Model Predictive Control of a Hybrid Rooftop Unit System Integrated with Phase Change Material for Load Shifting

Rohit Chintala[1*], Huang Ransisi[1], Xin Jin [1]

[1] National Renewable Energy Laboratory
Golden, Colorado, USA
rohit.chintala@nrel.gov

* Corresponding Author

ABSTRACT

Thermal energy storage (TES) offers a promising solution for load shifting and reducing peak demand in buildings. The TES system under consideration in this paper is a hybrid rooftop unit (RTU) from Emerson with 5-ton cooling capacity, integrated with a patented phase change material (PCM) from NETEnergy that functions as TES. The hybrid RTU system with PCM is currently in pre-commercialization stage, and offers important advantages over equivalent ice storage systems, namely significant reduction in added weight and cost, and a lot fewer moving parts and machinery. The hybrid RTU system with PCM has the potential to reduce peak demand by up to 40% and leverage time-of-use (TOU) utility rates to reduce operational costs for the customer.

The default controller used to regulate the operation of TES is rule-based control (RBC). A set of rules govern the charging and dishcarging of the PCM. The utility of the hybrid RTU + PCM can be further improved by replacing rule-based approach of RBC with an optimization-based approach of model predictive control (MPC). This paper provides the algorithm development, analysis, and results of implementing model predictive control (MPC) on the hybrid RTU with PCM using high-fidelity physics-based models as the simulation test-bed. The goal of the paper is to show how replacing RBC with MPC can provide greater benefits with regards to load shifting and peak-demand reduction. Reduced order models (ROMs) are developed for each component of the hybrid RTU and PCM, and mathematical optimization is performed to find the optimal charging/discharging trajectory of the PCM. MPC is then tested on high-fidelity physics-based models of the RTU and PCM components, and compared with an RBC strategy in a co-simulation environment. A high demand and low demand use case are tested under a TOU rate structure, and MPC performs better than RBC when the demand is low, yielding 4.5% savings.

1. INTRODUCTION

One of the main challenges the electric grid is facing in the era of rapid electrification and decarbonization is to manage the balance between supply from the power generation resources and demand from electric power consumers. Historically this balance was achieved through the use of operating reserves. The increased adoption of renewable energy resources such as wind, solar, hydroelectricity etc. into the power mix, however, has resulted in greater stress on the grid. The intermittent nature of renewable energy resources makes the prediction of power generation difficult, thereby making it more challenging to achieve this balance (Li, Wang, Hong, & Piette, 2021). A more dynamic way of achieving power balance from both the supply and demand side is required (Li et al., 2021).

Buildings—which consume 75% of the total electricity and 80% of the total demand in the United States (Center, 2020)—are still an untapped resource that can facilitate load balancing from the demand side by leveraging flexible loads. Load flexibility can be defined as a temporary (over the course of a few hours or days) change in the energy demand of the building to provide assistance to the grid without compromising the comfort requirements of the occupants, or the technical capabilities of the building (Le Dréau et al., 2023). The scheduling of building equipment such as lighting, miscellaneous electric loads, and heating ventilation and air-conditioning (HVAC) can be altered to provide load flexibility while still satisfying occupants' requirements (Beccali et al., 2020; Chen et al., 2019; Powells, Bulkeley, Bell, & Judson, 2014).

Among major building equipment, the load flexibility that can be potentially provided by an HVAC system is the most significant. The load flexibility can be further enhanced by incorporating thermal energy storage (TES) systems. A

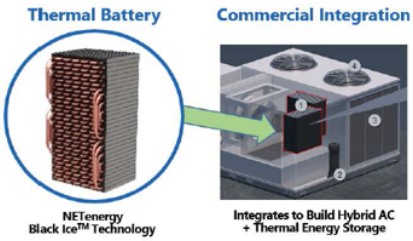

Figure 1: Hybrid RTU integrated with a PCM that serves as a thermal battery.

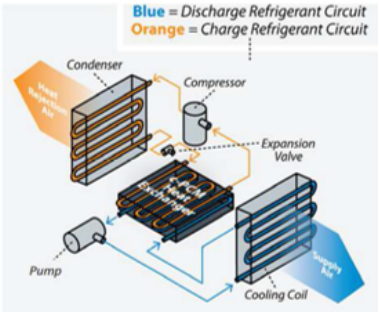

Figure 2: Charging and discharging circuits of the hybrid RTU.

TES system is able provide load flexibility by shifting a portion of the heating or cooling electric demand to a period in the day when the electricity prices are lower under a time-of-use (TOU) rate structure or by lowering the peak electric demand. A summary of the state of the art of thermal storage is provided in (Arteconi, Hewitt, & Polonara, 2012). In this paper we focus on a 5-ton hybrid rooftop unit (RTU) from Emerson that has been incorporated with a patented phase change material (PCM) from NETEnergy. The PCM serves as a TES system that stores the cooling energy from the hybrid RTU and discharges it to the conditioned space when required.

