DCcluster-Opt: Benchmarking Dynamic Multi-Objective Optimization for Geo-Distributed Data Center Workloads

Antonio Guillen-Perez*, Avisek Naug*, Vineet Gundecha, Sahand Ghorbanpour, Ricardo Luna Gutierrez, Ashwin Ramesh Babu, Munther Salim, Shubhanker Banerjee, Eoin H. Oude Essink, Damien Fay, Soumyendu Sarkar*

HPE Labs @ Hewlett Packard Enterprise

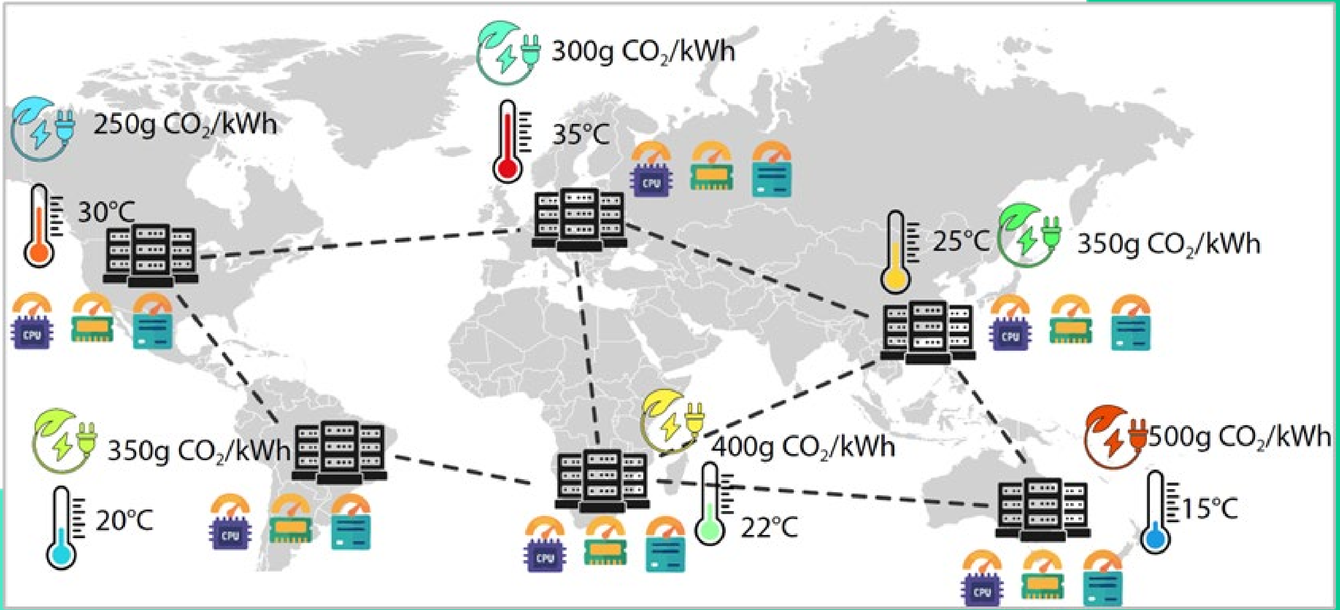

NEURAL INFORMATION PROCESSING SYSTEMS 2025

Dec 4, 2025

DCClusterOpt: HPC Cloud Optimization

With Multi-Agent Reinforcement Learning and LLMs

NeurIPS 2025 – Datasets & Benchmarks Track

The Challenge: AI's Unsustainable Footprint

- The computational demand of foundation models is growing exponentially, forcing a shift to vast, geo-distributed HPC ecosystems.
- This creates an unsustainable rise in **energy demands, carbon footprint, and operational costs**

