Enhancing Sustainability in Liquid-Cooled Data Centers with Reinforcement Learning Control

Avisek Naug[†], **Antonio Guillen Perez**[†], **Vineet Gundecha, Ricardo Luna Gutierrez, Ashwin Ramesh Babu, Sajad Mousavi, Paolo Faraboschi, Cullen Bash, Soumyendu Sarkar**[†][*]

Hewlett Packard Enterprise

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

Abstract

The growing energy demands of machine learning workloads require sustainable data centers with lower carbon footprints and reduced energy consumption. Supercomputing and many high-performance computing (HPC) data centers, which use liquid cooling for greater efficiency than traditional air cooling systems, can significantly benefit from advanced optimization techniques to control liquid cooling. We present RL-LC, a novel Reinforcement Learning (RL) based approach designed to enhance the efficiency of liquid cooling in these environments. RL-LC integrates a customizable analytical liquid cooling model suitable for simulations or digital twins of data centers, focusing on minimizing energy consumption and carbon emissions. Our method achieves an average reduction of approximately 4% compared to industry-standard ASHRAE guidelines, contributing to more sustainable data center management and offering valuable insights for reducing the environmental impact of HPC operations.

1 Introduction

As digital infrastructure continues to expand, the need for sustainable and energy-efficient data centers has become increasingly critical. Liquid cooling systems have emerged as a particularly promising solution for high-performance computing (HPC) applications with dense accelerator embeddings due to their superior thermal conductivity, allowing for more effective heat dissipation than traditional air-cooling methods. To derive the full potential of these liquid-cooled systems, we need sophisticated control strategies with a framework for real-time decision-making and dynamic adjustments based on changing workloads and environmental conditions. Our research introduces RL-LC, a Reinforcement Learning-based approach designed to optimize the liquid cooling process in data centers. By controlling the coolant flow rate and supply liquid temperature, RL-LC creates an adaptive system that continuously optimizes energy efficiency relative to varying workloads and environmental conditions. This approach contributes to achieving energy and carbon efficiency goals by minimizing resource consumption in data centers. Our method has demonstrated significant results, achieving an average reduction of around 4% in energy consumption and carbon emissions, illustrating its potential to advance sustainable data center operations.

2 Literature Review

Liquid cooling is increasingly favored in data centers for its energy efficiency and scalability, especially in high-density environments, reducing cooling energy consumption by $29 - 50\%$, improving

[*]Corresponding author. †These authors contributed equally.

Tackling Climate Change with Machine Learning: workshop at NeurIPS 2024.

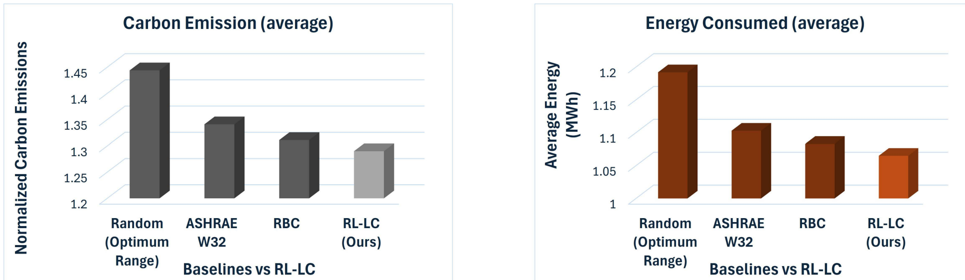

Figure 1: Comparison of Average Energy Consumption and Carbon Emissions between different baseline agents and our approach RL-LC.

Power Usage Effectiveness (PUE), and managing thermal loads as data centers grow Habibi Khalaj and Halgamuge (2017), Patterson et al. (2016), Gowdra (2017). Its potential to lower carbon footprints by enabling high heat density cabinet operations highlights the importance of accurate modeling Hernon et al. (2009). Machine learning (ML) algorithms, integrated with detailed cooling models, enhance predictive accuracy for energy consumption and thermal management Herring et al. (2022), with techniques like support vector regression and Gaussian Process Regression optimizing system performance Tang et al. (2021); Khan et al. (2022). Reinforcement learning (RL), unlike traditional control methods, handles non-linear dynamics and real-time adaptation to fluctuating workloads, optimizing long-term energy efficiency and multiple objectives like temperature control and energy consumption Roijers et al. (2013); Che et al. (2023); Humfeld et al. (2021); Sarkar et al. (2024a,b, 2023); Naug et al. (2023a,b). Early RL implementations in liquid cooling systems have shown significant energy savings, suggesting a major step towards sustainable and efficient data center operations Ran et al. (2019, 2022); Li et al. (2019). As data centers grow in scale and complexity, RL in liquid cooling systems represents a significant step toward achieving both operational efficiency and environmental sustainability.