The energy flexibility offered by the TES system can be exploited by activating a control strategy that seeks to minimize the peak demand or utility price while maintaining occupancy comfort. The most commonly used strategy is referred to as rule-based control (RBC), where a certain set of rules are developed from intuition or expert knowledge to dictate the charging and discharging strategy of the TES. An alternative to the RBC strategy, namely model-predictive control (MPC), is explored in this paper. MPC leverages mathematical models to compute optimal setpoints for the hybrid RTU to lower the utility costs of the consumer, while maintaining the cooling load requirements.

In this paper a co-simulation platform is developed using a high-fidelity physics-based and neural network model of the vapor compression cycle and the PCM of the hybrid RTU. The high-fidelity model was developed in a prior project. In this paper, reduced order models (ROMs) are developed for each component of the hybrid RTU to design and evaluate the performance of MPC and compare it to an RBC strategy. The paper lays down the steps for developing the component ROMs, steps taken to integrate it with the high-fidelity model in a cosimulation environment, and designing an MPC to compute optimal PCM setpoints under a TOU rate structure. Comparison is provided of the RBC and MPC strategies for two use cases: 1) high cooling load and 2) low cooling load. Although peak demand reduction from MPC can reduce demand charge and thereby lower utility bills, the focus of our paper is to demonstrate the load shifting capability and the energy charge reduction benefit for the MPC with a hybrid RTU system.

2. INTEGRATED HYBRID RTU AND PHASE CHANGE MATERIAL

In this paper we focus on a hybrid RTU system, which has been integrated with PCM as shown in Figure 1. The PCM serves as a thermal battery that is charged through a vapor compression circuit with R410-A as the refrigerant, and discharged through an airstream that is coupled to a water-glycol circuit. The charging and discharging circuits of the PCM are shown in Figure 2. When there is a cooling requirement in the conditioned space, the PCM is discharged by pumping air through the heat exchanger in the glycol circuit. The mass flow rate of the conditioned air is governed by

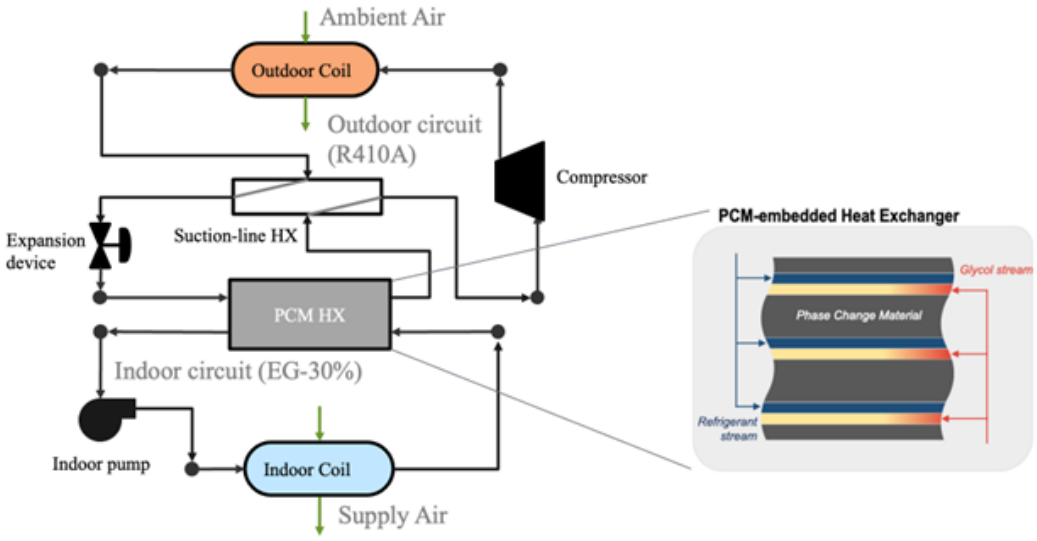

Figure 3: System schematic of the TES-integrated rooftop unit.

a local controller to meet the cooling demands of the space.

2.1 Physics-based models of vapor compression systems

A high-fidelity model of the hybrid RTU with physics-based modeling for the vapor compression system, and neural-network for PCM serves as the simulation test-bed in this paper. Figure 3 shows the schematic of the TES-integrated rooftop unit. It consists of a vapor compression refrigerant circuit (R410A) and a glycol circuit (30% ethylene glycol solution). A suction line heat exchanger is used in the refrigerant circuit to ensure sufficient compressor suction superheat. The two circuits are thermally coupled by a dual-circuit PCM embedded heat exchanger. The low pressure, two-phase refrigerant evaporates in the device and thermally charges TES (solidifying the PCM). The warm glycol solution returning from the indoor coil discharges TES (melting the PCM). As it is cooled down by the PCM, the glycol solution is pumped to the indoor coil and provides space cooling to the building. The dual circuit configuration allows the system to control the TES discharge rate and charge rate independently.

