieving practical environmental sustainability of an HPC system is extremely challenging for facility designers – different HPC components rank differently on different sustainability metrics (carbon and water).*

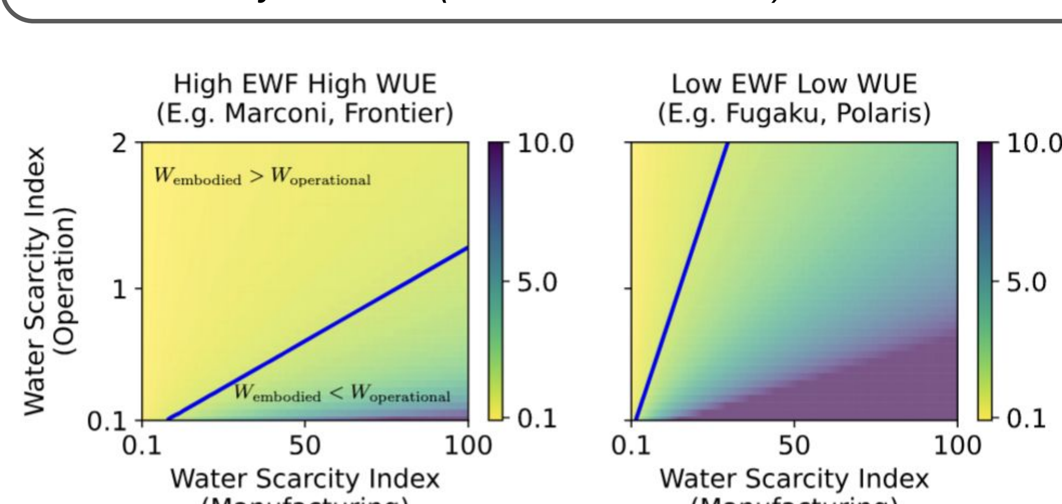

Observation 2. Both the geographic location of the hardware manufacturer and the HPC center play a critical role in the system's overall water footprint. Building fabrication facilities in water-scarce regions can lead to high embodied water footprints, even if operational water use remains low.

Insight: *Careful consideration of manufacturing sites is critical for HPC systems, in addition to operational site selection for HPC systems. Even for operational site selection, we highlight the need for modeling and accounting for water intensity of energy generation, year-round weather conditions, and water scarcity of the local region.*

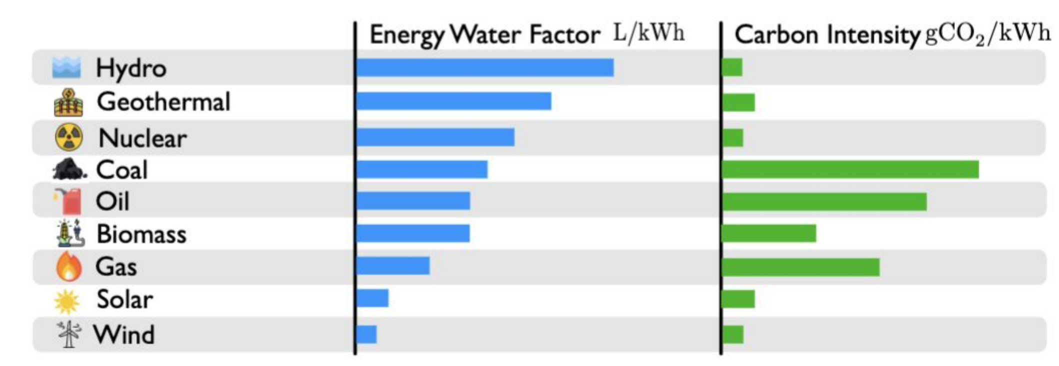

Observation 3. Water consumption during electricity generation can be significant, and energy sources that are typically environment-friendly are not necessarily water-friendly.

