Energy & Buildings 360 (2026) 117324

Contents lists available at ScienceDirect

Energy & Buildings

journal homepage: www.elsevier.com/locate/enbuild

Machine learning-enhanced MPC for demand flexibility in small commercial buildings: An experimental study

Sang Woo Ham [a,*], Donghun Kim [a], Lazlo Paul [a], Armando Casillas [a], Anand Krishnan Prakash [a], Richard Brown [a,b], Marco Pritoni [a]

[a] Building & Industrial Energy Systems Division, Lawrence Berkeley National Laboratory, Berkeley, CA, USA
[b] California Institute for Energy and Environment (CIEE), University of California, Berkeley, CA, USA

ARTICLE INFO

Keywords:
Model Predictive Control
Hybrid Modeling
Unmeasured Disturbance Forecasting
Load Shifting
Demand Flexibility
Machine Learning

ABSTRACT

Small- and medium-sized commercial buildings (SMCBs) represent the majority of U.S. commercial building stock and a significant share of peak electricity demand, yet they often lack centralized building automation systems, representing a significant untapped resource for urban energy management. This infrastructure gap makes advanced control implementation challenging, limiting the potential for widespread demand flexibility. Model Predictive Control (MPC) has shown strong potential for load shifting, peak demand reduction, and cost savings, but its effectiveness is hindered by unmeasured disturbances such as internal heat gains. This paper presents a Hybrid MPC framework that integrates a physics-based gray-box building thermal model, identified using a lumped disturbance (LD) approach, with a machine learning (ML) model for forecasting unmeasured disturbances. The hybrid approach is designed for buildings with multiple individually controlled heat pump and thermostat pairs, common in SMCBs, and aims to optimize coordinated scheduling of multiple heat pumps under dynamic electricity pricing while respecting comfort constraints. The methodology is validated through both simulations of case study buildings and experimental studies at a highly-instrumented test facility. Simulation results show that the Hybrid MPC achieves substantial load shifting and peak demand reduction, approaching the performance of an ideal MPC with perfect disturbance knowledge, and outperforming a conventional MPC without disturbance forecasting. In experiments, the Hybrid MPC reduced daily HVAC energy costs by 8.7%, peak-price time load (load shifting) by 41.7%, and peak demand by 29.2% compared to baseline control, demonstrating comparable benefits to the 11.6% cost savings, 42.9% load shifting, and 23.2% peak reduction of the ideal MPC. These results demonstrate that the proposed hybrid modeling approach can significantly improve MPC performance in real-world SMCB applications without requiring additional disturbance measurements.

1. Introduction

The U.S. electrical grid is increasingly challenged by rising demand from data center growth, the adoption of new electric technologies, and the retirement of conventional generation resources, all of which heighten risks to grid reliability and resilience [1,2]. Consequently, managing the timing of energy use via demand flexibility becomes as critical as energy efficiency. Buildings consume approximately 74% of the nation's electricity, with heating, ventilation, and air conditioning (HVAC) systems as major contributors to peak demand [3,4]. While this large footprint exacerbates grid stress, it also presents a significant opportunity for providing demand flexibility-particularly through load shifting and peak demand reduction-by leveraging the inherent thermal mass of building structures. For example, pre-heating or pre-cooling spaces before high-price or high-demand periods can shift electrical load to off-peak hours, reducing stress on the grid.

Model Predictive Control (MPC) has emerged as one of the most effective advanced control strategies for enabling such flexibility. Originating from the process industry, MPC has been applied to buildings since the early 1990s [5], with numerous studies demonstrating its ability to reduce energy costs [6], lower peak demand [7], integrate renewable generation [8–10], and participate in demand response (DR) programs [11–13]. Comprehensive reviews [14–16] have summarized MPC applications, challenges, and future research needs, including the high effort required for modeling, the lack of automated modeling tools, and the shortage of HVAC domain expertise.

Despite this extensive research base [19], specific applications for small- and medium-sized commercial buildings (SMCBs)-which

* Corresponding author.
E-mail address: sham@lbl.gov (S.W. Ham).

https://doi.org/10.1016/j.enbuild.2026.117324
Received 16 December 2025; Received in revised form 22 February 2026; Accepted 12 March 2026
Available online 17 March 2026

Nomenclature

Variables

BAS	Building Automation System
CFM	Cubic Feet per Minute
DR	Demand Response
HDP	Highly Dynamic Price
HP	Heat Pump
HP-RTU	Heat Pump Rooftop Unit
HVAC	Heating, Ventilation, and Air Conditioning
ID	Input Disturbance
IoT	Internet of Things
LD	Lumped Disturbance
LSTM	Long Short-Term Memory (neural network)
MPC	Model Predictive Control
OD	Output Disturbance
PRBS	Pseudo-Random Binary Signal
RTF	Runtime Fraction
RTU	Rooftop Unit
SMCB	Small- and Medium-Sized Commercial Building

Variables

$\mathbf{A}$	State transition matrix [-]
A_{win}	Effective window area [m^2]
$\mathbf{C}$	Output matrix of the state-space model [-]
C_{w}	Thermal capacitance of building thermal mass [J/K]
C_{za}	Thermal capacitance of zone air [J/K]
Δ_{h}	Thermostat hysteresis band [°C]
D_k	Available measurement data up to timestep k
δ	Peak power limit in MPC optimization [kW]
$\dot{Q}_{\text{g}}$	Lumped internal heat gains (unmeasured disturbance) [kW]
$\dot{Q}_{\text{h}}$, $\dot{Q}_{\text{c}}$	Full-capacity heating/cooling rate [kW]
$\dot{Q}_{\text{hc}}$	Full-capacity heating/cooling rate [kW]
$\dot{q}_{\text{sol,win}}$	Incident solar radiation on windows [kW/m^2]
ER	Electricity price [$/kWh]
f	Convective fraction of incident solar radiation [-]
Γ_{l}, Γ_{u}	Lower/upper comfort violation slack variables [°C]
$\hat{\bar{y}}_i(k+j\|\hat{\mathbf{x}}_i(k))$	Predicted zone air temperature for zone i at step $k+j$, conditioned on filtered state [°C]
$\hat{\bar{y}}_i^*(k+1\|\hat{\mathbf{x}}_i(k))$	Optimal predicted zone air temperature corresponding to optimal control sequence [°C]
$\hat{\mathbf{x}}_i(k)$	Kalman-filtered state estimate for zone i at timestep k
j	Prediction step index in MPC [-]
k	Current control timestep index [-]
n	Number of heat pumps or thermal zones [-]
N_p	Prediction horizon in MPC [-]
$P_{\text{HP,h},i}$, $P_{\text{HP,c},i}$	Electric power of ith heat pump heating or cooling operation at full capacity [kW]
R_{zo}	Thermal resistance between zone air and outdoor air [K/kW]
R_{zw}	Thermal resistance between zone air and thermal mass [K/kW]
$T_{\text{csp},i}$, $T_{\text{hsp},i}$	Cooling/heating setpoint temperature for zone i [°C]
$T_{\text{l},i}$, $T_{\text{u},i}$	Lower/upper comfort temperature bounds for zone i [°C]
T_{oa}	Outdoor air temperature [°C]
T_{w}	Thermal mass temperature [°C]
T_{za}	Zone air temperature [°C]
$\mathbf{u}$	Control input vector [-]
$\bar{u}_{\text{h},i}$, $\bar{u}_{\text{c},i}$	Heating and cooling runtime fractions for zone i [-]
u_{h}, u_{c}	Heating/cooling stage signal (0–1 normalized) [-]
ω_{d}	Weight on peak demand reduction term in MPC objective [-]
ω_{l}, ω_{u}	Weights on lower/upper comfort violation penalties [-]
y_{za}	Measured zone air temperature [°C]

comprise 95% of U.S. commercial buildings [17]-remain underrepresented. While grey-box resistance-capacitance (RC) models are theoretically well-suited for these applications due to their balance of physical interpretability and computational efficiency [18], practical implementation is hindered by the heterogeneity of SMCB systems. Unlike large commercial facilities, SMCBs are typically served by multiple independent HVAC units (e.g., rooftop units, packaged heat pumps, ductless mini-splits) controlled by individual thermostats-often without a building automation system (BAS) [20]. In some cases, control is limited to infrared (IR) remotes, further constraining integration with conventional supervisory control methods [21,22]. These characteristics present both an opportunity and a challenge: there is potential for a universal, scalable control solution for SMCBs. However, the absence of centralized control infrastructure and the discrete ON/OFF nature of equipment pose significant challenges to implementing MPC.

The high cost and complexity of retrofitting SMCBs with advanced controls have limited broader adoption. However, the recent proliferation of Internet of Things (IoT) technologies-such as WiFi-enabled thermostats-has opened new pathways for cost-effective, data-driven, and remotely deployable advanced controls [21,23]. These devices enable both data collection and remote setpoint control without major hardware changes, making scalable deployment of MPC in SMCBs increasingly feasible. This technological shift is a prerequisite for integrating the vast stock of SMCBs into smart city energy management frameworks.

Previous field demonstrations have begun to validate the efficacy of MPC in residential and small commercial contexts, particularly for economic optimization under Time-of-Use (TOU) or dynamic pricing. Pedersen et al. [24] pioneered this approach by achieving 9% cost savings through shifting heat pump operations based on day-ahead market prices. Similarly, Afram and Janabi-Sharifi [25] reported substantial savings between 8.5% and 56.1% using TOU pricing in a residential setting. Thorsteinsson et al. [26] further demonstrated this potential in a single-family house, achieving up to 32.8% cost savings by shifting heating loads based on spot prices. Jiang et al. [27] also showed substantial savings through optimal precooling using existing smart thermostats. In terms of algorithmic diversity, Bunning et al. [28] utilized machine learning-based Data Predictive Control (DPC) to strictly maintain comfort boundaries, achieving nearly 25% cooling energy savings. Building on the concept of coordinating multiple units, Kim and Braun [29] implemented MPC for multiple ON/OFF RTUs in a small retail store, reducing peak demand by 18%. These studies highlight the potential for MPC to deliver significant grid-interactive benefits in smaller buildings without major hardware retrofits.

However, most prior MPC studies for SMCBs have focused on peak demand reduction or traditional event-based DR, rather than predictive load shifting under dynamic pricing. A key technical barrier to successful MPC implementation is the accurate prediction of indoor temperature trajectories in the presence of significant *unmeasured disturbances*-such as internal heat gains from occupants, equipment, and lighting-which are costly or impractical to measure in SMCBs. Gray-box models

are attractive for MPC because they contain building thermal dynamics information and can be embedded in linear optimization formulations [31]. Yet their accuracy deteriorates when unmeasured disturbances are large and unpredictable [32–34], leading to suboptimal control decisions that can cause excessive pre-heating or pre-cooling, eroding cost savings. This challenge mirrors findings in residential energy management, where Blonsky et al. [35] demonstrated through a comparative benchmarking framework that while deterministic MPC with perfect forecasts is optimal, uncertainties significantly degrade performance, necessitating robust strategies to mitigate these errors. Regarding unmeasured disturbances specifically, Dong et al. [36] showed that internal heat gains are a fundamental bottleneck in load forecasting, requiring hybrid approaches that estimate these gains from consumption patterns. Similarly, Dhaliwal et al. [37] found that disturbance forecasting accuracy for unmeasured gains is critical for MPC, noting that misclassifying solar and occupancy gains can lead to severe overheating risks.

To address this challenge, we build upon our earlier low-cost MPC framework for SMCBs [30,38] and propose a machine-learning-enhanced MPC (Hybrid MPC) approach. The key innovation lies in systematically integrating a gray-box thermal model-identified using a lumped disturbance (LD) approach to handle unmeasured disturbances during system identification-with a data-driven machine learning model that forecasts future unmeasured disturbances during MPC optimization.

Despite the extensive body of research on MPC for buildings [18], a critical gap remains in how unmeasured disturbances are handled throughout the control pipeline. Many system identification studies for building thermal models do not explicitly address the unmeasured disturbance problem. Some studies do incorporate measurements of internal heat gains-such as plug loads, occupancy, and lighting-during model identification and MPC optimization. However, such detailed monitoring is typically feasible only in lab-level facilities or highly instrumented testbeds, and is impractical for cost-effective deployment in real SMCBs where metering infrastructure is limited. While some hybrid modeling approaches combining physics-based and data-driven methods have been proposed to correct systematic grey-box errors [39], these have not been systematically validated for MPC applications through experimental studies in real buildings. Furthermore, the performance gap between practical MPC implementations (without disturbance measurements) and research-level scenarios (with measured plug load settings, as in lab-level settings) has not been well characterized, making it difficult to assess the value of disturbance forecasting approaches.

This work makes three main contributions. First, we present a systematic framework that addresses unmeasured disturbances throughout the entire MPC pipeline-from identification to optimization-rather than treating them in isolation or assuming perfect knowledge. Second, we validate the Hybrid MPC through both simulation and experimental studies in an actual SMCB, comparing its performance against baseline control, standard MPC without disturbance forecasting ($\text{MPC}_{\text{vanilla}}$), and an idealized MPC with perfect disturbance knowledge ($\text{MPC}_{\text{ideal}}$). This comparison quantifies the practical value of disturbance forecasting by showing how closely the Hybrid MPC approaches theoretical best performance. Third, we design the controller specifically for SMCBs with multiple independent HVAC units and thermostat-based control, demonstrating its capability to provide both load shifting under dynamic electricity pricing and peak demand reduction while maintaining occupant comfort. Simulation results for both heating and cooling seasons show that the Hybrid MPC achieves cost savings and load flexibility approaching the idealized controller while significantly outperforming standard MPC, confirming the importance of systematic disturbance forecasting in practical building control applications.

The remainder of this paper is organized as follows: Section 2 describes the Hybrid MPC formulation, target building types, and the hybrid modeling approach. Section 3 provides how we evaluate our proposed approach based on simulation and experimental studies. Section 4 presents the simulation setup, scenarios, and results. Section 5 describes the experiment setup, scenarios, and results. Finally, conclusions and future research directions are discussed in Section 7 alongside the Discussion 6.

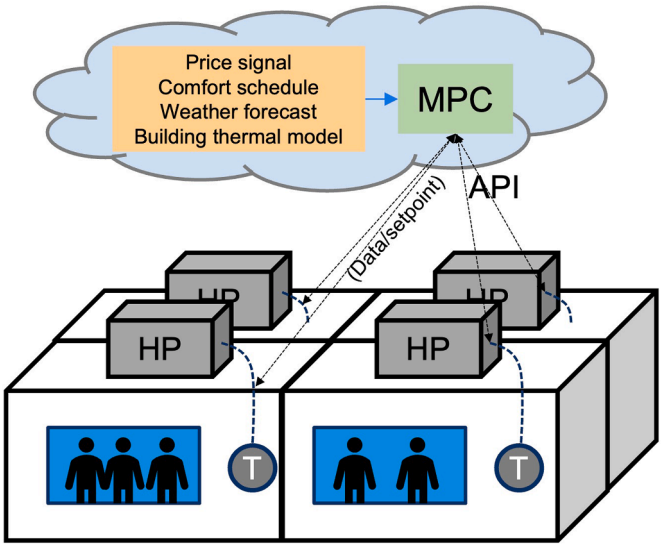

Fig. 1. Schematic diagram of MPC controlling multiple zones in SMCBs.

2. Hybrid MPC Framework

This section summarizes the hybrid MPC framework, which integrates a previously developed MPC formulation with a hybrid modeling approach for building thermal dynamics. We briefly describe the target building infrastructure (Section 2.1), summarize the MPC formulation from prior work (Section 2.2), and present the hybrid modeling approach that combines gray-box models with disturbance prediction (Section 2.3).

2.1. Target SMCBs for MPC

Fig. 1 illustrates the MPC deployment architecture. The MPC targets SMCBs with multiple thermal zones, each equipped with dedicated thermostats and HVAC equipment (RTUs, heat pumps, or mini-splits). The MPC communicates with thermostats via local/cloud APIs or building automation protocols (BACnet, Modbus) to read real-time data and send optimal temperature setpoints based on electricity prices, occupancy schedules, and weather forecasts.

2.2. MPC Formulation

The MPC formulation builds upon the previous work [30,38]. The controller minimizes total energy cost while limiting peak demand and maintaining thermal comfort:

$$\begin{aligned}\min \quad & \sum_{j=1}^{N_p}\sum_{i=1}^{n} ER(k+j-1)\Big(P_{\text{HP,h},i}(k+j-1)\bar{u}_{\text{h},i}(k+j-1) \\ & +P_{\text{HP,c},i}(k+j-1)\bar{u}_{\text{c},i}(k+j-1)\Big) \\ & +\omega_{\text{d}}\delta+\omega_{\text{l}}\Gamma_{\text{l}}+\omega_{\text{u}}\Gamma_{\text{u}} \qquad (1)\\ s.t. \quad & T_{\text{l},i}^{\text{MPC}}(k+j)-\Gamma_{\text{l}} \le \hat{y}_i(k+j|\hat{\mathbf{x}}_i(k)) \le T_{\text{u},i}^{\text{MPC}}(k+j)+\Gamma_{\text{u}} \\ & \sum_{i=1}^{n}\Big(P_{\text{HP,h},i}(k+j-1)\bar{u}_{\text{h},i}(k+j-1) \\ & +P_{\text{HP,c},i}(k+j-1)\bar{u}_{\text{c},i}(k+j-1)\Big) \le \delta \qquad (2)\\ & 0 \le \bar{u}_{\text{h},i}(k+j-1) \le 1, \quad 0 \le \bar{u}_{\text{c},i}(k+j-1) \le 1,\end{aligned}$$

where $\bar{u}_{\mathrm{h},i}$ and $\bar{u}_{\mathrm{c},i}$ are the heating and cooling runtime fractions during the given timestep based on the ON/OFF signals for heat pump i, $P_{\mathrm{HP,h},i}$ and $P_{\mathrm{HP,c},i}$ are the electric power of the ith heat pump at full capacity during heating and cooling operation [kW], which are computed assuming a constant coefficient of performance (COP), such that power is the product of the rated heat pump power and the runtime fraction (i.e., $P_{\mathrm{HP,h},i}\bar{u}_{\mathrm{h},i}$ or $P_{\mathrm{HP,c},i}\bar{u}_{\mathrm{c},i}$). This simplification is reasonable for mild climates where COP variations within the prediction horizon are small, and it enables practical deployment by eliminating the need to collect equipment power data under various outdoor air conditions (See more in Section 6.1). ER is the electricity price [$/kWh], N_p is the prediction horizon, n is the number of heat pumps, $\hat{\bar{y}}_i(k+j|\hat{\mathbf{x}}_i(k))$ is the predicted zone air temperature for zone i, conditioned on the Kalman-filtered state $\hat{\mathbf{x}}_i(k)$, δ is the peak power limit [kW], and Γ_{l}, Γ_{u} are the lower and upper comfort violation slack variables [°C], and ω_{d}, ω_{l}, ω_{u} are the penalty weights on peak demand, lower comfort violation, and upper comfort violation, respectively. At each control timestep k, a Kalman filter estimates the current state $\hat{\mathbf{x}}_i(k)$ from the available measurement data D_k, accounting for process and measurement noise. The predicted temperature is then obtained as $\hat{\bar{y}}_i(k+j|\hat{\mathbf{x}}_i(k)) = \mathbf{C}\hat{\mathbf{x}}_i(k+j)$, where $\mathbf{C}$ is the output matrix of the state-space model and $\hat{\mathbf{x}}_i(k+j)$ is propagated from $\hat{\mathbf{x}}_i(k)$ using the state transition equation. Since the measurement noise is assumed to be white Gaussian, the expectation of the future predicted output equals this deterministic state propagation from the filtered state, i.e., $\mathbb{E}(\bar{y}_i(k+j)|D_k) = \hat{\bar{y}}_i(k+j|\hat{\mathbf{x}}_i(k))$.

$P_{\mathrm{HP,h},i}$ and $P_{\mathrm{HP,c},i}$ are the electric power of the ith heat pump at full capacity during heating and cooling operation [kW],

The comfort boundaries $T_{\mathrm{l},i}^{\mathrm{MPC}}$ and $T_{\mathrm{u},i}^{\mathrm{MPC}}$ are adjusted from the thermostat heating and cooling setpoints ($T_{\mathrm{l},i}$ and $T_{\mathrm{u},i}$) to account for the ON/OFF hysteresis band Δ_{h}:

$$\begin{aligned} T_{\mathrm{l},i}^{\mathrm{MPC}}(k+j) &= T_{\mathrm{l},i}(k+j) - \Delta_{\mathrm{h}}/2 \\ T_{\mathrm{u},i}^{\mathrm{MPC}}(k+j) &= T_{\mathrm{u},i}(k+j) + \Delta_{\mathrm{h}}/2 \end{aligned} \tag{3}$$

This adjustment is necessary because the zone temperature fluctuates within the hysteresis band during ON/OFF cycling, so the time-averaged temperature over a control timestep settles at approximately $T_{\mathrm{l},i}^{\mathrm{MPC}}$ or $T_{\mathrm{u},i}^{\mathrm{MPC}}$ during continuous heating or cooling operation, respectively [30,40].