The PCM embedded heat exchanger (HX) (Figure 3, right) contains the refrigerant stream, the glycol stream, and a composite material. The refrigerant and the glycol solution flow through multi-port extruded aluminum channels. These channels are stacked directly on top of each other to ensure good thermal contact. Each channel pair is sandwiched between the composite material that contains a graphite matrix impregnated PCM. Table 1 summarizes the key parameters of the PCM used in the analysis. We define the fully charged state (state of charge, or SOC, = 100%) as an enthalpy state where the PCM is isothermal at 0°C and fully discharged (SOC = 0%) as isothermal at 10°C.

Table 1: Key properties of the phase change material.

Material Property	Value
Total energy capacity (kWh)	17.6
PCM Tt($.^{\circ}C$)	4.6
Latent heat ($\frac{kJ}{kg}$)	168
Density ($\frac{kg}{m^3}$)	836

The system is simulated via a quasi-steady state, component-based framework that primarily consists of two parts: the component models and the system solver. We formulated a system of non-linear residual equations with the smallest possible set of equations to describe the steady state of the system. The formulation method is adapted from the tripartite graph method (Huang, Aute, & Ling, 2021). The system solver executes the components in the flow directions of the working fluids to obtain the residual values, and used quasi-Newton method to iteratively solve the thermodynamic states of the refrigerant and the glycol solution. The following assumptions are adopted for the system

simulation.

- Expansion device
- Pressure drop neglected
- Fan power and fan power neglected.

Each component is a module on its own represented by a set of engineering equations or a performance map to predict the component steady-state performance under given model inputs. The compressor is represented with a 20-coefficient performance map. The outdoor coil is modeled with a three-zone approach (Huang, Ling, & Aute, 2020) (also referred to as moving boundary approach), in which the heat exchanger is divided into two or three zones, and refrigerant heat transfer in each zone is characterized by a lumped heat transfer coefficient. The indoor coil is modeled with an effectiveness-NTU approach. The suction line heat exchanger is also modeled with the moving boundary approach.

Finally, we used a surrogate model to represent the PCM HX embedded HX. It predicts the refrigerant pressure and glycol solution temperatures required to deliver a specific discharge and charge rate. The surrogate model is trained by MATLAB Deep Learning Toolbox using a detailed 2D transient finite-difference model. The detailed PCM HX model is structured similar to the ones by (Woods et al., 2021) and (Huang, Mahvi, James, Kozubal, & Woods, 2024). It discretizes the device into three types of nodes containing the PCM composite, working fluids, and aluminum channel walls. The conduction heat transfer rates between PCM composite nodes are calculated with a discretized version of Fourier's law, accounting for directional differences in the thermal conductivity of the PCM composite. Heat transfer rates between different types of nodes are calculated by considering all relevant resistances, including fluid convection, contact, and conduction resistances. The model iteratively solves for the refrigerant pressure in each timestep so that the evaporating stream exits as a saturated vapor.

The surrogate model approximates the refrigerant pressure and glycol solution outlet and inlet temperatures given 1) refrigerant and glycol mass flow rates, 2) glycol solution discharge heat rate (equivalent to the required cooling load), and 3) PCM state of charge (at the current timestep). The surrogate model then uses the approximation outputs to update the PCM state of charge to be used for the next timestep as shown in equation 1.

$$
\begin{aligned}
r_{soc}^{i+1} &= \frac{(r_{soc}^{i} \cdot C_{tes} + \dot{Q_{net}} \cdot \Delta \cdot t)}{C_{tes}} \\
\dot{Q_{net}} &= \frac{(h_{o,R410} - h_{in,R410}) \cdot \dot{m}_{R410}}{\dot{Q}_{chargedtoTES}} - \frac{(T_{in,EG} - T_{o,EG}) \cdot c_p \cdot \dot{m}_{EG}}{\dot{Q}_{dischargedfromTES}}
\end{aligned}
\tag{1}
$$

2.2 Reduced order models

MPC is a strategy in which an optimal trajectory of the control variables is computed over a future time period known as the control horizon. Only the first input of the control trajectory is implemented, and the optimization is recomputed at the next time step. This allows the control system to re-calibrate at each timestep owing to improper forecast of disturbance inputs and discrepancies between model prediction and actual system behavior. The optimization is generally facilitated through reduced order models (ROMs) that abstract the system dynamics into simpler mathematical equations. The optimal control trajectory over the control horizon is then computed using prediction of the system behavior from the ROM, and an optimization solver that seeks to minimize an objective function. For the hybrid RTU system under consideration in this paper, the control input is the SOC trajectory of the PCM. In order to compute the optimal trajectory of the PCM SOC, the following ROMs have been developed using simulation data from the high fidelity model of the hybrid RTU.