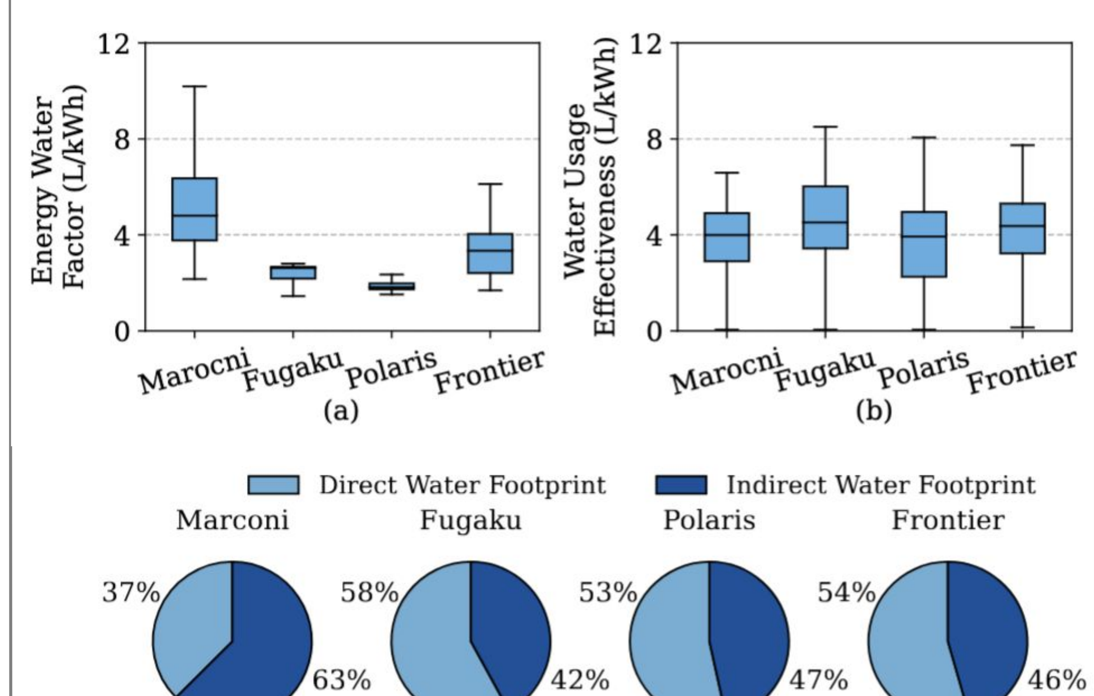

Observation 4. Indirect operational water footprint can often be comparable to direct operational water footprint. Unfortunately, the factors that affect these components show strong temporal and spatial variation. Favorable climate conditions for cooling the datacenter do not necessarily mean overall lower water consumption -- since water consumed during energy generation may become dominant.

Insight: The water consumed to generate electricity (besides water needed to cool the HPC system) must be taken into account toward overall water-optimized HPC system operations.

Insight: 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.

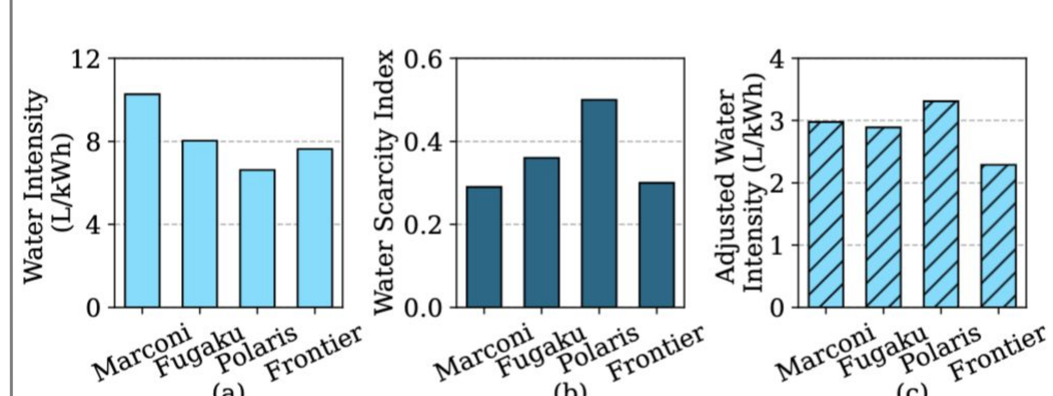

Observation 5. Geographical dependence of the water scarcity index affects the effective water footprint consumption of different HPC centers. Interestingly, the water scarcity index can vary significantly even at a kilometer scale, and hence, the operational water footprint is greatly impacted by which nearby power grids are being used for electricity generation.