The optimization yields optimal heat pump schedules that are converted to thermostat setpoints via Eq. 4, where $\hat{\bar{y}}_i^*(k+1|\hat{\mathbf{x}}_i(k))$ denotes the optimal predicted temperature corresponding to the optimal control sequence $\bar{u}_{\mathrm{h},i}^*$ and $\bar{u}_{\mathrm{c},i}^*$:

$$\begin{aligned} T_{\mathrm{csp},i}(k+1) &= \begin{cases} \hat{\bar{y}}_i^*(k+1|\hat{\mathbf{x}}_i(k)), & \text{if } \bar{u}_{\mathrm{c},i}^*(k+1) > 0 \\ T_{\mathrm{u},i}(k+1), & \text{otherwise} \end{cases} \\ T_{\mathrm{hsp},i}(k+1) &= \begin{cases} \hat{\bar{y}}_i^*(k+1|\hat{\mathbf{x}}_i(k)), & \text{if } \bar{u}_{\mathrm{h},i}^*(k+1) > 0 \\ T_{\mathrm{l},i}(k+1), & \text{otherwise} \end{cases} \end{aligned} \tag{4}$$

The MPC solves a linear program where $\bar{u}_{\mathrm{h},i}$ and $\bar{u}_{\mathrm{c},i}$ are continuous variables between 0 and 1, representing the runtime fraction within each control timestep (e.g., 15 or 30 minutes). The MPC does not directly command the ON/OFF signal; instead, the optimal runtime fraction $\bar{u}_{\mathrm{h},i}^*$ or $\bar{u}_{\mathrm{c},i}^*$ determines the expected zone temperature $\hat{\bar{y}}_i^*(k+1|\hat{\mathbf{x}}_i(k))$, which is then assigned as the thermostat setpoint via Eq. 4. The thermostat's built-in deadband controller then cycles the heat pump ON/OFF to maintain this setpoint, naturally achieving the intended runtime fraction. For a variable speed heat pump, the compressor modulates its speed to maintain the setpoint, corresponding to the same runtime fraction. Further details on this setpoint-based MPC implementation are provided in [38].

Fig. 2 illustrates the MPC strategy: pre-heating/cooling during off-peak periods with staggered operation to reduce peak demand, then allowing temperature drift during peak-price periods.

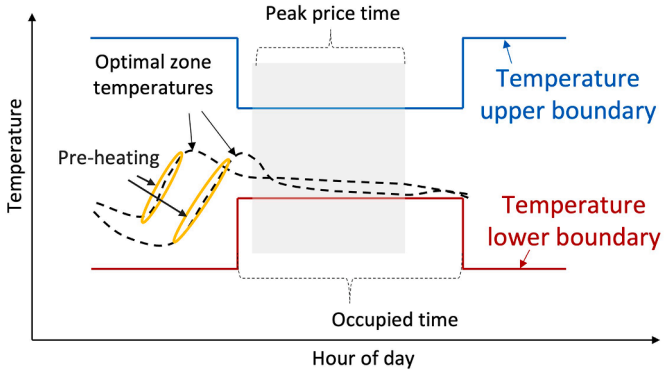

Fig. 2. Conceptual MPC operation showing pre-heating and load shifting (adapted from [38]).

2.3. Hybrid Modeling Approach

2.3.1. Overview

The hybrid modeling approach addresses the challenge of unmeasured disturbances in MPC by combining a physics-based gray-box model with a neural network-based disturbance predictor. The workflow consists of five steps: (1) building thermal parameters are identified offline using a robust system identification method that accounts for unmeasured disturbances (Section 2.3.3), (2) these parameters are used to construct a disturbance model that represents unmeasured heat gains as direct inputs (Section 2.3.4), (3) historical disturbance profiles are inferred from operational data using Kalman filtering (Section 2.3.4), (4) an LSTM neural network is trained to predict future disturbance profiles (Section 2.3.5), and (5) predicted disturbances are substituted into the gray-box model for MPC temperature prediction (Section 2.3.6). Fig. 3 illustrates the comparison between a conventional gray-box-only approach and the proposed hybrid approach.

2.3.2. Gray-Box Building Model

Accurate temperature prediction ($\mathbb{E}(\bar{y}_i(k+j)|D_k)$) is critical for MPC performance. Conventional gray-box models (Fig. 4) capture building thermal dynamics through physical parameters (thermal resistances R, capacitances C) but struggle when unmeasured disturbances (occupancy, plug loads, infiltration) are significant [32]. The gray-box model can be formulated as a system of ordinary differential equations (Eqs. 5-6).

Accurate temperature prediction is critical for MPC performance. The building's thermal dynamics are modeled using a gray-box approach (Fig. 4), which captures the dominant heat transfer physics through thermal resistances (R) and capacitances (C). The model is formulated as a system of ordinary differential equations (Eqs. 5–6).

$$\begin{aligned} \underbrace{\begin{bmatrix} \dot{T}_{\mathrm{w}}(t) \\ \dot{T}_{\mathrm{za}}(t) \end{bmatrix}}_{\dot{\mathbf{x}}} &= \underbrace{\begin{bmatrix} \frac{-1}{C_{\mathrm{w}}R_{\mathrm{zw}}} & \frac{1}{C_{\mathrm{w}}R_{\mathrm{zw}}} \\ \frac{-1}{C_{\mathrm{za}}R_{\mathrm{zw}}} & \frac{-1}{C_{\mathrm{za}}R_{\mathrm{zw}}} + \frac{-1}{C_{\mathrm{za}}R_{\mathrm{zo}}} \end{bmatrix}}_{\mathbf{A}} \underbrace{\begin{bmatrix} T_{\mathrm{w}}(t) \\ T_{\mathrm{za}}(t) \end{bmatrix}}_{\mathbf{x}} \\ &+ \underbrace{\begin{bmatrix} 0 & \frac{(1-f)A_{\mathrm{win}}}{C_{\mathrm{w}}} \\ \frac{1}{C_{\mathrm{za}}R_{\mathrm{zo}}} & \frac{fA_{\mathrm{win}}}{C_{\mathrm{za}}} \end{bmatrix}}_{\mathbf{B}_{\mathrm{w}}} \underbrace{\begin{bmatrix} T_{\mathrm{oa}}(t) \\ \dot{q}_{\mathrm{sol,win}}(t) \end{bmatrix}}_{\mathbf{w}} \\ &+ \underbrace{\begin{bmatrix} 0 & 0 \\ \frac{\dot{Q}_{\mathrm{h}}}{C_{\mathrm{za}}} & \frac{\dot{Q}_{\mathrm{c}}}{C_{\mathrm{za}}} \end{bmatrix}}_{\mathbf{B}_{\mathrm{u}}} \underbrace{\begin{bmatrix} u_{\mathrm{h}}(t) \\ u_{\mathrm{c}}(t) \end{bmatrix}}_{\mathbf{u}} + \underbrace{\begin{bmatrix} 0 \\ \frac{1}{C_{\mathrm{za}}} \end{bmatrix}}_{\mathbf{B}_{\mathrm{g}}} \underbrace{\begin{bmatrix} \dot{Q}_{\mathrm{g}}(t) \end{bmatrix}}_{\dot{\mathbf{Q}}_{\mathrm{g}}} \end{aligned} \tag{5}$$

$$y_{\mathrm{za}} = \underbrace{\begin{bmatrix} 0 & 1 \end{bmatrix}}_{\mathbf{C}} \mathbf{x} \tag{6}$$

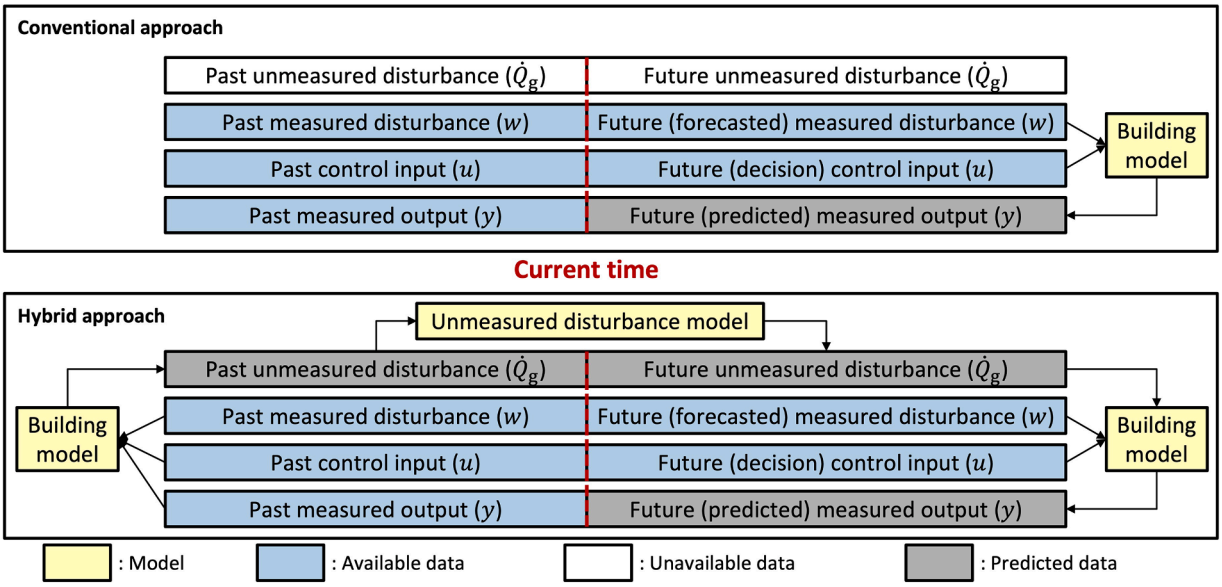

Fig. 3. Comparison of conventional (gray-box only) versus hybrid approach (gray-box + disturbance prediction).

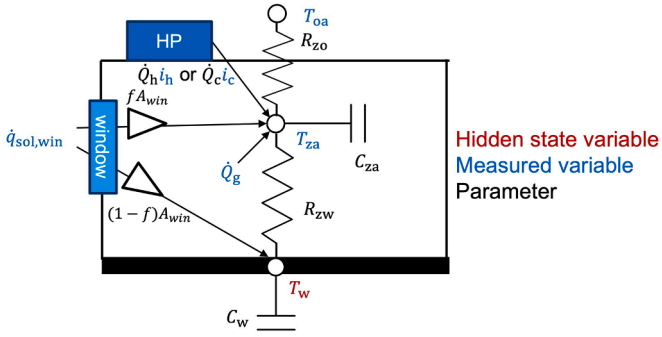

Fig. 4. Gray-box thermal model with two temperature nodes and disturbance terms.

where T_w, T_{za}, and T_{oa} are temperature nodes of large thermal mass (w), zone air (za), and outdoor air (oa), R_{zw} and R_{zo} are thermal resistances between temperature nodes [K/kW], C_w and C_{za} are thermal capacitances [J/K], $\dot{q}_{sol,win}$ is incident solar radiation per area on windows [kW/m^2], A_{win} is effective window area [m^2], f is convective fraction of the incident solar radiation, $\dot{Q}_h$ and $\dot{Q}_c$ are heating and cooling rates at full speed [kW], u_h and u_c are heating and cooling control signals, $\dot{Q}_g$ is a lumped term for all heat gains (e.g., summation of plug, lighting, and occupancy loads) [kW], and y_{za} is measured zone air temperature [°C]. A multi-zone building can be modeled either as an aggregation of multiple single-zone gray-box models or as a single multi-zone gray-box model. Both approaches are applicable to the objective function in Eq. 1. When thermal coupling between zones is negligible, independent single-zone models suffice. However, when inter-zone heat transfer is significant, a thermally coupled multi-zone gray-box model is necessary, as presented in A.

While the gray-box model effectively captures the building's thermal dynamics, even well-calibrated models produce inaccurate future predictions without knowledge of future unmeasured disturbances ($\dot{Q}_g$) [33,34]. The hybrid approach (Fig. 3) addresses this limitation by combining the physics-based model with a data-driven disturbance predictor.

The hybrid approach (Fig. 3) addresses this limitation by combining: (1) a physics-based gray-box model with system parameters identified via output disturbance (OD) method, and (2) a neural network-based model that predicts future input disturbances (ID). Here, ID is used instead of OD for the hybrid model because ID directly represents internal heat gains (easier to interpret and predict). Complete mathematical details and validation results are provided in [33,34].

2.3.3. System Identification via Output Disturbance

To identify the building's thermal parameters (R, C), the unmeasured disturbance $\dot{Q}_g$ must be accounted for during system identification. Fig. 5 illustrates the relationship between the inputs and output of the building envelope transfer function, $G_{bldg} : (\mathbf{w}, \mathbf{u}, \dot{Q}_g) \rightarrow y$. Since $\dot{Q}_g$ is not directly measured in most buildings, system identification methods differ in how they treat this unmeasured disturbance. The output disturbance (OD) approach represents the cumulative effect of unmeasured disturbances through the measurement process (i.e., as a disturbance on the output y), which has been shown to provide more robust parameter estimates compared to conventional approaches that treat unmeasured disturbances as white noise [32,41].

The OD formulation is given by:

$$\begin{aligned} \mathbf{x}(k+1) &= \mathbf{A}_d\mathbf{x}(k) + \mathbf{B}_{w,d}\mathbf{w}(k) + \mathbf{B}_{u,d}\mathbf{u}(k) \\ y(k) &= \mathbf{C}_d\mathbf{x}(k) + \zeta_{OD}(k) + e_{OD}(k) \\ \zeta_{OD}(k+1) &= \rho_1\zeta_{OD}(k) + \rho_2 e_{OD}(k) \end{aligned} \tag{7}$$

where ζ_{OD} captures the cumulative effect of unmeasured disturbances on temperature, and $\mathbf{w}(k) = [T_{oa}(k), \dot{q}_{sol,win}(k)]$ represents measured weather disturbances. The OD approach provides robust building parameters (thermal resistances and capacitances in $\mathbf{A}_d$, $\mathbf{B}_{w,d}$, $\mathbf{B}_{u,d}$) even under significant unmeasured disturbances. Parameters are identified using prediction error minimization with data from active testing experiments (PRBS setpoint variations over unoccupied weekends) [41].

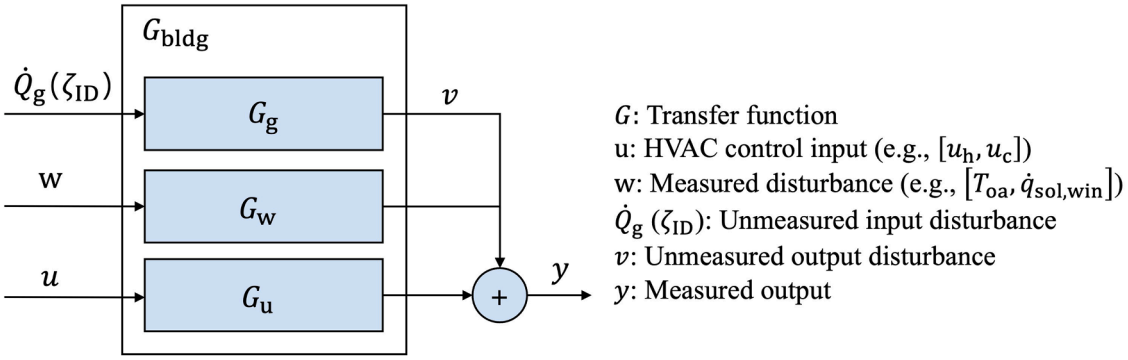

Fig. 5. Relationship between inputs and output of the building envelope transfer function, showing the treatment of unmeasured disturbance ($\dot{Q}_g$) as input disturbance (ID, ζ_{ID}) or output disturbance (OD, ζ_{OD}).

2.3.4. Input Disturbance Model for Prediction

As shown in Fig. 5, the unmeasured disturbance can alternatively be represented as an input disturbance (ID), where it appears before the building system as a direct heat input (ζ_{ID}). While OD is preferred for system identification due to its robustness, the input disturbance (ID) model is more suitable for the hybrid modeling approach because: (1) ID directly represents internal heat gains ($\dot{Q}_g$), making it physically interpretable and easier to predict, and (2) ID profiles are non-cumulative, whereas OD profiles include both the direct impact of heat gains on the temperature and the indirect impact of previous heat gains stored in the thermal mass. Essentially, ID and OD capture the same unmeasured disturbance but at different locations relative to the building system: ID before the system (in units of heat gains [kW]) and OD after the system (in units of temperature [°C]). The ID model uses the same building parameters from OD identification but represents unmeasured disturbances as direct heat inputs:

$$\begin{aligned}\mathbf{x}_{ID}(k+1) &= \mathbf{A}_{d,ID}\mathbf{x}_{ID}(k) + \mathbf{B}_{w,d}\mathbf{w}(k) + \mathbf{B}_{u,d}\mathbf{u}(k)\\ y(k) &= \mathbf{C}_{d,ID}\mathbf{x}_{ID}(k) + e_{ID}(k)\end{aligned} \tag{8}$$

where $\mathbf{x}_{ID}(k) = [\mathbf{x}(k), \zeta_{ID}(k)]$ is the augmented state vector including the ID term ζ_{ID}, which represents direct internal heat gains [kW]. The building-related parameters in $\mathbf{A}_{d,ID}$ and $\mathbf{B}$ matrices are identical to those from OD identification. Historical ID profiles are then inferred from measured operational data by applying a Kalman filter to Eq. 8, which extracts the disturbance signal ζ_{ID} from the augmented state estimate.

2.3.5. Disturbance Prediction Model

To predict future unmeasured disturbances, we train an LSTM neural network that forecasts the ID term (ζ_{ID}). The LSTM then learns weekly disturbance patterns:

$$\hat{\zeta}_{ID}(k+1:k+24) = \text{LSTM}(\text{hour\_of\_week}(k), \mathbf{w}(k), \hat{\zeta}_{ID}(k-23:k)) \tag{9}$$

where the model predicts the next 24 hours of internal heat gains based on time of week, weather conditions, and past disturbance patterns. Training uses at least one month of operational data to capture typical occupancy and usage patterns. To ensure reproducibility, the final disturbance forecasting model employs a specific Long Short-Term Memory (LSTM) network architecture selected via grid search (designated as `DEEP_AR_LSTM_FORECAST-case01` in [34]). The network consists of a single hidden layer ($n_{layer} = 1$) with 20 hidden units ($n_z = 20$) and a ReLU activation function. The input features include one day (24 hours) of historical pattern features (past ζ_{ID}), along with past and future measured weather disturbances (outdoor air temperature T_{oa} and solar radiation $q_{sol,win}$) and the hour of the week (how). Operational data is resampled to hourly intervals, standardized using z-score normalization, and partitioned chronologically: one month is used for training and the subsequent months for testing. The model is trained using the Adam optimizer with a batch size of 32, a Mean Squared Error (MSE) loss function, and an initial learning rate of 0.001 (decaying exponentially by a factor of 0.95 every 50 epochs). To prevent overfitting, the network utilizes dropout layers and an early stopping criterion with a patience of 100 epochs, capped at a maximum of 1000 epochs. By successfully forecasting these unmeasured disturbances, the hybrid model achieves a one-day-ahead temperature prediction root-mean-square error (RMSE) of 0.2 to 0.35°C in simulations, and 0.3 to 0.6°C in the experimental validation on test dataset.

2.3.6. Integration for MPC

The hybrid modeling framework operates through the following workflow: (1) building parameters are identified offline using the robust OD approach (Eq. 7), (2) these parameters are used to construct the ID model (Eq. 8) for prediction, (3) historical ID profiles (ζ_{ID}) are inferred from operational data using Kalman filtering, (4) the LSTM disturbance predictor is trained on these ID profiles, and (5) future temperatures are predicted by substituting the predicted unmeasured disturbance as $\dot{Q}_g$ in the original gray-box formulation (Eq. 5). Complete mathematical details and validation results are provided in [33,34].

3. Evaluation Methodology

To validate the proposed Hybrid MPC framework, we conduct both simulation studies and field experiments. This section describes the rationale for our evaluation methodology, the selection of test systems, and the performance metrics used.

3.1. Evaluation Approach and Test System Selection

We evaluate the Hybrid MPC through both simulation and experimental studies to provide complementary insights. The simulation study enables systematic comparison under controlled conditions with perfect knowledge of ground truth, allowing us to isolate the impact of disturbance prediction accuracy on MPC performance. The experimental study validates real-world applicability and reveals practical implementation challenges not captured in simulation.

For the simulation study (Section 4), we developed a building model representing a typical SMCB with nine independent thermal zones. This configuration was selected to: (1) demonstrate the scalability of the approach to buildings with multiple HVAC units, (2) evaluate peak demand reduction through coordinated control, and (3) test the framework under realistic internal load variability across zones.