- Compressor power
- Cooling power delivered
- State of charge of the PCM.

2.2.1 **Compressor Power ROM:** The compressor power ROM is utilized by the MPC to compute the cost of increasing the PCM SOC. While there are other components of the vapor compression system such as the refrigerant pump that consume power, the compressor power is the most significant. There is always a compromise that exists between the

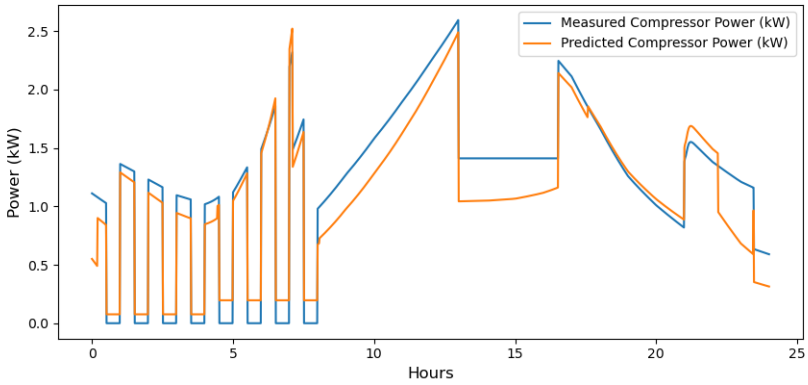

Figure 4: Comparison between compressor power simulated by ROM and physics-based model.

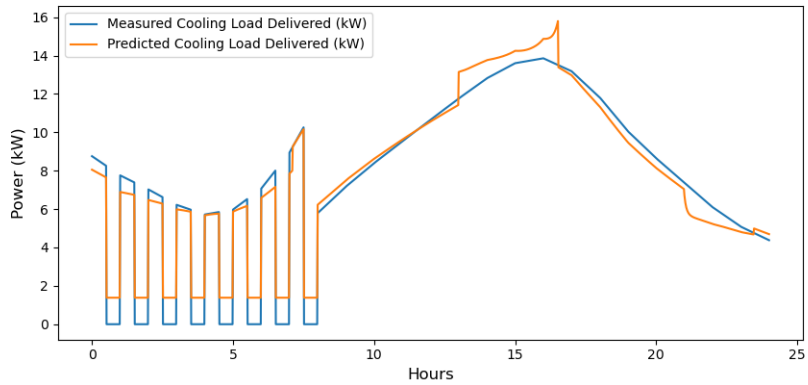

Figure 5: Comparison between cooling load delivered by ROM and physics-based model.

accuracy of the ROMs for better MPC performance, and the time and effort required to develop the models. Only the compressor power of the vapor of the charging cycle is considered in the paper. Future iterations would investigate the added value of a more accurate ROM.

The main control variable governing the compressor power and the PCM SOC is its RPM. Ramping up the compressor speed accelerates the charging of the PCM SOC, but also results in greater electrical demand. Hence, the compressor power ROM model structure is selected as shown in equation 2.

$$P_{comp} = k^i \cdot \omega_{comp}^2 + c^i \qquad i \in 1, 2, 3, 4 \tag{2}$$

A quadratic relationship is assumed between the compressor power P_{comp} and its RPM ω_{comp}, with k as the proportionality constant and c as an interpolation constant. The relationship, however, is dependent on the operating conditions. The superscript i corresponds to the current operating condition of the hybrid RTU. Four different operating conditions are identified based on the ambient temperature T_{oa} and the PCM SOC. A comparison between the ROM predictions and high-fidelity simulations is shown in Figure 4.

2.2.2 **Cooling Power Model:** The hybrid RTU system has to meet the cooling demand requirements of the conditioned space. In this section we model the amount of cooling power that is delivered by the hybrid RTU to the space. The cooling delivered depends on the mass flow rate ($\dot{m}_{fra}$) of the discharge circuit and the current PCM SOC (r_{soc}). A cooling power model is also created by using a linear parametric approach with a model structure as shown in

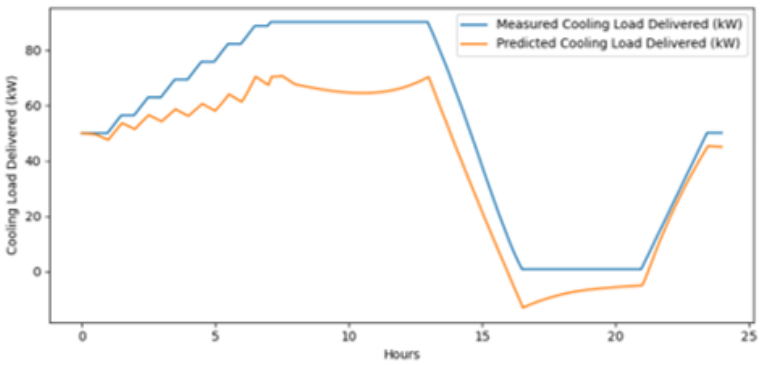

Figure 6: Comparison between the PCM SOC by ROM and neural network model.