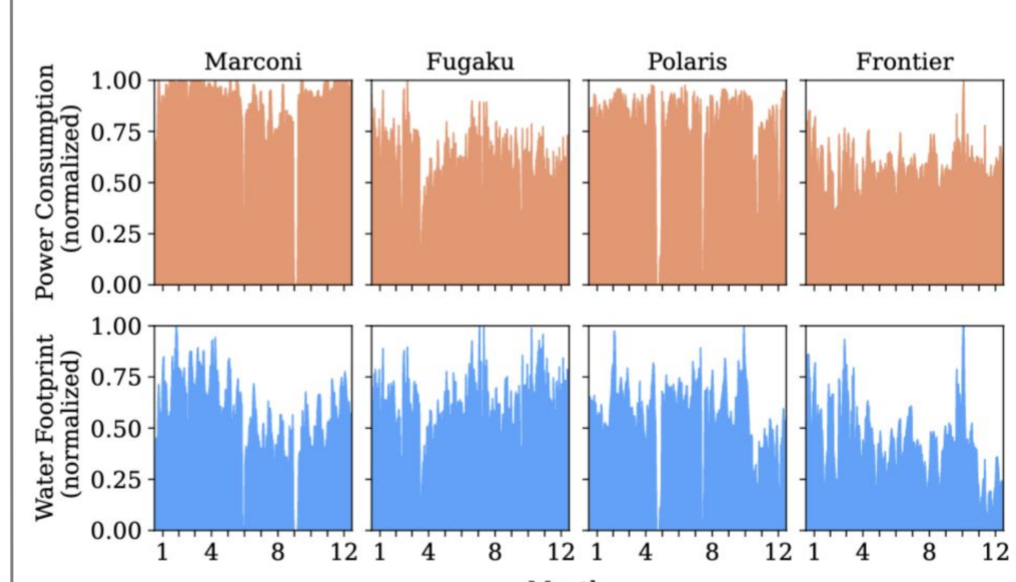

Observation 6. energy-aware HPC system operation does not necessarily mean water-optimal operation. Existing popular research strategies that attempt to minimize energy consumption (e.g., workload shifting among HPC centers purely based on energy consumption) may still lead to disproportionately high water use if regional water constraints and patterns are not carefully considered.

