Simulator-Based Reinforcement Learning for Data Center Cooling Optimization

Chi Zhou
chizhou@meta.com
Meta Platforms

Doris Gao
dorisgao@meta.com
Meta Platforms

Lisa Rivalin
lisarivalin@meta.com
Meta Platforms

Andrew Grier
agrier@meta.com
Meta Platforms

Gerson Arteaga Ramirez
garteaga@meta.com
Meta Platforms

John Fabian
jfabian@meta.com
Meta Platforms

Abstract

Data centers power the internet, making digital communication and connection possible. In 2022, 1-1.3% of the total energy consumption in the world went to data centers. They also consume a lot of water in their cooling systems. Optimizing their energy and water usage is a priority in the industry. This paper presents one of our approaches to optimize energy and water consumption in data center cooling by leveraging a Simulator Based Reinforcement Learning method. First we developed a physics-based simulation model that can predict the thermal behavior within 1°F of MAE (mean absolute error) for cold aisles. Then an RL model is trained offline resulting in a better policy for controlling the supply airflow setpoint. The production model at one of our data center regions has shown reduction of supply fan energy consumption by 20% and water usage by 4% on average across various weather conditions.

1 Introduction

Efficiency is one of the key components of designing, building and operating sustainable data centers. Besides the IT load, cooling is the primary consumer of energy and water in the data center environment. By improving the cooling efficiency, it will help reduce our energy use, water use and greenhouse gas (GHG) emissions and address one of the biggest challenges of all - climate change.

One type of efficient data center design is the penthouse cooling system, which uses outdoor air and evaporative cooling systems to maintain environmental conditions within the envelope of temperature between 65°F and 85°F (18°C and 30°C) and relative humidity between 13 and 80 percent. As water and energy are consumed in the conditioning of this air, optimizing the amount of supply airflow that has to be conditioned is a high priority in terms of improving operational efficiency.

The idea of using AI for large scale data center cooling optimization is not new. Others have started their trial with neural network models back in 2016 (Evans & Gao, 2016). There are also various Reinforcement Learning approaches reported (Li et al., 2020; Lazic et al., 2018; Luo et al., 2022). However, applying the control policy determined by an online RL model may result in various risks including breaches of service requirements and even thermal unsafety. To address this challenge, in this paper we detail our experience using an offline simulator based RL approach to optimize the amount of airflow supplied into data centers for cooling purposes. As a result, at one of the pilot regions we have reduced the supply fan energy consumption by 20% and water usage by 4% on average across various weather conditions.

2 Data Center Penthouse Cooling System

Many of our data centers adopt a two-tiered penthouse design that utilizes 100% outside air for cooling. As shown in Fig. 1, the air enters the facility through louvers on the second-floor "penthouse," with modulating dampers regulating the volume of outside air. The air passes through a mixing room, where outdoor air if too cold can be mixed with heat from server exhaust when needed to regulate the temperature. The air then passes through a series of air filters and a misting chamber where the Evaporative Cooling and Humidification (ECH) system is used to further control the temperature and humidity. The air continues through a fan wall that pushes the air through openings in the floor that serve as an air shaft leading into the server area on the first-floor. The hot air coming out from the server exhaust will be contained in the hot aisle, through exhaust shafts, then recycled back to the economizer or released out of the building with the help of relief fans. Water is mainly used in two ways: evaporative cooling and humidification. The evaporative cooling system converts water into vapor to lower the temperature when the outside air is too hot, while the humidification process maintains the humidity level if the air is too dry.

