Hybrid MPC for Building Heat Pump Demand Flexibility under Unmeasured Disturbances

Sang woo Ham, Donghun Kim
Building Technology & Urban Systems Division
Lawrence Berkeley National Laboratory
Berkeley, CA, USA
{sham, donghunkim}@lbl.gov

Abstract

Implementing advanced controls like Model Predictive Control (MPC) to unlock the grid flexibility of small- and medium-sized commercial buildings (SMCBs) is often prohibitively expensive due to the extensive sensors required to measure internal heat gains. This paper introduces a Hybrid MPC framework that overcomes this challenge without additional sensors by integrating a physics-informed gray-box model with a neural network (NN) that forecasts these unmeasured disturbances. Evaluated in a multi-building case study, the Hybrid MPC achieves substantial load shifting and peak demand reduction, with performance nearly matching an ideal controller with perfect foresight.

1 Introduction

The increasing strain on the U.S. electrical grid [6, 7], driven by electrification and the integration of intermittent renewables, necessitates new sources of demand flexibility. Buildings, consuming nearly 74% of the nation's electricity, are a prime resource. By leveraging their inherent thermal mass, buildings can shift their heating and cooling loads, reducing electricity consumption during periods of high grid stress or high prices [9, 8]. Model Predictive Control (MPC) is a powerful strategy for organizing this load shifting, with numerous studies demonstrating its ability to reduce energy costs and peak demand while maintaining occupant comfort [2].

Despite its potential, MPC has seen limited adoption in small- and medium-sized commercial buildings (SMCBs), which constitute about 95% of U.S. commercial building stock [10]. A critical barrier is the difficulty of accurately modeling building thermal dynamics in the presence of significant unmeasured disturbances. Internal heat gains from occupants, lighting, and plug loads can drastically alter a building's thermal trajectory [5], yet they are rarely measured directly because doing so requires a commercial-grade Building Automation System (BAS), an infrastructure that is both costly and infeasible in most SMCBs. Standard MPC approaches attempt to circumvent this by using approximations based on total plug load data [1] and periodic recalibration, but these methods often fail to capture the true dynamics of the disturbances. Consequently, models that rely on these inaccurate estimations may make suboptimal control decisions, leading to poor performance, occupant discomfort, and wasted energy [3].

To address this challenge, we propose a Hybrid MPC framework designed specifically for SMCBs. Our approach combines the strengths of physics-informed and data-driven models:

1. A lightweight gray-box thermal model captures the building's fundamental thermal dynamics (e.g., thermal mass, envelope resistance). This model is identified from operational data using a robust Lumped Disturbance (LD) method that isolates the effects of unknown heat gains [5].

39th Conference on Neural Information Processing Systems (NeurIPS 2025) Workshop: UrbanAI: Harnessing Artificial Intelligence for Smart Cities.

2. A data-driven neural network (NN) is trained to forecast the future profile of the unmeasured internal heat gains by using the historical data and a gray-box model.

By integrating the NN disturbance forecast into the gray-box model's predictions (i.e., hybrid modeling approach), the MPC can optimize HVAC operations with a much more accurate view of future conditions. This paper first introduces this hybrid modeling methodology. We then present a simulation-based case study to evaluate its performance, comparing several MPC strategies, including Hybrid MPC. The paper closes with a conclusion and a discussion of future research directions.

2 Hybrid MPC Framework

The proposed Hybrid MPC is designed for typical SMCBs, which often feature multiple independent thermal zones, each served by its own thermostat and HVAC unit (e.g., a rooftop unit or heat pump (HP)). The MPC coordinates these units to achieve building-level objectives for load shifting and peak demand reduction.

2.1 MPC Formulation

The MPC solves an optimization problem at each control timestep k to find the optimal sequence of HVAC operations over a prediction horizon N_p. The objective is to minimize the total electricity cost, a penalty on peak power demand, and any comfort violations. The core optimization problem is:

$$
\begin{aligned}
\min \quad & \sum_{j=1}^{N_p}\sum_{i=1}^{n} ER(k+j-1)P_{\mathrm{HP},i}(k+j-1)\bar{u}_i(k+j-1)+\omega_{\mathrm{d}}\delta+\omega_{\mathrm{l}}\Gamma_{\mathrm{l}}+\omega_{\mathrm{u}}\Gamma_{\mathrm{u}} \qquad (1)\\
s.t. \quad & T_{\mathrm{l},i}(k+j)-\Gamma_{\mathrm{l}} \le \mathbb{E}[\bar{y}_i(k+j)|\mathcal{D}_k] \le T_{\mathrm{u},i}(k+j)+\Gamma_{\mathrm{u}}\ (\forall i \in \{1,\cdots,n\})\\
& \sum_{i=1}^{n} P_{\mathrm{HP},i}(k+j-1)\bar{u}_i(k+j-1) \le \delta \quad \text{and} \quad 0 \le \bar{u}_i(k+j-1) \le 1
\end{aligned}
$$

where for each zone i and future step j, ER is the electricity price ($/kWh), $P_{\mathrm{HP},i}$ is the HP power draw, and $\bar{u}_i$ is the control input, representing the normalized HP runtime fraction (RTF) for $(k+j)^{\mathrm{th}}$ timestep of the MPC. The decision variables are the RTF trajectories $\bar{u}_i$, the peak power limit over the horizon δ, and a slack variable for comfort violations Γ_c. The constraints enforce that the predicted zone air temperature given the historic and forecast data $\mathbb{E}[\bar{y}_i(k+j)|\mathcal{D}_k]$ stays within the lower and upper comfort bounds ($T_{\mathrm{l},i}, T_{\mathrm{u},i}$) and that the aggregate power does not exceed δ. The weights ω_d and ω_c penalize peak demand and comfort violations, respectively. The output of the optimization is a sequence of optimal RTFs, which are then translated into temperature setpoints for the individual thermostats.

2.2 Hybrid Thermal Modeling Approach

Accurate prediction of the future zone temperature, $\mathbb{E}[\bar{y}_i(k+j)|\mathcal{D}_k]$, is critical to the MPC's success. Our hybrid approach achieves this in three steps, as illustrated in Figure 1.

Step 1: Gray-Box Model Identification. We model the thermal dynamics of each zone using a standard 2R2C gray-box model structure [4], which represents the zone with two temperature states (zone air and thermal mass) and a set of thermal resistances and capacitances. A key challenge is identifying the parameters of this model (R and C values) from data that is corrupted by unmeasured heat gains $\dot{Q}_{\mathrm{g}}$. We use an LD system identification technique [5], which models the unmeasured heat gains as an additional disturbance term during the parameter estimation process. This allows us to obtain an unbiased estimate of the building's physical thermal parameters, even without measuring $\dot{Q}_{\mathrm{g}}$.

Step 2: Disturbance Model Training. With the identified gray-box model, we can use a Kalman filter on historical operational data to infer the past time-series of the unmeasured heat gain, $\dot{Q}_{\mathrm{g}}$. This inferred historical profile serves as the training data for a data-driven disturbance model. We use an NN model. A long short-term memory (LSTM) structure is chosen from the previous research [3],

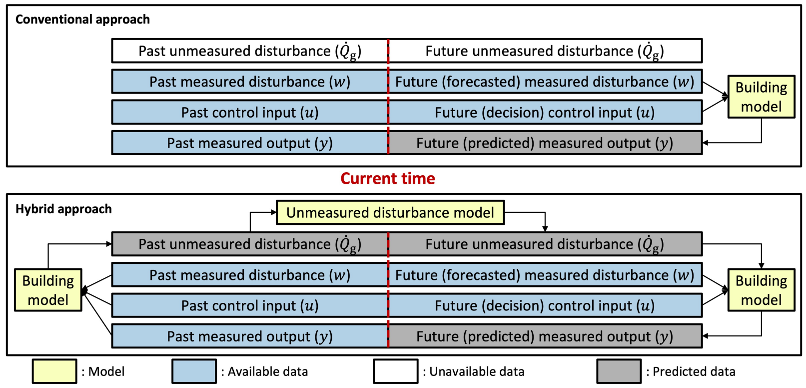

Figure 1: Comparison of the prediction process for a conventional MPC, which ignores future unmeasured disturbances, and our Hybrid MPC, which forecasts and incorporates them.

a type of recurrent neural network well-suited for time-series forecasting. Specifically, the LSTM is trained to predict the future hourly profile of $\dot{Q}_\mathrm{g}$ based on inputs such as time of week, outdoor air temperature, solar radiation, and the recent history of $\dot{Q}_\mathrm{g}$ itself. The model structure and its required inputs are chosen based on the building's thermal system transfer function. Since system identification can be imperfect, a disturbance model needs to account for these inaccuracies and the effects of inputs on the building thermal system.