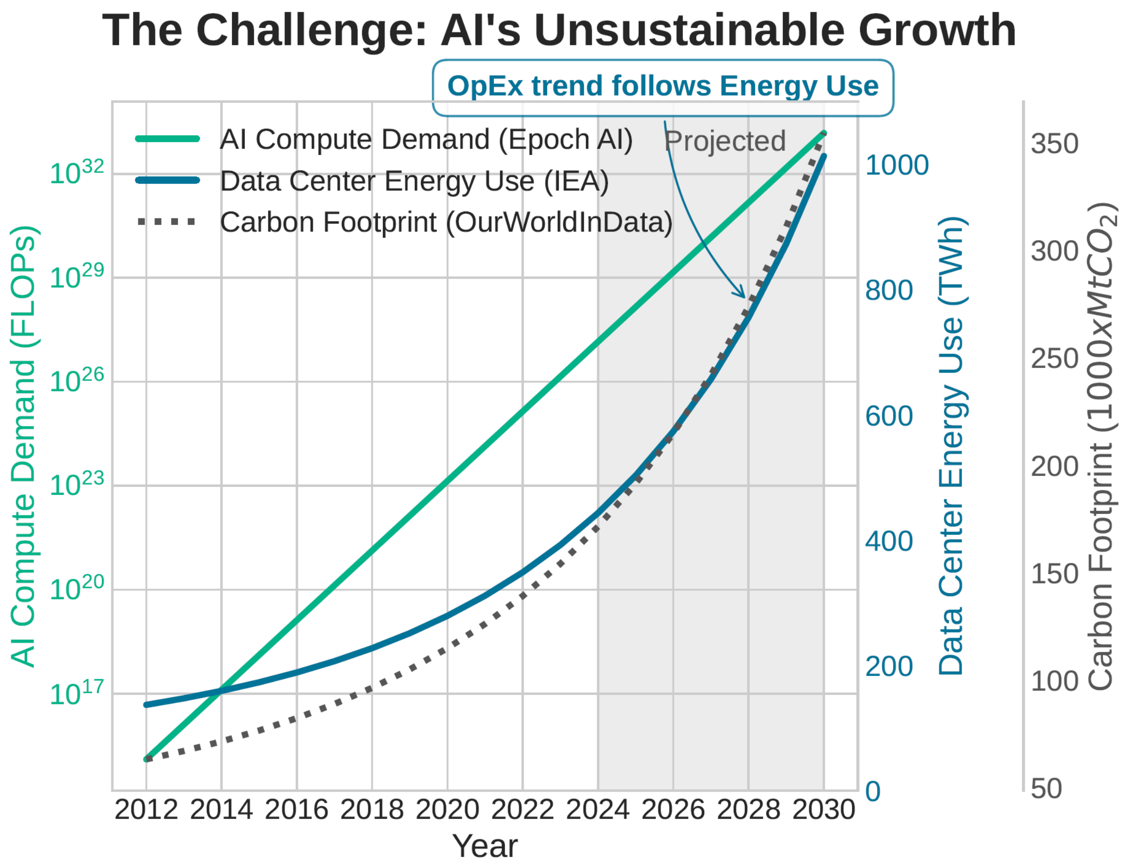

[1] https://epoch.ai/blog/training-compute-of-frontier-ai-models-grows-by-4-5x-per-year
[2] https://www.iea.org/data-and-statistics/charts/global-data-centre-electricity-consumption-by-equipment-base-case-2020-2030
[3] https://ourworldindata.org/grapher/carbon-intensity-electricity?tab=line&time=earliest..2024&country=EU-27~OWID_WRL

Motivation & The Gap

The Problem: AI's Footprint vs. The Reality Gap

Managing global HPC infrastructure is a massive operations research problem, requiring trade-offs between conflicting goals:

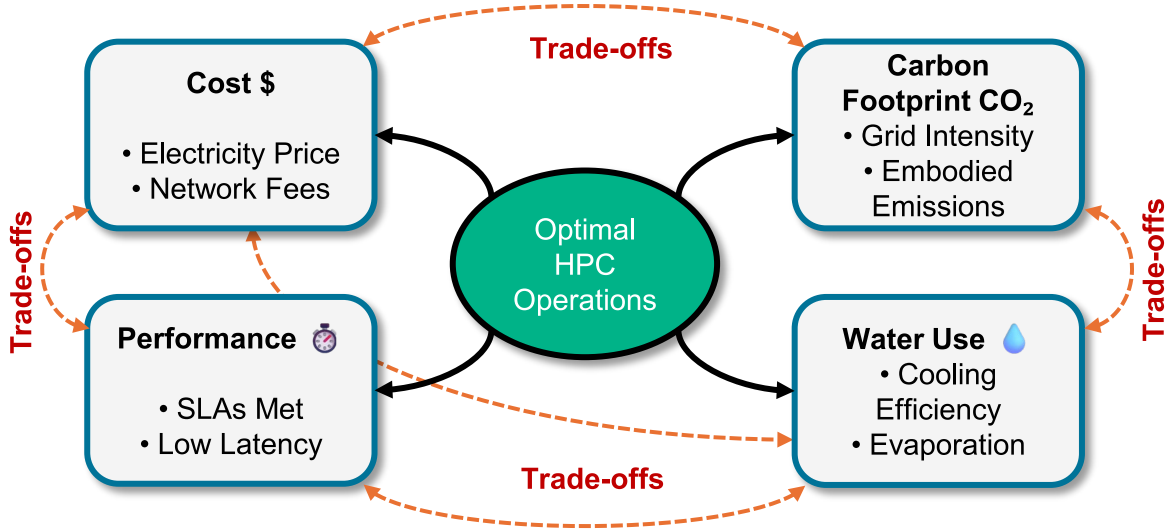

The Problem: AI's Footprint vs. The Reality Gap

Progress is blocked by a lack of realistic testbeds. Existing research is siloed and often:

- Uses abstract or simplified physical models.
- Optimizes either scheduling or cooling, but not both.
- Lacks real-world, dynamic environmental data.

We can't deploy trustworthy AI controllers without a realistic, high-fidelity digital twin to train and validate them.

Core Innovation: DCcluster-Opt

NEURAL INFORMATION PROCESSING SYSTEMS 2025

Our Contribution: The DCcluster-Opt Benchmark

An open-source, high-fidelity simulation benchmark for sustainable, geo-temporal task scheduling. It is the first to holistically integrate four crucial elements at a global scale:

1. Geo-Distributed Cluster
Models a global cluster across 20+ regions.

2. Real-World Data
Driven by real AI workload traces, grid carbon, electricity prices, and weather data.

3. Physics-Informed Models
High-fidelity models of all DC physics: IT Load, Network, and Complex Cooling (HVAC, Chillers & Liquid Cooling).

