

Sustainability of Data Center Digital Twins with Reinforcement Learning

Soumyendu Sarkar^{*†}, Avisek Naug[†], Antonio Guillen[†], Ricardo Luna[†], Vineet Gundecha[†],
Ashwin Ramesh Babu, Sajad Mousavi

Hewlett Packard Enterprise

{soumyendu.sarkar, avisek.naug, antonio.guillen, rluna, vineet.gundecha, sajad.mousavi, ashwin.ramesh-babu}@hpe.com

Abstract

The rapid growth of machine learning (ML) has led to an increased demand for computational power, resulting in larger data centers (DCs) and higher energy consumption. To address this issue and reduce carbon emissions, intelligent design and control of DC components such as IT servers, cabinets, HVAC cooling, flexible load shifting, and battery energy storage are essential. However, the complexity of designing and controlling them in tandem presents a significant challenge. While some individual components like CFD-based design and Reinforcement Learning (RL) based HVAC control have been researched, there's a gap in the holistic design and optimization covering all elements simultaneously. To tackle this, we've developed DCRL-Green, a multi-agent RL environment that empowers the ML community to design data centers and research, develop, and refine RL controllers for carbon footprint reduction in DCs. It is a flexible, modular, scalable, and configurable platform that can handle large High Performance Computing (HPC) clusters. Furthermore, in its default setup, DCRL-Green provides a benchmark for evaluating single as well as multi-agent RL algorithms. It easily allows users to subclass the default implementations and design their own control approaches, encouraging community development for sustainable data centers. Open Source Link: <https://github.com/HewlettPackard/dc-rl>

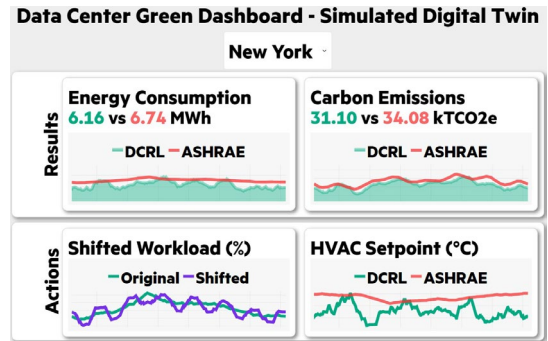


Figure 2: Sustainability Dashboard for Data Centers

Introduction

In order to reduce carbon emissions, data centers have to be efficient with respect to design and control. Recent advances in Reinforcement Learning (RL) have led to the creation of key RL environments for optimizing energy and resource allocation in building and data center management (Va'zquez-Canteli et al. 2019; Scharnhorst et al. 2021; Moriyama et al. 2018; Crawley et al. 2000; Acun et al. 2023; Radovanovic' et al. 2023). While they focus on individual approaches like flexible load shifting, HVAC setpoint optimization and battery scheduling, none of them jointly optimize the behavior of these individual approaches in real-time due to the difficulty in modeling the complex interdependencies (Figure 1). Our work achieves this through the DCRL-Green framework (Figure 2).

DCRL-Green

DCRL-Green is a tool for evaluating designing carbon-efficient data centers and single as well as multi-agent reinforcement learning strategies aimed at reducing data center carbon footprints. Built on the OpenAI Gym, this comprehensive framework highlights:

Simulation Framework (Data Center Digital Twin in Figure 3): DCRL models multizone zone IT Rooms, where the user can provide custom IT Cabinet geometry, server power specifications, and an HVAC system comprising chillers, pumps, and cooling towers (Naug et al. 2023a,b). It also has default as well as easily customizable models for




Figure 1: Internal and External System Dependencies

*Corresponding Author

†These authors contributed equally.