Step 3: Integrated Prediction for MPC. During real-time operation, the MPC uses the combined hybrid model for prediction. At each step, the LSTM provides a 24-hour forecast of the unmeasured heat gains, $\hat{\dot{Q}}_\mathrm{g}$. This forecast is fed as a known input into the gray-box model, along with weather forecasts and the planned HP control actions ($\bar{u}_i$). This provides a comprehensive and accurate prediction of future zone temperatures, enabling the MPC to make informed, proactive decisions.

3 Case study

We evaluated the Hybrid MPC framework in a simulation study using hypothetical 9 buildings. Each building is thermally independent and served by its own HP, with thermal dynamics governed by a 2R2C model. Internal loads representing occupants and equipment were applied on a typical weekday schedule.

3.1 Simulation Scenario

We compared the performance of four control strategies over a one-month heating period and a one-month cooling period. A dynamic electricity price signal with a high morning peak (6-8 AM) was used to incentivize load shifting. Four types of scenarios are investigated: (a) **Baseline:** A standard programmable thermostat with fixed occupied and unoccupied setpoints, (b) **MPC$_\mathbf{vanilla}$:** A conventional MPC using only the gray-box model identified in Step 1 (with LD approach, i.e., without knowing unmeasured disturbances). It assumes future unmeasured disturbances are zero, representing a typical MPC implementation without disturbance forecasting, (c) **MPC$_\mathbf{ideal}$:** An idealized MPC that has perfect knowledge of future unmeasured disturbances. The unmeasured disturbances are used for system identification. It serves as a theoretical performance benchmark, (d) **MPC$_\mathbf{hybrid}$:** Our proposed controller, using the gray-box model (with LD approach, i.e., without knowing unmeasured disturbances), integrated with the LSTM-based disturbance forecast.

3.2 Results

Due to the space limitation, only the heating season results are presented. Figure 2 shows the average daily temperature and power profiles for each controller. The Baseline controller reacts to the 6 AM occupancy schedule, causing a large, simultaneous power spike that coincides directly with the high-price period.

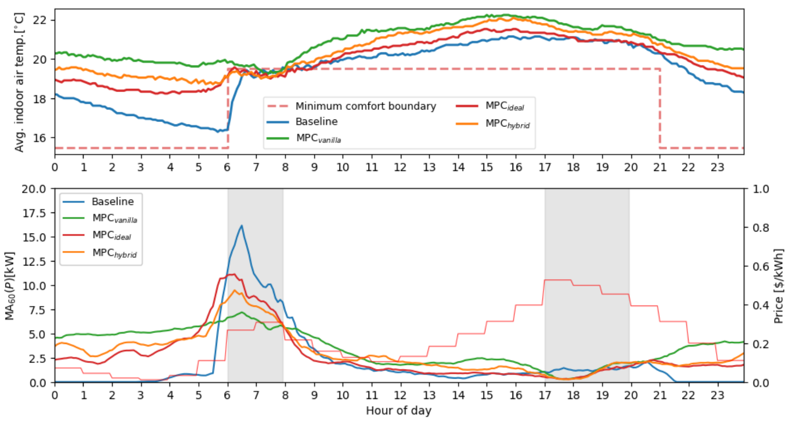

Figure 2: Average operational profiles during the heating season simulation. **Top:** The Hybrid and Ideal MPCs maintain comfort without the significant overheating caused by the Vanilla MPC. **Bottom:** The Hybrid MPC effectively shifts load away from the 6-8 AM peak price period, closely mimicking the Ideal MPC's behavior.

All MPC strategies successfully shift this load by pre-heating the building during the cheaper, early morning hours. However, their effectiveness varies significantly. The $\mathbf{MPC_{vanilla}}$, unaware of the impending internal gains from occupancy, pre-heats excessively. This leads to temperature overshoot and wasted energy, resulting in a 15.5% *increase* in monthly cost compared to the Baseline (Table 1). In contrast, the $\mathbf{MPC_{ideal}}$ optimally pre-heats just enough to coast through the morning peak, achieving a 3.9% cost savings.

