Reinforcement Learning for Building Control and its Real-World Implementation

Multi-Objective Deep Reinforcement Learning for Grid-Interactive Energy-Efficient Buildings (MODRLC)

Principal Investigator: Andrey Bernstein

PI Email: Andrey.Bernstein@nrel.gov

Additional Project Contributors: NREL: Xiangyu Zhang, Rohit Chintala, Yue Chen, Xin Jin, and Peter Graf

CU Boulder: Gregor Henze, Emiliano Dall'Anese, Sourav Dey, Thibault Marzullo, and Ana M. Ospina

Industrial Partner: QCoefficient

BACKGROUND

- Building HVAC control offers vast untapped potential for:
  - Buildings to provide grid services such as demand response (DR) and bulk power signal tracking.
  - Reducing on-site energy use intensity.
- Traditional controllers such as model predictive control (MPC) can be challenging in mass deployment because it requires:
  - Carefully developed, building-specific models.
  - On-demand computation to implement real-time control.
- Deep reinforcement learning (DRL) can be trained offline, is adaptive to different building types, and has very limited online computational requirement, making it ideal for grid-interactive efficient building (GEB) technologies.

PROJECT OVERVIEW / OBJECTIVES

- Algorithm/methodology/simulation platform development for training multi-objective DRL (MODRL) controllers in building applications.
- Implement MODRL controller on a real building in New York City and achieve the following objectives:
  - Save 20% on the utility bills with time of use (TOU) pricing, when compared to a default baseline.
  - Reduce peak demand by 20%.
  - Perform within 10% of an MPC trained specifically on the building.

APPROACH

- Develop and test Deep Reinforcement Learning policy using a continuous action space RL formulation and compare performance with MPC in 3 incremental phases :
  - Reduced 3-resistor-2-capacitance model where the building thermodynamics are abstracted into an equivalent electrical circuit .
  - EnergyPlus™ model calibrated on the real building.
  - Real building in New York having 40 floors, and equipped with 4 electric chillers and 8 cooling towers.

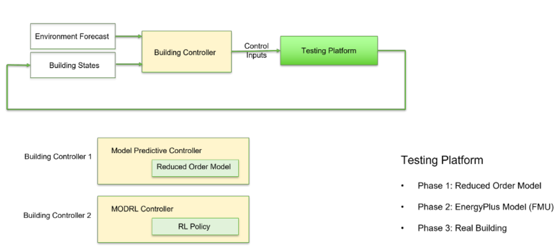

Line diagram of DRL and MPC implementation in incremental phases

FUTURE WORK

- Implement Reinforcement Learning controller on the real New York high-rise building.

Research Outcomes

We developed a Deep Reinforcement Learning controller that serves grid objectives, lowers energy costs, maintains occupant comfort in unseen testing scenarios, and performs comparably to a baseline MPC for a 5-zone EnergyPlus model.

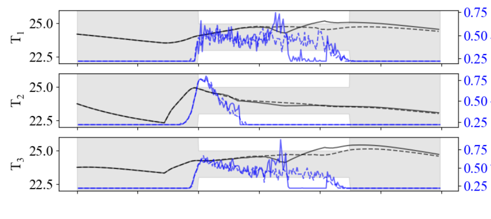

Indoor air temperatures and mass flow rates($\dot{m}_i$) for 3 of the 5 zones. Shaded areas show the temperature comfort band.

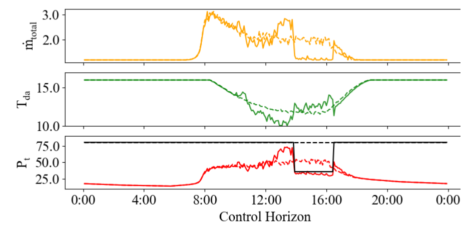

Total mass flow rate discharge air temperature and building HVAC power. In all plots, dashed lines represent the non-demand response (DR) day profiles and solid ones show those of the DR days.

Our team trained the DRL for real-world implementation on a high-rise commercial building:

- We designed the multi-objective function to reduce peak-demand charge, TOU electricity charges, and carbon emissions while strictly adhering to constraints on indoor air temperatures and chiller operation using a 24-hr-ahead pricing signal and weather forecast.
- We performed extensive testing on a high-fidelity simulation platform calibrated using data from the real building.

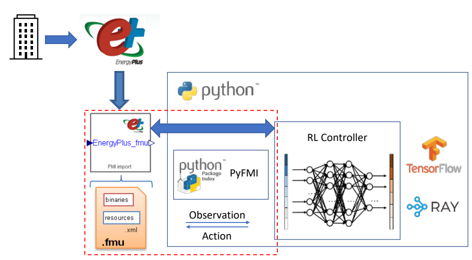

Line diagram of DRL implementation on calibrated EnergyPlus commercial building model.

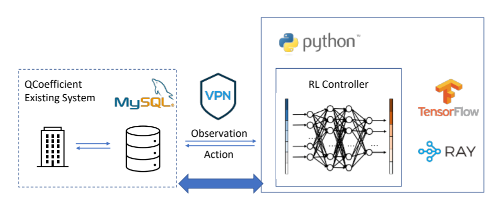

Line diagram of DRL implementation on real commercial building.

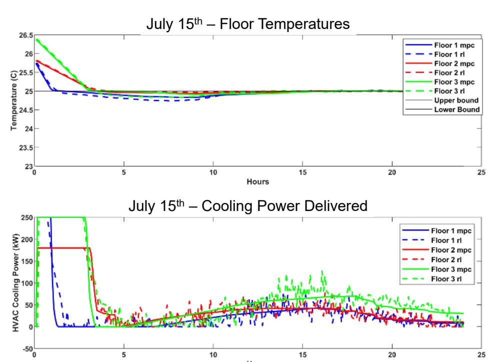

DRL implementation on 3-Resistance-2-Capacitance model on a warm summer day.

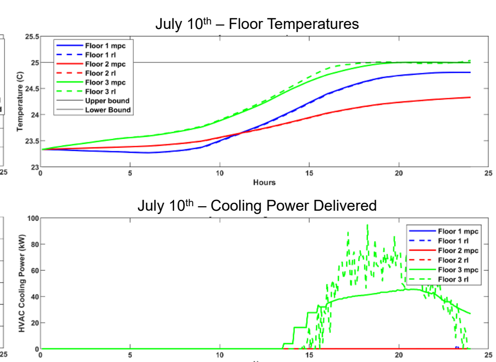

DRL implementation on 3-Resistance-2-Capacitance model with different initial conditions

Impact

- **Large-scale implementation** because Reinforcement Learning has the potential to bypass the need for building-specific models.
- **Lower the utility bills of building occupants** while also maintaining occupant comfort and providing grid services.
- **Equip buildings to provide demand side management** grid services.

This work was authored [in part] by the National Renewable Energy Laboratory, operated by Alliance for Sustainable Energy, LLC, for the U.S. Department of Energy (DOE) under Contract No. DE-AC36-08GO28308. Funding provided by the U.S. Department of Energy Office of Energy Efficiency and Renewable Energy Building Technologies Office. The views expressed in the article do not necessarily represent the views of the DOE or the U.S. Government. The U.S. Government retains and the publisher, by accepting the article for publication, acknowledges that the U.S. Government retains a nonexclusive, paid-up, irrevocable, worldwide license to publish or reproduce the published form of this work, or allow others to do so, for U.S. Government purposes.

BTO Peer Review 2023
Arlington, VA
April 24-28, 2023
NREL/PO-5500-85670