Real-time Carbon Footprint Minimization in Sustainable Data Centers with Reinforcement Learning

Soumyendu Sarkar, Avisek Naug, Ricardo Luna Gutierrez, Antonio Guillen, Vineet Gundecha, Ashwin Ramesh Babu, Cullen Bash

Hewlett Packard Labs @
Hewlett Packard Enterprise

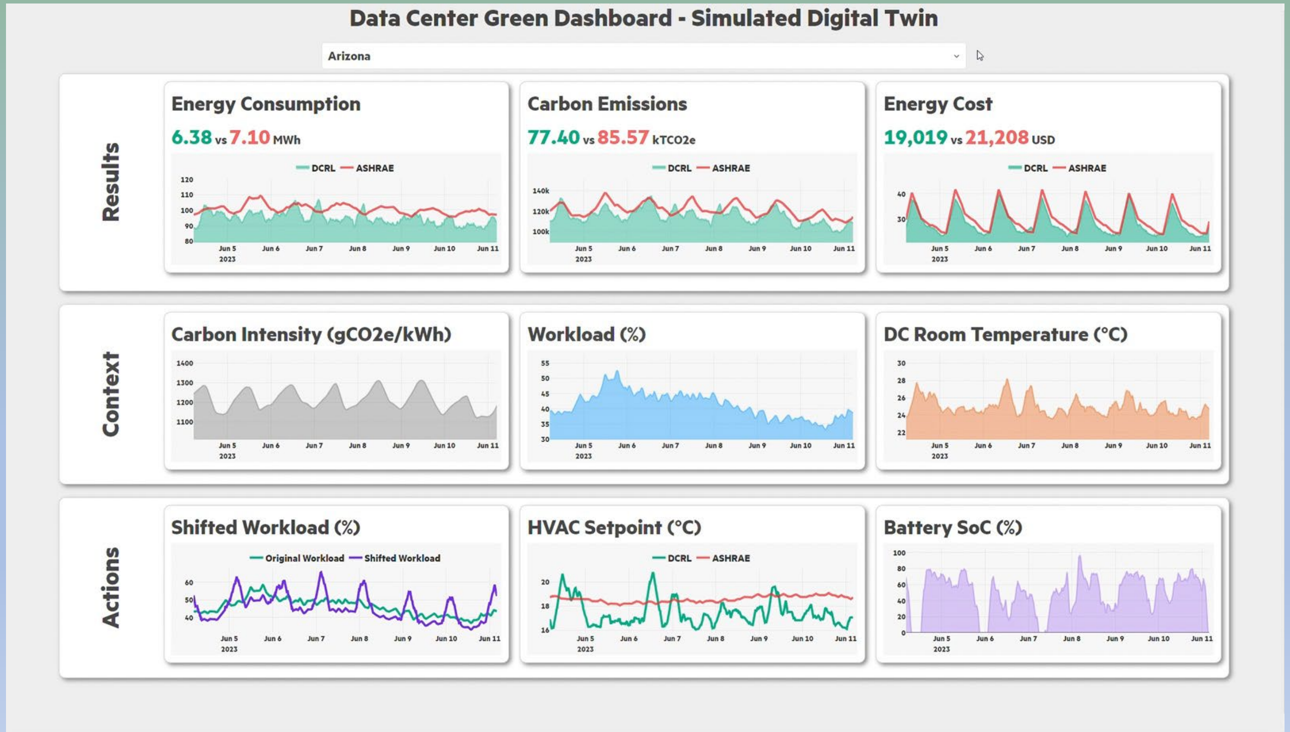

Motivation

- ❑ Sustainable data centers with
  - ❑ Lower Carbon Emissions
  - ❑ Lower Energy Consumption
  - ❑ Lower Energy Cost
- ❑ Paradigm shift in optimizing Cooling and IT loads
  - ❑ Schedule flexible loads
  - ❑ Leverage battery storage
- ❑ Real-time controller to optimize all these goals is lacking.