equation 3.

$$P_{cooling} = a_0 \cdot \dot{m}_{fra} \cdot T_{oa} + a_1 \cdot \dot{m}_{fra} \cdot T_{room} + a_2 \cdot \dot{m}_{fra} \cdot r_{soc} \tag{3}$$

In the equation above, T_{room} is the temperature of the conditioned air, and parameters a_0, a_1, and a_2 are identified through a ridge regression. The right-hand side of the equation contains product terms to capture the enthalpy of the discharging circuit mass flow rate before and after passing through the heat exchanger. Figure 5 shows a comparison between the cooling power delivered as calculated by the high-fidelity model and the ROM, respectively.

2.2.3 **PCM SOC Model:** The charging and discharging process of the hybrid RTU system are heavily dependent on the current SOC of the PCM. If the SOC is low, there is a possibility that a cooling demand requirement in the near future may not be met since charging the PCM takes time. On the other hand, keeping the PCM at a high SOC all the time is highly inefficient. Hence for MPC to plan the optimal trajectory of the SOC, it is critical to model the impacts of the charging and discharging models described in the previous sections. A linear relationship is again assumed for modeling the PCM SOC with the model structure as shown in equation 4.

$$r_{soc}^{i+1} = a_{soc} \cdot r_{soc}^{i} + b_0 \cdot P_{comp} + b_1 \cdot P_{comp} \cdot T_{oa} + b_2 \cdot P_{cooling} \tag{4}$$

The parameters $\theta_{soc} = [a_{soc}, b_0, b_1, b_2]$ are computed using ARX system identification process. The features of the model are selected based on some physics intuition, and feature engineering. Figure 6 shows a comparison between the evolution of the SOC of the PCM using the high-fidelity model and the linear parametric model.

The three ROMs described above are then combined to form a single hybrid RTU ROM with compressor speed ω_{comp} and conditioned air mass flow rate m_{fra} as inputs to the model. A perfect local controller is also designed that makes use of the analytical equations shown above to translate a PCM SOC setpoint to an equivalent compressor speed and conditioned air mass flow rate.

3. CO-SIMULATION SETUP

A co-simulation environment was developed to design and test the MPC strategy on the hybrid RTU system, and compare it against an RBC. In this paper, MPC is tested with only the hybrid RTU system in the feedback loop. The building load is not part of the closed-loop control. A flow-diagram of the co-simulation setup is shown in Figure 7. As can be seen from the figure, the only feedback that the MPC receives is the PCM SOC. In this particular use case, the hybrid RTU system always tries to meet a predefined building cooling load profile. This significantly reduces the amount of flexibility that can be leveraged from the home, but we wanted to test and demonstrate the capabilities of MPC just on the hybrid RTU system. Adding the building model into the closed loop of the MPC would require a house thermodynamic ROM, which may not always be available or easy to obtain. Developing a house thermodynamic ROM would require historical data of the weather conditions, the cooling provided by the HVAC system, and the temperature of the conditioned space. Thus in this paper we wanted to investigate the benefits of MPC if we just have knowledge about the building cooling load profile and the weather forecast.

The co-simulation setup can be divided into three components:

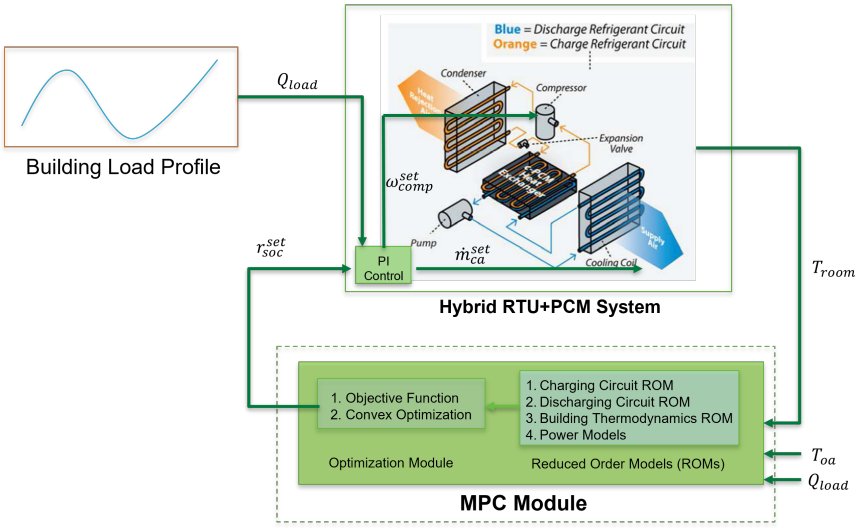

Figure 7: Flow diagram of the co-simulation setup between the high-fidelity physics-based and neural network models, and the ROMs.