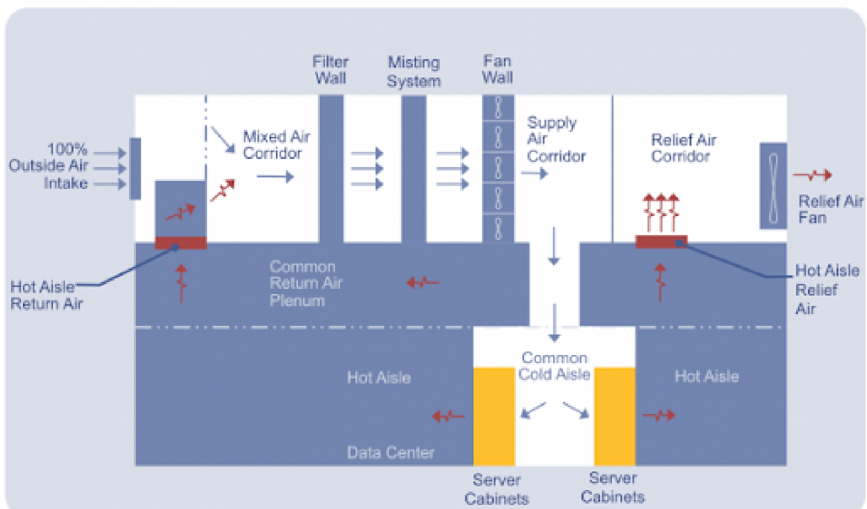

Figure 1: Penthouse Cooling System Design

In order to supply air within the defined operating envelope, the penthouse relies on the Building Management System (BMS) to monitor and control different components of the mechanical system to perform the task of conditioning the intake air from outside by mixing, humidifying/dehumidifying, evaporative cooling, or a combination of these operations. There are three major control loops responsible for adjusting setpoints for supply air temperature, humidity, and airflow respectively. The airflow setpoint is typically calculated based on a small set of input variables like current IT load, cold aisle temperature and differential pressure between cold aisle and hot aisle, and the logic is often very simple at a linear scale, but becomes very difficult to accurately model as these values at different locations in the datacenter are coupled to one another and highly dependent on complex local boundary conditions. However the amount of airflow will largely dictate the energy used by the supply fan arrays and water consumption when cooling or humidification is required. Therefore, optimizing the airflow setpoint would have the biggest impact on further improving the cooling efficiency given the fact that the temperature and humidity boundary of the operating envelope is fixed.

3 Methodology

3.1 Physical Simulator

To quantify and visualize the fluid dynamics of data centers, companies typically rely on detailed physics-based models called computational fluid dynamics (CFD) simulations. We can study previously unseen conditions and designs with these models, but they demand a significant amount of computational power, and take a long time to run. Alternatively using measurements from the many sensors in data centers, we can create numerical models that predict the energy consumption and thermal response of the building. These models, although highly accurate, may not generalize well about situations outside the normal operating conditions reflected in the training data.

Recently gray box models open a realm of new opportunities beyond data science approaches and CFD. They solve thermal balance equations to describe the conditions in an individual room or even row in a facility, producing results quickly and with little input data. Our simulator adopts this approach which combines first-principle physics with building-modeling language Modelica (Fritzson, 2010). The Modelica Buildings library is an open-source library with component and system models for building energy and control systems. It is also accompanied by Python modules that can be used to automate simulations and post-processing of simulation results, which is crucial to large scale model productionisation.

To simulate a particular data center building, we first need static, site-specific details about the facility's geometry, construction materials, and HVAC (heating, ventilation and air conditioning), as well as its system configurations and component efficiencies. Then we numerically re-create the control strategies (using the Control Description Language) governing the behavior of all the HVAC and water equipment as functions of the indoor and outdoor conditions.