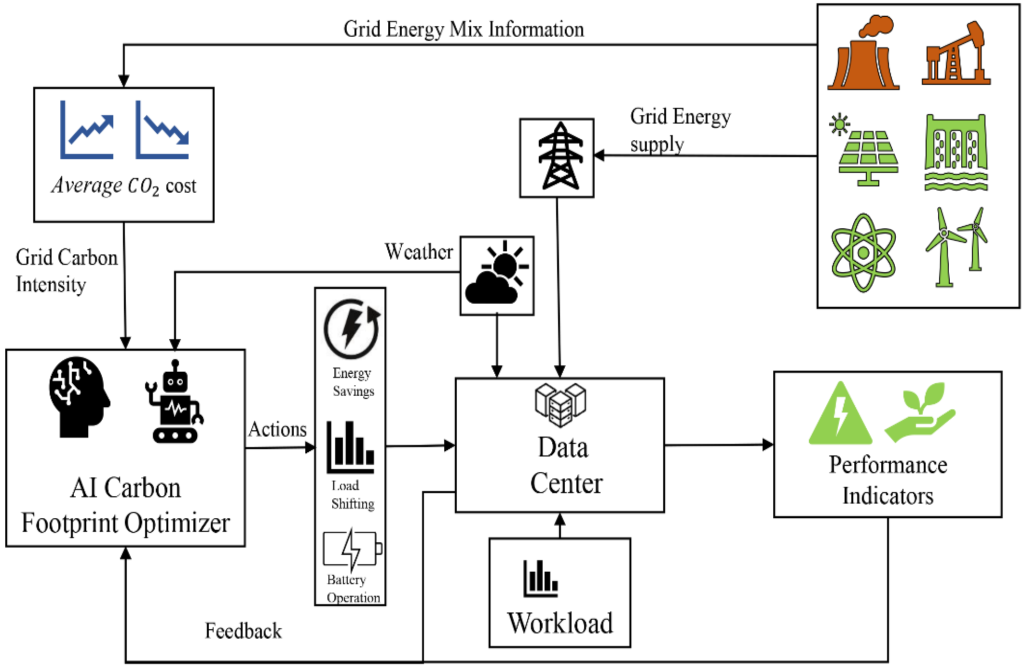

System Diagram

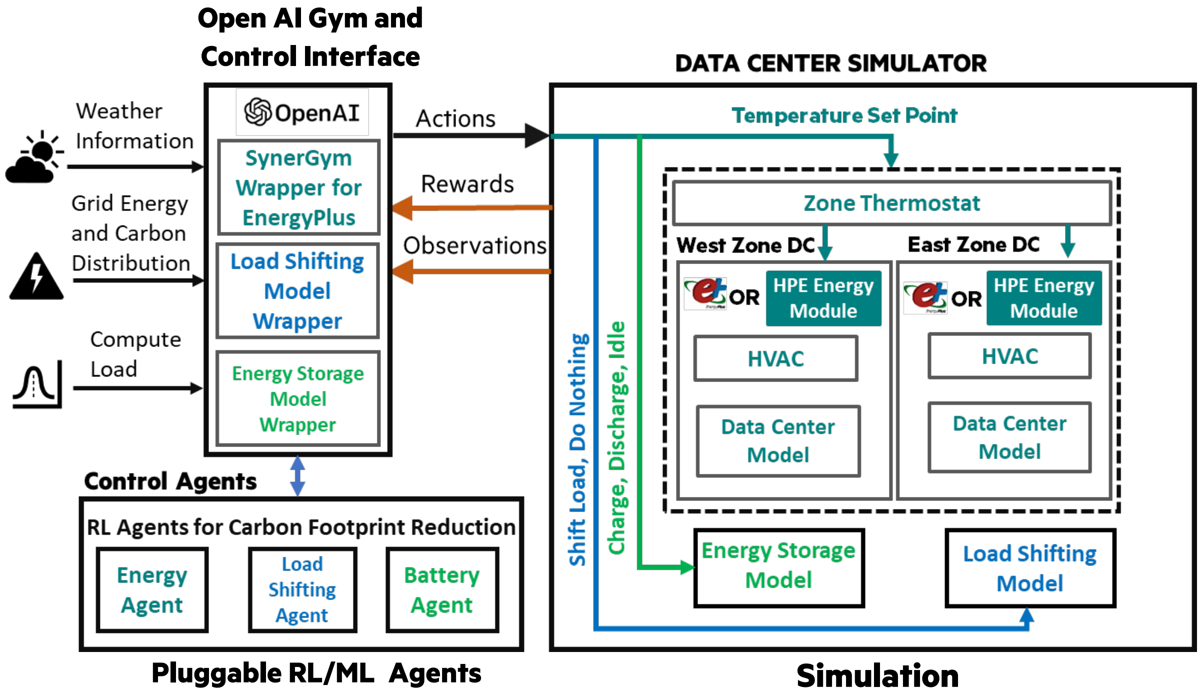

Digital System

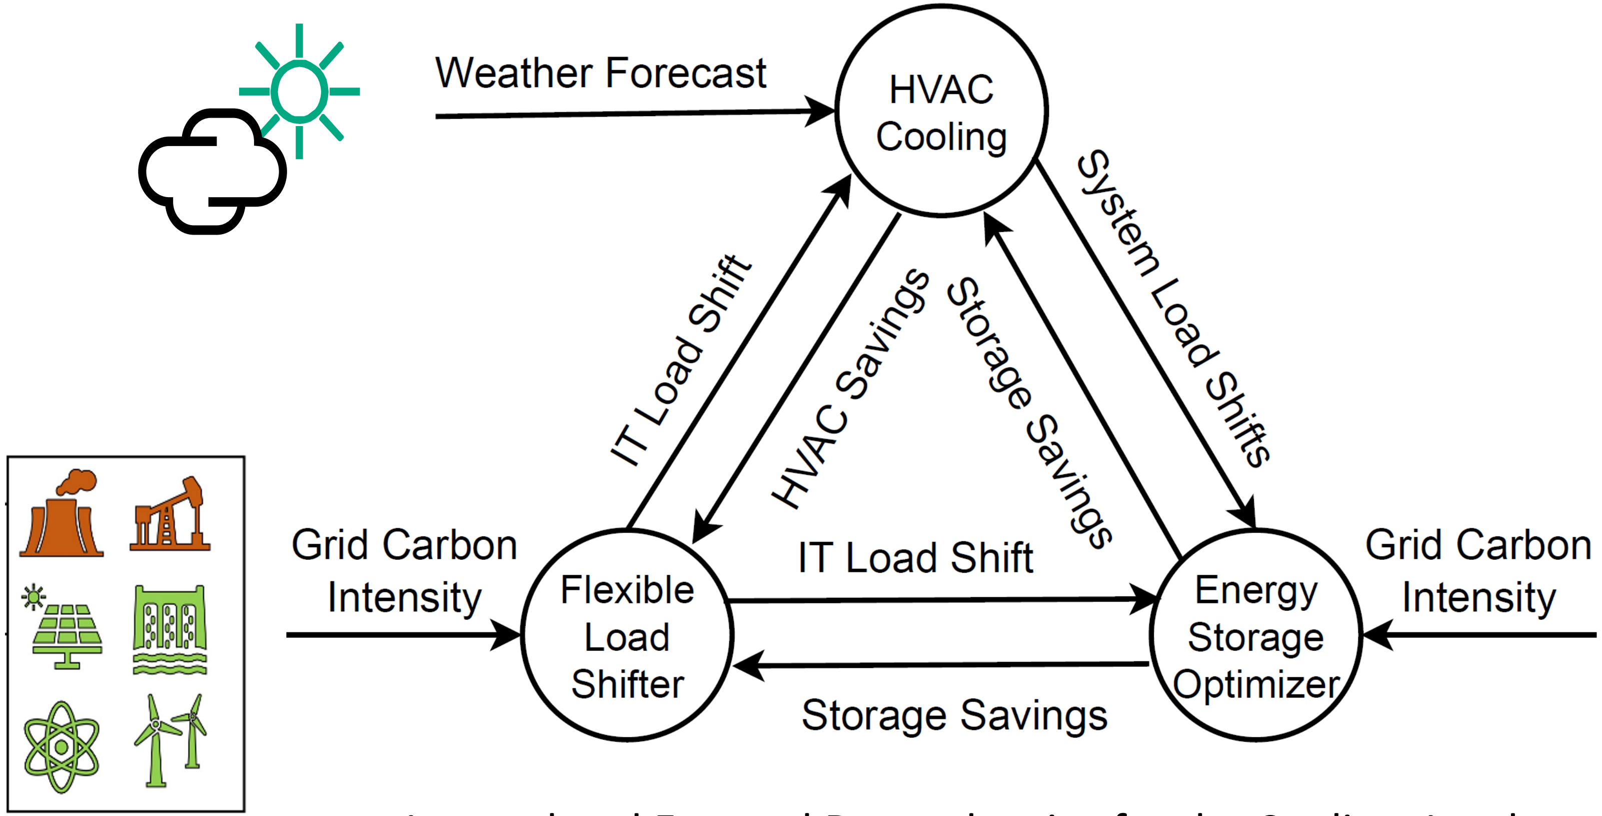

Internal and External Dependencies for the Cooling, Load Shifting and Battery agents