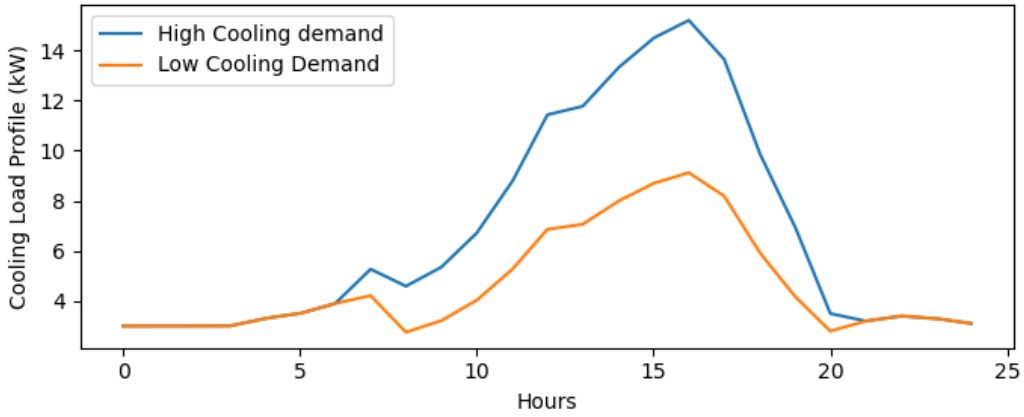

Figure 8: Building cooling load profile for the high demand and low demand use cases.

- High-fidelity simulation model described in Section 2.1.
- MPC module comprising ROMs described in Section 2.2, and the objective function and optimization solver (discussed in Section 3.1).
- Disturbance input forecast comprising 1) building cooling load profile and 2) outside air drybulb temperature.

Using the forecast of the disturbance inputs and feedback of the current PCM SOC, MPC computes an optimal trajectory of the PCM SOC over the control horizon. MPC uses the hybrid RTU ROM and an optimization solver to compute this trajectory. The first SOC setpoint in this trajectory is then sent to the high-fidelity model described in Section 2.1. A local controller using a method of trial and error converts the SOC setpoint from MPC to an RPM signal to the compressor, and a mass flow rate input to the discharging circuit of the high-fidelity model. The updated PCM SOC is fed back to the MPC, and the process is repeated at the next time step. A further description of the co-simulation use case is provided below.

3.1 MPC and RBC under TOU Rate Structure

As discussed above, a static building cooling load profile is provided to the co-simulation environment for the hybrid RTU to track. Two synthetic load profiles are generated, as shown in Figure 8 representing the high load and the low

load use cases. Each of the two use cases are run with RBC as well as MPC strategies in the co-simulation framework. The control performance is evaluated based on the utility cost under a TOU rate structure shown in Table2. For both control strategies, the minimum and maximum SOCs are restricted between 10% and 90%, respectively.

Table 2: TOU Rate Structure.

Period	Time	Rate
Peak	13:00 - 17:00	1.35 $/kWh
Off-Peak	00:00 - 13:00 and 17:00 to 24:00	0.45 $/kWh

3.1.1 **RBC Strategy:**The design of an RBC strategy is generally subjective but there are a few guiding principles that can be applied to devise the rules. For the case of the hybrid RTU, to reduce the cost under a TOU rate structure, a strategy that can be adopted is to charge the PCM completely just before the start of the peak period, and discharge completely by the end of the peak period. This strategy is adopted for RBC in this paper. But since this strategy doesn't adapt to the cooling load requirements, the same SOC setpoint profile is sent to the high-fidelity model for both the high-demand and low-demand use cases.

3.1.2 **MPC Strategy:**Unlike RBC, an MPC strategy can adapt to the changes in the ambient weather and cooling load requirements. A control horizon of 6 hours is selected in this paper as a compromise between planning for a more distant future and computational time. MPC leverages the hybrid RTU ROM to compute an optimal trajectory of the PCM SOC for each hour of the control horizon as shown in equation 5.

$$\mathbf{u_{soc}} = \arg\min_{r_{soc}} J^{mpc}(r_{soc}) \tag{5}$$

where $\mathbf{u_{soc}}$ is the set of 6 optimal PCM SOC setpoints over the control horizon, and $J^{mpc}(r_{soc})$ is the objective function to be minimized. The objective function designed for the study is the utility cost associated with the compressor power ω_{comp} and the fan energy to deliver the conditioned air, under the TOU rate structure provided above. The objective function shown in equation 6.

$$\mathbf{J^{mpc}} = \sum_{k=1}^{n_{hrzn}} r_k^{elec} \cdot t_{sim-hr} \cdot (P_{comp} + P_{m_{fra}}) \tag{6}$$

where k is the simulation timestep, n_{hrzn} is the number of timesteps over the control horizon, r_{elec} is the utility rate, and t_{sim-hr} is the simulation timestep in hours. The hybrid RTU system is, however, constrained to track the cooling load requirements described above.