4. Multi-Objective Optimization
A modular reward system to study complex trade-offs (Cost vs. Carbon vs. SLA vs. Water).

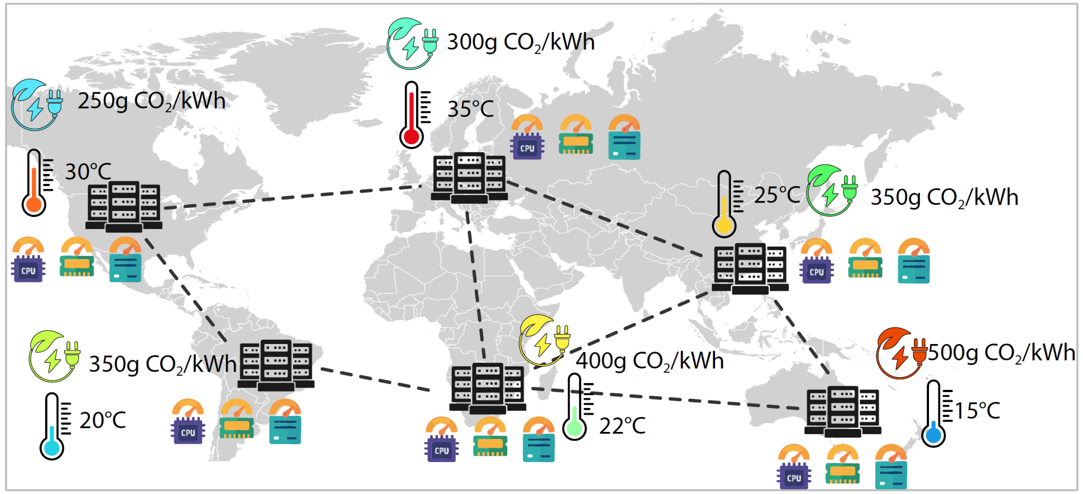

Varying Carbon Intensity

Dynamic Workloads

Volatile Energy Prices

Network Latency & Costs

Why it matters: For the first time, we can rigorously test and compare complex AI scheduling strategies on a realistic, globally-distributed, and physically-accurate simulation.

The DCcluster-Opt MDP: Geo-Temporal Task Scheduling

NEURAL INFORMATION PROCESSING SYSTEMS 2025

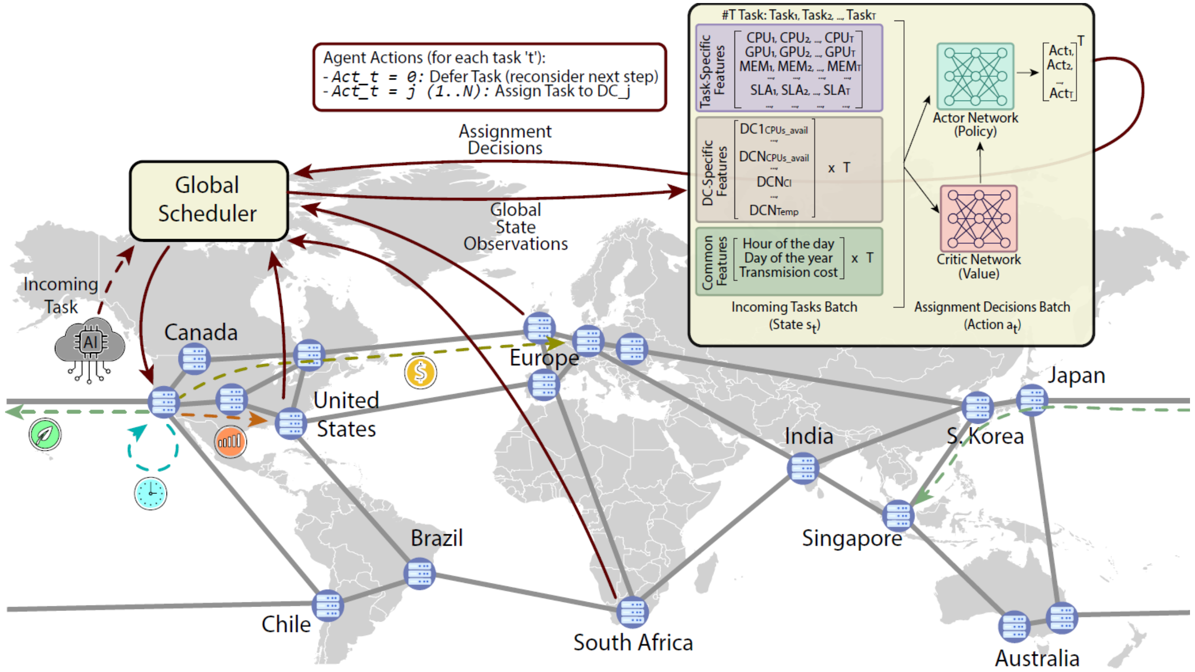

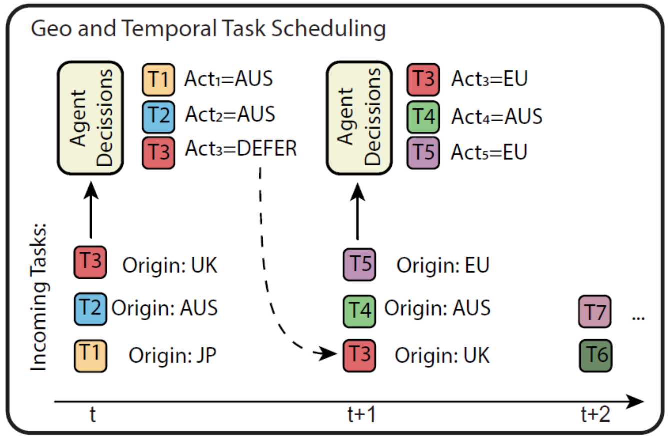

At each 15-minute timestep, the **Global Scheduler** solves a complex Markov Decision Process:

1. Observe State (S)

The agent receives a rich, variable-length state for a batch of `k` pending tasks. For each task, the state includes:

- **Task-Specific Features:** CPU, GPU, Memory, and SLA requirements.
- **DC-Specific Features:** Real-time load, energy price, and carbon intensity from all `N` data centers.
- **Common Features:** Global time of day and year.