Multi-agent Reinforcement Learning

Multi-Agent Reinforcement Learning (MARL)

- Achieve individual goals
- Collaborate through a shared reward

Explored various MARL methods:

- Multi-Agent Deep Deterministic Policy Gradient (MADDPG), (centralized critic)
- Independent Proximal Policy Algorithm (IPPO) (independent critics)

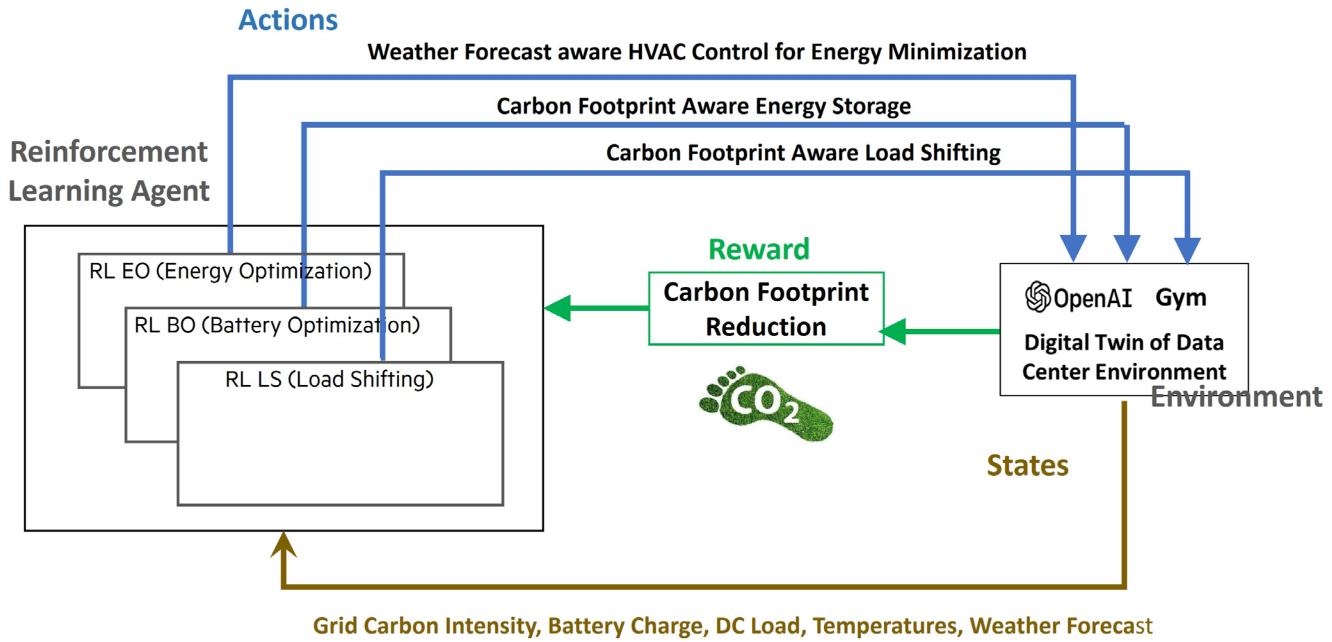

Concurrent Carbon Footprint Reduction Reinforcement Learning Control with Multi-agent Proximal Policy Optimization (PPO)

Multi-agent Reinforcement Learning

	$\mathbf{MDP_{LS}}$ Flexible Load Shifting	$\mathbf{MDP_{E}}$ Energy HVAC Optimizer	$\mathbf{MDP_{BAT}}$ Battery Agent
State: S_t	Time, DC temperature, IT Load, Unassigned Flexible Load, DC Energy, Carbon Intensity, Battery Charge	Time, DC temperature, Weather, DC Energy, IT Load, HVAC Setpoint	Time, DC Energy, Battery Charge, Carbon Intensity
Action: A_t	Assign Flexible Load, Idle	HVAC Setpoint	Charge, Supply, Idle
Reward: $R_{t+1}(S_t, A_t)$	$0.8 * r_{LS} + 0.1 * r_E + 0.1 * r_{BAT}$	$0.1 * r_{LS} + 0.8 * r_E + 0.1 * r_{BAT}$	$0.1 * r_{LS} + 0.1 * r_E + 0.8 * r_{BAT}$

Table 1: MDPs for Load Shifting, HVAC Energy Optimization, and Battery Operation. Here $r_{LS} = -(CO_2\ Footprint + LS_{Penalty})$, $r_E = -(Total\ Energy\ Consumption \times Cost\ per\ kWh)$, and $r_{BAT} = -(CO_2\ Footprint)$, where $LS_{Penalty}$ is the scalar value of the unassigned flexible IT workload.