4. RESULTS AND DISCUSSION

In this section we provide the results of the comparison of the RBC and MPC strategies when applied to the hybrid RTU system in a co-simulation framework. Two use cases are investigated: 1) high cooling demand and 2) low cooling demand. For a fair comparison between the two control strategies, it is important to take into consideration both the resultant utility cost and the cooling load tracking errors. Figure 9 shows how the two control strategies track the cooling load and the PCM SOC trajectories of the high-fidelity simulation. The dashed lines correspond to the low demand use case.

We see that both control strategies track the cooling load profile with relatively small errors. The PCM SOC trajectories, however, are vastly different. In the RBC strategy, the PCM is fully charged just before the start of the peak period, and fully discharged at the end of the peak period. No discharge occurs prior to the start of the peak period. MPC, however, discharges the PCM at the start of the simulation, as the cooling load over the control horizon can be met with a much smaller SOC during the beginning of the day. MPC sees the forecast of higher TOU prices beginning at 7 am (6 hours prior). It, however, doesn't start charging until 10 am, as the predictions from the ROM suggest that cooling demand during the peak period can be met without having to reach the maximum SOC. As with RBC, MPC also completely discharges during the peak hour period. The same behavior is observed for the low-demand use case, with MPC having to charge to a smaller SOC level before the start of the peak period.

Table 3 shows the summary of the utility costs incurred in the two use cases and the corresponding cooling load tracking error. The cooling load tracking error is counted only when the cooling provided is less than the load requirement. We

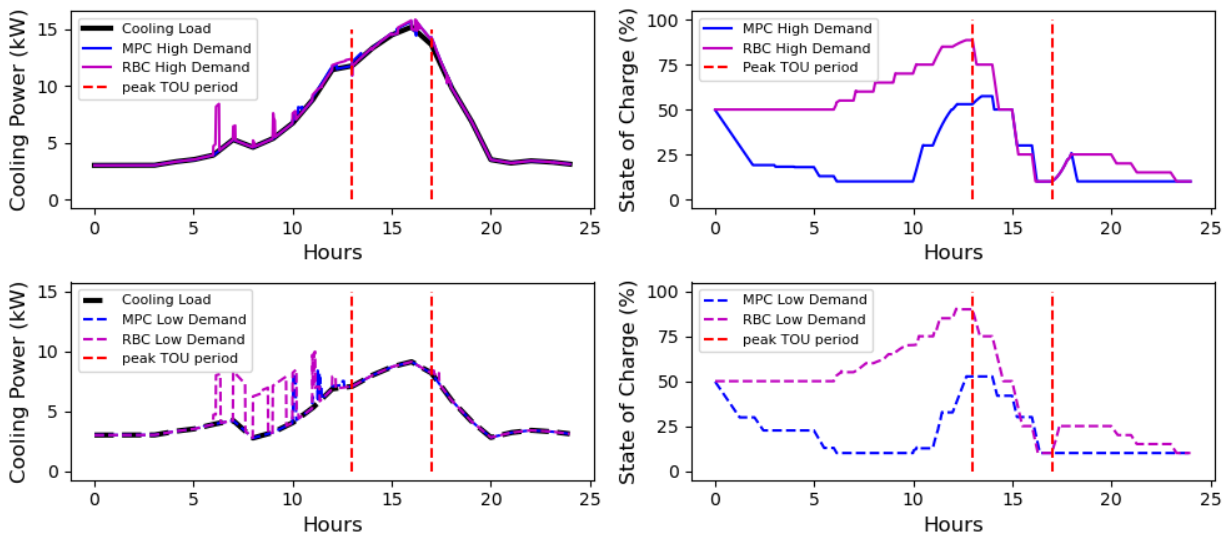

Figure 9: Cooling load tracking and PCM SOC trajectories of RBC and MPC control strategies for high demand and low demand use cases.

see that for the high demand use case there is very little difference in performance between the two control strategies. Although the results are not favorable for MPC for the high demand case where the savings are -3.8% on utility costs, MPC does better than RBC in the low demand case with savings of 4.5%. One explanation for the results is that it does make sense to charge the PCM completely just before the start of the peak period when the demand is high. But this strategy is not as fruitful on medium or low demand use cases. There are still, however, many more avenues that warrant further investigation. Current research is focusing on studying the impacts of increasing MPC control horizon, improving ROM accuracy, alternate optimization solvers and initial guesses for the MPC solution, and different TOU rate structures. we are also currently investigating the impact of adding the building load into the closed loop of the MPC.