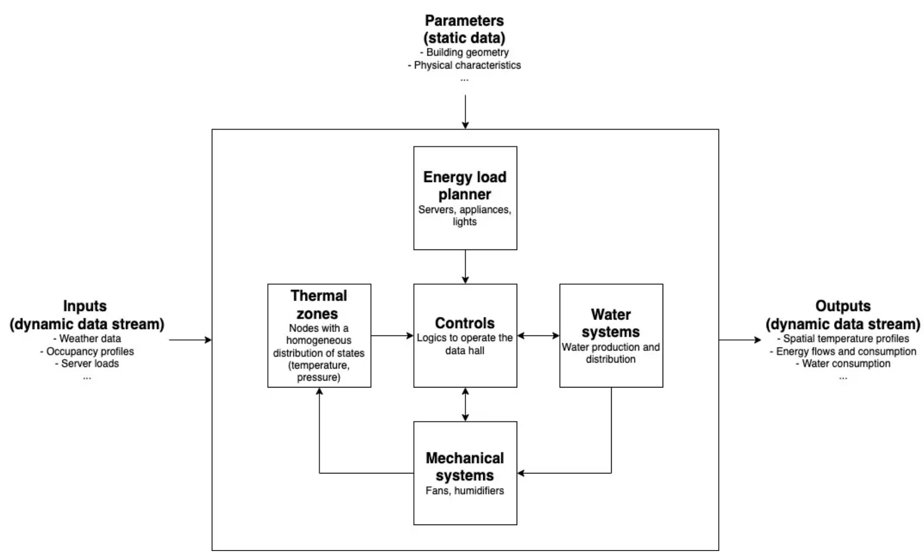

Figure 2: Components of Data Center Physical Model Simulator

Once those tasks are complete, we can begin to input data such as weather data, IT load, and setpoint schedules. It uses differential equations to output the dynamic system response, such as the thermal load and resulting energy use, along with related metrics like cold aisle temperature and differential pressure profiles. Fig. 2 illustrates the components of our physical model simulator (Rivalin et al., 2022).

Using the Modelica simulation in the data center environment is not new either, for example, a previous study has shown potential energy savings by resolving control-related issues and optimizing the supply air temperature (Fu et al., 2019). However we found out that one important source of parasitic heat in data centers is the hot-air recirculation, flowing back from the hot to the cold aisle at the rack level and having multiple negative consequences, such as the creation of hot spots damaging the hardware and increasing the data center cooling needs. To alleviate this phenomenon, the ability to predictively model conditions and events, even those not seen before, in data centers is increasingly important. The main novelty of our approach is that we have identified three main reasons leading to recirculation: baseline recirculation characterized by design Hot Aisle Containment and racks leakage, recirculation caused by negative differential pressure locally at the rack, and flow deficit at the end of an aisle when there is not enough supply air to cool the racks. Our model can quantify the three types of recirculation airflow and predict the thermal behavior within 1°F of MAE (mean absolute error) for cold aisles in normal conditions. It achieves more accurate predictions than a data-science-only-based model in under-provisioned conditions. Please refer to (Rivalin et al., 2023) for more technical details if the reader is interested.

3.2 RL Approach

Reinforcement Learning (RL) is good at modeling control systems as sequential state machines. It functions as a software agent that determines what action to take at each state based on some transition model, which leads to a different state, and constantly gets feedback from the environment in terms of reward. In the end, the agent would learn the best policy model (typically parameterized by a deep neural network) to achieve the optimal accumulated reward. The data center cooling control can be naturally modeled under this paradigm.

At any given time, the state of a data center can be represented by a set of environmental variables monitored by many different sensors for outside air, supply air, cold aisle and hot aisle plus IT load, i.e., power consumption by servers, etc. The action is to control setpoints, for example, the supply airflow setpoint determining how fast the supply fans run to meet the demand. The policy is the function mapping from the state space to action space, i.e., determining the appropriate airflow setpoint based on current state conditions. Now the task is to leverage historical data we have collected from thousands of sensors in our data centers, augmented with simulated data of potential, but not yet experienced conditions, and train a better policy model that gives us better reward in terms of energy or water usage efficiency.

As illustrated in Fig. 3, our RL agent operates in a simulated environment by starting from real-life historical observations S, exploring the action space, and feeding into the simulator to predict the anticipated new state S' and reward R given each sampled action A, and collect the (S, A) pairs having the best reward to form a new training data set to update the parameterized policy model.