2. Take Action (A)

The policy network outputs `k` individual decisions, one for each task:

- **Geographical Action:** Assign task to a specific data center `DC_j`.
- **Temporal Action:** Defer task to be reconsidered in the next timestep.

3. Receive Reward (R)

The agent's goal is to learn a policy that maximizes a global reward signal, which is a weighted sum of:

- Total Energy Cost
- Total Carbon Emissions
- Water usage
- SLA Violations

AI Approach: MARL + LLMs

The AI Control Plane: Agents Optimizing Infrastructure

We move beyond simple heuristics to a hierarchical, AI-driven control system.

Multi-Agent Reinforcement Learning (MARL)

Global Scheduler (RL Agent): A coordinating agent (SAC, PPO) makes high-level decisions.
- **Action 1:** *Where* to run a task? (Geo-placement)
- **Action 2:** *When* to run a task? (Temporal deferral)

Local Controllers (RL Agents): Agents inside each DC optimize local physics.
- **Action 3:** *How* to cool? (Dynamic setpoints for HVAC / Liquid Cooling)

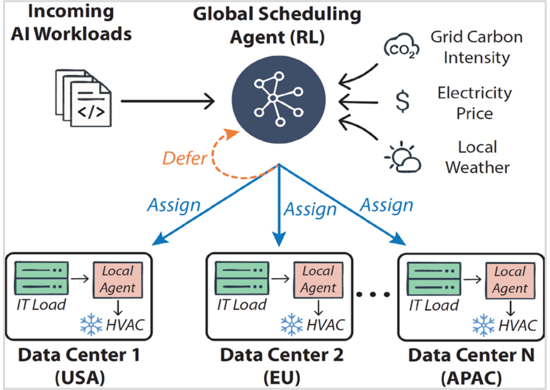

LLM-based Reasoning Controller

Moves beyond "black box" RL to an auditable, "agentic" system.

We use **Policy Distillation** to transfer the expert RL strategy into a flexible LLM controller.

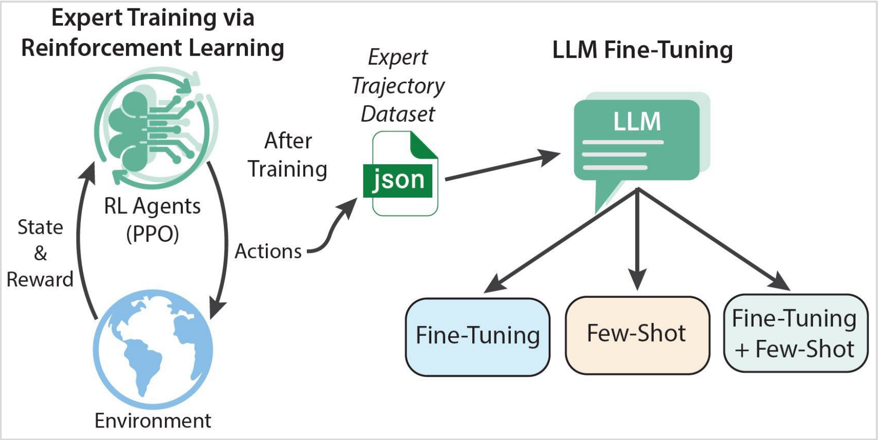

Why it matters: This shows a clear path from complex dynamic optimization (**RL**) to trustworthy, auditable, and adaptable control (**LLMs**) required for production HPC systems.

Agentic Design

Architecture: The Digital Twin Feedback Loop

Our agentic system mimics an expert human operations team, creating a transparent and auditable control plane.