Sustainability Dashboard for Data Center

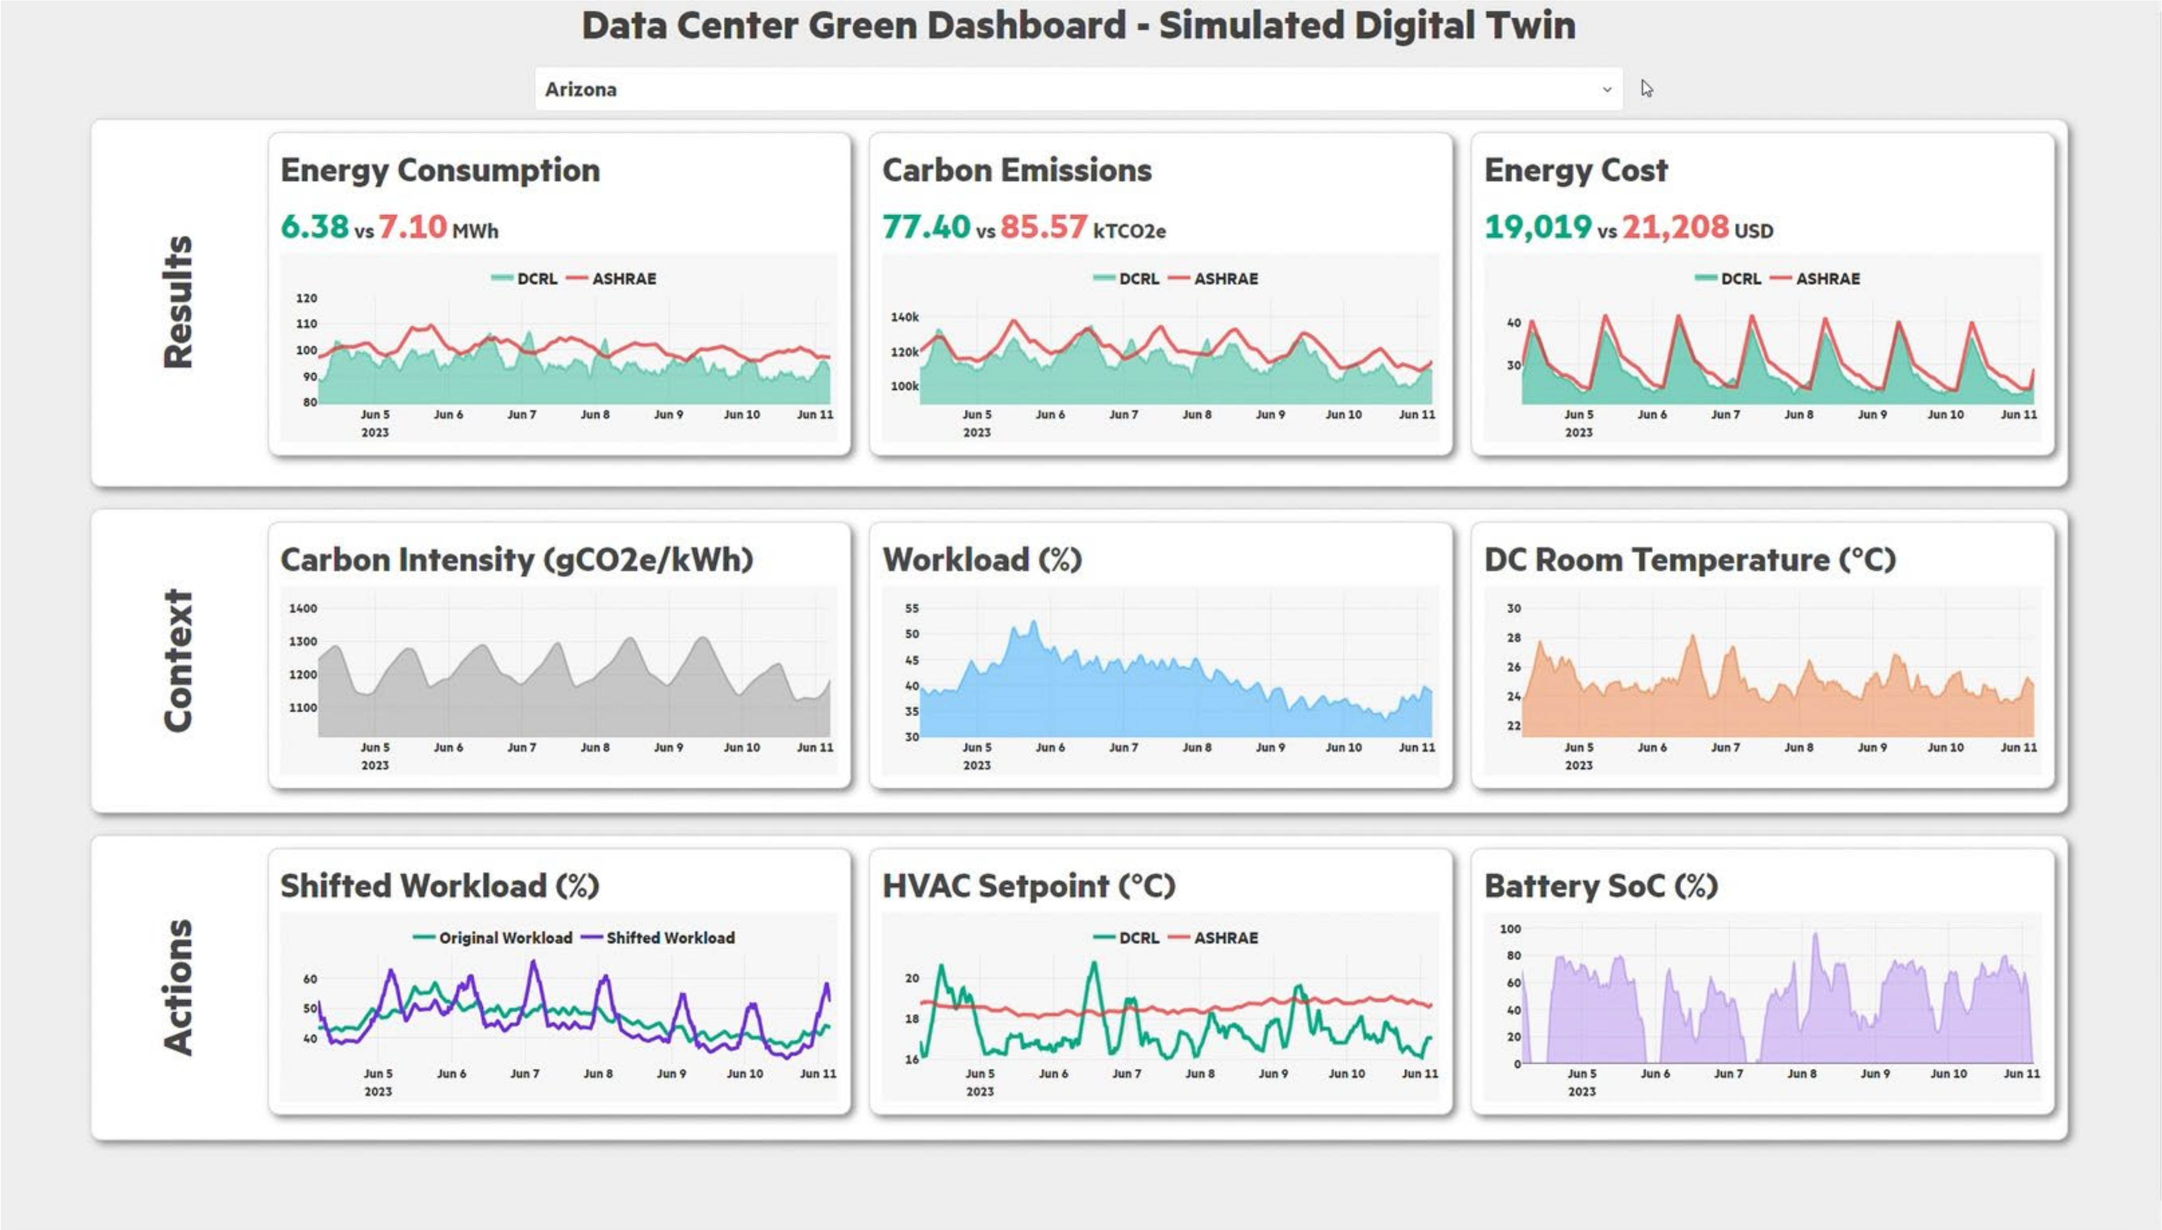

Results: CO2 emissions across