For the experimental study (Section 5), we selected the FLEXLAB testbed at Lawrence Berkeley National Laboratory [42]. FLEXLAB provides a controlled yet realistic environment with: (1) the ability to directly measure internal loads (plug, lighting), enabling us to provide values similar to the ground truth of unmeasured disturbances, (2) precise control over experimental conditions while maintaining realistic building characteristics, and (3) well-instrumented mechanical systems

Table 1
Specifications of the Building, HVAC System, and Baseline Schedule.

Category	Parameter	Value(s)
Building	Location	Berkeley, CA, USA (9 rooms)
	Floor Area (per building)	1076 ft^2 (100 m^2)
	Internal Loads	5 occupants (75 W/person), 2 kW peak internal load [a](a)
	Fenestration	One South-facing window (3 m^2) per building
HVAC System	Type	Split-system Heat Pump
	Control	Thermostat (ON/OFF)
	Rated Heating Capacity/Power	10 kW ($\dot{Q}_h$) / 3.4 kW (P_h)
	Rated Cooling Capacity/Power	-7.5 kW ($\dot{Q}_c$) / 3.0 kW (P_c)
Baseline Schedule	Occupied Hours	6:00 - 21:00
	Occupied Temperature Range	20.0 - 23.3 °C (68 - 74 °F)
	Unoccupied Temperature Range	15.6 - 30.0 °C (60 - 86 °F)

[a] (a) Internal loads include all lighting, equipment, occupant heat gains, and infiltration. For the simulation environment, we added random noise ranging from 0 to 20% of the total internal loads.

that provide high-quality data for validation of the proposed approach. More details can be found in Section 5.

3.2. Performance Metrics

We use the following metrics to quantify MPC performance:

- **Energy Cost [$/day or $/month]:** Total electricity cost under the dynamic price signal, calculated as $\sum_t P(t) \cdot ER(t) \cdot \Delta t$ where $P(t)$ is power demand, $ER(t)$ is electricity price, and Δt is the time interval.
- **Cost Savings [%]:** Percentage reduction in energy cost relative to the Baseline control.
- **Load Shifting [%]:** Percentage reduction in energy consumption during peak-price periods (6-8 AM for heating, 5-9 PM for cooling), calculated as: $\frac{E_{\text{baseline,peak}} - E_{\text{MPC,peak}}}{E_{\text{baseline,peak}}} \times 100\%$.
- **Peak Demand Reduction** [%]: Percentage reduction in maximum building power demand, calculated as: $\frac{P_{\text{max, baseline}} - P_{\text{max,MPC}}}{P_{\text{max,baseline}}} \times 100\%$.
- **HVAC Cycling [cycles/day]:** Total number of on/off state transitions across all HVAC units and operating modes (heating and cooling) per day, used to assess the potential for equipment degradation due to increased cycling under MPC control.
- **Comfort Violation [°C·h/day]:** Daily average cumulative temperature deviation beyond the comfort boundaries (with a 0.5°C hysteresis band), summed across all zones: $\sum_z \sum_t \left[\Delta T^-_{z,t} + \Delta T^+_{z,t}\right] \cdot \Delta t$, where $\Delta T^-_{z,t} = \max(0,\ T_l - \Delta_h - T_{za})$ and $\Delta T^+_{z,t} = \max(0,\ T_{za} - T_u - \Delta_h)$ with $\Delta_h = 0.5$°C.

All metrics are computed while maintaining thermal comfort constraints (temperatures within the specified comfort bounds during occupied hours). The comparison between $\text{MPC}_{\text{vanilla}}$, $\text{MPC}_{\text{ideal}}$, and $\text{MPC}_{\text{hybrid}}$ isolates the contribution of disturbance prediction to overall performance.

4. Simulation study

4.1. Case study building

For the simulation study, a hypothetical building consisting of 9 thermally disconnected rooms was created. Each room is served by a heat pump (HP) and a thermostat, and its thermal behavior is modeled using the 2R2C model described in Section 2.3. The parameter values for each room are: $C_w = 4$ kWh/K, $C_z = 1$ kWh/K, $R_{zw} = 1.2$ K/kW, $R_{zo} = 6$ K/kW, $f = 0.3$, and $A_{win} = 3$ m^2. HVAC-related parameters and other building specifications are summarized in Table 1.

Fig. 6 illustrates the building's heating operation on a typical winter day. The top plot displays temperature profiles for individual rooms alongside the lower comfort boundary (heating setpoint minus the thermostat deadband). The outdoor air temperature is also shown on the left y-axis. The bottom plot visualizes the total number of heat pump (HP) operations along with the dynamic electricity price signal (see Section 4.2). Individual HP ON/OFF signals are also shown on separate lines. During unoccupied nighttime hours, the building cools down due to cold weather, and HP operation begins at the start of the occupied period (i.e., 6 AM). After 9 AM, internal gains maintain the building temperature within the comfort zone without heating operation.

4.2. MPC simulation scenario

The performance of the Baseline control, summarized in Table 1, is compared with that of MPC for both heating and cooling months using the simulated building model. To evaluate the impact of the hybrid modeling approach, we develop and analyze three different MPC variants: $\text{MPC}_{\text{vanilla}}$, $\text{MPC}_{\text{ideal}}$, and $\text{MPC}_{\text{hybrid}}$. All three MPCs share the same underlying building model structure but differ in how they identify the building model and handle future unmeasured disturbances during the optimization process.

$\text{MPC}_{\text{vanilla}}$ employs the building model identified through the SYSID procedure described in Section 2.3. This identification is based on data collected during active testing but excludes any information about unmeasured disturbances. Consequently, $\text{MPC}_{\text{vanilla}}$ solves the optimization problem using only the building model dynamics, considering external inputs such as weather conditions, solar radiation, and building envelope characteristics, but without accounting for future unmeasured disturbances.

$\text{MPC}_{\text{ideal}}$ represents an idealized version of the controller that utilizes unmeasured disturbance information both during SYSID and in the MPC optimization. The primary purpose of $\text{MPC}_{\text{ideal}}$ is to provide a performance benchmark, indicating the upper bound achievable by $\text{MPC}_{\text{vanilla}}$ and $\text{MPC}_{\text{hybrid}}$. It is important to note that unmeasured disturbance information is generally unavailable in practice; however, since this study is simulation-based, the disturbance data are accessible and used directly. As a result, the building model obtained from SYSID closely matches the true simulated building model described in Section 4.1. Moreover, during optimization, $\text{MPC}_{\text{ideal}}$ incorporates the actual unmeasured disturbance forecasts.

$\text{MPC}_{\text{hybrid}}$ leverages the hybrid modeling approach. It uses the same building model as $\text{MPC}_{\text{vanilla}}$ but enhances the optimization by estimating future unmeasured disturbances via the hybrid model. Unlike $\text{MPC}_{\text{ideal}}$, which has direct access to true disturbance data, $\text{MPC}_{\text{hybrid}}$ predicts future unmeasured disturbances at each MPC timestep based solely on available historical information.

Several tuning parameters related to the MPC influence its performance. The primary objective of the MPC is to achieve load shifting-shifting consumption from high-price periods to low-price periods-with peak demand reduction as a secondary goal. The parameter (ω_d) in Eq. 1 determines the weight assigned to peak demand reduction over each

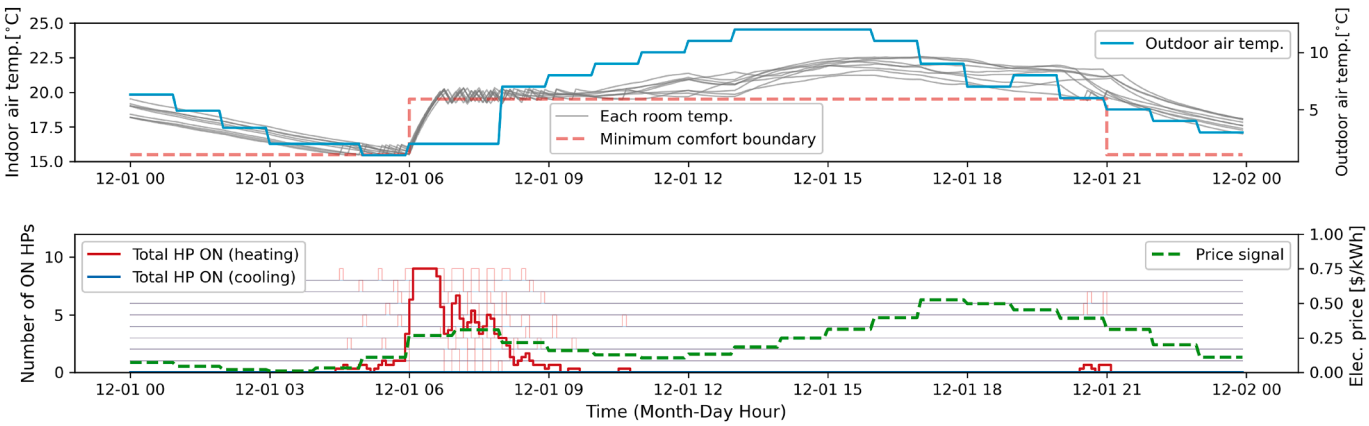

Fig. 6. Typical heating day operation (Baseline).

prediction horizon. A higher value of (ω_d) encourages the MPC to reduce the simultaneous operation of multiple heat pumps (HPs) through pre-heating or pre-cooling. However, excessive pre-heating or pre-cooling can lead to increased energy costs. Therefore, ω_d is typically chosen empirically to balance energy cost and peak demand reduction; a small value is generally sufficient when load shifting is the primary goal, since larger values aggressively curtail simultaneous device operation at the expense of cost savings (see B for a systematic sensitivity analysis and Section 6.3). In this study, ω_d is set to 0.1 based on the sensitivity analysis, which shows that this value preserves the majority of peak-time and cost savings while still providing meaningful peak demand reduction.

Additionally, the parameter (δ), which controls the allowable fraction of devices that can be ON during each timestep, is set to 0.7. This limits heating or cooling operations for each timestep to approximately 70% of all devices' operation. The sensitivity analysis (B) confirms that δ has a relatively minor effect on peak-time and cost savings but influences peak demand reduction; $\delta = 0.7$ offers a balanced trade-off without noticeably degrading cost performance The parameters ω_l and ω_u represent weights for enforcing comfort boundaries, and are assigned a high value (e.g., 1000) to ensure the MPC strictly maintains temperatures within the comfort range.

The MPC's comfort boundaries are defined based on the heating and cooling setpoints used in the Baseline control. Setpoints are updated every 15 minutes, a sampling interval chosen to align with typical times required to reach setpoints and the demand charge time window. For the electricity price signal, this study uses the Highly Dynamic Price (HDP) signal developed by Lawrence Berkeley National Laboratory (LBNL) [43] (see Fig. 6). The HDP signal provides a realistic, though simulated, representation of how electricity prices may fluctuate based on grid conditions, encouraging responsive load management.

4.3. MPC simulation results

In Fig. 7, we illustrate the typical heating operation scenarios for four different control algorithms: Baseline, $MPC_{vanilla}$, MPC_{ideal}, and MPC_{hybrid}. This figure provides a comparative overview of how each control strategy manages heating over a specified period during the heating season. Note that the simulation is run continuously across the entire season, so each day's initial zone temperatures are determined by the previous day's operation under each respective control strategy. The snapshots shown therefore reflect genuine differences in thermal state resulting from each controller's behavior. The y-axis bounds have been made consistent across all subplots for easier comparison.

Fig. 7 (a) depicts the Baseline control strategy, which operates all heat pumps (HPs) at the same time during morning peak hours (i.e., when prices are high), specifically from 6:00 to 8:00. This is done to meet the rise in the heating setpoint as the building transitions from an unoccupied to an occupied state. Occasional cooling operations are also observed during the afternoon when solar and internal heat gains cause select zones to exceed the upper comfort boundary.

In contrast to the Baseline control strategy, the MPC approaches-specifically $MPC_{vanilla}$, MPC_{ideal}, and MPC_{hybrid}-utilize pre-heating to avoid the abrupt and simultaneous heating operation during the morning peak hours. Overall, the three MPC variants exhibit qualitatively similar load-shifting behavior: all begin pre-heating during the early morning off-peak hours and successfully avoid the simultaneous HP operation observed in the baseline around 6:00–9:00. However, important differences are observed in the details. The $MPC_{vanilla}$ initiates the most aggressive pre-heating, starting around midnight, because its prediction model does not account for unmeasured internal heat gains and therefore overestimates the heating demand. As a result, zone temperatures are driven higher than necessary before occupancy begins. The MPC_{ideal} maintains zone temperatures most strictly within the comfort boundaries by leveraging an accurate building model and disturbance predictions; it applies slightly more pre-heating than MPC_{hybrid} to ensure that all zones reach the occupied setpoint precisely by 6:00. The MPC_{hybrid}, which relies on a data-driven prediction that may over-predict internal gains on certain days, applies somewhat less pre-heating, and consequently a few zones fall slightly below the comfort boundary at the start of occupied hours. While MPC_{ideal} and $MPC_{vanilla}$ achieve tighter temperature regulation, this comes at the cost of higher compressor cycling, as both controllers frequently switch HPs on and off to maintain temperatures close to the setpoint throughout the day. MPC_{hybrid}, by contrast, exhibits smoother HP operation with fewer on/off transitions, reflecting the smoothing effect of its less precise predictions.

Fig. 8 illustrates the comparison of operational performance of the Baseline controller alongside three MPC strategies-vanilla, ideal, and hybrid-representing the average performance during the heating period. The Baseline controller exhibits a schedule-based approach, initiating the heat pump precisely at 6 AM to meet the scheduled increase in the heating setpoint. This results in a significant power demand spike that directly coincides with the morning peak electricity price period from 6 AM to 8 AM. In contrast, all MPC strategies proactively leverage the building's thermal inertia by pre-heating the space during the cheaper, off-peak hours before 6 AM. This strategy effectively shifts a substantial portion of the heating load away from the high-cost period to the earlier low-cost period.

An analysis of the economic and load-shifting performance reveals that all three MPC strategies reduce both peak-time load and daily energy cost relative to the Baseline. The $MPC_{vanilla}$ controller achieved the most aggressive peak-time load reduction (85.0%) and the highest cost saving (34.3%, $9.20/day); unlike in the cooling season where its performance degrades significantly, the vanilla controller performs competitively here because the unmeasured internal heat gains are

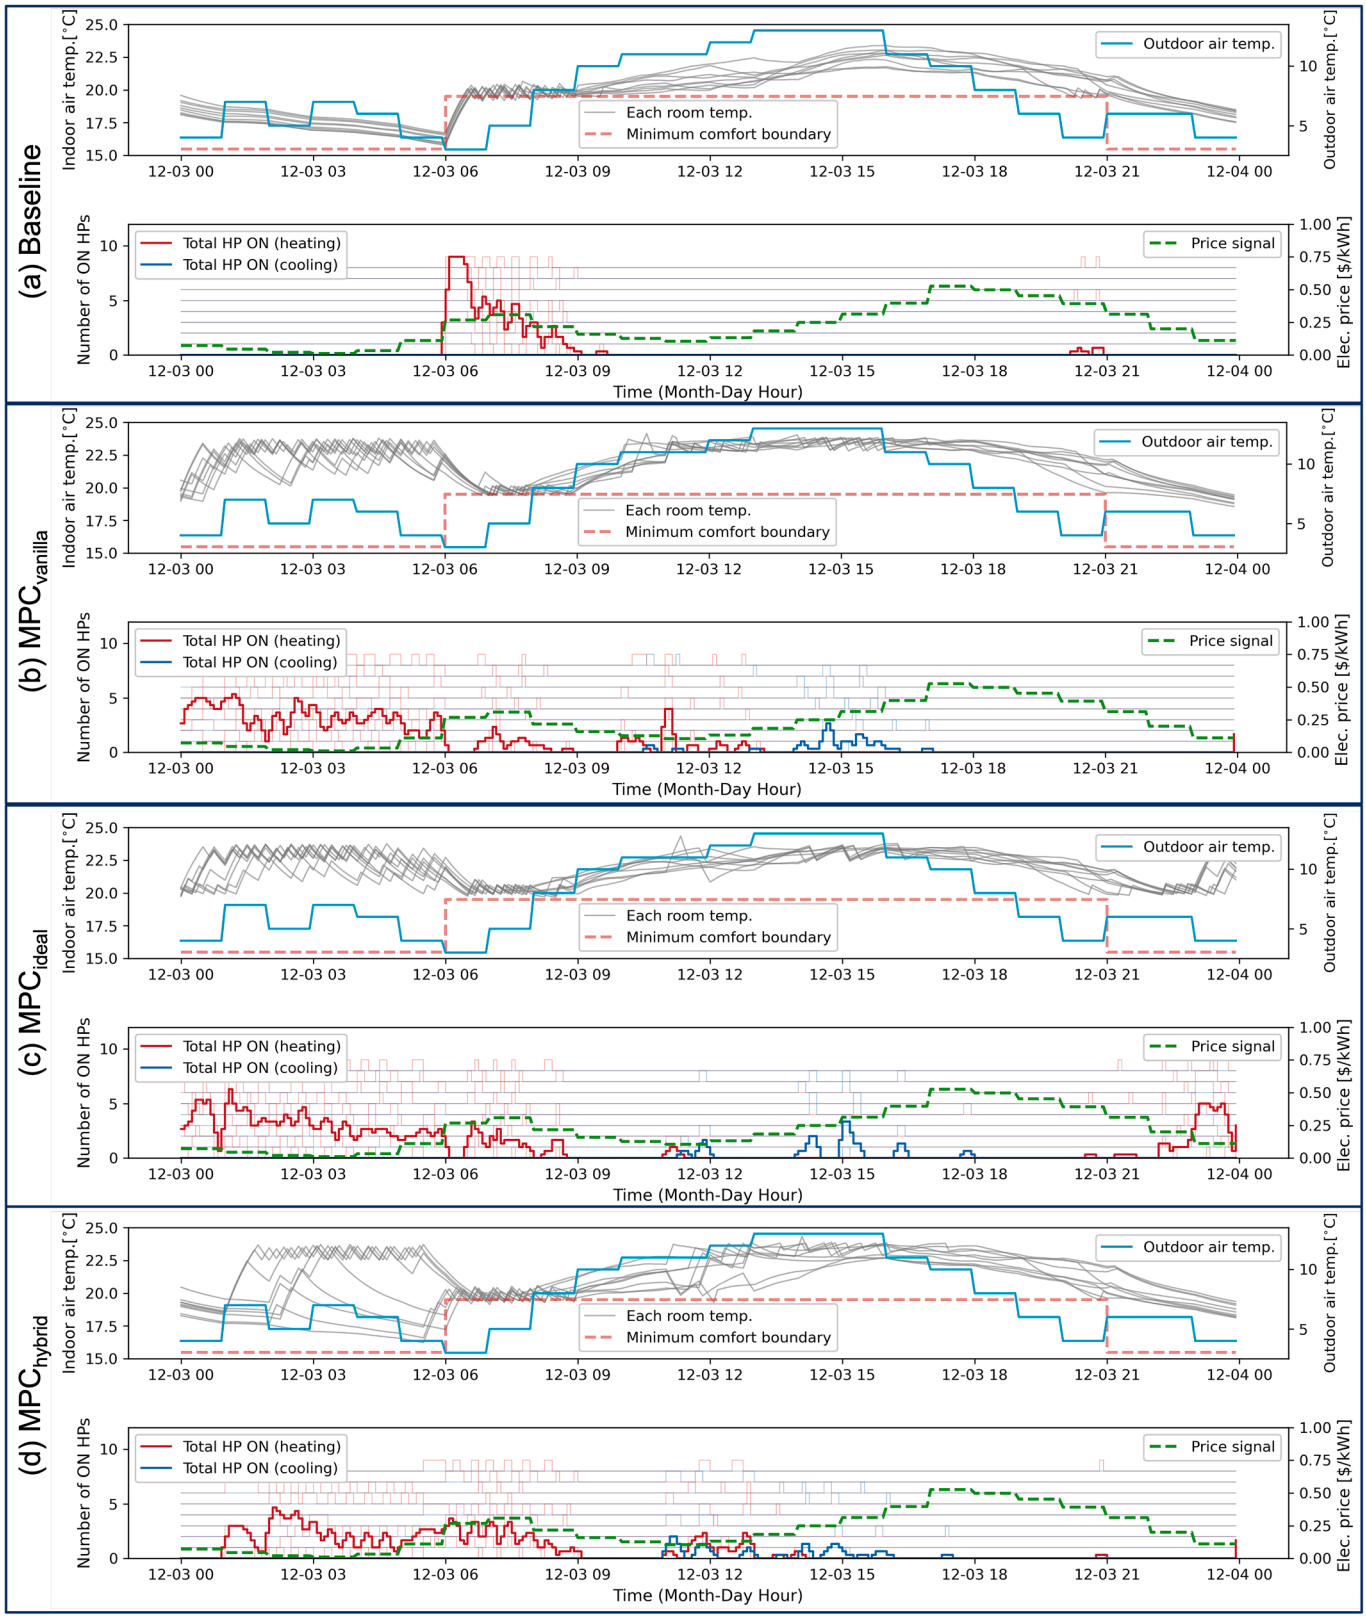

Fig. 7. Typical heating day operations of (a) Baseline, (b) $MPC_{vanilla}$, (c) MPC_{ideal}, (d) MPC_{hybrid}.