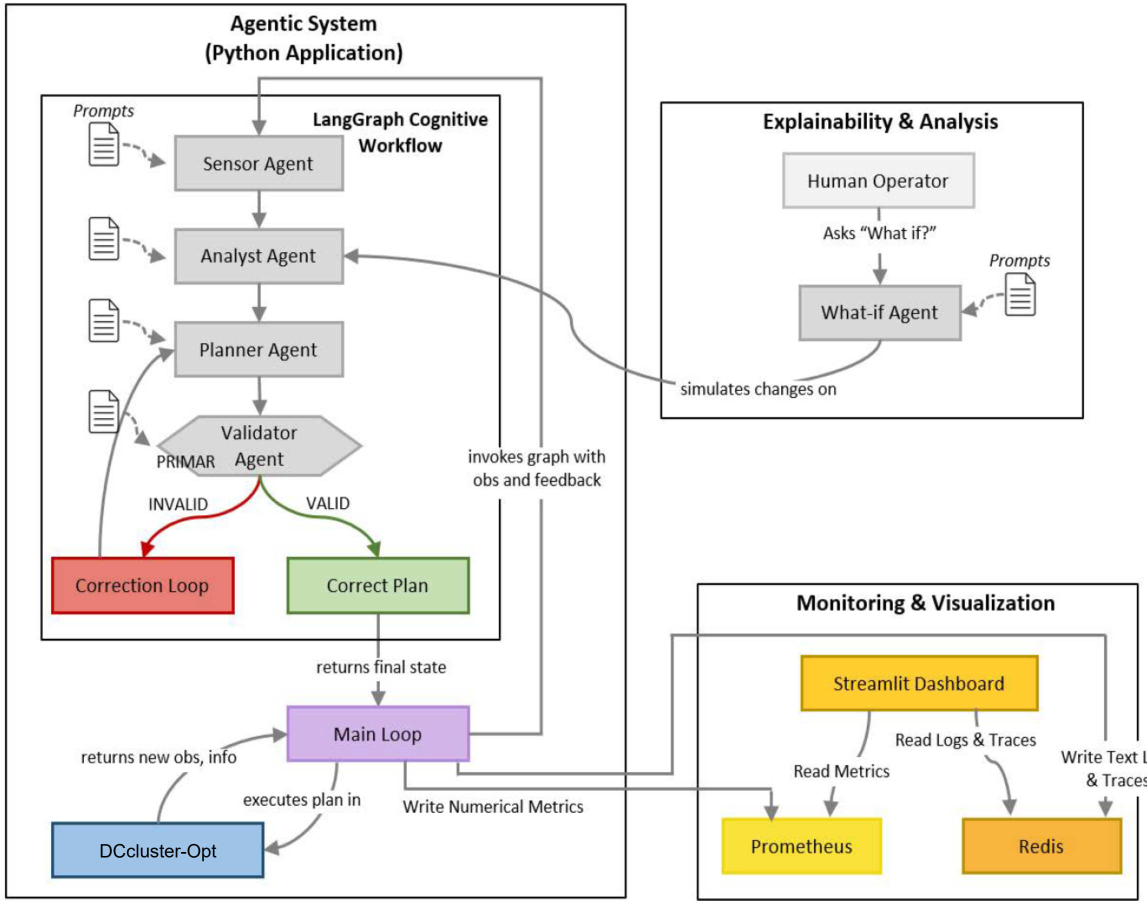

1. **DCcluster-Opt (Digital Twin):** Emits complex, high-dimensional state (IT load, carbon, price, weather, cooling temps).
2. **Sensor Agent:** Translates raw state into semantic JSON for the LLM.
3. **Analyst & Planner Agents (LLM):** Formulates high-level strategy and a concrete, low-level action plan.
4. **Validator Agent (Guardrail):** Inspects the plan for safety, compliance, and correctness.
5. **Monitor Agent:** Observes outcomes and provides qualitative feedback for continuous adaptation.
6. **What-if Agent:** Allows to create counterfactual situations.

Putting it all together

NEURAL INFORMATION PROCESSING SYSTEMS 2025

Problem: Sustainable Geo-Distribution of Workloads in HPC Data Center Cluster

RL: Multi-agent Multi-objective scheduler resolving external and DC dependencies

LLM Reasoning Controller: RL trajectories textualized and **distilled into fine-tuned LLMs**

Agentic Mesh: Framework with guard rails and transparency for real-time deployment and analytics