locations and seasons

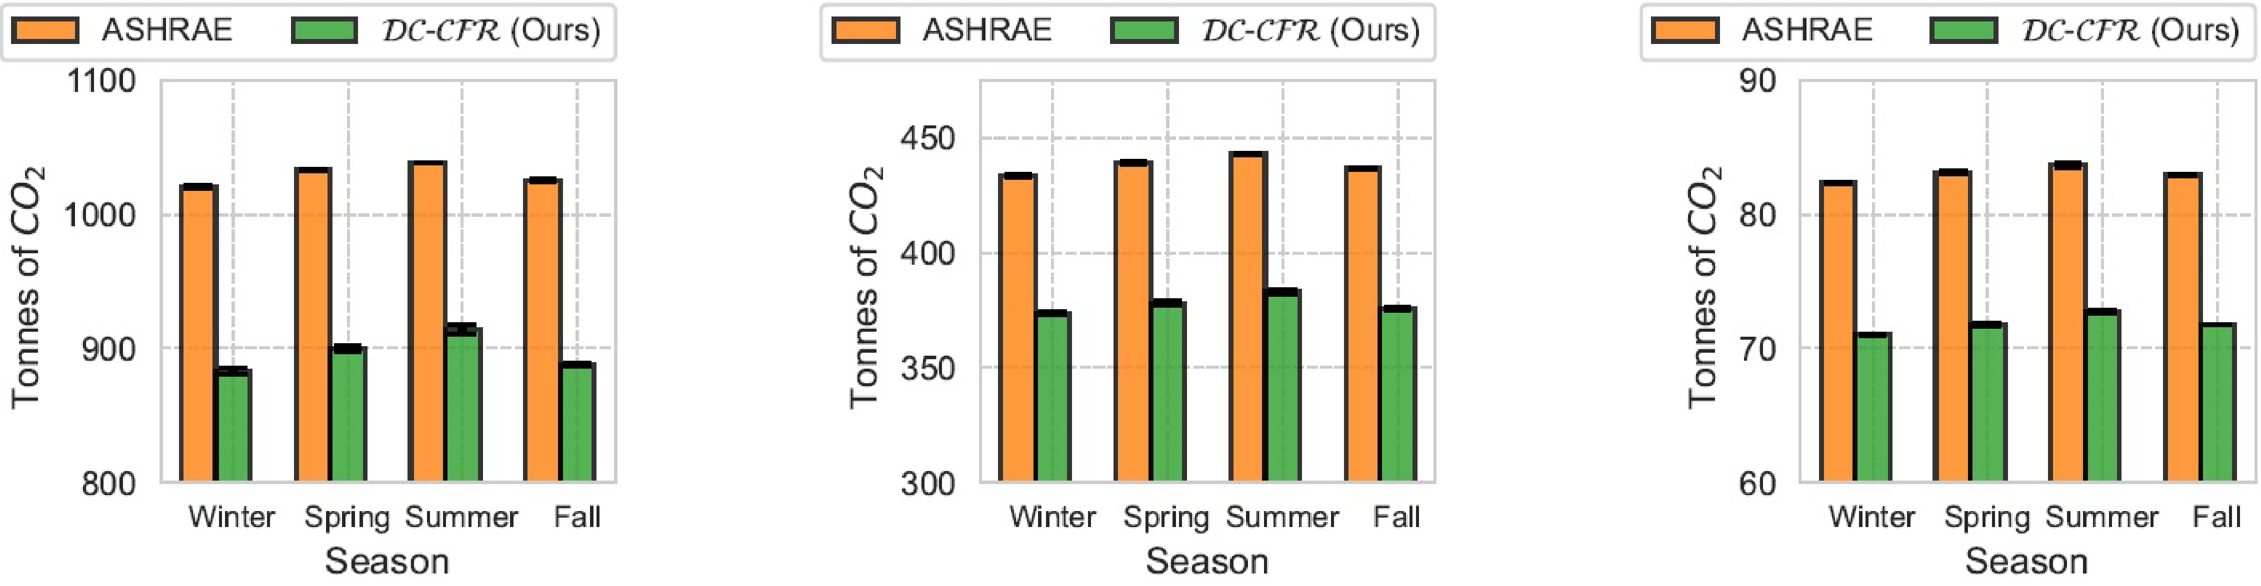

(a) Arizona, USA. (b) New York, USA. (c) Washington, USA.