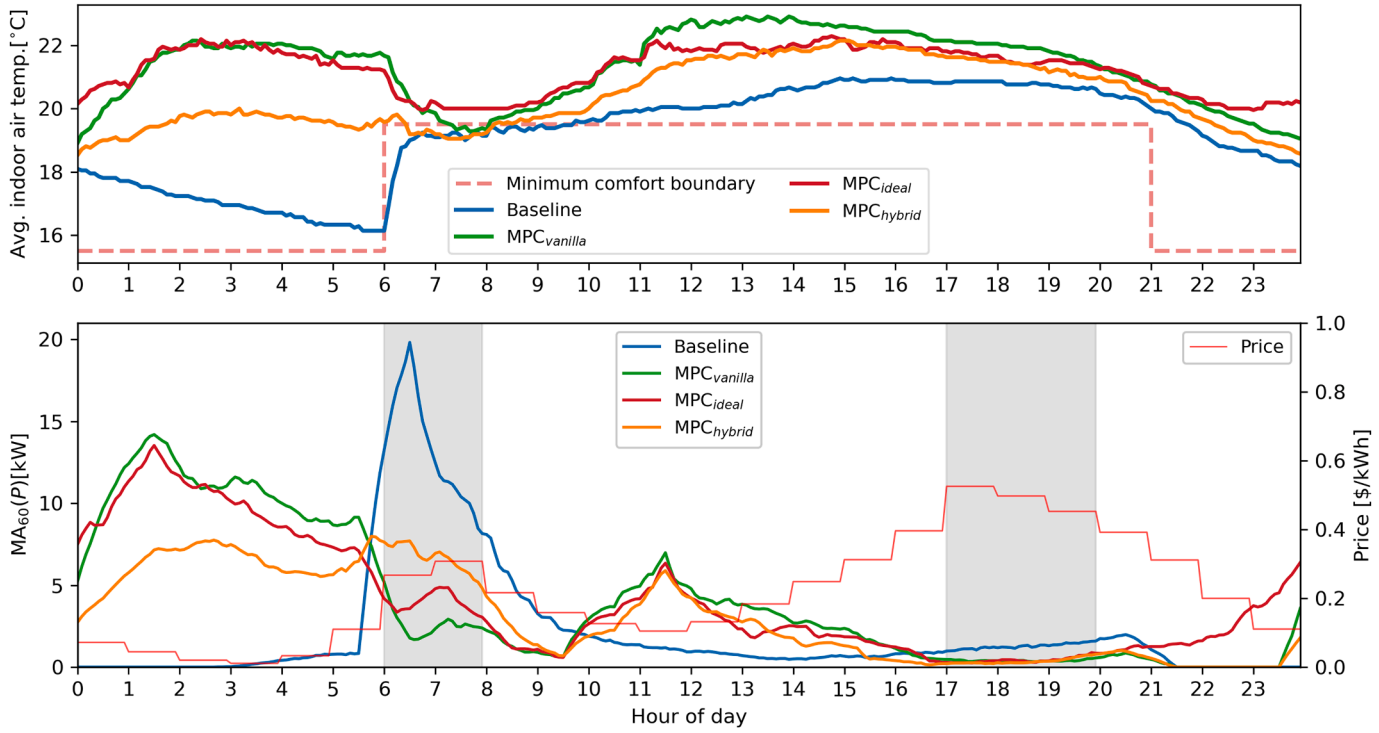

Fig. 8. Average operational profiles during the heating period. **Top:** Indoor air temperature profiles for each control strategy compared to the minimum comfort boundary. **Bottom:** Heat pump power consumption and the hourly electricity price (the power consumption is 60-minute averaged to visualize hourly load characteristics.)

relatively small during the early-morning pre-heating hours. Nevertheless, $MPC_{vanilla}$ overheats the building during the afternoon when occupant and equipment heat gains increase, resulting in the highest average indoor temperature. The MPC_{ideal} reduced peak-time load by 74.4% and daily cost by 24.4% ($10.59/day), while achieving the best comfort performance with only 0.10 °C·h/day of comfort violation—a 98.6% reduction relative to the Baseline (7.0 °C·h/day). Its temperature profile remains above the minimum comfort boundary during the morning transition and stays moderate during the afternoon, reflecting the benefit of accurate disturbance predictions. The MPC_{hybrid} achieved a comparable cost saving of 31.5% ($9.59/day) with a 53.9% peak-time reduction. However, due to limited prediction accuracy, its pre-heating was the smallest among the three MPC variants, causing some zones to slightly violate the minimum comfort boundary at the onset of occupancy (0.85 °C·h/day). On the other hand, MPC_{hybrid} avoided the afternoon over-heating observed in the vanilla case. Regarding HVAC cycling, all MPC strategies increased the number of daily on/off transitions relative to the Baseline (207 cycles/day). This is because the MPC controllers set the zone temperature targets near the upper allowable setpoint to maximize pre-heating, which causes more frequent ON/OFF switching around that setpoint. MPC_{ideal} exhibited the largest increase (65.1%, 342 cycles/day), followed by $MPC_{vanilla}$ (47.4%, 306 cycles/day) and MPC_{hybrid} (18.8%, 246 cycles/day), consistent with their respective degrees of pre-heating aggressiveness. These findings confirm that prediction accuracy—particularly of unmeasured disturbances—governs the trade-off between comfort maintenance, over-heating avoidance, and load-shifting aggressiveness. The comparison metrics are all summarized in Table 2.

Beyond average performance, examining the peak demand during the most extreme heating day, as shown in Fig. 9, highlights the effectiveness of MPC in mitigating grid strain. The Baseline controller showed a peak 15-minute demand of 30.6 kW, driven by the simultaneous activation of all heat pumps at 6 AM. All three MPC strategies reduced this peak by distributing the heating load across off-peak hours. The $MPC_{vanilla}$ strategy, due to its extensive pre-heating, achieved a 14.8% reduction (26.1 kW). The MPC_{ideal} achieved a 16.7% reduction (25.5 kW), while MPC_{hybrid} achieved the largest peak demand reduction (22.2%, 23.8 kW). The comparison metrics are all summarized in Table 2.

In Fig. 10, we illustrate the typical cooling operation scenarios for four different control algorithms: Baseline, $MPC_{vanilla}$, MPC_{ideal}, and MPC_{hybrid}. This figure provides a comparative overview of how each control strategy manages cooling over a specified period during the cooling season.

Fig. 10 (a) depicts the Baseline control strategy, which operates reactively, activating the cooling system as indoor temperatures rise based on the scheduled cooling setpoints. With increased outdoor air temperature and solar radiation, it showed higher cooling energy consumption during the afternoon periods. Unlike the heating season, which has a significant morning peak, there is no specific peak operation time during the occupied period because there are no abrupt setpoint changes that would cause a simultaneous operation of all cooling units.

In contrast to the Baseline's reactive approach, the MPC strategies attempt to pre-cool the space to shift load away from the afternoon peak price hours to some extent. However, the potential for significant load shifting is inherently limited. The air-based system cannot store thermal energy for extended periods as effectively as a water-based system can, making it difficult to shift cooling loads for long durations (e.g., from 5-9 p.m. to 9-11 a.m.). In contrast, shifting a heating load is easier as the required time shift is much shorter (e.g., from 6-8 a.m. to 4-6 a.m.). Consequently, the benefits of pre-cooling are less pronounced in visual inspection, though all MPC strategies performed some pre-cooling, as shown by the temperature fluctuations between 9 a.m. and 2 p.m.

Fig. 11 illustrates the comparison of the average operational performance of the four control strategies during the cooling season. Unlike the heating season, the peak electricity prices occur in the afternoon (5 PM to 9 PM), when cooling loads are mild. An analysis of the daily average performance shows that all MPC strategies successfully

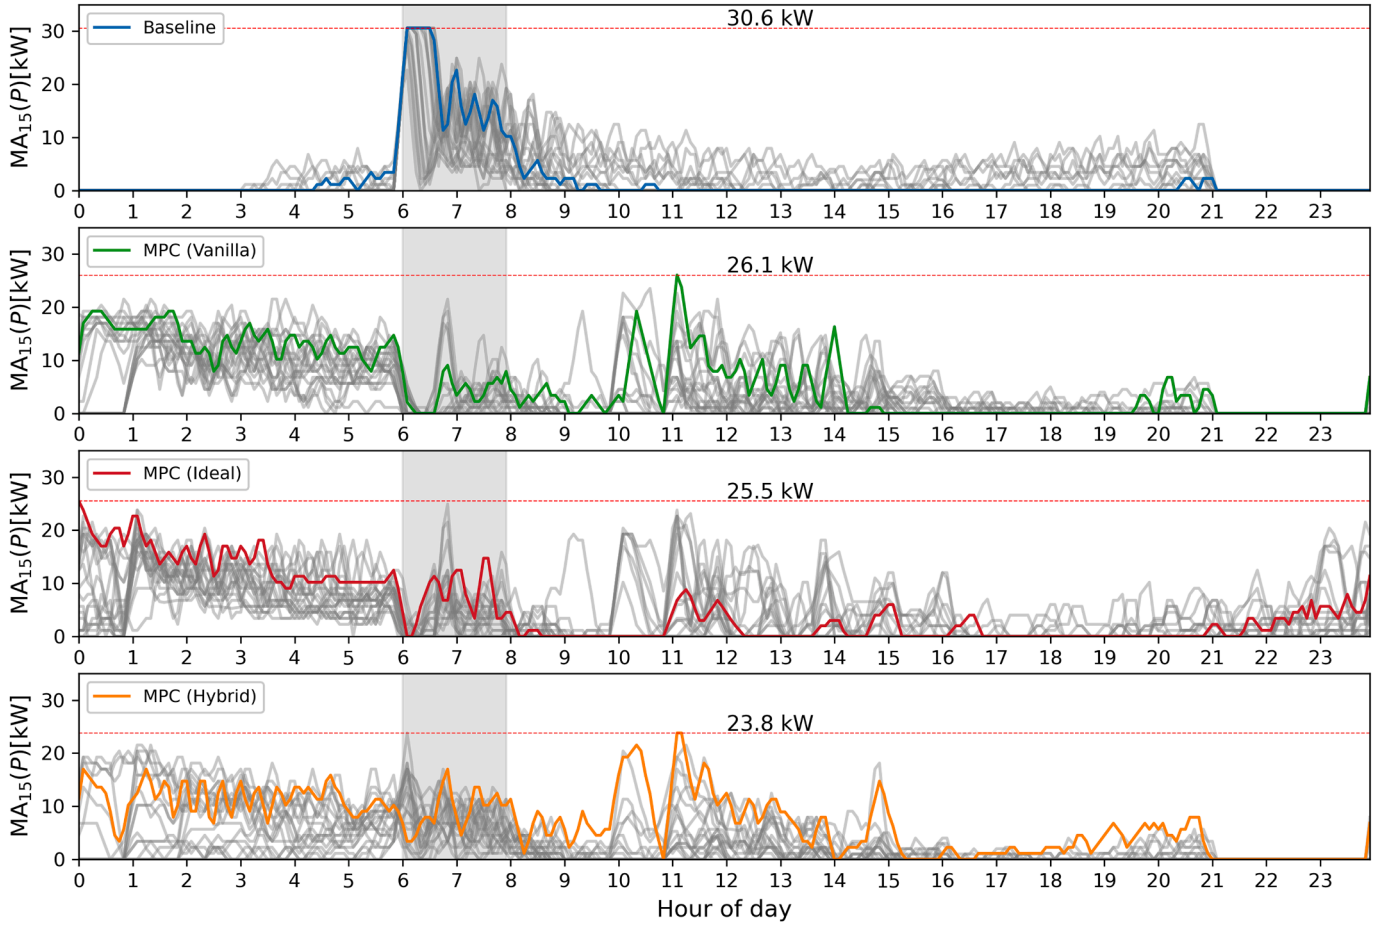

Fig. 9. Heat pump power demand profiles for all individual days during the heating season. The day with the highest peak demand for each strategy is highlighted in color.

reduced both energy costs and peak-period load compared to the Baseline. The Baseline controller incurred a daily cost of $18.88. The MPC_{ideal} strategy was the most effective in comfort performance, achieving only 0.10 °C·h/day of comfort violation—a 90.3% reduction relative to the Baseline (1.03 °C·h/day)—while, reducing the daily cost to $16.22 (14.1% saving) and shifting 30.5% of the cooling load away from the 5-9 PM peak price window. The MPC_{hybrid} controller performed similarly, with a daily cost of $13.85 (26.6% saving) and a peak-time load reduction of 38.4%, while maintaining a comfort violation of 0.43 °C·h/day. While the $MPC_{vanilla}$ also provided savings with a daily cost of $17.25 (8.7% saving), it was the least effective among the three MPC variants, achieving only a 10.0% peak-time load reduction and a comfort violation of 0.66 °C·h/day; this degraded performance relative to the heating season is attributed to the larger influence of unmeasured disturbances (solar radiation, occupant and equipment heat gains) during the cooling period, which the vanilla model cannot capture. Regarding HVAC cycling, all MPC strategies increased the daily on/off transitions relative to the Baseline (232 cycles/day), with MPC_{ideal} showing the largest increase (54.4%, 359 cycles/day), followed by MPC_{hybrid} (17.4%, 273 cycles/day) and $MPC_{vanilla}$ (12.3%, 261 cycles/day). The comparison metrics are all summarized in Table 2.

An analysis of the peak demand, illustrated in Fig. 12, highlights some of the benefits of the MPC controllers. The Baseline strategy resulted in a maximum peak demand of 30.6 kW during the cooling season. In comparison, all MPC strategies successfully lowered this peak. The $MPC_{vanilla}$ controller achieved a 31.4% reduction to 21.0 kW. The MPC_{ideal} controller reduced the peak to 23.0 kW (24.8% reduction), while MPC_{hybrid} achieved a 28.1% reduction to 22.0 kW. The modest nature of this peak reduction is rooted in the fundamental characteristics of cooling operations. Unlike heating demand, which often has a predictable, sharp peak due to synchronized system start-ups in the morning, the cooling load is distributed more broadly throughout the afternoon. As a result, periods of peak simultaneous operation occur more stochastically. While the MPC controller can mitigate these random peaks, the lack of a distinct, universal trigger makes it more challenging to achieve the same level of substantial, targeted peak shaving seen in heating scenarios. The comparison metrics are all summarized in Table 2.

In Fig. 10, we illustrate the typical cooling operation scenarios for four different control algorithms: Baseline, $MPC_{vanilla}$, MPC_{ideal}, and MPC_{hybrid}. This figure provides a comparative overview of how each control strategy manages cooling over a specified period during the cooling season.

Fig. 10 (a) depicts the Baseline control strategy, which operates reactively, activating the cooling system as indoor temperatures rise based on the scheduled cooling setpoints. With increased outdoor air temperature and solar radiation, it showed higher cooling energy consumption during the afternoon periods. Unlike the heating season, which has a significant morning peak, there is no specific peak operation time during the occupied period because there are no abrupt setpoint changes that would cause a simultaneous operation of all cooling units.

In contrast to the Baseline's reactive approach, the MPC strategies attempt to pre-cool the space to shift load away from the afternoon peak price hours to some extent. However, the potential for significant load shifting is inherently limited. The air-based system cannot store

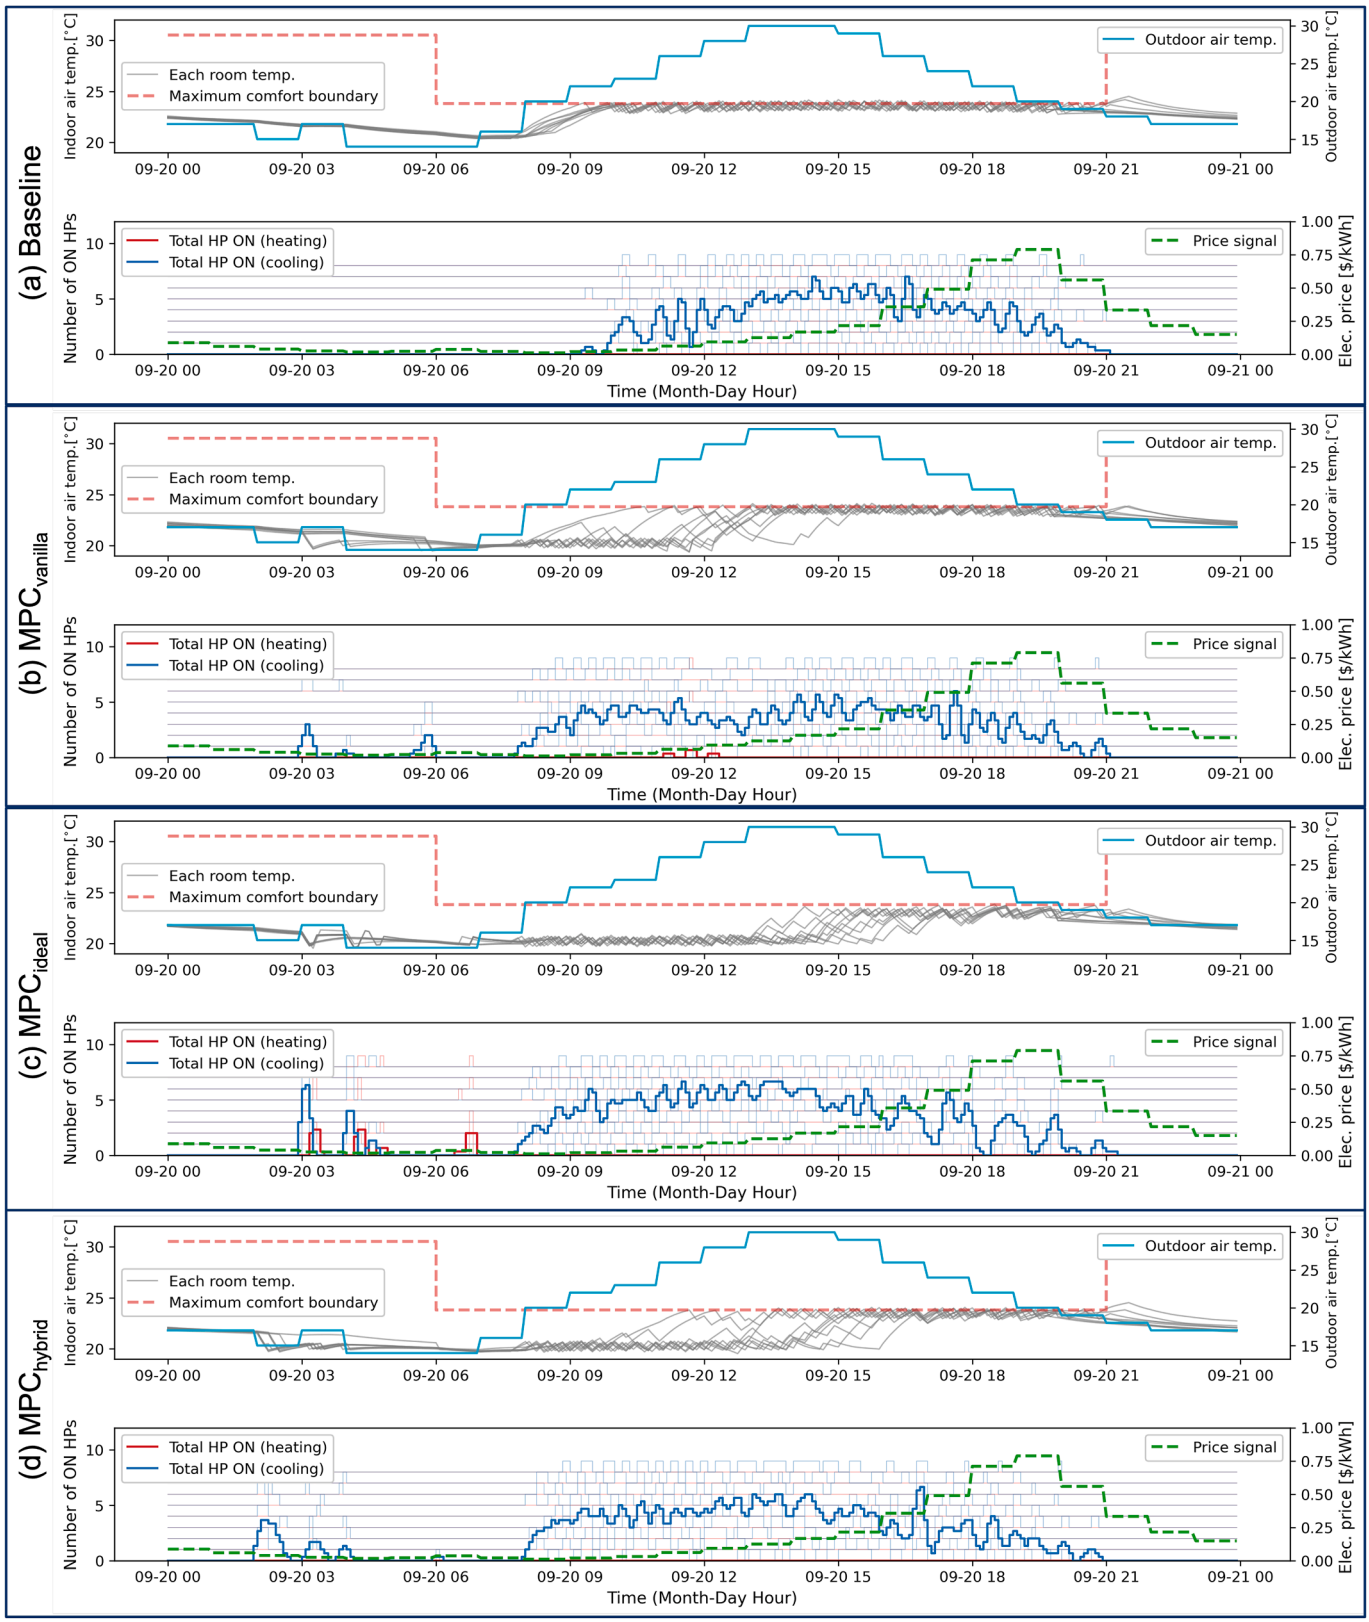

Fig. 10. Typical cooling day operations of (a) Baseline, (b) $MPC_{vanilla}$, (c) MPC_{ideal}, (d) MPC_{hybrid}.

