Enhancing Energy Efficiency in Cooling Systems with Physics-Informed Reinforcement Learning

Introduction

The consumption of energy in buildings is very critical. Buildings account for 40% of the total energy consumption and HVAC represents roughly 40% of a building's consumption. My master's thesis explores an innovative approach to optimize the energy efficiency of these cooling systems through Physics-Informed Reinforcement Learning (PIRL). This method integrates physical principles with advanced machine learning techniques to achieve significant energy savings. Experiments have been performed on a simulated model from a data center in Florida.

Main Text

I worked on the development and implementation of a PIRL-based controller for a cooling system from a data center. The results were promising, as the PIRL controller may be able to save energy by around 3% compared to the baseline controller. This improvement is achieved by optimizing the operational parameters of the cooling system. However, the optimal parameters could vary over time, and the main challenge is to develop a controller that adapts to these variations and learns new control laws online.

The main problem addressed was the high energy consumption of data center cooling systems. Traditional controllers follow simple predefined rules to control the cooling equipment. PIRL can find optimal control rules by trial and error with a simulated model of the plant, which lead to less energy consumption.

These optimal laws are usually nontrivial. For example, it was found that when the cooling demands are higher, it is better to lower the temperature in the chilled water circuit. This is somehow a surprising result, since one would think that this is inefficient because the chiller must cool the water to a lower temperature. However, if this temperature is too high, the water flow rate needs to be higher to maintain the heat transfer. Since the power consumption of the pumps has a quadratic relation with the flow rate, its power consumption grows quick for large flow rates, increasing the aggregated power consumption of all the equipment.

The need to reduce energy consumption in buildings is essential from both an economic and an environmental point of view. Data centers are growing rapidly, and their energy demands are becoming a significant part of global energy consumption. Optimizing their cooling systems can lead to substantial cost savings and a reduction in carbon footprint, contributing to more sustainable and efficient data center operations. Moreover, having a controller that "learns" continuously would reduce the need for controller retuning, which would save costs for the company.

My thesis proposes a method to design an RL controller for the cooling system that can be trained in a simulated environment and then deployed to a cooling plant. Then it also suggests future research paths towards crafting a controller that can "learn" continuously on the real plant, being continuously adapted to variations in the plant or equipment aging.