3 System Description

In a liquid cooling system for data centers, heat is removed from servers using a cold plate with high thermal conductivity, such as copper or aluminum Nada et al. (2021). A simplified schematic of the primary liquid cooling loop with three server cabinets is presented in Figure 2. The liquid coolant, typically water or a dielectric fluid, flows through channels within the cold plate, absorbing heat from the server components. The heated liquid is then transported to a heat exchanger, where it is cooled before being recirculated. The design of the cooling system, including the flow channel configuration (e.g., serpentine or parallel), plays a crucial role in cooling efficiency by reducing thermal resistance, maintaining lower server temperatures, and minimizing pressure drop. We want to highlight here that the model has been adapted from a larger data center framework in Naug et al. (2024); Sarkar et al. (2024), which considers the default operation of the secondary chiller loops, pump operations and battery storage management as shown in Figure 3. It simulates a data center with 1 MW of computing power and 50 cabinets. Also, this design is scalable for much larger HPC data centers.

This paper focuses on modeling and optimizing the primary cooling loop, crucial for maintaining server temperatures and ensuring energy-efficient operations. RL-LC aims to enhance overall cooling efficiency, reduce energy consumption, and contribute to more sustainable data center operations by improving the liquid cooling process within this primary loop.

4 Model Formulation

We formulate the optimization problem for liquid-cooled data centers, focusing on modeling the liquid cooling system as shown in Figure 2. A differential equation model for server liquid cooling dynamics is implemented, considering two key components: the heat capacitor and the convection model. The heat capacitor C represents the server's thermal storage capacity, $C \cdot \frac{dT}{dt} = Q_{port}$ where T is the server temperature and Q_{port} is the net heat flow. The convection model represents heat transfer between the server plate and the cooling liquid: $Q_{flow} = G_c \cdot (T_{solid} - T_{fluid})$ where G_c is the effective convection thermal conductance, dependent on coolant properties and flow rate

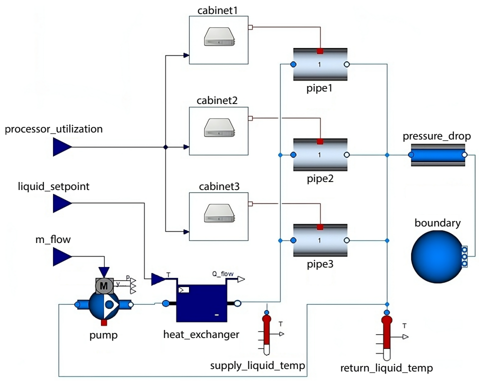

Figure 2: Schematic of the primary liquid cooling loop in a data center with three server cabinets, highlighting key components like the pump, heat exchanger, and coolant flow.

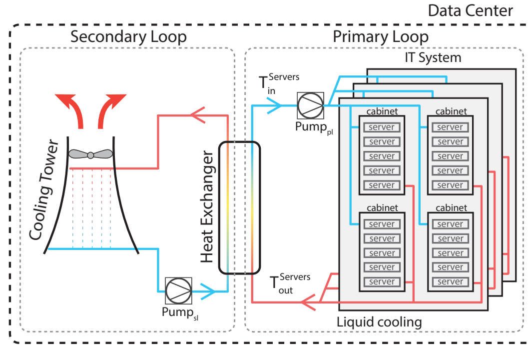

Figure 3: Representation of a data center's liquid cooling system, showing the primary and secondary loops. The focus is on optimizing the primary loop by controlling pump speed m_{flow} and liquid supply temperature ($liquid\_setpoint$).

(m_{flow}) and the temperature of the incoming fluid(T_{fluid}). The overall heat transfer problem is described by: $Q_{prescribed} = P_{util} \cdot P_{full\_load}$ and $Q_{prescribed} + Q_{flow} = C \cdot \frac{dT_{server}}{dt}$ where P_{util} is the server utilization percentage and $P_{fullload}$ is the full load power. The model simulates thermal behavior by accounting for energy storage in the server's plate and convective heat transfer. Server utilization (P_{util}) is an exogenous variable determining heat generation at different temperatures (T). The optimization goal is to minimize energy consumption by controlling the cooling liquid flow rate m_{flow} and its temperature entering the servers T_{fluid}.