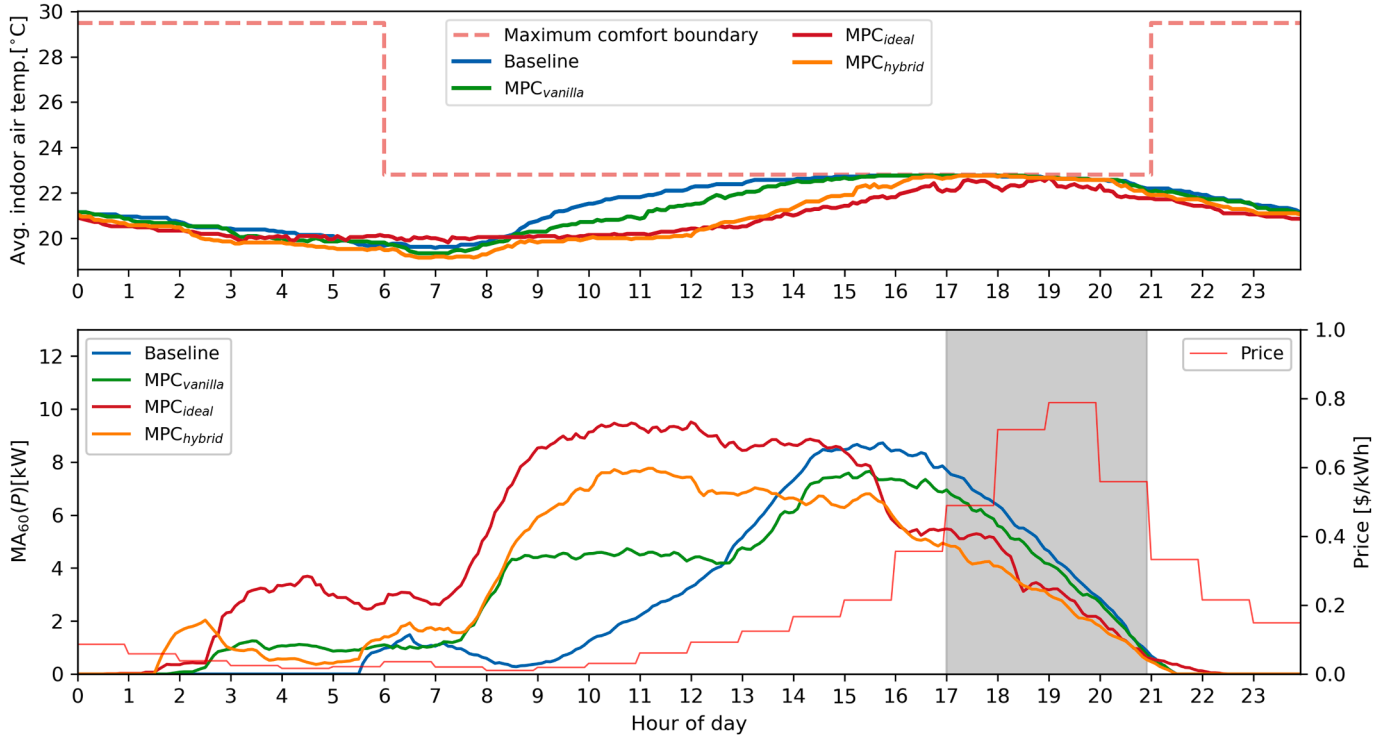

Fig. 11. Average operational profiles during the cooling period. **Top:** Indoor air temperature profiles for each control strategy compared to the maximum comfort boundary. **Bottom:** Heat pump power consumption and the hourly electricity price (the power consumption is 60-minute averaged to visualize hourly load characteristics.)

thermal energy for extended periods as effectively as a water-based system can, making it difficult to shift cooling loads for long durations (e.g., from 5-9 p.m. to 9-11 a.m.). In contrast, shifting a heating load is easier as the required time shift is much shorter (e.g., from 6-8 a.m. to 4-6 a.m.). Consequently, the benefits of pre-cooling are less pronounced in visual inspection, though all MPC strategies performed some pre-cooling, as shown by the temperature fluctuations between 9 a.m. and 2 p.m.

Fig. 11 illustrates the comparison of the average operational performance of the four control strategies during the cooling season. Unlike the heating season, the peak electricity prices occur in the afternoon (5 PM to 9 PM), when cooling loads are mild. An analysis of the daily average performance shows that all MPC strategies successfully reduced both energy costs and peak-period load compared to the Baseline. The Baseline controller incurred a daily cost of $18.88. The MPC_{ideal} strategy was the most effective in comfort performance, achieving only 0.10 °C·h/day of comfort violation—a 90.3% reduction relative to the Baseline (1.03 °C·h/day)—while, reducing the daily cost to $16.22 (14.1% saving) and shifting 30.5% of the cooling load away from the 5-9 PM peak price window. The MPC_{hybrid} controller performed similarly, with a daily cost of $13.85 (26.6% saving) and a peak-time load reduction of 38.4%, while maintaining a comfort violation of 0.43 °C·h/day. While the $MPC_{vanilla}$ also provided savings with a daily cost of $17.25 (8.7% saving), it was the least effective among the three MPC variants, achieving only a 10.0% peak-time load reduction and a comfort violation of 0.66 °C·h/day; this degraded performance relative to the heating season is attributed to the larger influence of unmeasured disturbances (solar radiation, occupant and equipment heat gains) during the cooling period, which the vanilla model cannot capture. Regarding HVAC cycling, all MPC strategies increased the daily on/off transitions relative to the Baseline (232 cycles/day), with MPC_{ideal} showing the largest increase (54.4%, 359 cycles/day), followed by MPC_{hybrid} (17.4%, 273 cycles/day) and $MPC_{vanilla}$ (12.3%, 261 cycles/day). The comparison metrics are all summarized in Table 2.

An analysis of the peak demand, illustrated in Fig. 12, highlights some of the benefits of the MPC controllers. The Baseline strategy resulted in a maximum peak demand of 30.6 kW during the cooling season. In comparison, all MPC strategies successfully lowered this peak. The $MPC_{vanilla}$ controller achieved a 31.4% reduction to 21.0 kW. The MPC_{ideal} controller reduced the peak to 23.0 kW (24.8% reduction), while MPC_{hybrid} achieved a 28.1% reduction to 22.0 kW. The modest nature of this peak reduction is rooted in the fundamental characteristics of cooling operations. Unlike heating demand, which often has a predictable, sharp peak due to synchronized system start-ups in the morning, the cooling load is distributed more broadly throughout the afternoon. As a result, periods of peak simultaneous operation occur more stochastically. While the MPC controller can mitigate these random peaks, the lack of a distinct, universal trigger makes it more challenging to achieve the same level of substantial, targeted peak shaving seen in heating scenarios. The comparison metrics are all summarized in Table 2.

Table 2 shows that all three MPC strategies achieve statistically significant cost savings and peak load reductions compared to the baseline, confirming their effectiveness for demand-side management. Among them, MPC_{hybrid} and MPC_{ideal} generally deliver the largest benefits. However, MPC_{ideal} occasionally yields smaller cost savings than MPC_{hybrid} because its perfect disturbance forecast enables tighter temperature control near the comfort boundaries, which in turn leads to substantially higher cycling counts (e.g., +65.1% in heating and +54.4% in cooling). Ironically, MPC_{hybrid} exhibits fewer cycling events because its less accurate disturbance predictions result in slightly less aggressive temperature regulation, providing a practical advantage for equipment longevity. $MPC_{vanilla}$ performs well during the heating season but shows diminished benefits during cooling, where the effect of unmeasured disturbances such as solar radiation and internal heat gains becomes more significant. Peak demand reduction, while not the primary optimization objective, is achieved as a supplementary benefit across all strategies, with reductions of up to 22.2% in heating and 31.4% in cooling.

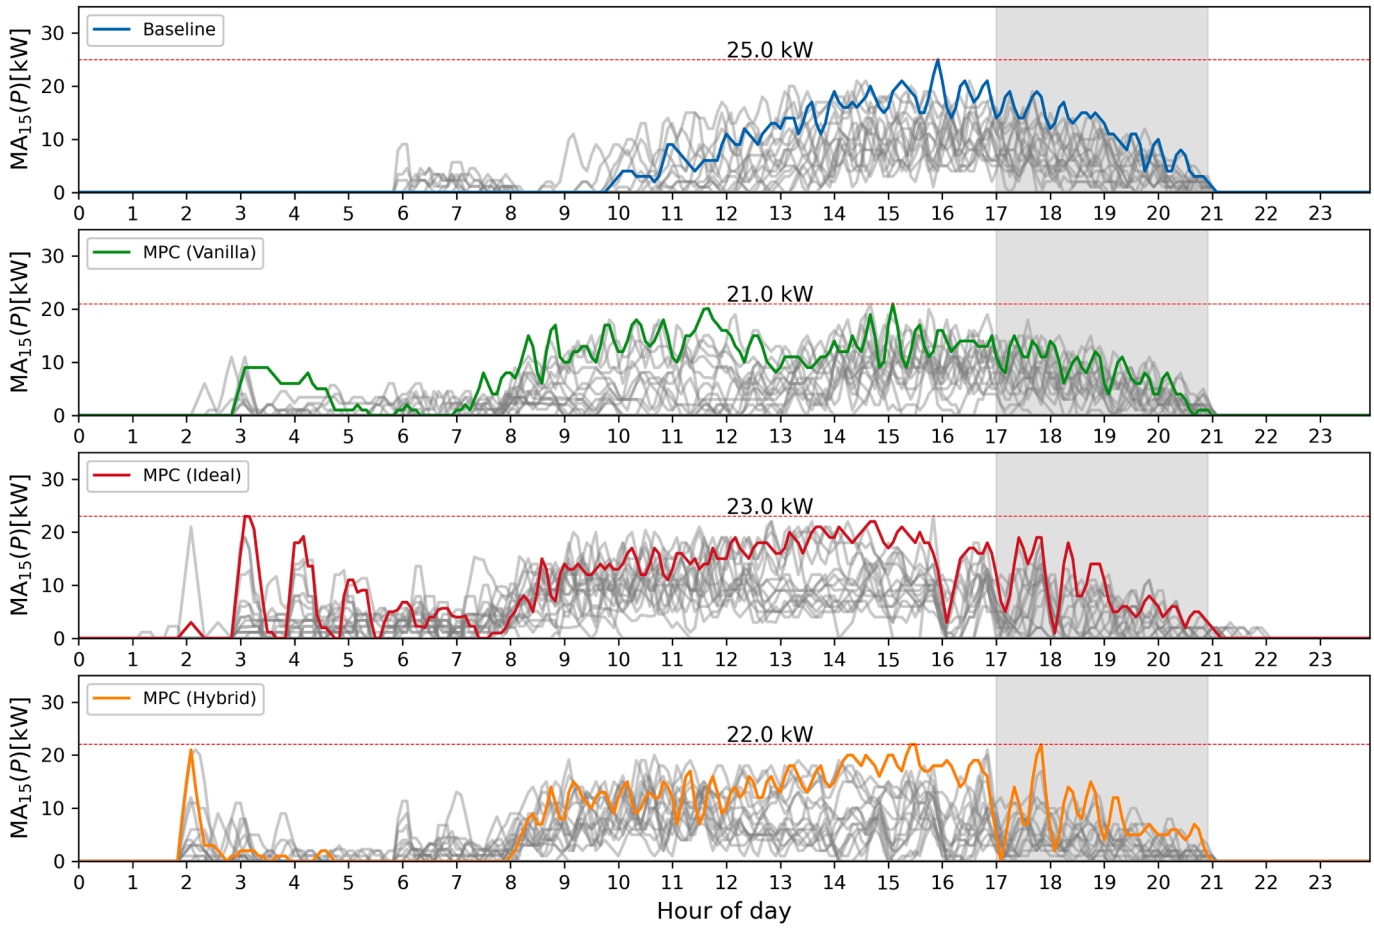

Fig. 12. Heat pump power demand profiles for all individual days during the cooling season. The day with the highest peak demand for each strategy is highlighted in color.

Table 2
Summary of Control Strategy Performance in simulation case study. Values are daily means ± 95% confidence intervals (CI) over 21 days. *p*-values are from paired *t*-tests comparing each MPC strategy against the baseline. All reported differences are statistically significant at $\alpha = 0.05$.

			$MPC_{vanilla}$		MPC_{ideal}		MPC_{hybrid}	
Season	Performance Metric	Baseline	Value	*p*	Value	*p*	Value	*p*
Heating	Daily Cost ($/day)	14.00 ± 4.18	9.20 ± 1.92	<0.001	10.59 ± 2.67	0.002	9.59 ± 2.20	<0.001
	Change[a] (%)	—	*-34.3%*		*-24.4%*		*-31.5%*	
	Morning Peak Load Shift (6–8 AM) (kW)	14.70 ± 2.29	2.20 ± 0.87	<0.001	3.76 ± 0.93	<0.001	6.78 ± 1.19	<0.001
	Change[a] (%)[a]	—	*-85.0%*		*-74.4%*		*-53.9%*	
	Peak Demand (kW)	27.0	17.0	—	22.7	—	22.7	—
	Change[a] (%)[a]	—	*-37.0%*		*-16.0%*		*-16.0%*	
	Comfort Violation (°C·h/day)	7.00 ± 1.89	0.24 ± 0.05	<0.001	0.10 ± 0.08	<0.001	0.85 ± 0.21	<0.001
	Change[a] (%)[a]	—	*-96.6%*		*-98.6%*		*-87.8%*	
	Cycling Events[b] (count/day)	207 ± 74	306 ± 43	<0.001	342 ± 69	<0.001	246 ± 58	0.008
	Change[a] (%)[a]	—	*+47.4%*		*+65.1%*		*+18.8%*	
Cooling	Daily Cost ($/day)	18.88 ± 6.47	17.25 ± 5.63	0.011	16.22 ± 5.23	<0.001	13.85 ± 5.20	<0.001
	Change[a] (%)[a]	—	*-8.7%*		*-14.1%*		*-26.6%*	
	Afternoon Peak Load Shift (5–9 PM) (kW)	4.48 ± 1.58	4.04 ± 1.38	0.009	3.12 ± 1.13	<0.001	2.76 ± 1.12	<0.001
	Change[a] (%)[a]	—	*-10.0%*		*-30.5%*		*-38.4%*	
	Peak Demand (kW)	24.9	23.0	—	23.8	—	23.8	—
	Change[a] (%)[a]	—	*-7.6%*		*-4.4%*		*-4.4%*	
	Comfort Violation (°C·h/day)	1.03 ± 0.39	0.66 ± 0.21	0.004	0.10 ± 0.04	<0.001	0.43 ± 0.15	<0.001
	Change[a] (%)[a]	—	*-36.1%*		*-90.5%*		*-57.6%*	
	Cycling Events[b] (count/day)[a]	232 ± 60	261 ± 65	<0.001	359 ± 64	<0.001	273 ± 65	<0.001
	Change[a] (%)[a]	—	*+12.3%*		*+54.4%*		*+17.4%*	

[a] Percentage change compared to baseline; - indicates reduction, + indicates increase.
[b] Cycling events are aggregated across all 9 heat pumps.

Table 3
Baseline Control Specifications.

Category[a]	Parameter	Value(s)
Setpoints	Occupied Hours Unoccupied Hours	Heating: 70°F, Cooling: 74°F (06:00-18:00 PT) Heating: 60°F, Cooling: 80°F (18:00-06:00 PT)
HVAC Schedule	Weekdays Weekends	06:00-18:00 PT (Mon-Fri) Same as weekdays
Climate	Test Location	Berkeley, CA (typical weather)
Fan Operation	Fixed Speed Min. Ventilation	950 CFM (85% fan speed) 150-160 CFM (fixed OA damper)
Heating Supply Air Temp	Control Limit	Thermostat controls HP-RTU stages; supply air not directly controlled Max 100°F to avoid hot air short-circuiting
Economizer	Mode	Disabled; minimum position maintained
Occupancy Loads	Fully Occupied Partially Occupied	600 W: 08:00-12:00, 13:00-17:00 300 W: 07:00-08:00, 12:00-13:00, 17:00-18:00
Lighting Loads	Fully Occupied Partially Occupied	350 W: 07:00-16:00 175 W: 06:00-07:00, 16:00-20:00

[a] HP-RTU: Heat Pump Rooftop Unit; PT: Pacific Time.

5. Experimental Study

5.1. Experiment overview

To validate the performance of the proposed Hybrid MPC under real-world conditions, we conducted an experimental study at the FLEXLAB Buildings [42]. A unique advantage of the FLEXLAB is its ability to directly measure internal loads (e.g., plug and lighting loads), which are typically treated as unmeasured disturbances in SMCBs. This enables a direct comparison between the proposed Hybrid MPC (MPC_{hybrid}) and an idealized MPC (MPC_{ideal}) that has perfect knowledge of these disturbances, as well as a Baseline (i.e., fixed setpoint control).

The Baseline control maintains fixed heating and cooling setpoints during occupied hours, with setback setpoints during unoccupied hours (Table 3). Both MPC strategies aim to minimize utility costs in response to a dynamic price signal while maintaining zone temperatures within the comfort range. The key difference between MPC_{ideal} and MPC_{hybrid} lies in their treatment of unmeasured disturbances. MPC_{ideal} assumes that both measured disturbances (w) and unmeasured disturbances ($\dot{Q}_{gain}$) are available for system identification and forecasting. In contrast, MPC_{hybrid} has access only to measured disturbances for system identification and estimates $\dot{Q}_{gain}$ using a black-box neural network model, which is combined with the gray-box building model to form the hybrid model.

5.2. Test Facility and Building Characteristics

The experiment was conducted in Cell B of the rotating test cells at FLEXLAB, located at Lawrence Berkeley National Laboratory in Berkeley, California (Fig. 13). The cell represents a small office space with a floor area of 57 m^2 and a large north-facing window. It is served by a packaged rooftop heat pump unit (HP-RTU; AAON RQ 2-ton) with two-stage heating and cooling control. The nominal heating and cooling capacities are 6.53 kW and 6.16 kW, respectively, at 700 CFM (1190 CMH). The HP-RTU is controlled by a Schneider Electric SE8600 thermostat, which determines the heating/cooling stage based on internal logic.