Table 3: Cost and cooling load tracking error summary of the co-simulation study.

Use case	Utility cost ($)	Load tracking error (kWh)
MPC High Demand	37.88	0.10
RBC High Demand	36.48	0.10
MPC Low Demand	18.74	0.05
RBC Low Demand	17.90	0.06

5. CONCLUSIONS

The full potential of the flexibility that buildings can provide to the grid is far from realized. HVAC systems are the most prominent source of this flexibility within buildings. The flexibility can be further enhanced with the help of TES devices, which can store heating or cooling energy to be used at a later time and assist the grid by following a more desirable energy use profile.

A hybrid RTU system integrated with a PCM that serves as a TES system is the subject of this paper. To maximize the flexibility potential of the hybrid RTU we designed and tested an advanced control strategy—MPC in a co-simulation environment. A high-fidelity physics-based vapor compression system and neural network model trained from finite-element analysis of the PCM serves as the simulation testbed. Reduced order models are developed for each component of the hybrid RTU enabling MPC to compute an optimal SOC trajectory for the PCM. These setpoints are sent to the high-fidelity model, and feedback is received via the PCM SOC.

This paper presents the steps for the component ROM development, integration with the high-fidelity model in a co-

simulation framework, and the design of an MPC strategy for the hybrid RTU system. The co-simulation platform is then used to evaluate the performance of MPC by comparing it with an RBC strategy applicable for a TOU rate structure. RBC does slightly better in the high demand use case, where is MPC is slightly better for the low demand use case. The results make intuitive sense, since charging the PCM SOC to a high value on high demand days would require less mass flow rate and compressor speed to meet the cooling demand requirements. The strategy, however, may not be as fruitful on medium or low cooling demand days. There are still plenty of other factors that that need to be investigated to test the MPC strategy: 1) the control horizon, 2) the accuracy of the component ROMs, 3) TOU and demand utility rate structures, etc., and these are currently ongoing areas of research.

ACKNOWLEDGEMENT

This work was authored 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 was provided by the U.S. DOE 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.

Please note that specific commercial products discussed in this report are used as research examples only—NREL does not endorse or support any particular brand, product, service, or manufacturer.

REFERENCES

Arteconi, A., Hewitt, N., & Polonara, F. (2012). State of the art of thermal storage for demand-side management. *Applied Energy*, *93*, 371-389.

Beccali, M., Bellia, L., Fragliasso, F., Bonomolo, M., Zizzo, G., & Spada, G. (2020). Assessing the lighting systems flexibility for reducing and managing the power peaks in smart grids. *Applied Energy*, *268*, 114924.

Center, B. P. (2020). Annual energy outlook 2020. *Energy Information Administration, Washington, DC*, *12*, 1672–1679.

Chen, Y., Chen, Z., Xu, P., Li, W., Sha, H., Yang, Z., … Hu, C. (2019). Quantification of electricity flexibility in demand response: Office building case study. *Energy*, *188*, 116054.

Huang, R., Aute, V., & Ling, J. (2021). A tripartite graph based methodology for steady-state solution of generalized multi-mode vapor compression systems. *Applied Thermal Engineering*, *185*, 116385.

Huang, R., Ling, J., & Aute, V. (2020). Comparison of approximation-assisted heat exchanger models for steady-state simulation of vapor compression system. *Applied Thermal Engineering*, *166*, 114691.

Huang, R., Mahvi, A., James, N., Kozubal, E., & Woods, J. (2024). Evaluation of phase change thermal storage in a cascade heat pump. *Applied Energy*, *359*, 122654.

Le Dréau, J., Lopes, R. A., O'Connell, S., Finn, D., Hu, M., Queiroz, H., … Saeed, M. H. (2023). Developing energy flexibility in clusters of buildings: A critical analysis of barriers from planning to operation. *Energy and Buildings*, *300*, 113608.

Li, H., Wang, Z., Hong, T., & Piette, M. A. (2021). Energy flexibility of residential buildings: A systematic review of characterization and quantification methods and applications. *Advances in Applied Energy*, *3*, 100054.

Powells, G., Bulkeley, H., Bell, S., & Judson, E. (2014). Peak electricity demand and the flexibility of everyday life. *Geoforum*, *55*, 43–52.

Woods, J., Mahvi, A., Goyal, A., Kozubal, E., Odukomaiya, A., & Jackson, R. (2021). Rate capability and ragone plots for phase change thermal energy storage. *Nature Energy*, *6*(3), 295–302.