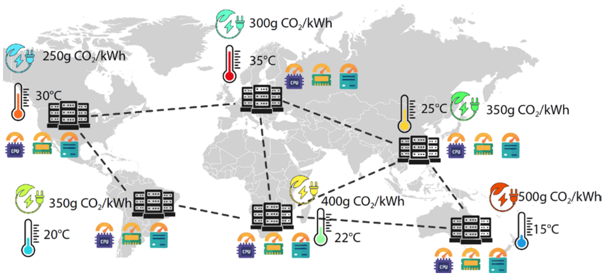

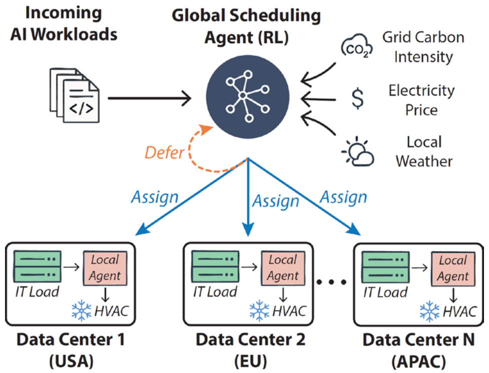

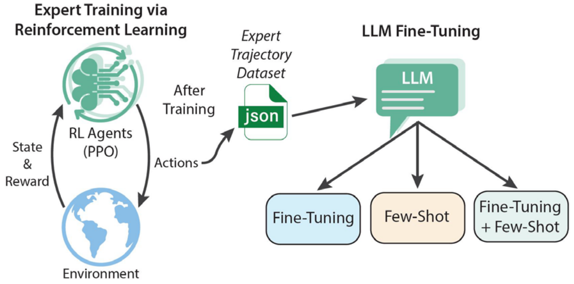

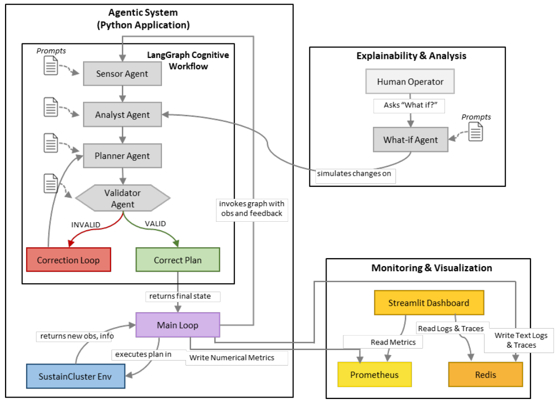

Explainability, Trust, Guidance & Insights

Priority Switching & Generalizability

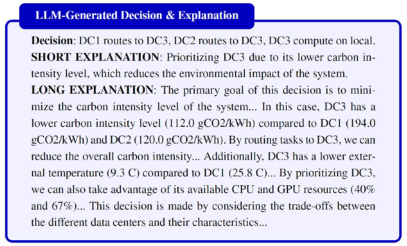

LLM-Generated Decision & Explanation

Decision: DC1 routes to DC3, DC2 routes to DC3, DC3 compute on local.
SHORT EXPLANATION: Prioritizing DC3 due to its lower carbon intensity level, which reduces the environmental impact of the system.
LONG EXPLANATION: The primary goal of this decision is to minimize the carbon intensity level of the system... In this case, DC3 has a lower carbon intensity level (112.0 gCO2/kWh) compared to DC1 (194.0 gCO2/kWh) and DC2 (120.0 gCO2/kWh). By routing tasks to DC3, we can reduce the overall carbon intensity... Additionally, DC3 has a lower external temperature (9.3 C) compared to DC1 (25.8 C)... By prioritizing DC3, we can also take advantage of its available CPU and GPU resources (40% and 67%)... This decision is made by considering the trade-offs between the different data centers and their characteristics...

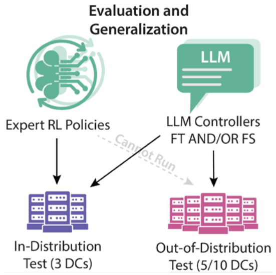

Grounded in Reality: Real-World Datasets

DCcluster-Opt integrates diverse, real-world data streams to drive the simulation. **This is not a synthetic problem.**

Data Component	Source(s)	Role & Coverage
AI Workloads	Alibaba Cluster Trace 2020/ MIT*/Marconi Supercomputer*	Real GPU training/inference jobs (>6.5k GPUs)
Carbon Intensity	Electricity Maps	Real-time grid emissions (gCO_2eq/kWh) for 20+ global regions.
Electricity Prices	GridStatus.io, ISOs	Dynamic spot market prices ($/kWh).
Weather	Open-Meteo	Ambient temperatures that drive cooling physics.
Network	Cloud Providers, Persico et al. [2]	Real monetary transfer costs ($/GB) and empirical latency (ms).

* WIP

Why it matters: The AI controller learns to exploit real, complex patterns, such as scheduling tasks to "follow the renewables" or avoid peak pricing.

[2] V. Persico, et al., "A First Look at Public-Cloud Inter-Datacenter Network Performance," 2016 IEEE Global Communications Conference (GLOBECOM), Washington, DC, USA, 2016, pp. 1-7, doi: 10.1109/GLOCOM.2016.7841498.

Results & Insights

Holistic AI Control Delivers >11% Savings

Comparing our Global Scheduler (SAC) + Local RL HVAC Controller against a baseline with fixed cooling setpoints. *(30-day simulation)*

Scenario	Total Energy (MWh) ↓	Total CO_2 (t) ↓	Total Cost ($) ↓	Water usage (m^3)
Fixed Baseline	1087.1	328.7	$101,401	7394.1
Holistic AI Control	**921.3**	**273.1**	**$81,300**	**7253.6**

-15.27% Total Energy (MWh)

-16.78% Total CO_2 Emissions (t)

-19.82% Total Cost ($)

140,500 Liters of Water Saved

Multi-Objective Trade-offs

Rule-based methods (e.g., "Lowest Carbon") excel in one metric but fail in others (e.g., high SLA violations).

Our **RL (Geo+Time)** agent learns a balanced policy that navigates the trade-offs defined in the reward function.

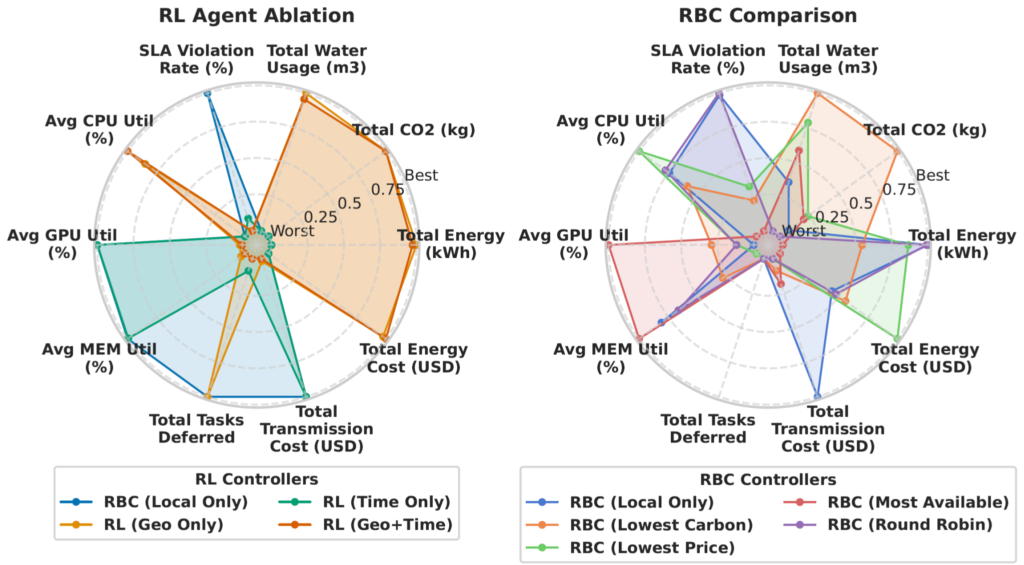

Why it matters: Holistic optimization wins. Optimizing scheduling and cooling physics together unlocks double-digit efficiency gains that siloed approaches miss.