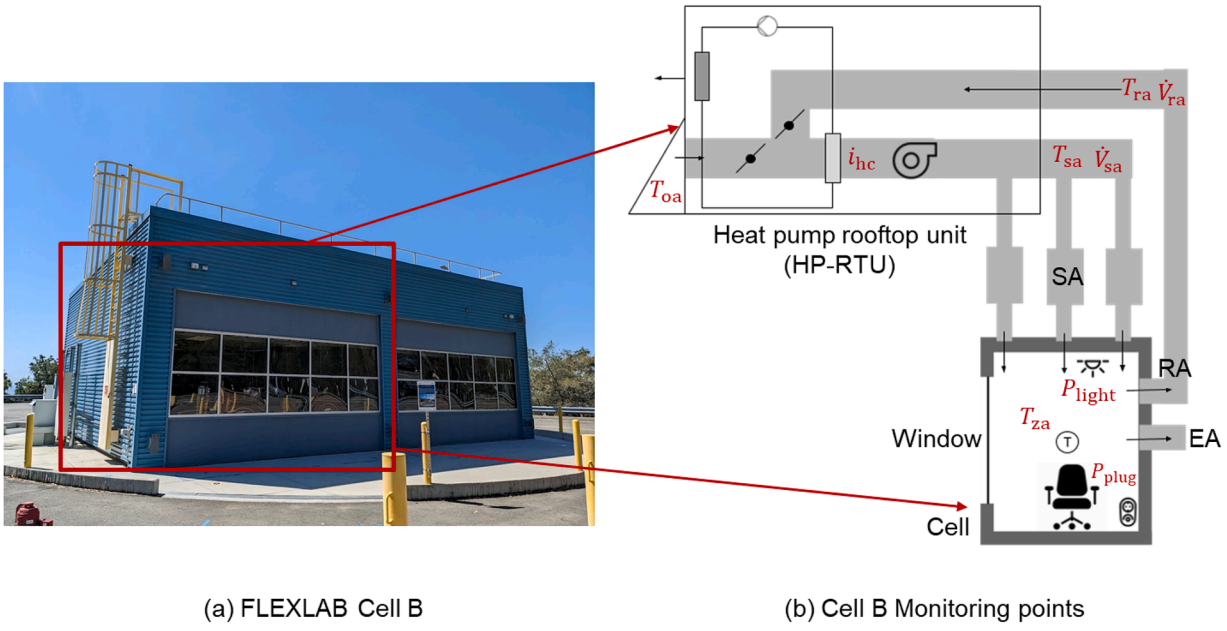

Fig. 13. FLEXLAB Cell B and its mechanical schematic diagram.

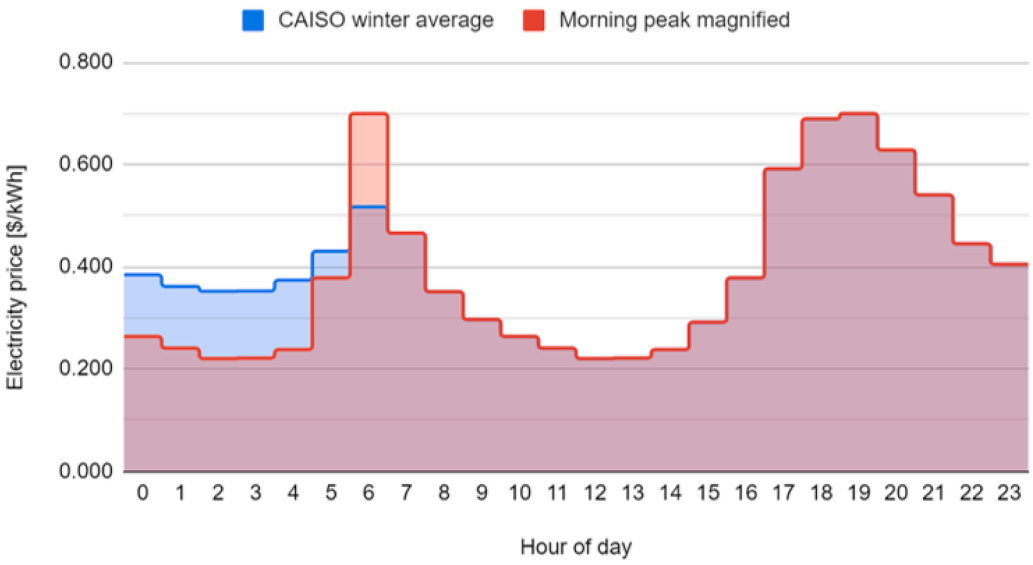

Fig. 14. Price signal.

Table 4
Experiment Schedule.

Period	Description
Weeks 1-4 of March	Commissioning; Baseline operation and System Identification
Apr. 1st week	Experiments with Baseline and MPC_{ideal}
Apr. 2nd-4th weeks	Experiments with Baseline, MPC_{ideal}, and MPC_{hybrid}

Conditioned air is delivered via ceiling supply and return grilles, with natural exhaust through the plenum.

The monitoring points used in the experiment are shown in Fig. 13(b) and include outdoor air temperature T_{oa} [°C], supply air temperature T_{sa} [°C], return air temperature T_{ra} [°C], zone air temperature from thermostat T_{za} [°C], heating/cooling stage signal i_{hc} [-], supply air flow rate $\dot{V}_{sa}$ [m^3/s], return air flow rate $\dot{V}_{ra}$ [m^3/s], plug load power P_{plug} [kW], and lighting load power P_{light} [kW].

Data were collected at one-minute intervals and stored in the FLEXLAB database. The Baseline control schedule, HVAC operation parameters, and internal load schedules are summarized in Table 3.

5.3. Experimental Setup

The dynamic price signal used for the MPC tests is shown in Fig. 14. It was adapted from the average CAISO winter price profile to (1) increase heating load incentives during early-morning hours and (2) provide a larger peak-to-off-peak price ratio to encourage load shifting.

The experimental procedure consisted of three phases (Table 4):

1. **Commissioning:** Verification of HP-RTU and thermostat performance.
2. **Model Calibration:** Collection of operational data for system identification of the gray-box building model and, for MPC_{hybrid}, the black-box disturbance model. This included Baseline operation and manual setpoint perturbations during unoccupied periods to excite building thermal dynamics.
3. **Experiment:** Alternating daily operation between Baseline, MPC_{ideal}, and MPC_{hybrid} to ensure comparable weather conditions across control strategies. Due to mild weather, controller assignments were occasionally adjusted to capture data under a range of outdoor air temperatures.

The control and data flow architecture is illustrated in Fig. 15. All sensor data were collected in the FLEXLAB datastore and passed to a VOLTTRON-based middleware [44], which communicated with the MPC server. More details on the middleware can be found in previous work [22,45]. The MPC server computed optimal setpoints every control interval and returned them to the middleware for implementation. A Grafana dashboard was used for real-time monitoring.

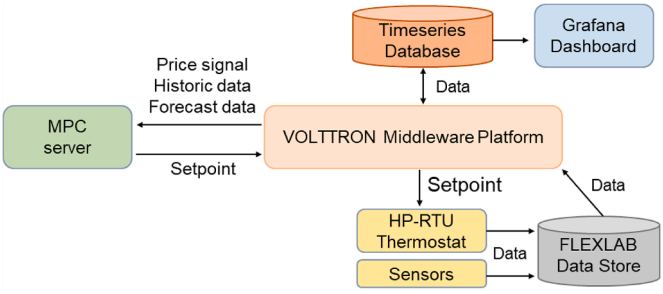

Fig. 15. Schematic diagram of testing environment.

5.4. Regression-Based Performance Comparison

In the simulation study, different control strategies can be evaluated under identical outdoor conditions, enabling direct comparison of performance metrics. In the experimental study, however, each strategy is tested on different days with different weather conditions. Because heating and cooling energy consumption depends strongly on outdoor air temperature, raw daily metrics cannot be compared directly across strategies without accounting for this confounding variable.

To enable a fair comparison, we employ a dummy variable regression approach (also known as analysis of covariance). For each performance metric Y (e.g., daily energy cost, peak demand), we pool the daily observations from the Baseline and MPC strategies and fit the following model:

$$Y = \beta_0 + \beta_1 T_{oa} + \beta_2 D + \beta_3 (T_{oa} D) + \epsilon, \quad (10)$$

where T_{oa} is the daily mean outdoor air temperature, D is a dummy variable that equals 1 for MPC days and 0 for Baseline days, and ϵ is the residual error. In this formulation, β_2 captures the shift in the intercept due to MPC (i.e., a level change in the metric at a reference temperature), and β_3 captures any change in the slope with respect to outdoor temperature (i.e., whether the MPC effect varies with weather

conditions). When $D = 0$, the model reduces to the Baseline relationship $Y = \beta_0 + \beta_1 T_{oa}$; when $D = 1$, it becomes $Y = (\beta_0 + \beta_2) + (\beta_1 + \beta_3) T_{oa}$.

To test the overall statistical significance of the MPC effect, we perform an F-test comparing the full model (Eq. (10)) against a reduced model that omits the MPC terms ($Y = \beta_0 + \beta_1 T_{oa} + \epsilon$). The null hypothesis is $H_0 : \beta_2 = \beta_3 = 0$, meaning MPC has no effect on the metric after controlling for outdoor temperature.

To quantify the practical significance of the MPC effect, we evaluate the fitted Baseline and MPC regression lines at every observed outdoor temperature in the combined dataset and compute their means:

$$\overline{Y}_{\text{Baseline}} = \frac{1}{n}\sum_{i=1}^{n}(\beta_0 + \beta_1 T_{\text{oa},i}), \quad \overline{Y}_{\text{MPC}} = \frac{1}{n}\sum_{i=1}^{n}[(\beta_0 + \beta_2) + (\beta_1 + \beta_3) T_{\text{oa},i}], \tag{11}$$

where the summation runs over all n daily observations from both strategies. The weather-normalized percentage reduction is then:

$$\text{Reduction}(\%) = \frac{\overline{Y}_{\text{Baseline}} - \overline{Y}_{\text{MPC}}}{\overline{Y}_{\text{Baseline}}} \times 100\%. \tag{12}$$

This approach averages the gap between the two regression lines over the actually observed range of outdoor temperatures, providing a single weather-normalized summary of the MPC benefit that accounts for the distribution of conditions encountered during the experiment.

5.5. MPC experiment results

5.5.1. Prediction comparison

Fig. 16 compares zone air temperature predictions for $\text{MPC}_{\text{ideal}}$ and $\text{MPC}_{\text{hybrid}}$ against measurements with disturbances ($\dot{q}_{\text{sol,win}}$, $\dot{Q}_{\text{gain}}$, u_h, and u_c). Both models exhibit good agreement with measured temperatures, with $\text{MPC}_{\text{hybrid}}$ showing significantly improved accuracy compared to the gray-box-only model, denoted $T_{za,\text{hybrid0}}$ in Fig. 16, which uses the gray-box model without the disturbance prediction component (no disturbance forecast). This confirms the value of the disturbance prediction component in the hybrid approach. Minor deviations in $\text{MPC}_{\text{ideal}}$ predictions are attributed to unmeasured disturbances such as infiltration and differences between measured plug loads and actual heat gains. The bottom subplot of Fig. 16 shows the measured HVAC control signals at the original data resolution to illustrate the on/off switching behavior of the heat pumps. The heating command u_h represents the compressor stage and can take the values 0, 1, or 2; these values have been normalized to [0, 0.5, 1] in the figure for visual clarity.

Fig. 16 compares zone air temperature predictions for $\text{MPC}_{\text{ideal}}$ and $\text{MPC}_{\text{hybrid}}$ against measurements with disturbances ($\dot{q}_{\text{sol,win}}$, $\dot{Q}_{\text{gain}}$, u_h, and u_c). Both models exhibit good agreement with measured temperatures, with $\text{MPC}_{\text{hybrid}}$ showing significantly improved accuracy compared to the gray-box-only model (no disturbance forecast). This confirms the value of the disturbance prediction component in the hybrid approach. Minor deviations in $\text{MPC}_{\text{ideal}}$ predictions are attributed to unmeasured disturbances such as infiltration and differences between measured plug loads and actual heat gains.

5.5.2. Load shifting and energy cost reduction

Fig. 17 shows average daily heating power profiles for load shifting performances. The Baseline control exhibits a sharp morning peak (6-8 AM) coinciding with high electricity prices. Both MPC strategies reduce heating during this period through pre-heating in the early morning (1–5 AM), achieving substantial reductions in peak-period heating energy. Since the three control strategies were tested sequentially in a single building under different weather conditions, quantitative effects are evaluated via the dummy variable regression approach described in Section 5.4, with results summarised in Tables 5 and 6.

Both MPC strategies reduce heating during this period through pre-heating in the early morning (1-5 AM), achieving substantial48% ($\text{MPC}_{\text{ideal}}$) and 45% ($\text{MPC}_{\text{hybrid}}$) reductions in peak-period heating energy(Table 5). Since the three control strategies were tested sequentially in a single building under different weather conditions, we evaluate the quantitative effects via the dummy variable regression approach described in Section 5.4, with results summarized in Tables 5 and 6.

Fig. 18 presents the morning peak-price average power (6–8 AM) and daily HVAC energy cost as functions of daily mean outdoor air temperature. Both metrics decrease with increasing T_{oa}, as expected, since warmer outdoor conditions reduce the heating load. Across the observed temperature range, both MPC strategies consistently achieve lower peak-period power and energy cost compared to the Baseline, with the regression lines for $\text{MPC}_{\text{ideal}}$ and $\text{MPC}_{\text{hybrid}}$ lying well below the Baseline line.

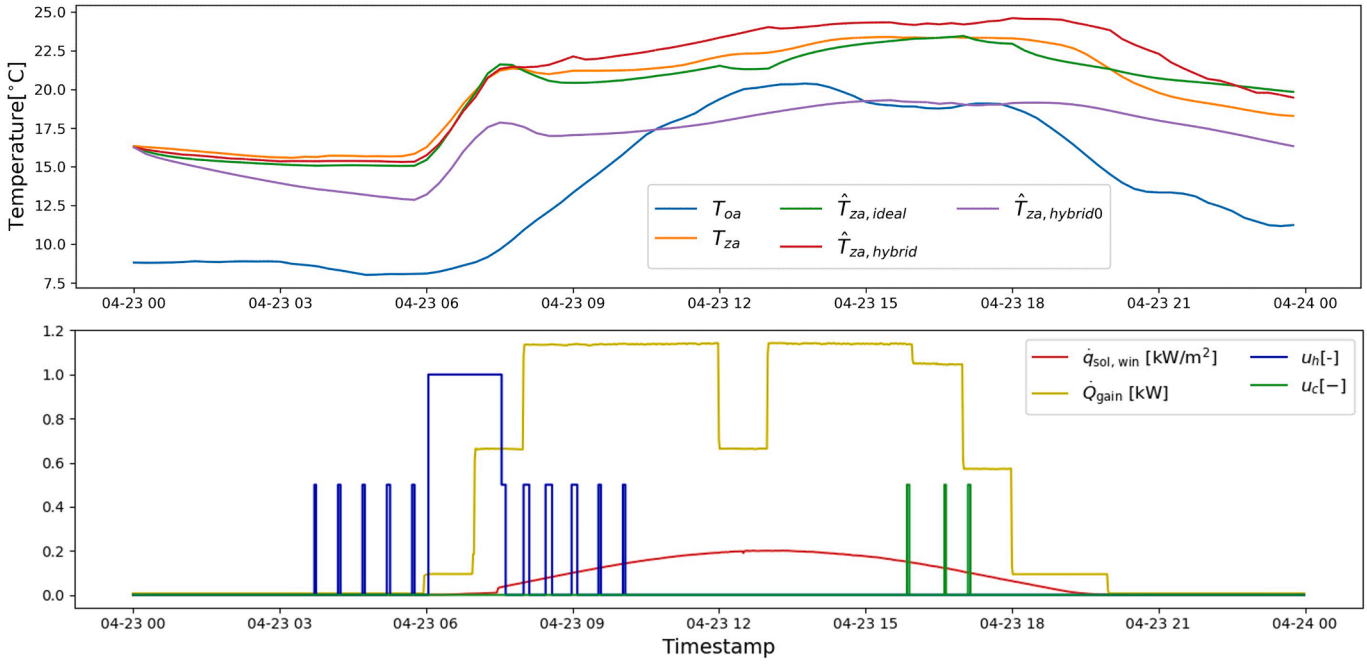

Fig. 16. Comparison of model prediction.

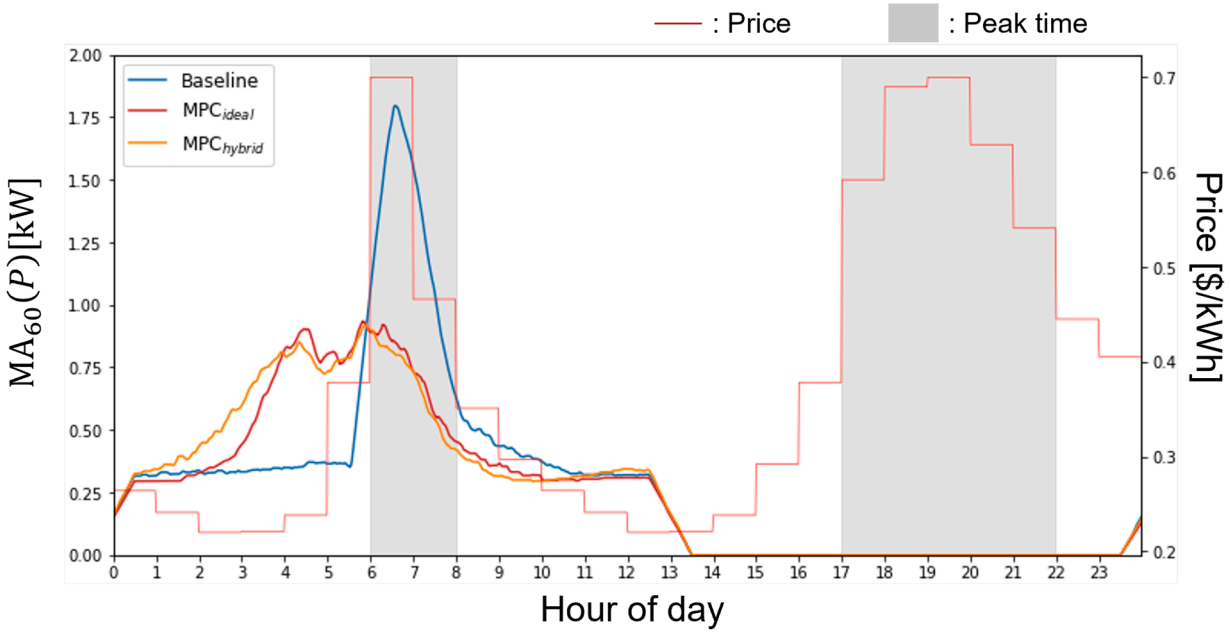

Fig. 17. Averaged daily heating energy use for Baseline and MPCs.

Table 5
Dummy Variable Regression Results.

Metric	MPC Case	β_2 (SE)	β_3 (SE)	F	R^2
Energy Cost [$/day]	MPC$_{ideal}$[a]	-0.610 (0.657)	0.017 (0.058)	10.79**	0.914
	MPC$_{hybrid}$[a]	-1.294 (0.467)	0.086 (0.040)	11.48**	0.944
Peak Load, 6–8 AM [kW]	MPC$_{ideal}$[a]	-1.618 (0.428)	0.094 (0.038)	46.75***	0.935
	MPC$_{hybrid}$[a]	-1.716 (0.399)	0.104 (0.035)	41.22***	0.938
Comfort [°C·h/day]	MPC$_{ideal}$[a]	-9.205 (1.297)	0.619 (0.114)	89.84***	0.960
	MPC$_{hybrid}$[a]	-10.050 (1.276)	0.678 (0.110)	87.14***	0.960
Cycling [events/day]	MPC$_{ideal}$[a]	-32.010 (28.447)	3.353 (2.493)	2.01	0.721
	MPC$_{hybrid}$[a]	-3.960 (32.726)	1.578 (2.833)	3.98	0.716

[a] Model: $Y = \beta_0 + \beta_1 T_{oa} + \beta_2 D + \beta_3 (T_{oa} \times D) + \epsilon$, where $D = 1$ for MPC and $D = 0$ for Baseline. F-test evaluates $H_0 : \beta_2 = \beta_3 = 0$.
[*] $p < 0.05$; ** $p < 0.01$; *** $p < 0.001$.

Table 6
Weather-Normalized MPC Performance: Mean Effect Size.

Metric	MPC Case	$\overline{Y}_{BL}$	$\overline{Y}_{MPC}$	ΔY	Reduction [%]
Energy Cost [$/day]	MPC$_{ideal}$[a]	3.67	3.24	-0.42	11.6
	MPC$_{hybrid}$[a]	3.64	3.32	-0.32	8.7
Peak Load, 6–8 AM [kW]	MPC$_{ideal}$[a]	1.30	0.74	-0.56	42.9
	MPC$_{hybrid}$[a]	1.28	0.75	-0.54	41.7
Peak Demand, 15-min [kW]	MPC$_{ideal}$[a]	2.13	1.64	-0.50	23.2
	MPC$_{hybrid}$[a]	2.13	1.51	-0.62	29.2
Comfort [°C·h/day]	MPC$_{ideal}$[a]	3.03	0.79	-2.24	73.8
	MPC$_{hybrid}$[a]	2.95	0.61	-2.35	79.4
Cycling [events/day]	MPC$_{ideal}$[a]	22.3	28.0	+ 5.7	-25.7
	MPC$_{hybrid}$[a]	21.5	35.5	+ 14.0	-65.0

[a] $\overline{Y}_{BL}$ and $\overline{Y}_{MPC}$ are the mean predicted values from the fitted Baseline and MPC regression lines, evaluated over all observed outdoor temperatures. $\Delta Y = \overline{Y}_{MPC} - \overline{Y}_{BL}$. Negative reduction indicates an increase relative to Baseline.