5 RL Control Problem

Overall, the goal of the paper is to solve the RL-agent control-problem for liquid cooled data centers, which allows us to optimize the Cooling Energy consumption $Q_{cooling}$ by adjusting the actions A_{liq} comprising m_{flow} and T_{fluid}. Let Env_{LIQ} denote the data center liquid cooling model. $Agent_{LIQ}$ maps the state s_t to the action a_t and Env_{LIQ} represents the transition model mapping states s_t and action a_t to the next state and resultant energy consumption. Then, the sequence of operations for the problem can be represented as:

$$Agent_{LIQ} :(t_{amb} \times m_{flow\_prev} \times T_{liquid\_return} \times P_{util}) \rightarrow (m_{flow}, T_{fluid}) \tag{1}$$

$$Env_{LIQ} :(t_{amb} \times m_{flow} \times T_{liquid\_return} \times P_{util} \times T_{fluid}) \rightarrow (Q_{cooling}, Q_{it}) \tag{2}$$

where t_{amb} indicates the ambient temperature at time t, which influences cooling requirements; m_{flow_prev} represents the actual pump speed at time t, affecting the liquid flow rate in the primary loop cooling system; T_{liquid_return} is the return (mixed) liquid temperature at time t in the primary loop, reflecting the heat removed from the servers; $Q_{cooling}$ and Q_{it} represent the energy consumption of the cooling system and IT equipment at time t, respectively; $m_{flow,t}$ and $T_{fluid,t}$ denotes the action taken by the Liquid Cooling agent at time t, adjusting cooling system parameters.

We train the agent to find the optimal θ_{LIQ} that parameterizes the RL policy $Agent_{LIQ}$ which minimizes the total energy consumption over a specified horizon N. Here we choose N to be $(31 \times 24 \times 4)$ i.e. a horizon of 31 days, where we assume a step duration of 15 minutes i.e.: $\sum_{t=0}^{t=N} EnergyCons_t$ is minimized. We consider the following reward for the agent A_{LIQ}: $r(s_t, a_t) = -(Q_{cooling,t} + Q_{it,t})$. The reward was normalized using the mean and standard deviation of historical energy consumption data of the model.

6 Experimental Setup and Results

We compared our RL-based controller with three baselines. ASHRAE W32 follows industry guidelines by maintaining a fixed pump speed and a constant supply temperature of 32 °C, reflecting the standard and conservative W32 control strategy defined by ASHRAE (2021). The RBC rule-based

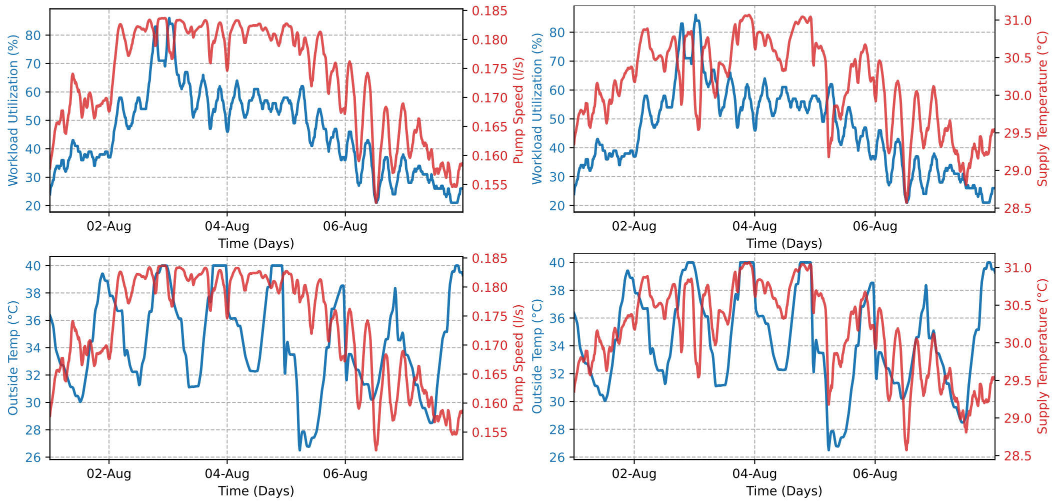

Figure 4: Pump speed (m_{flow}) and liquid supply temperature (T_{fluid}) actions in relation to workload utilization and outside temperature variations over a one-week period.

controller adjusts the pump speed in response to workload fluctuations while keeping the supply temperature fixed at 32°C. This baseline attempts to balance cooling capacity with computational demands, but lacks the dynamic optimization offered by RL-LC. We also provide a Random baseline that randomly sets pump speed and supply temperature within optimal ranges, representing a non-optimized approach to system control. Visual representations of these baselines are shown in Figure **??**.