From 'What' to 'Why': Trustworthy LLM Controllers

Policy distillation creates a transparent agent that can justify its complex scheduling decisions in natural language.

LLM-Generated Decision & Explanation

```
DECISION: Assign Task T-12 (Origin: US-CAL) to DC3 (DE-LU). Defer
Task T-13.

SHORT EXPLANATION: Prioritizing DC3 due to its significantly lower
carbon intensity and favorable electricity price, which aligns with
our global objective of minimizing environmental impact and cost.

LONG EXPLANATION: The primary goal is to minimize carbon. DC3 (DE-
LU) has a current carbon intensity of 112.0 gCO2/kWh, compared to
194.0 gCO2/kWh at the origin DC1 (US-CAL)... By routing this task,
we accept a small network cost to gain larger carbon and energy
savings...
```

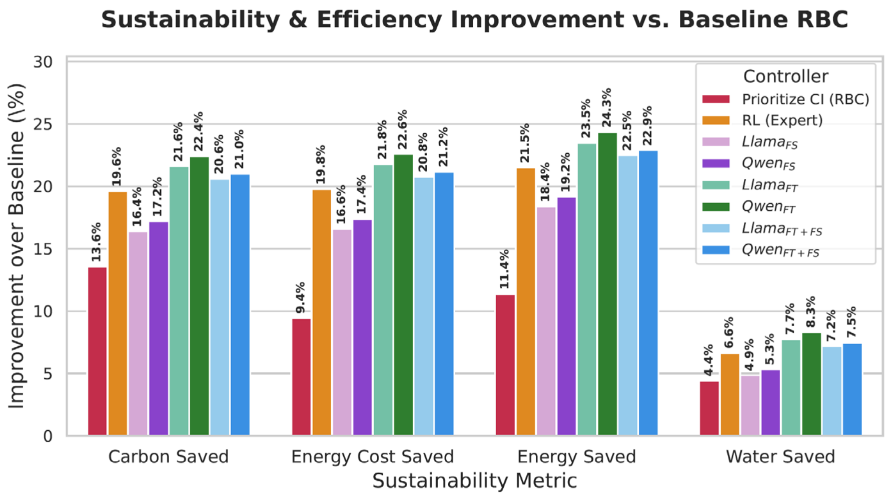

Key Benefits:

- **Explainability & Trust**: Creates auditable, human-readable logs for operators.
- **Adaptability**: Objectives can be changed "on-the-fly" with natural language prompts.
- **Scalability**: Reasoning generalizes to new data centers without costly retraining

Scenario	Controller	Total CO_2 (t) ↓	SLA (%) ↓
5-Datacenters	RBC	62.76	1.48
	LLM	59.55	1.08
10-Datacenters	RBC	131.81	1.95
	LLM	125.10	1.19

Table 2. Generalization (Out-of-Distribution) results

Impact & Business Value

NEURAL INFORMATION PROCESSING SYSTEMS 2025

From Simulation to Sustainable Operations

This research provides a direct, actionable blueprint for enterprise and national lab sustainability.

Reduce Operational Costs (OpEx)

Demonstrated **>12% reduction** in total operational energy costs from holistic AI control.

Meet Enterprise ESG Goals

Quantifiably reduce total CO_2 emissions by **>11.5%** and enable granular carbon-aware scheduling.

Optimize Resource Efficiency

Physics-informed models for water usage and Heat Recovery Units (HRUs) allow for minimizing waste.

De-Risk Deployment

Develop, audit, and validate AI controllers before deploying them on mission-critical HPC systems.

Broader Vision and Future Work

Towards Self-Optimizing, Net-Zero HPC Clouds

The next generation of scientific infrastructure (like the **American Science Cloud - AmSC**) will be too complex for human operators alone. They require agentic, trustworthy AI control planes.