The performance of our $\mathbf{MPC_{hybrid}}$ closely tracks that of the ideal controller. By accurately forecasting the upcoming internal gains, it avoids the excessive pre-heating of the vanilla MPC. It achieves a substantial 32.5% load shift away from the morning peak period and a 16% reduction in the season's highest peak demand, with only a marginal 2.5% increase in monthly cost over the baseline. The cooling season results showed a similar trend, with the Hybrid MPC achieving an 18.4% cost reduction, nearly matching the 19.5% savings of the Ideal MPC.

Table 1: Summary of Heating Season Simulation Performance

Performance Metric	**Baseline**	$\mathbf{MPC_{vanilla}}$	$\mathbf{MPC_{ideal}}$	$\mathbf{MPC_{hybrid}}$
Monthly Cost ($)	149.70	172.98 (+15.5%)	143.81 (-3.9%)	153.45 (+2.5%)
Morning Load Shift (%)	—	45.5	23.9	32.5
Peak Demand Reduction (%)	—	37.0	16.0	16.0

4 Conclusion

This paper presents a Hybrid MPC framework that enables effective load shifting and peak demand reduction in commercial buildings without requiring direct measurement of internal heat gains. By combining a physics-informed gray-box model with a data-driven LSTM network for disturbance forecasting, our controller makes intelligent, proactive decisions that account for these critical unmeasured dynamics. Simulation results demonstrate that this hybrid approach is essential for achieving high performance; whereas a conventional MPC that ignores future disturbances may shift load inefficiently and increase energy costs, our Hybrid MPC achieves significant load shifting and peak reduction while closely matching the economic performance of an ideal controller. Ultimately, this work demonstrates a practical and scalable pathway for deploying advanced building controls to unlock the vast demand flexibility potential of the SMCB sector. Future work will explore the integration of stochastic methods to explicitly account for uncertainty in disturbance forecasts.

Acknowledgments

This work was supported by the Assistant Secretary for Energy Efficiency and Renewable Energy, Building Technologies Office, of the U.S. Department of Energy under Contract No. DE-AC02-05CH11231, by the New York State Energy Research & Development Authority (NYSERDA), and by the California Energy Commission through grant EPC-19-013.

References

[1] D Blum, Z Wang, C Weyandt, D Kim, M Wetter, T Hong, and M A Piette. Field demonstration and implementation analysis of model predictive control in an office HVAC system. *Applied energy*, 318:119104, 2022.

[2] Ján Drgoňa, Javier Arroyo, Iago Cupeiro Figueroa, David Blum, Krzysztof Arendt, Donghun Kim, Enric Perarnau Ollé, Juraj Oravec, Michael Wetter, Draguna L Vrabie, and Lieve Helsen. All you need to know about model predictive control for buildings. *Annual reviews in control*, 50:190–232, January 2020.

[3] Sang Woo Ham and Donghun Kim. Hybrid modeling approach for better identification of building thermal network model and improved prediction. page Paper 420, 2022.

[4] Sang Woo Ham, Donghun Kim, Tanya Barham, and Kent Ramseyer. The first field application of a low-cost MPC for grid-interactive K-12 schools: Lessons-learned and savings assessment. *Energy and Buildings*, 296:113351, 2023.

[5] Donghun Kim, Jie Cai, Kartik B Ariyur, and James E Braun. System identification for building thermal systems under the presence of unmeasured disturbances in closed loop operation: Lumped disturbance modeling approach. *Building and environment*, 107:169–180, October 2016.

[6] NERC. 2024 long-term reliability assessment. Technical report, North American Electric Reliability Corporation, 2024.

[7] NERC. Urgent need for resources over 10-year horizon as electricity demand growth accelerates, 2024 LTRA finds. Technical report, North American Electric Reliability Corporation, 2024.

[8] M Neukomm, Valerie Nubbe, and R Fares. Grid-interactive efficient buildings technical report series: Overview of research challenges and gaps. Technical report, Office of Energy Efficiency & Renewable Energy, U.S. Department of Energy, 2019.

[9] U.S. EIA. U.S. energy consumption by source and sector, 2023. Technical report, U.S. Energy Information Administration, 2024.

[10] U.S. Energy Information Administration (EIA). Commercial buildings energy consumption survey (CBECS) data. Technical report, U.S. Department of Energy, Washington, DC, 2018.