Figure 1 and Table 1 present the performance comparison between RL-LC and the three baselines. This illustrates that our method not only reduces energy consumption and carbon emissions by approximately 4% over AHRAE on average, but also achieves a similar reduction in carbon emissions, highlighting its effectiveness in optimizing liquid cooling systems for data centers.

Controller →	Random Opt Range	ASHRAE W32	RBC	RL-LC (Ours)
Energy MWh (mean)	1.192	1.103	1.083	1.065
Energy (SD)	0.018	0.017	0.017	0.016
CO2 Emissions Norm (mean)	1.444	1.341	1.311	1.29
CO2 (SD)	0.022	0.021	0.021	0.02

Table 1: Energy and CO2 emission reduction with various controllers. Evaluation with 20 seeds.

Figure 4 demonstrates how RL-LC adapts pump speed and liquid supply temperature to varying workload utilization and outside temperature over a week. The top two graphs highlight that the pump speed (left) and the supply temperature (right) generally increase during periods of high workload, reflecting the system's response to increased cooling demands. The bottom two plots indicate that both pump speed (left) and supply temperature (right) are also influenced by fluctuations in outside temperature. Higher outside temperatures prompt an increased cooling effort to maintain server stability. In contrast, when the outside temperature is lower, the system reduces the supply temperature, allowing more efficient cooling with less energy. This adaptability highlights RL-LC's ability to maintain server stability under different environmental conditions while optimizing energy use.

7 Conclusions

This paper introduced RL-LC, a novel reinforcement learning-based method for optimizing liquid cooling in data centers, achieving a notable 4% reduction in energy consumption and carbon emissions compared to ASHRAE standards. RL-LC can scale for large HPC data centers. By dynamically adjusting pump speed and supply liquid temperature in response to workload and environmental changes, RL-LC consistently outperformed traditional control methods, demonstrating significant potential for enhancing sustainability, especially as liquid cooling becomes more prevalent beyond

HPC with the increasing use of GPU accelerators. Future work includes adapting and deploying RL-LC to supercomputing digital twins. We also plan to extend RL-LC to more complex operations, such as managing secondary cooling loops and integrating with holistic energy optimization controls. This further enhances its scalability and sustainability impact in larger data centers.

References

A. Habibi Khalaj, S. K. Halgamuge, A Review on efficient thermal management of air- and liquid-cooled data centers: From chip to the cooling system, Appl. Energy 205 (2017) 1165–1188. doi:10.1016/j.apenergy.2017.08.037.

M. K. Patterson, S. Krishnan, J. M. Walters, On energy efficiency of liquid cooled HPC datacenters, in: 2016 15th IEEE Intersociety Conference on Thermal and Thermomechanical Phenomena in Electronic Systems (ITherm), IEEE, 2016, pp. 2016–03. doi:10.1109/ITHERM.2016.7517615.

N. Gowdra, DynaCool - Simulating Efficient Liquid Cooling for Current and Next Generation Large Scale Data Centres, 2017. URL: https://openrepository.aut.ac.nz/items/d08063db-08b7-4bf5-a796-791dd5d3732c, [Online; accessed 8. Aug. 2024].

D. Hernon, T. Salamon, R. Kempers, S. Krishnan, A. Lyons, M. Hodes, P. Kolodner, J. Mullins, L. McGarry, Thermal management: Enabling enhanced functionality and reduced carbon footprint, Bell Labs Tech. J. 14 (2009) 7–19. doi:10.1002/bltj.20385.

J. Herring, P. Smith, J. Lamotte-Dawaghreh, P. Bansode, S. Saini, R. Bhandari, D. Agonafer, Machine Learning-Based Heat Sink Optimization Model for Single-Phase Immersion Cooling, ASME Digital Collection (2022). doi:10.1115/IPACK2022-97481.

X. Tang, Q. Guo, M. Li, C. Wei, Z. Pan, Y. Wang, Performance analysis on liquid-cooled battery thermal management for electric vehicles based on machine learning, J. Power Sources 494 (2021) 229727. doi:10.1016/j.jpowsour.2021.229727.

S. A. Khan, C. Eze, K. Dong, A. R. Shahid, M. S. Patil, S. Ahmad, I. Hussain, J. Zhao, Design of a new optimized U-shaped lightweight liquid-cooled battery thermal management system for electric vehicles: A machine learning approach, Int. Commun. Heat Mass Transfer 136 (2022) 106209. doi:10.1016/j.icheatmasstransfer.2022.106209.

D. M. Roijers, P. Vamplew, S. Whiteson, R. Dazeley, A survey of multi-objective sequential decision-making, Journal of Artificial Intelligence Research 48 (2013) 67–113.