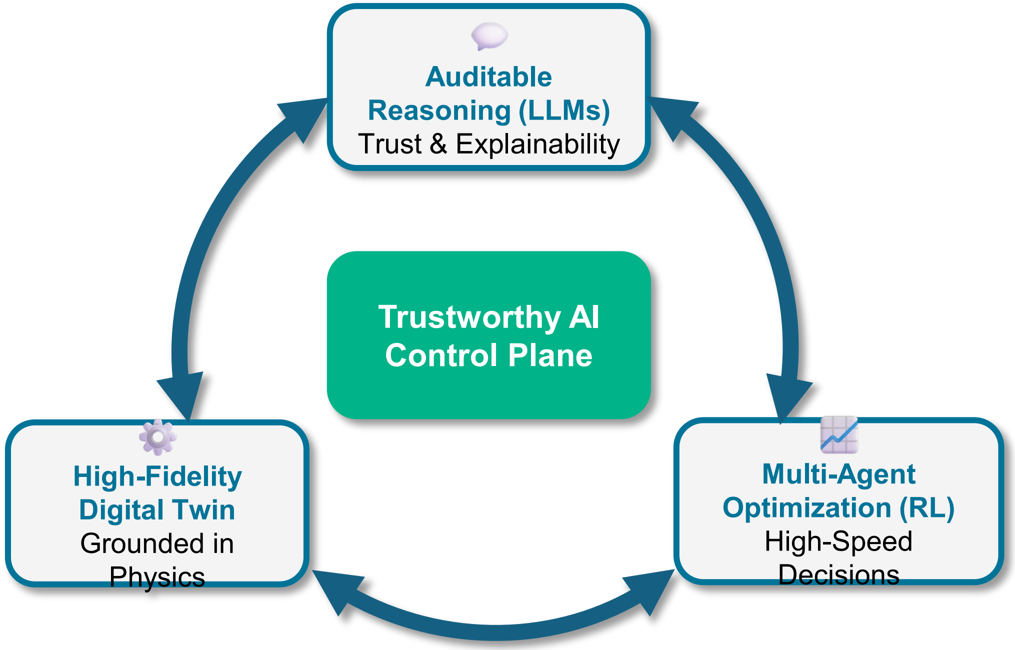

Addressing Core AmSC Requirements:

- **Managing a Geo-Distributed, Heterogeneous System:** DCcluster-Opt is purpose-built to model the AmSC's federated topology, simulating diverse hardware across locations with unique electricity grids (price & carbon) and network links.
- **Sustainable and Cost-Effective Operation:** Our benchmark provides the essential multi-objective framework (Cost, Carbon, Water, SLA) to train an AmSC scheduler to verifiably minimize both operational cost and environmental impact.
- **Trustworthy and Auditable Control:** DCcluster-Opt serves as a **high-fidelity digital twin**, allowing AI controllers to be developed, audited, and de-risked before deployment on mission-critical hardware. Our LLM agents show a path to the required explainability.

The Goal: Autonomous, resilient, and sustainable infrastructure that self-optimizes for cost, carbon, and science, 24/7.

Conclusion & Call to Action

NEURAL INFORMATION PROCESSING SYSTEMS 2025

AI Optimizing AI: From Simulation to Sustainable Operations

We are releasing the entire Dccluster-Opt ecosystem to provide the robust, standardized sandbox the field has been missing.

Our Contributions:

- **DCcluster-Opt:** A high-fidelity, open-source benchmark integrating global scheduling, real-world data, and complex cooling physics.
- **AI Control Plane:** A novel hybrid MARL + LLM agentic system for multi-objective optimization.
- **Quantifiable Impact:** Demonstrated >11% reductions in Cost, Carbon, and Energy.
- **Gymnasium API and Ray RLlib Integration:** Compatibility with SOTA libraries (SAC, PPO, APPO, IMPALA).
- **Discoverable Datasets:** Fully documented with Croissant metadata, a new standard for ML-ready datasets.
- **Community-Ready:** Open-source and organized to encourage contributions.

Join Us

"We provided the testbed to solve a hard operations problem. Now, let's go play with it.“

Paper:

Code:

Code

github.com/HewlettPackard/sustain-cluster

Docs

hewlettpackard.github.io/sustain-cluster/

NeurIPS 2025

Thank You

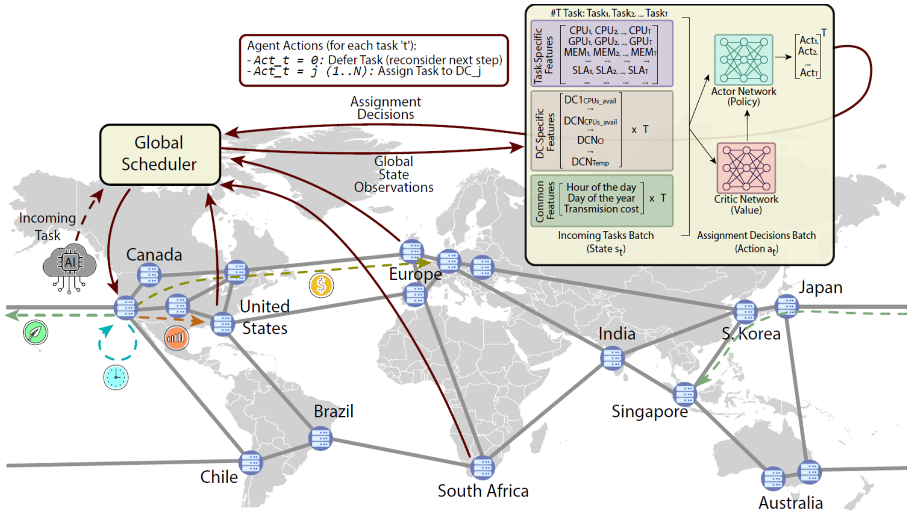

Paper:

Code:

Antonio Guillen

Avisek Naug

Vineet Gundecha

Sahand Ghorbanpour

Ricardo Luna

Ashwin R Babu

Munther Salim

Damien Fay

Soumyendu Sarkar

ML Research
@HPE Labs

DCcluster-Opt: Benchmarking Dynamic Multi-Objective Optimization for Geo-Distributed Data Center Workloads

soumyendu.sarkar@hpe.com