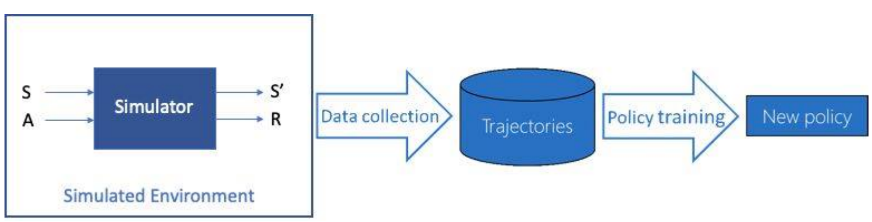

Figure 3: Simulated based Offline Reinforcement Learning Approach

The simulator plays a very important role here as our goal is to optimize energy or water usage while maintaining the data center condition under specs so hardware performance won't be affected. More specifically we want to keep the cold aisle temperature rise below a certain threshold or a positive pressurization from cold aisle to hot aisle to minimize the parasitic heat caused by recirculation. Additionally, the physics-based simulator enables us to train the reinforcement learning model with all possible scenarios, not only those present in the historical data. This increases reliability during outlier events as well as allows for rapid deployment in newly commissioned data centers.

4 Results

We started a pilot at one of our data center regions in 2021 having the RL model directly controlling the supply airflow setpoint. Fig. 4 shows a comparison of the new setpoint, in the unit of CFM (cubic feet per minute) as the red line to the original BMS setpoint (as the dotted blue line) over one week's duration for illustration purposes (The Y-axis display of setpoint values has been intentionally obscured for reasons of confidentiality). The fluctuation is mainly determined by the supply air temperature and server load cycles at different times of day. More importantly as shown in Fig. 5, the data center temperature conditions never went out of spec with reduced airflow supply with respect to both cold aisle average and maximum temperature compared against the supply air temperature.

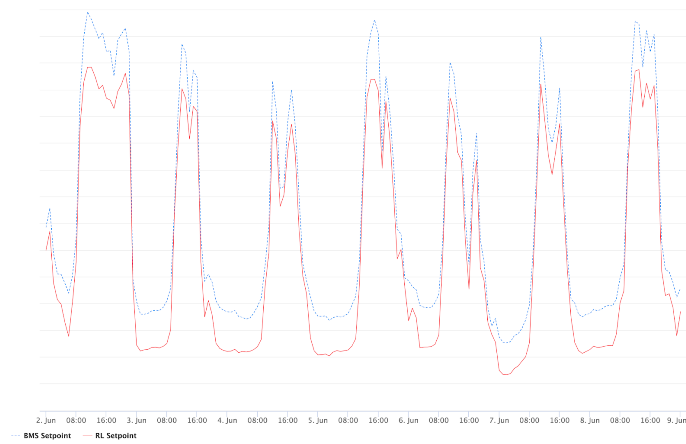

Figure 4: Comparison of the RL model versus original BMS setpoint

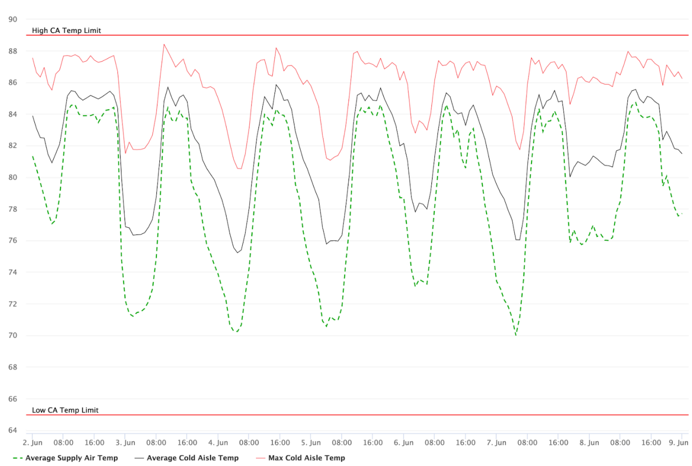

Figure 5: Data Center Temperature Profile Under RL Model Control

It is noticeable that the CFM savings vary under different SA temperatures as the univariate chart in Fig. 6 has shown. The CFM savings can easily be converted to energy savings used by the supply fans. Under hot and dry conditions when evaporative cooling or humidification is required, using less air will result in less water usage as well. Over the past couple years of pilot, on average across various weather conditions, we are able to reduce the supply fan energy consumption by 20% and water usage by 4%.