	Percentage Reduction of Carbon Footprint with IPPO compared to ASHRAE Data Center Max Load 1.2MWh Experiment with EnergyPlus for a period of 1 year; Lookahead N = 4 hours						
	Algorithms						
	LS	EO	BAT	LS+EO	LS+BAT	EO+BAT	(Our proposal)
Arizona	7.72 ± 0.18	8.16 ± 0.05	0.25 ± 0.08	13.26 ± 0.07	7.98 ± 0.1	8.46 ± 0.05	14.36 ± 0.09
New York	7.13 ± 0.19	8.02 ± 0.06	0.41 ± 0.03	14.39 ± 0.08	7.68 ± 0.20	8.21 ± 0.07	15.08 ± 0.11
Washington	4.27 ± 0.20	7.54 ± 0.11	0.46 ± 0.05	13.62 ± 0.08	4.53 ± 0.17	7.78 ± 0.08	13.96 ± 0.06

Table 2: **Carbon Footprint Reduction Percentages** compared to industry standard ASHRAE: Performance of the individual approaches over a period of one year.

	Percentage Reduction of Energy Consumption with IPPO compared to ASHRAE Data Center Max Load 1.2MWh Experiment with EnergyPlus for a period of 1 year; Lookahead N = 4 hours						
	Algorithms						
	LS	EO	BAT	LS+EO	LS+BAT	EO+BAT	(Our proposal)
Arizona	7.11 ± 0.17	8.32 ± 0.04	0.00 ± 0.00	14.28 ± 0.07	7.15 ± 0.09	8.41 ± 0.05	14.54 ± 0.33
New York	7.05 ± 0.18	8.07 ± 0.06	0.00 ± 0.00	14.35 ± 0.08	7.12 ± 0.20	8.28 ± 0.08	14.62 ± 0.07
Washington	4.38 ± 0.21	7.42 ± 0.11	0.00 ± 0.00	13.78 ± 0.06	4.46 ± 0.18	7.31 ± 0.04	13.85 ± 0.07

Table 3: **Energy Reduction Percentages** compared to industry standard ASHRAE evaluated over one year.