Copyright © 2024, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

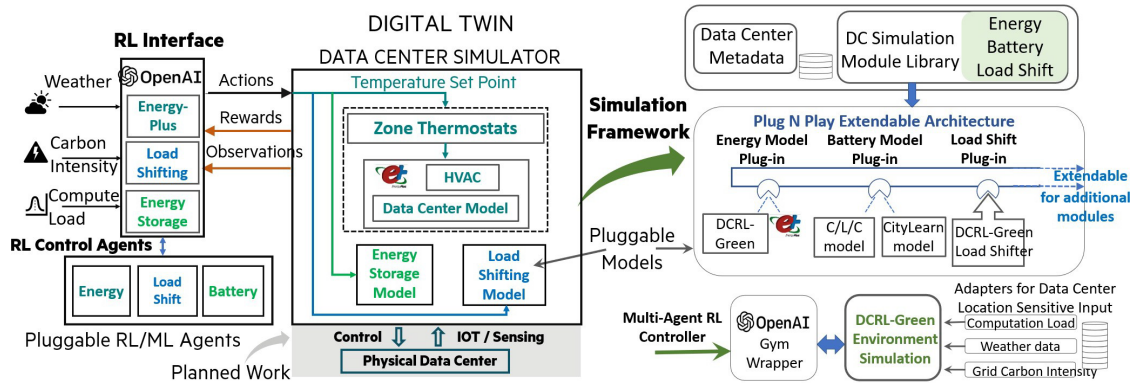


Figure 3: DCRL-Green Data Center Digital Twin Configurable Modeling.

grid carbon intensity aware Load Shifting Model and auxiliary Battery supply. The models seamlessly integrate with external software like EnergyPlus or directly in Python.

Configurability (Plug N Play Extendable Architecture in Figure 3): Users can extensively tailor data center designs, adjusting elements from server specifics to HVAC details via a JSON object file. Moreover, parameters such as workload profiles and weather data can be modified. These allow the users to rapidly prototype different data center designs, monitor potential temperature hotspots. It also accommodates variations in geography and sustainability metrics.

Interface (Open AI Gym Controller Interface in Figure 3) To facilitate single as well as holistic multiagent optimization of Data Center Carbon footprint, DCRL-Green provides interfaces for applying reinforcement learning-based control (Control Agents in Figure 3). It uses the massively scalable RLLib (Liang et al. 2018) that allows users to readily try an array of single as well as multiagent algorithms for optimizing flexible load scheduling, HVAC setpoint control, and battery supply (Sarkar et al. 2023b,a,c,d).

Overall, DCRL-Green fosters sustainable data center operations, promoting collaborative green computing research within the ML community.

Applications of DCRL

Single and Multi-Agent Optimization

DCRL has also been used to optimize HVAC setpoint control using classical single-agent reinforcement learning Table 1 leading to 7% carbon emission reduction and multi-agent reinforcement learning Table 2 leading to 13% carbon emission reduction.

Conclusion

The paper presents DCRL-Green, an OpenAI Gym environment for reinforcement learning in data centers (DCs), offering customizable DC configurations and targeting various sustainability goals, thereby enabling ML researchers to mitigate climate change effects from rising DC workloads. As part of future work, CFD neural surrogates can automate parameter generation for a custom data center configuration.




Figure 4: Temperature Distribution generated from a given Configuration of Data Center

Controller	Energy Footprint (MWh)	Carbon Footprint (TonnesCO ₂)
PPO	26.42	13726.56
A2C	26.74	13878.30
ASHRAE	28.52	14839.00
Guideline36		

Table 1: Energy and Carbon Footprint Comparison for single agent (HVAC Control setpoint)

Region	Energy	Carbon Footprint
Arizona	13.4 ± 0.48	13.38 ± 0.62
New York	13.01 ± 0.12	12.77 ± 0.11
Washington	13.27 ± 0.06	13.51 ± 0.05

Table 2: Percentage Reduction in Energy and Carbon footprint using a Multiagent approach (MADDPG (Lowe et al. 2017)) when compared against the ASHRAE G36 control

Lowering the energy usage of data centers and shifting the energy consumption to periods when more green energy is available on the power grid can significantly impact the reduction of the carbon footprint of the data centers and help fight climate change. Also this open source implementation democratizes access to machine learning researchers to contribute to a greener planet.