The regression analysis reveals that the MPC effect on morning peak load is highly significant for both strategies (F-test $p < 0.001$; Table 5), with weather-normalized reductions of 42.9% for MPC$_{ideal}$ and 41.7% for MPC$_{hybrid}$ (Table 6). The interaction terms (β_3) are also significant ($p = 0.027$ and $p = 0.012$, respectively), indicating that the MPC load shifting effect is slightly larger on colder days when heating loads are higher. For daily energy cost, both MPC strategies achieve statistically significant reductions (F-test $p \approx 0.002$), with weather-normalized savings of 11.6% for MPC$_{ideal}$ and 8.7% for MPC$_{hybrid}$. The peak demand (15-min maximum) is also significantly reduced, by 23.2% for MPC$_{ideal}$ and 29.2% for MPC$_{hybrid}$. The close performance between the two MPC strategies confirms that the black-box disturbance estimator in MPC$_{hybrid}$ provides a practical substitute for direct measurement of internal loads.

5.5.3. Peak demand reduction

Fig. 19 presents daily HVAC power profiles (15-minute moving average) for each control strategy, with the day exhibiting the highest peak demand highlighted. The Baseline controller concentrates heating at the 6–8 AM setpoint transition, producing a sharp HVAC peak of 2.13 kW. Both MPC strategies distribute the heating load more evenly by pre-heating during the low-price early morning hours, reducing the HVAC peak demand to 1.64 kW for MPC$_{ideal}$ and 1.51 kW for MPC$_{hybrid}$. The dummy variable regression analysis confirms that peak demand reductions are statistically significant (F-test $p < 0.001$ for both; Table 5), with weather-normalized reductions of 23.2% for MPC$_{ideal}$ and 29.2% for MPC$_{hybrid}$ (Table 6). Notably, MPC$_{hybrid}$ achieves a slightly larger peak demand reduction than MPC$_{ideal}$, which can be attributed to its more distributed pre-heating pattern visible in the early morning hours (0–6 AM). The gray traces in Fig. 19 show the variability across individual experimental days, illustrating that the peak demand reduction is consistent rather than driven by a single favorable day.

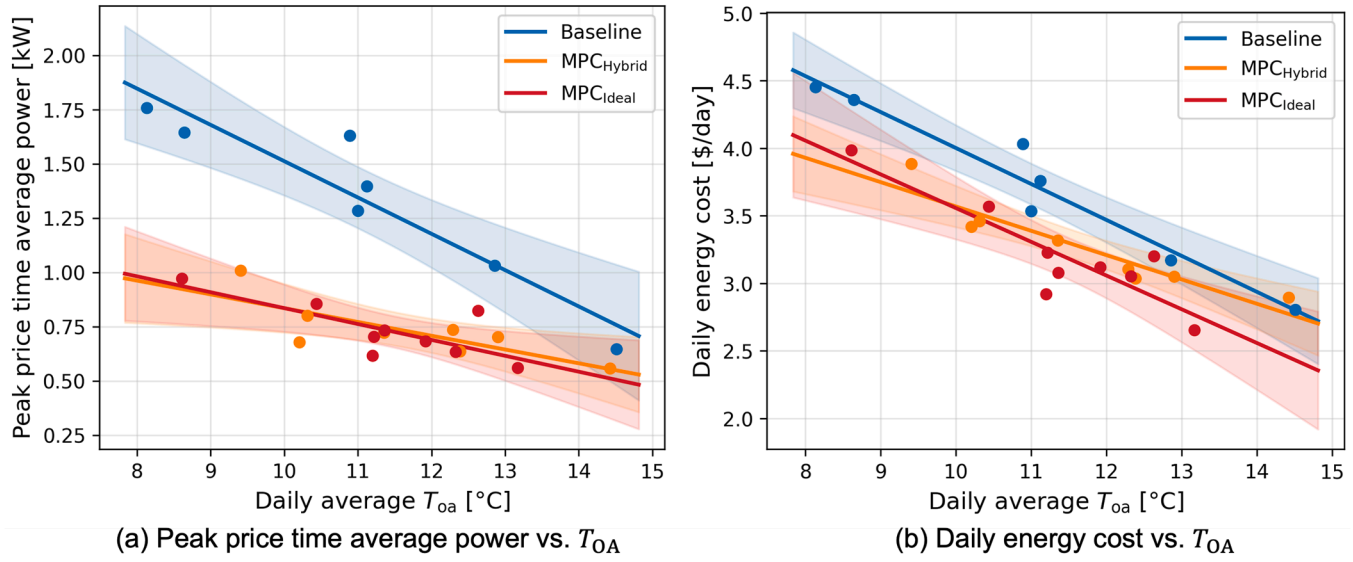

Fig. 18. (a) Peak-price time average power (6–8 AM) and (b) daily energy cost as functions of daily mean outdoor air temperature T_{oa}. Solid lines show fitted regression models; shaded regions indicate 95% confidence intervals.

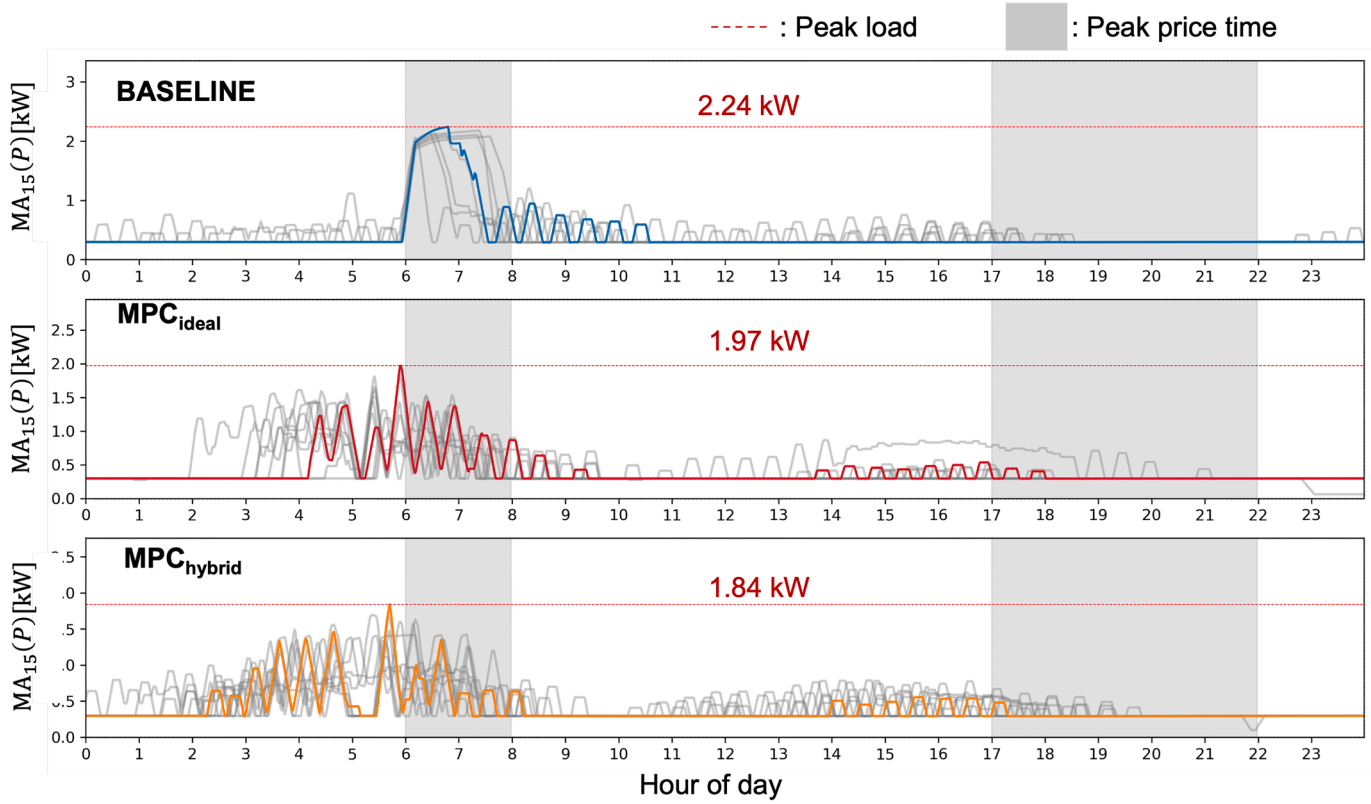

Fig. 19. HVAC peak demand comparison between control strategies. Each panel shows the 15-minute moving average of heat pump power $\mathrm{MA}_{15}(P)$ for all experimental days (gray), with the peak-demand day highlighted in color. The dashed red line indicates the observed peak load, and the shaded region marks the peak electricity price period (6–8 AM).

5.5.4. *Comfort and HVAC cycling*

Fig. 20 examines comfort violation and HVAC cycling behavior. As shown in Fig. 20(a), the Baseline exhibits substantially higher comfort violations, particularly on colder days where the fixed-setpoint controller struggles to maintain zone temperatures within the comfort band during the morning warm-up period. Both MPC strategies achieve dramatic reductions in comfort violation degree-hours: 73.8% for $\mathrm{MPC}_{\mathrm{ideal}}$ and 79.4% for $\mathrm{MPC}_{\mathrm{hybrid}}$ (Table 6). These reductions are highly significant (F-test $p < 0.001$; Table 5) with large effect sizes (Cohen's $f^2 > 14$). The improvement occurs because MPC anticipates the occupied-period comfort requirements and begins pre-heating before the setpoint transition, whereas the Baseline reacts only after the setpoint change, resulting in temperature undershoot.

However, this improved comfort performance comes at the cost of increased HVAC cycling, as shown in Fig. 20(b). $\mathrm{MPC}_{\mathrm{ideal}}$ increases daily cycling events by 25.7% and $\mathrm{MPC}_{\mathrm{hybrid}}$ by 65.0% relative to the Base-

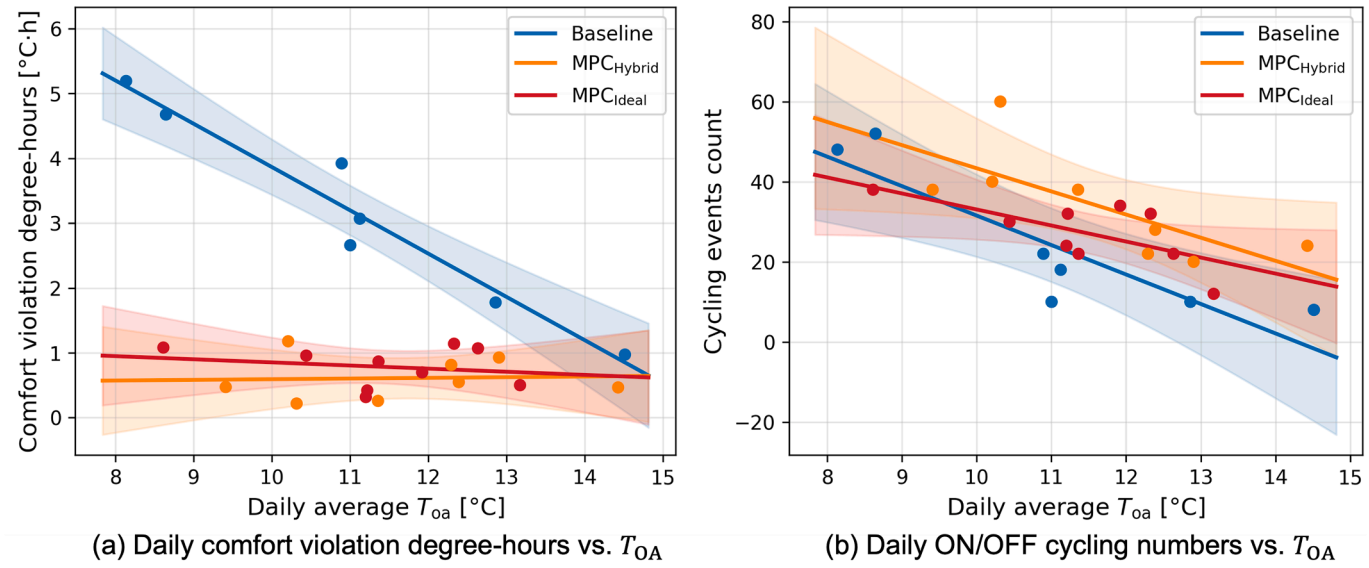

Fig. 20. (a) Daily comfort violation degree-hours and (b) daily HVAC on/off cycling events as functions of daily mean outdoor air temperature T_{oa}. Solid lines show fitted regression models; shaded regions indicate 95% confidence intervals.

line (Table 6). This increase is primarily because MPC maintains zone temperatures near the comfort boundary to minimize energy cost, resulting in more frequent on/off transitions of the heat pump as the controller modulates heating to track the optimal temperature trajectory. The F-tests for cycling are not statistically significant at the 0.05 level ($p = 0.177$ for $\text{MPC}_{\text{ideal}}$, $p = 0.050$ for $\text{MPC}_{\text{hybrid}}$; Table 5), partly due to the high day-to-day variability in cycling counts evident in the wide confidence bands in Fig. 20(b). Nevertheless, the practical increase in cycling warrants consideration for long-term equipment wear and could be mitigated through explicit cycling penalties in the MPC cost function.

5.5.5. Summary of regression analysis

Tables 5 and 6 consolidate the dummy variable regression results across all performance metrics. Table 5 reports the regression coefficients, F-test statistics, and effect sizes. The MPC dummy variable (β_2) and interaction term (β_3) jointly capture the level shift and slope change attributable to MPC after controlling for outdoor temperature. The F-test evaluates the null hypothesis $H_0 : \beta_2 = \beta_3 = 0$, and Cohen's f^2 quantifies the practical effect size. Table 6 presents the weather-normalized mean performance for each strategy and the corresponding percentage reductions, computed by evaluating the fitted regression lines over the observed range of outdoor temperatures as described in Section 5.4.

6. Discussion

The proposed hybrid MPC framework combines first-principles process models with data-driven disturbance modeling to improve prediction accuracy and control performance under realistic operating conditions. While the hybrid approach demonstrates promising results in both simulation and experimental tests, several limitations and practical considerations remain.

6.1. Constant COP assumption and variable-efficiency equipment

The MPC formulation assumes a constant COP for each heat pump, making the electrical power consumption a linear function of the runtime fraction. This assumption is reasonable for the demonstration sites on the U.S. West Coast, where mild weather results in relatively small COP variations within the prediction horizon. However, in climates with large outdoor air temperature swings, the COP of air-source heat pumps can vary significantly—particularly during heating operation at low ambient temperatures, where defrost cycles and auxiliary electric resistance heating further reduce effective COP. In such cases, the constant-COP assumption introduces modeling error in the cost objective, potentially leading to suboptimal pre-heating or pre-cooling schedules. To address this, the power terms $P_{\text{HP,h},i}$ and $P_{\text{HP,c},i}$ in Eq. 1 can be replaced with outdoor-air-temperature-dependent functions, or auxiliary heating modes can be treated as separate operations with distinct power characteristics, as described in [38]. The impact of variable COP on MPC performance warrants further investigation, particularly for cold-climate applications where COP degradation is most pronounced.

6.2. Model Accuracy and Data Requirements

The quality of the hybrid model depends heavily on the quality and representativeness of the training dataset. When the available training data is highly correlated or lacks sufficient excitation, the system parameters become unidentifiable [41]. This can lead to biased disturbance models and degraded prediction accuracy. In the simulation study, insufficiently diverse training data led to noticeable performance degradation compared to the ideal MPC benchmark, while the experimental study, which contains a certain amount of unmeasured disturbances, shows similar performance to the ideal MPC benchmark.

The unmeasured disturbance in our case study exhibits stochastic behavior, making it inherently difficult to model accurately with deterministic approaches. While the hybrid model partially compensates for this disturbance, its accuracy is still lower than that of an ideal MPC with perfect disturbance knowledge. This is consistent with the expectation that stochastic, time-varying disturbances cannot be fully captured without explicit stochastic modeling. Our results suggest that at least one month of operational data containing the unmeasured disturbance is necessary to achieve satisfactory disturbance rejection in practice [34]. Shorter training periods may lead to underfitting and poor generalization to unseen disturbance patterns.

6.3. Practical guidelines for MPC parameter tuning

The sensitivity analysis in B reveals that the MPC performance is governed primarily by the demand penalty weight ω_d, while the demand target fraction δ has a comparatively secondary effect. Fig. 21 illustrates this effect by comparing the optimal prediction trajectories at two representative values of ω_d.

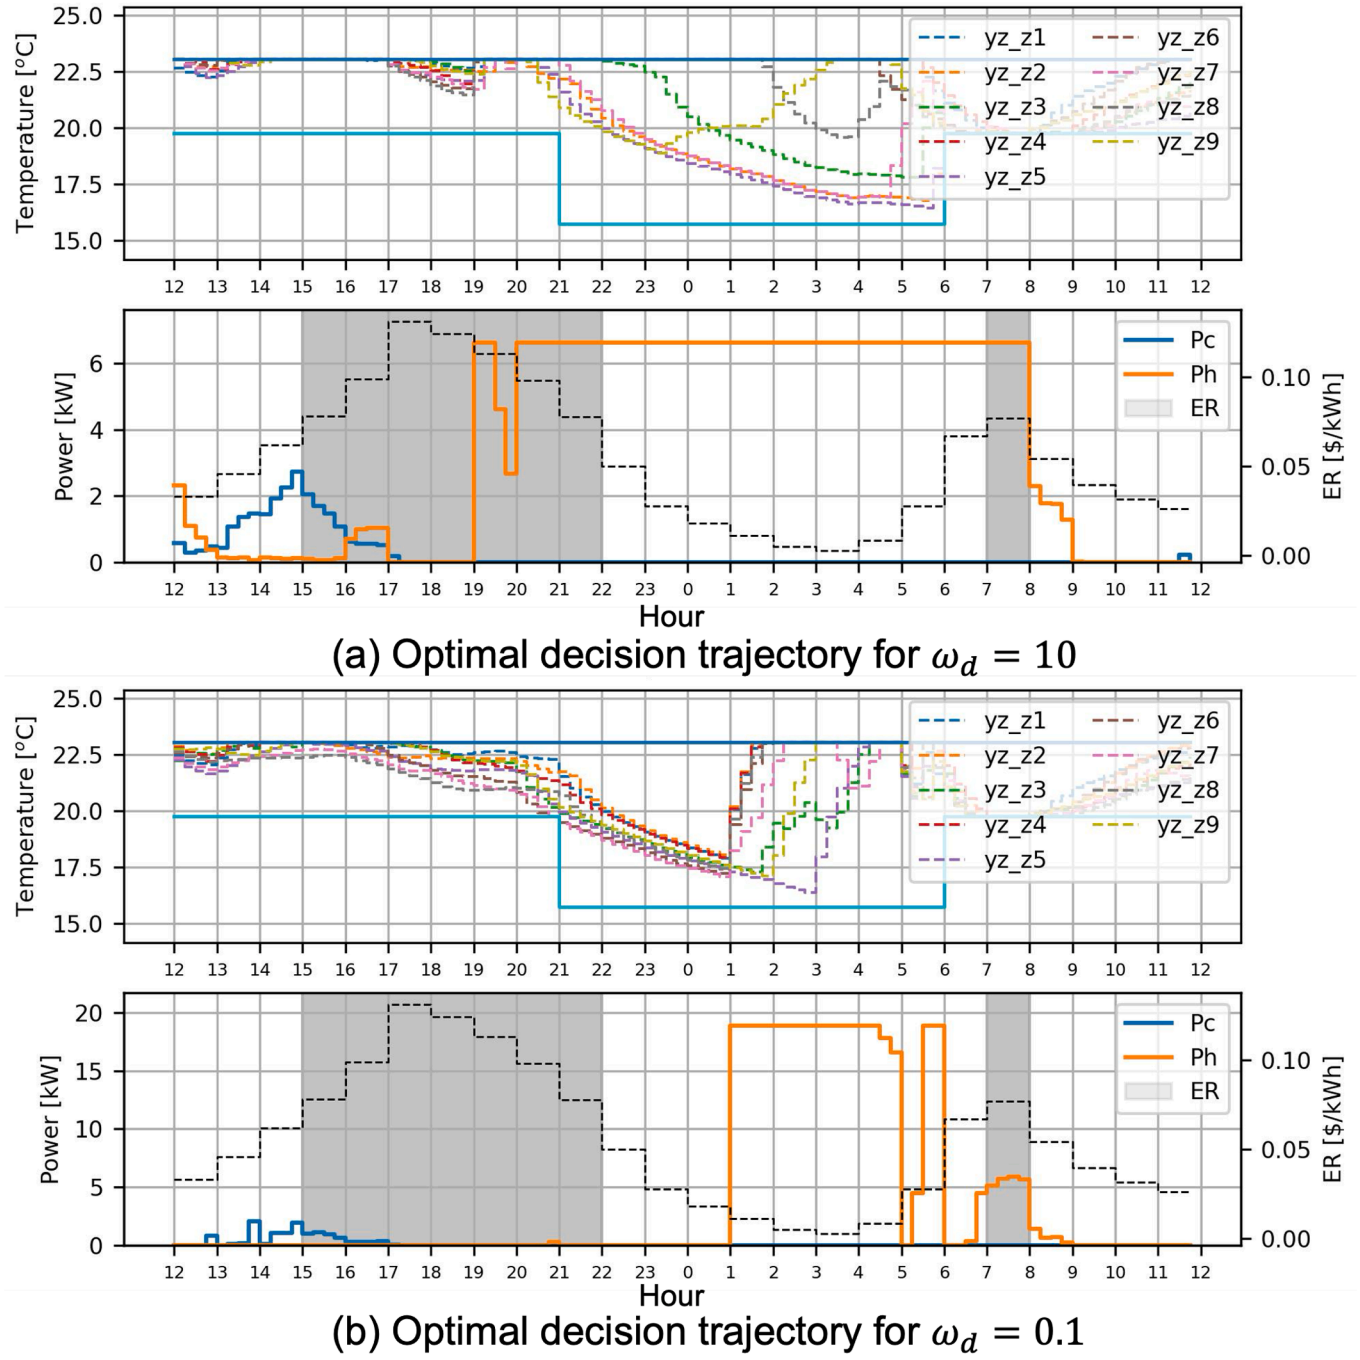

Fig. 21. Optimal MPC prediction trajectories for (a) $\omega_d = 10$ and (b) $\omega_d = 0.1$. Upper panels show predicted zone temperatures for nine zones; lower panels show predicted heating power (P_h), cooling power (P_c), and the time-of-use electricity rate (ER, shaded). A high demand penalty forces sustained overnight pre-heating with a peak of approximately 6.5 kW, while a low demand penalty permits concentrated pre-heating with a peak of approximately 19 kW but lower total energy cost.