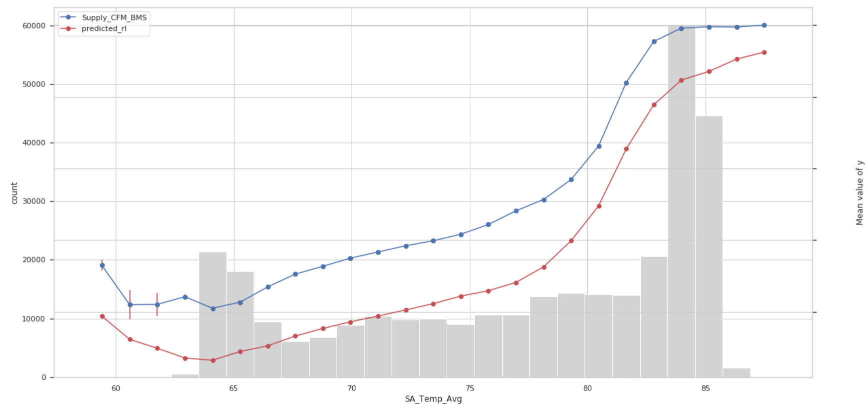

Figure 6: Airflow Savings Breakdown at Different Supply Air Temperature

5 Future Work

This effort opens the door to scale our data center facilities operations and transform how we operate by introducing automated predictions and continuous optimization to the tuning of environmental conditions, while bending the cost curve and reducing effort on labor intensive tasks. We are in the process of rolling out the program to our current data center fleet in the next couple of years, which would achieve significant energy and water usage savings and contribute to the long term sustainability goals.

Also as we are revamping the new data center design to optimize for artificial intelligence, we plan to apply the same methodology for future DC optimization at the design phase so our new type data centers are commissioned with the capability from day one of its operation.

References

Richard Evans and Jim Gao. Deepmind ai reduces energy used for cool- ing google data centers by 40%, 2016. URL https://deepmind.com/blog/article/deepmind-ai-reduces-google-data-centre-cooling-bill-40.

Peter Fritzson. *Principles of object-oriented modeling and simulation with Modelica 2.1*. John Wiley & Sons, 2010.

Yangyang Fu, Wangda Zuo, Michael Wetter, James W. VanGilder, and Peilin Yang. Equation-based object-oriented modeling and simulation of data center cooling systems. *Energy and Buildings*, 198:503–519, 2019. ISSN 0378-7788.

Nevena Lazic, Tyler Lu, Craig Boutilier, MK Ryu, Eehern Jay Wong, Binz Roy, and Greg Imwalle. Data center cooling using model-predictive control. In *Proceedings of the Thirty-second Conference on Neural Information Processing Systems (NeurIPS-18)*, pp. 3818–3827, Montreal, QC, 2018.

Yuanlong Li, Yonggang Wen, Dacheng Tao, and Kyle Guan. Transforming cooling optimization for green data center via deep reinforcement learning. *IEEE Transactions on Cybernetics*, 50(5): 2002–2013, 2020.

Jerry Luo, Cosmin Paduraru, Octavian Voicu, Yuri Chervonyi, Scott Munns, Jerry Li, Crystal Qian, Praneet Dutta, Jared Quincy Davis, Ningjia Wu, et al. Controlling commercial cooling systems using reinforcement learning, 2022.

Lisa Rivalin, Andrew Grier, Chi Zhou, and Doris Gao. How thermal simulation helps optimize meta's data centers, 2022. URL https://engineering.fb.com/2022/09/14/data-center-engineering/data-centers-meta-thermal-simulation-optimization/.

Lisa Rivalin, Andrew Grier, Tobias Tiecke, Chi Zhou, Doris Gao, Prakriti Choudhury, and John Fabian. Predicting temperature and differential pressure in data centers using physical modeling. In *Indoor Air Quality, Ventilation and Energy Conservation in Buildings (IAQVEC)*, 2023.