References

- Acun, B.; Lee, B.; Kazhamiaka, F.; Maeng, K.; Gupta, U.; Chakkaravarthy, M.; Brooks, D.; and Wu, C.-J. 2023. Carbon Explorer: A Holistic Framework for Designing Carbon Aware Datacenters. In *Proceedings of the 28th ACM International Conference on Architectural Support for Programming Languages and Operating Systems, Volume 2*. ACM.
- Crawley, D. B.; Lawrie, L. K.; Pedersen, C. O.; and Winkelmann, F. C. 2000. Energy plus: energy simulation program. *ASHRAE journal*, 42(4): 49–56.
- Liang, E.; Liaw, R.; Nishihara, R.; Moritz, P.; Fox, R.; Goldberg, K.; Gonzalez, J.; Jordan, M.; and Stoica, I. 2018. RLlib: Abstractions for distributed reinforcement learning. In *International Conference on Machine Learning*, 3053–3062. PMLR.
- Lowe, R.; Wu, Y. I.; Tamar, A.; Harb, J.; Pieter Abbeel, O.; and Mordatch, I. 2017. Multi-agent actor-critic for mixed cooperative-competitive environments. *Advances in neural information processing systems*, 30.
- Moriyama, T.; Magistris, G. D.; Tatsubori, M.; Pham, T.; Munawar, A.; and Tachibana, R. 2018. Reinforcement Learning Testbed for Power-Consumption Optimization. *CoRR*, abs/1808.10427.
- Naug, A.; Guillen, A.; Luna Gutierrez, R.; Gundecha, V.; Ghorbanpour, S.; Dheeraj Kashyap, L.; Markovikj, D.; Krause, L.; Mousavi, S.; Babu, A. R.; et al. 2023a. PyDCM: Custom Data Center Models with Reinforcement Learning for Sustainability. In *Proceedings of the 10th ACM International Conference on Systems for Energy-Efficient Buildings, Cities, and Transportation*, 232–235.
- Naug, A.; Guillen, A.; Luna Gutierrez, R.; Gundecha, V.; Ghorbanpour, S.; Mousavi, S.; Ramesh Babu, A.; and Sarkar, S. 2023b. A Configurable Pythonic Data Center Model for Sustainable Cooling and ML Integration. In *NeurIPS 2023 Workshop on Tackling Climate Change with Machine Learning*.
- Radovanovic, A.; Koningstein, R.; Schneider, I.; Chen, B.; Duarte, A.; Roy, B.; Xiao, D.; Haridasan, M.; Hung, P.; Care, N.; Talukdar, S.; Mullen, E.; Smith, K.; Cottman, M.; and Cirne, W. 2023. Carbon-Aware Computing for Datacenters. *IEEE Transactions on Power Systems*, 38(2): 1270–1280.
- Sarkar, S.; Guillen, A.; Carmichael, Z.; Gundecha, V.; Naug, A.; Ramesh Babu, A.; and Luna Gutierrez, R. 2023a. Enhancing Data Center Sustainability with a 3D CNN-Based CFD Surrogate Model. In *NeurIPS 2023 Workshop on Tackling Climate Change with Machine Learning*.
- Sarkar, S.; Naug, A.; Guillen, A.; Gutierrez, R. L.; Ghorbanpour, S.; Mousavi, S.; Babu, A. R.; and Gundecha, V. 2023b. Concurrent Carbon Footprint Reduction (C2FR) Reinforcement Learning Approach for Sustainable Data Center Digital Twin. In *2023 IEEE 19th International Conference on Automation Science and Engineering (CASE)*, 1–8.
- Sarkar, S.; Naug, A.; Guillen, A.; Luna Gutierrez, R.; Gundecha, V.; Ghorbanpour, S.; Mousavi, S.; and Ramesh Babu, A. 2023c. Sustainable Data Center Modeling: A Multi-Agent Reinforcement Learning Benchmark. In *NeurIPS 2023 Workshop on Tackling Climate Change with Machine Learning*.
- Sarkar, S.; Naug, A.; Luna Gutierrez, R.; Guillen, A.; Gundecha, V.; Ramesh Babu, A.; and Bash, C. 2023d. Real-time Carbon Footprint Minimization in Sustainable Data Centers with Reinforcement Learning. In *NeurIPS 2023 Workshop on Tackling Climate Change with Machine Learning*.
- Scharnhorst, P.; Schubnel, B.; Fernandez Bandera, C.; Salom, J.; Taddeo, P.; Boegli, M.; Gorecki, T.; Stauffer, Y.; Peppas, A.; and Politi, C. 2021. Energym: A Building Model Library for Controller Benchmarking. *Applied Sciences*, 11(8).
- Va'zquez-Canteli, J. R.; Ka'mpf, J.; Henze, G.; and Nagy, Z. 2019. CityLearn v1.0: An OpenAI Gym Environment for Demand Response with Deep Reinforcement Learning. In *Proceedings of the 6th ACM International Conference on Systems for Energy-Efficient Buildings, Cities, and Transportation*, BuildSys '19, 356–357. New York, NY, USA: Association for Computing Machinery. ISBN 9781450370059.