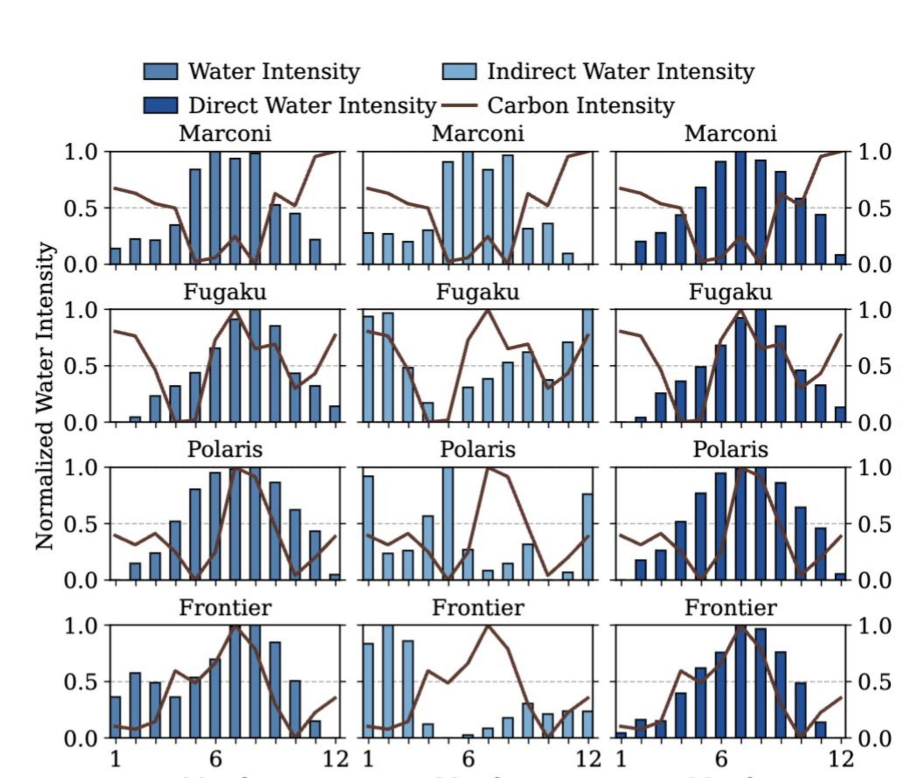

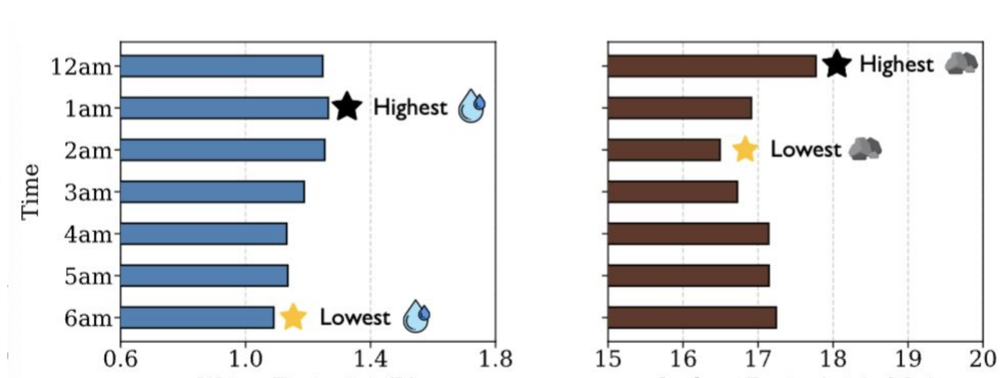

Insight: Both major sustainability metrics (carbon and water footprint) do not always compete, but the future HPC system design should carefully consider both metrics explicitly. New schedulers need to be developed if the HPC centers want to co-optimize for multiple sustainability metrics such as water, carbon, energy, etc.

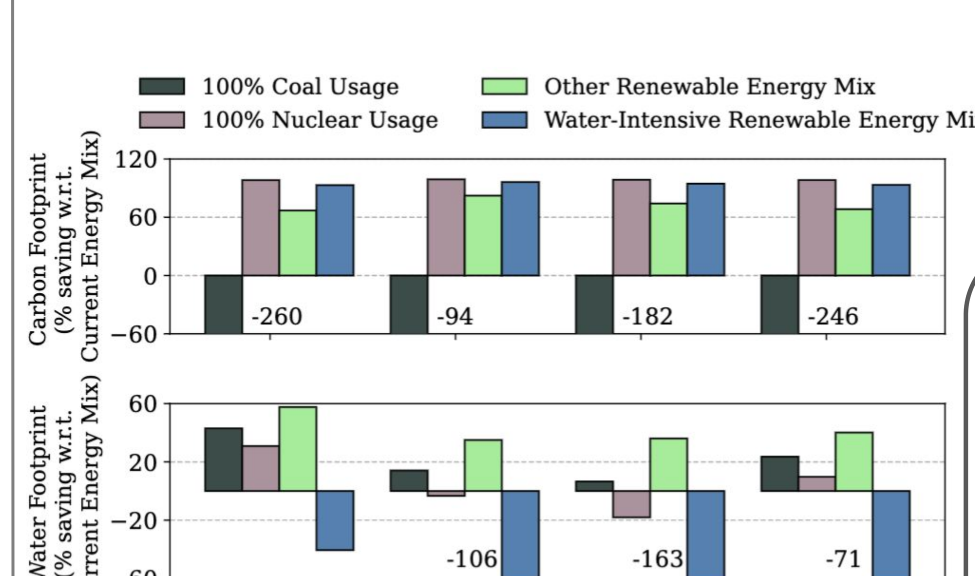

Observation 7. The water footprint of nuclear power reactors can be significant and the impact is location-dependent.

Insight: Small nuclear reactors are a promising option for mitigating the exploding energy needs of datacenters and their environmental carbon footprint. However, naively employing nuclear reactors to power HPC centers, to mitigate energy and carbon footprint concerns, can be significantly sub-optimal depending upon the location.