HVAC cooling optimization insights

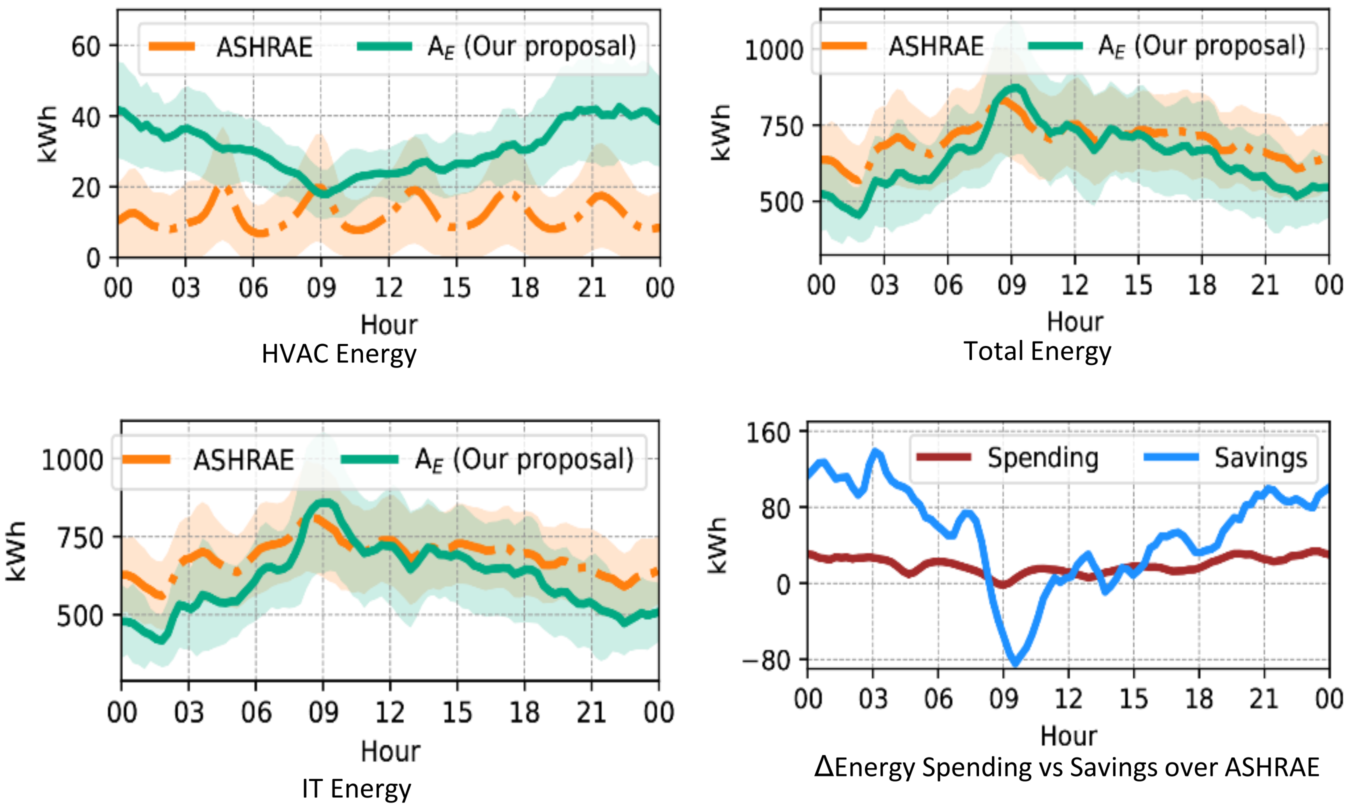

HVAC Cooling Optimization over ASHRAE Controller

Flexible Load Shifting optimization

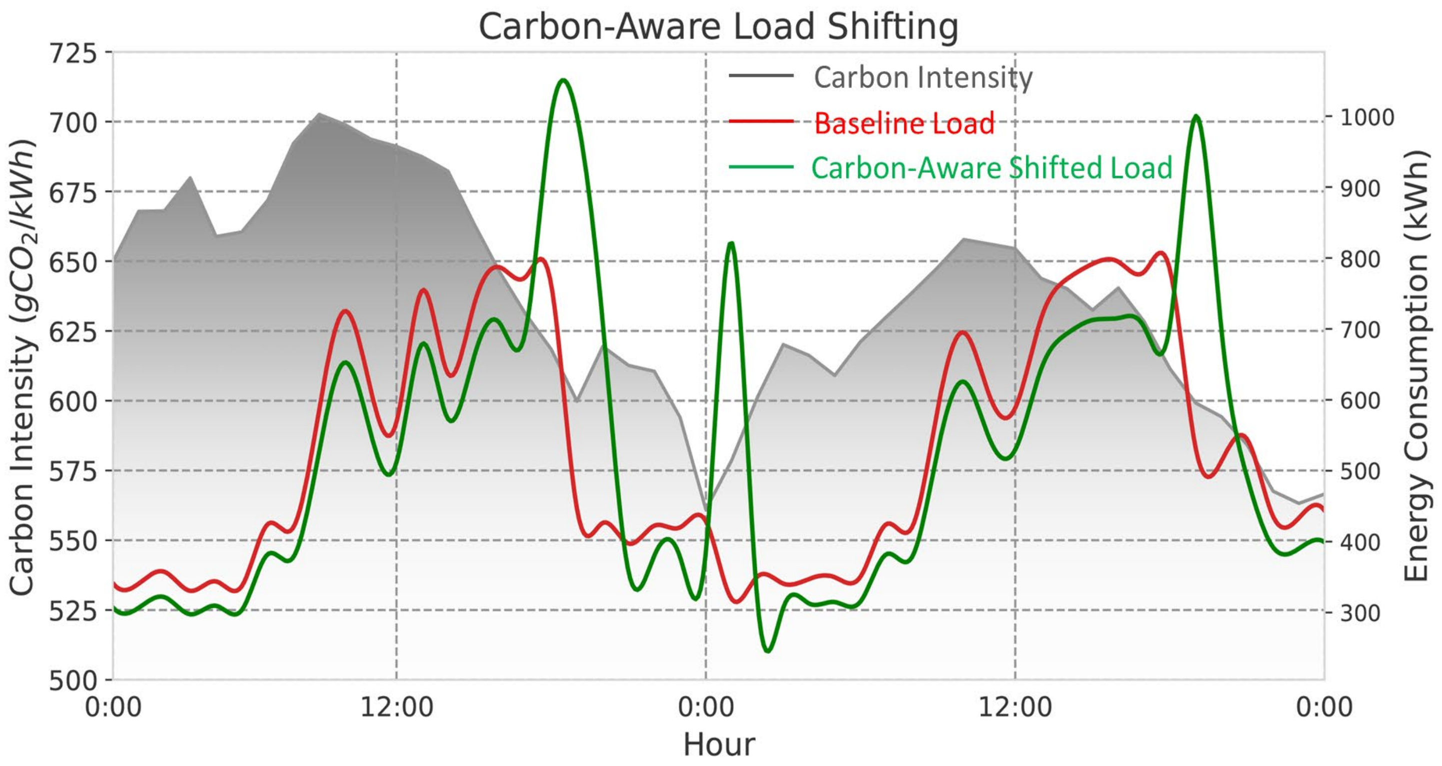

Carbon Intensity Aware Flexible Load Shifting

Energy Storage optimization

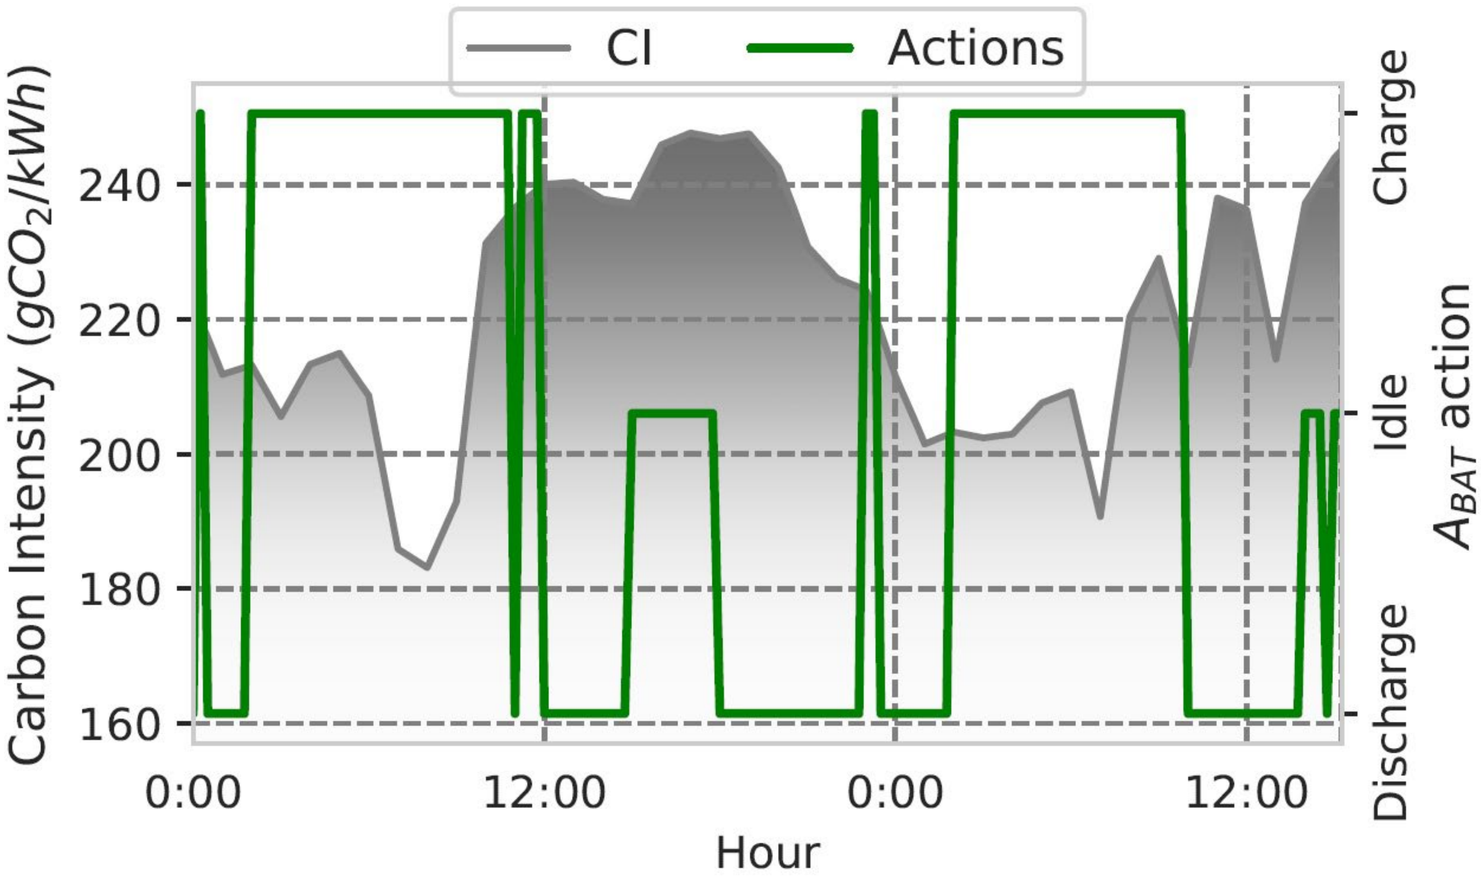

Sample time slice demonstrating the incremental carbon footprint savings using C2FR. We observe that the battery considers both the carbon intensity and the spikes in load to discharge and reduce carbon footprint.

Carbon Footprint reduction with the C2FR Framework

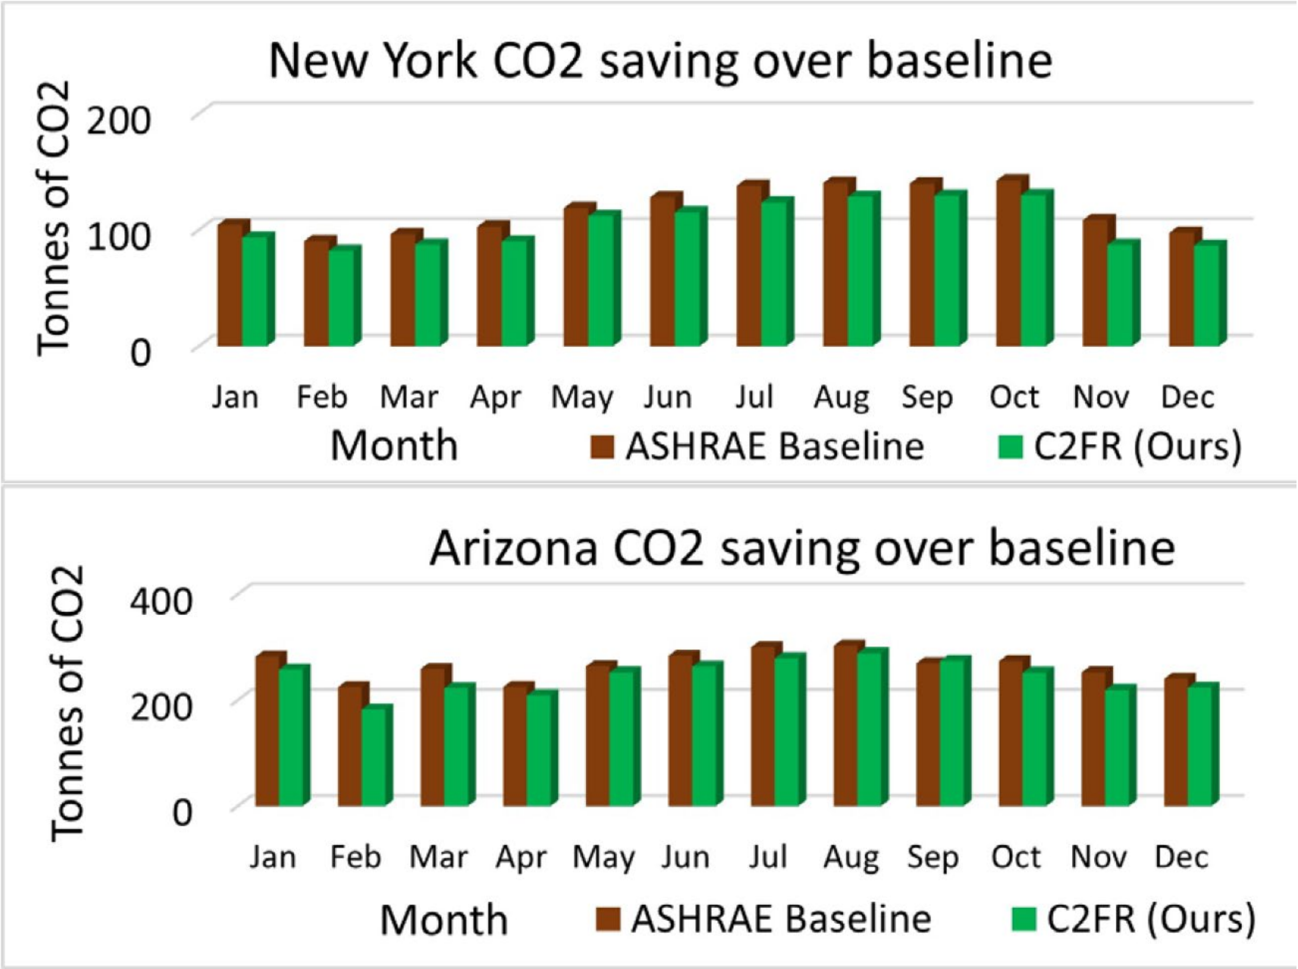

Monthly variations of Carbon footprint in the data centers in NY, and AZ states included in this study, controlled by ASHRAE baseline and C2FR (Ours)

Solution Highlights

Reduces data center carbon footprint and energy consumption

Reduces energy costs by shifting power load to lower-priced hours

Maximizes renewable energy usage through energy storage

Deep reinforcement learning (DRL) based

Implements carbon-aware load scheduling

Offers real-time control with all three optimizations, which has never been published

Optimizes HVAC control with weather forecasts

Modular, easily deployed, and integrated

Generalizable across multiple location/climate zones

Thank You

Soumyendu Sarkar
Hewlett Packard Labs @ Hewlett Packard Enterprise