G. Che, Y. Zhang, L. Tang, S. Zhao, A deep reinforcement learning based multi-objective optimization for the scheduling of oxygen production system in integrated iron and steel plants, Applied Energy 345 (2023) 121332.

K. D. Humfeld, D. Gu, G. A. Butler, K. Nelson, N. Zobeiry, A machine learning framework for real-time inverse modeling and multi-objective process optimization of composites for active manufacturing control, Composites Part B: Engineering 223 (2021) 109150.

S. Sarkar, A. Naug, R. Luna, A. Guillen, V. Gundecha, S. Ghorbanpour, S. Mousavi, D. Markovikj, A. Ramesh Babu, Carbon footprint reduction for sustainable data centers in real-time, Proceedings of the AAAI Conference on Artificial Intelligence 38 (2024a) 22322–22330. URL: https://ojs.aaai.org/index.php/AAAI/article/view/30238. doi:10.1609/aaai.v38i20.30238.

S. Sarkar, A. Naug, A. Guillen, R. Luna, V. Gundecha, A. Ramesh Babu, S. Mousavi, Sustainability of data center digital twins with reinforcement learning, Proceedings of the AAAI Conference on Artificial Intelligence 38 (2024b) 23832–23834. URL: https://ojs.aaai.org/index.php/AAAI/article/view/30580. doi:10.1609/aaai.v38i21.30580.

S. Sarkar, A. Naug, R. L. Gutierrez, A. Guillen, V. Gundecha, A. Ramesh Babu, C. Bash, Real-time carbon footprint minimization in sustainable data centers with reinforcement learning, in: NeurIPS 2023 Workshop on Tackling Climate Change with Machine Learning, 2023.

A. Naug, A. Guillen, R. Luna Gutiérrez, V. Gundecha, S. Ghorbanpour, L. Dheeraj Kashyap, D. Markovikj, L. Krause, S. Mousavi, A. R. Babu, S. Sarkar, 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, BuildSys '23, Association for Computing Machinery, New York, NY, USA, 2023a, p. 232–235. URL: https://doi.org/10.1145/3600100.3623732. doi:10.1145/3600100.3623732.

A. Naug, A. Guillen, R. Luna Gutierrez, V. Gundecha, S. Ghorbanpour, S. Mousavi, A. Ramesh Babu, S. Sarkar, A configurable pythonic data center model for sustainable cooling and ml integration, in: NeurIPS 2023 Workshop on Tackling Climate Change with Machine Learning, 2023b.

Y. Ran, H. Hu, X. Zhou, Y. Wen, Deepee: Joint optimization of job scheduling and cooling control for data center energy efficiency using deep reinforcement learning, in: 2019 IEEE 39th International Conference on Distributed Computing Systems (ICDCS), IEEE, 2019, pp. 645–655.

Y. Ran, X. Zhou, H. Hu, Y. Wen, Optimizing data center energy efficiency via event-driven deep reinforcement learning, IEEE Transactions on Services Computing 16 (2022) 1296–1309.

Y. Li, Y. Wen, D. Tao, K. Guan, Transforming cooling optimization for green data center via deep reinforcement learning, IEEE transactions on cybernetics 50 (2019) 2002–2013.

S. Nada, R. El-Zoheiry, M. Elsharnoby, O. Osman, Experimental investigation of hydrothermal characteristics of data center servers' liquid cooling system for different flow configurations and geometric conditions, Case Studies in Thermal Engineering 27 (2021) 101276.

A. Naug, A. Guillen, R. Luna, V. Gundecha, D. Rengarajan, S. Ghorbanpour, S. Mousavi, A. R. Babu, D. Markovikj, L. D. Kashyap, S. Sarkar, Sustaindc: Benchmarking for sustainable data center control, 2024. URL: https://arxiv.org/abs/2408.07841. arXiv:2408.07841.

S. Sarkar, A. Naug, R. Luna, A. Guillen, V. Gundecha, S. Ghorbanpour, S. Mousavi, D. Markovikj, A. R. Babu, Carbon footprint reduction for sustainable data centers in real-time, in: Proceedings of the AAAI Conference on Artificial Intelligence, volume 38, 2024, pp. 22322–22330.

ASHRAE, Emergence and Expansion of Liquid Cooling in Mainstream Data Centers, Technical Report, ASHRAE, Peachtree Corners, GA, 2021. URL: https://www.ashrae.org/file%20library/technical%20resources/bookstore/emergence-and-expansion-of-liquid-cooling-in-mainstream-data-centers_wp.pdf, white Paper.