When ω_d is set to a high value (e.g., 10), the optimizer aggressively curtails peak demand by restricting the maximum heating power to approximately 6.5 kW (Fig. 21a). To satisfy the comfort constraints under this tight demand cap, the controller is forced to distribute heating across the entire nighttime period, including the peak-price hours, resulting in sustained pre-heating from the evening through the early morning. While this strategy successfully limits instantaneous demand, the prolonged pre-heating increases thermal losses to the outdoor environment and raises overall energy cost, as the building must maintain elevated temperatures for an extended duration before occupancy. In contrast, when ω_d is reduced to 0.1 (Fig. 21b), the demand constraint is largely relaxed, permitting the optimizer to concentrate heating in a shorter window just before the occupied period. The resulting peak power reaches approximately 19 kW, but the total energy consumption and cost are lower because the pre-heating duration is shorter and thermal losses are minimized. This comparison demonstrates that unless the demand charge component of the utility tariff is substantial, excessive demand penalization leads to a counterproductive increase in energy cost.

When load shifting for cost reduction is the primary objective, a small ω_d (e.g., 0.1) is therefore sufficient; this preserves cost and peak-time savings while still providing indirect peak demand reduction through

the temporal redistribution of heating or cooling operations. Values in the range 0.1–5 serve as a practical rule of thumb for most building configurations and pricing structures. If explicit demand curtailment is desired—for instance, when demand charges constitute a significant fraction of the utility bill—the trade-off between demand charge savings and increased electricity cost can be quantified directly, and ω_d can be increased toward 5–10 accordingly. However, values beyond this threshold yield diminishing demand savings at the cost of rapidly deteriorating economic performance, particularly in heating-dominated seasons where pre-heating energy penalties are large, as illustrated in Fig. 21a.

It should be noted that the current single-level MPC formulation controls peak demand only indirectly: the optimizer determines runtime fractions per timestep, and if multiple units are activated simultaneously within a timestep, instantaneous demand may exceed the intended target even when the aggregate runtime fraction remains below δ. Hierarchical MPC architectures that decompose the problem into an upper-level demand scheduler and a unit-level MPC can enforce hard demand limits more reliably [30,40]. Nevertheless, when the MPC solution distributes heating or cooling operations across extended periods (e.g., overnight pre-conditioning), the resulting load profile naturally regulates peak demand, and aggressive demand penalization can be relaxed.

For practitioners deploying the controller without a preliminary simulation study, the following empirical guidelines are recommended: (i) initialize ω_d in the range 0.1–1 and δ at 0.7; (ii) if demand charges are a significant cost component, increase ω_d toward 5–10 and tighten δ toward 0.55; (iii) monitor cost, peak demand, and comfort metrics over the first billing cycle and adjust accordingly, noting that building load characteristics and weather conditions will influence the optimal balance. A formal re-tuning or simulation-based calibration is recommended when the controller is transferred to buildings with substantially different thermal mass, occupancy patterns, or utility rate structures.

6.4. Robustness to disturbance prediction mismatch

A key limitation is that the validation relies on a nine-zone simulated building and a single experimental test cell with relatively structured occupancy and load patterns, which raises questions about generalizability to other SMCB types (e.g., retail stores, restaurants, or schools) where internal heat gains may be more irregular. Because the disturbance model is deterministic and trained on recent historical profiles, its prediction accuracy degrades when actual disturbance patterns deviate substantially from the training data. To investigate the robustness of the hybrid MPC under such mismatch conditions, we conducted a sensitivity study in which the actual unmeasured disturbance level was scaled to 100% (design condition), 50%, and 0% of the original profile while the disturbance predictor remained trained on the original data. Figs. 22 and 23, and Table 7 summarize the results for both cooling and heating seasons.

In the cooling season (Fig. 22), the hybrid MPC consistently reduces peak-price-time average power relative to the Baseline across all disturbance levels, achieving reductions of 38% at the design condition, 85% at 50% disturbance, and 95% at 0% disturbance. Because the disturbance model over-predicts internal heat gains when actual gains are lower, the controller pre-cools more aggressively than necessary, which reduces cost savings from 27% at the design condition to 20% at 50% disturbance. Nevertheless, the reduced disturbance levels also decrease the overall cooling load relative to the Baseline, so the MPC still achieves substantial absolute cost reductions even under prediction mismatch. In the heating season (Fig. 23), the controller maintains large peak-price-time load reductions (54–65%) across all disturbance levels. However, when the disturbance model over-predicts nighttime heat losses (i.e., actual disturbance is lower than predicted), the MPC over-heats during the early morning pre-heating period. At 0% disturbance, this over-heating causes the daily energy cost to increase by 62% relative to the Baseline, as excess heating energy is consumed during non-peak hours to compen-

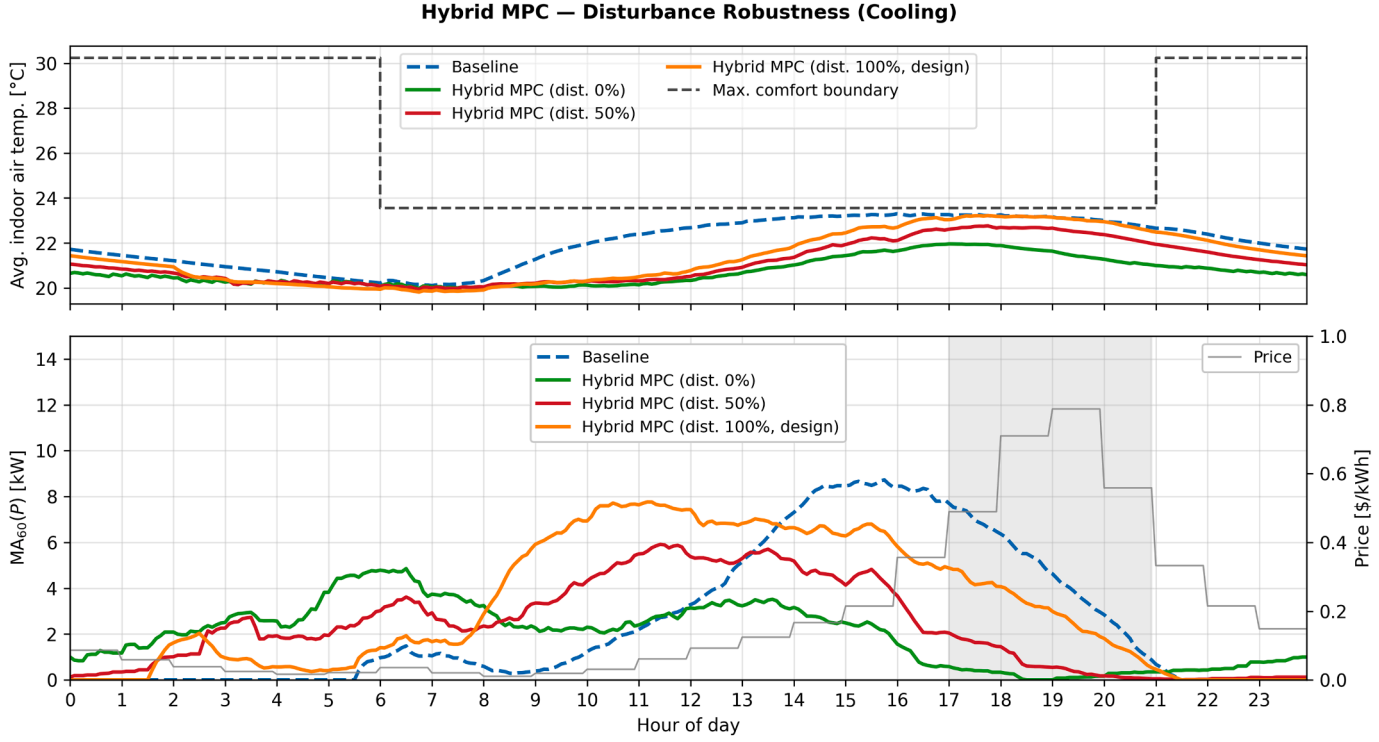

Fig. 22. Disturbance robustness analysis for the cooling season. Upper panel shows average indoor air temperature; lower panel shows 60-minute moving average power and time-of-use electricity price. The disturbance predictor is trained on the original (100%) profile while actual disturbances are scaled to 100%, 50%, and 0%.

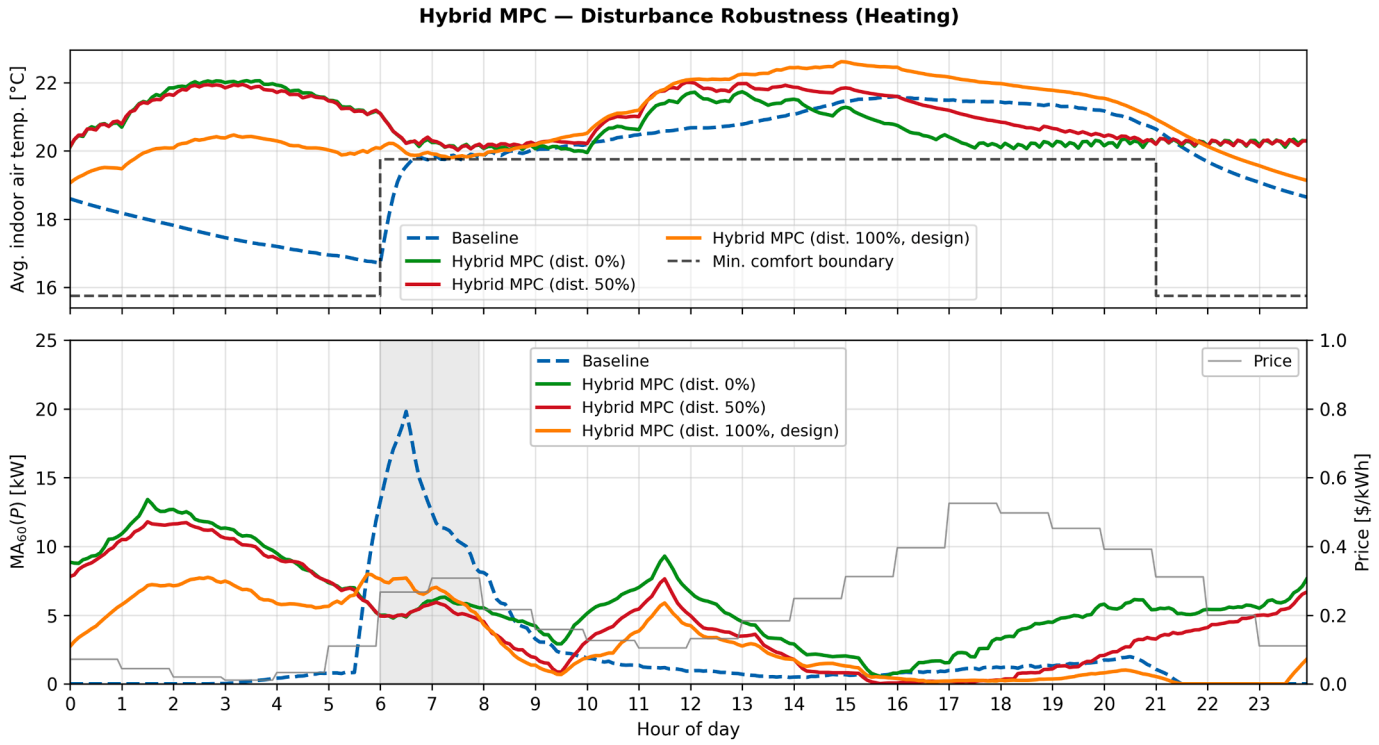

Fig. 23. Disturbance robustness analysis for the heating season. Upper panel shows average indoor air temperature; lower panel shows 60-minute moving average power and time-of-use electricity price. The disturbance predictor is trained on the original (100%) profile while actual disturbances are scaled to 100%, 50%, and 0%.

Table 7
Hybrid MPC performance under disturbance mismatch (percent reduction relative to Baseline (100%)).

Season[a]	Disturbance Level	Peak-Time Power	Energy Cost	Peak Demand	Comfort Viol.
Cooling	100% (design)	38%	27%	28%	58%
	50%	85%	66%	25%	95%
	0%	95%	80%	28%	91%
Heating	100% (design)	54%	32%	22%	88%
	50%	65%	3%	11%	97%
	0%	63%	-62%	11%	97%

[a] Peak-Time Power: average power during peak-price hours. Peak Demand: 15-minute maximum. Comfort Viol.: daily average comfort violation degree-hours. Positive values indicate reduction (improvement) over Baseline; negative values indicate increase (degradation).

sate for predicted losses that do not materialize. At 50% disturbance, the cost impact is modest (-3%), indicating that moderate prediction errors are well tolerated.

These results suggest that the hybrid MPC is robust to moderate disturbance mismatch (up to approximately 50% deviation from the training profile), maintaining both load shifting and cost reduction benefits. For buildings with highly irregular or rapidly changing occupancy patterns, this degradation can be mitigated through frequent re-calibration of the disturbance model at shorter intervals. When disturbance schedules change fundamentally—such as a change of business type—re-collection of baseline data and retraining of the disturbance model would be necessary. Future work could explore adaptive techniques such as sequential Bayesian updating, Gaussian process models, or transfer learning to improve disturbance prediction under non-stationary conditions. It is also worth noting that even under severe mismatch, the MPC consistently achieves peak-price-time load reduction through pre-conditioning, which remains a robust benefit of the predictive control strategy. For buildings where unmeasured disturbances are expected to be small or negligible, a standard gray-box MPC without the disturbance model may be sufficient.

6.5. *Context Dependency and Generalization of Results*

Our results indicate that MPC performance is also influenced by external factors such as building load profiles and electricity price signals. Different load patterns may amplify or attenuate the impact of unmeasured disturbances, while price signals can alter the optimal control trajectory. For instance, we observed significantly better load shifting performance in heating scenarios (which exhibit sharp morning peaks) compared to cooling scenarios (which have distributed afternoon loads). Additionally, in heating scenarios, the impact of unmeasured disturbances on short-term forecasts (e.g., a few hours in the early morning) is less critical for prediction accuracy.

This performance asymmetry is primarily driven by the temporal overlap between the building's thermal load profile and the peak-price windows, rather than by differences in thermal time constants alone. In the heating season, a large portion of the heating demand occurs during the early morning peak-price hours (6–8 AM), when the building must recover from nighttime setback before occupancy. The MPC can shift this concentrated load to pre-peak hours through overnight pre-heating, yielding substantial peak-time energy and cost reductions. In

contrast, the cooling season load is distributed across the afternoon and does not align as sharply with the evening peak-price window (5–9 PM), limiting the achievable load shifting benefit. More generally, when the thermal load profile and the price signal do not overlap—for example, a school building that closes near 4 PM will have negligible cooling load during the evening peak, or a building with a 9 AM occupancy start may already avoid the morning heating peak without MPC intervention—the load shifting benefit diminishes regardless of the building's thermal characteristics.

Therefore, the key determinant of MPC load shifting effectiveness is the degree of temporal coincidence between the building's load profile and the utility price structure. For air-based HVAC systems in lightweight commercial buildings, the thermal time constants are relatively short and do not critically differentiate performance across building types; the primary benefit arises instead from the supervisory coordination of multiple devices across time. When the load profile does not coincide with the price signal, or when a flat rate structure is used, the MPC can still achieve peak demand regulation through load coordination without temporal shifting, as demonstrated in our previous study [40].

Furthermore, building characteristics themselves affect the applicability of the hybrid approach. Buildings with higher thermal mass (e.g., radiant heating with concrete slabs) may be more suitable for load shifting strategies, as they can store thermal energy more effectively. In contrast, buildings with rapid thermal response (e.g., air-based systems in lightweight wooden construction) may demonstrate limited load shifting capability, potentially requiring excessive pre-heating or pre-cooling to achieve meaningful load shifts-particularly when the thermal load is already small. Therefore, careful model development and tuning are necessary to ensure satisfactory performance when applying the approach to new sites.

6.6. Scalability and Deployment Considerations

While our experimental validation focused on a single zone and simulation demonstrated nine-zone coordination, scaling to buildings with tens or hundreds of zones introduces significant challenges. Computational burden increases substantially, requiring an investigation into which optimization solvers can efficiently handle hundred-zone-scale problems. However, the MPC solution is designed for scale as it comes with a lightweight linear optimization solver. Also, the proposed architecture is inherently suited for aggregation into Virtual Power Plants (VPPs). By deploying this low-computational-cost hybrid approach across a city's stock of SMCBs, utilities could coordinate demand flexibility at an urban scale without requiring complex on-site hardware retrofits.

From a deployment perspective, the framework requires minimal hardware-only WiFi-enabled thermostats and internet connectivity. However, successful implementation depends on: (1) reliable data collection and communication infrastructure, (2) automated model identification and disturbance predictor training, and (3) robust fault detection and fallback to Baseline control when model accuracy degrades.

Cloud-based deployment also requires resilience mechanisms for communication outages, server failures, or thermostat malfunctions. Essential safeguards include: automatic reversion to baseline schedules during communication loss, local storage of backup schedules on thermostats, and watchdog timers to detect control system failures. These fault tolerance mechanisms are critical for maintaining occupant comfort and building owner confidence.

6.7. Future Research Directions

Several promising directions for future research emerge from this work:

Stochastic and Robust MPC: The stochastic nature of unmeasured disturbances points toward integrating stochastic or robust MPC formulations. Such approaches could explicitly account for uncertainty in disturbance forecasts through chance constraints or robust optimization, potentially improving closed-loop performance and reliability. Scenario-based MPC or tube-based MPC could provide probabilistic comfort guarantees even under uncertain disturbances. Additionally, incorporating stochastic weather forecast uncertainty into the MPC formulation could further improve robustness, as weather prediction errors propagate through the building model and affect temperature predictions.

Adaptive and Online Learning: The current hybrid model requires offline training on historical data. Future work should explore online learning methods that continuously update the disturbance predictor as new data becomes available. Adaptive control strategies could detect changes in building usage patterns or occupant behavior and automatically retrain models, improving long-term robustness.

Multi-Building Coordination and Thermally Coupled Multi-Zone Modeling: Extending the framework to coordinate multiple SMCBs could enable portfolio-level services such as frequency regulation or community-scale demand response. The current work demonstrates the coordination of multiple HPs serving thermally independent zones. By replacing the single-zone gray-box models with a coupled multi-zone gray-box model, the framework can be extended to account for inter-zone thermal interactions (see Appendix A for the multi-zone model formulation). This extension is left for future work. However, by design, the proposed framework targets buildings conditioned by HPs with dedicated thermostats for each unit; buildings served by centralized systems such as variable air volume (VAV) would require modifications to the optimization problem formulation.

Economic Optimization and Grid Services: Beyond time-of-use pricing, future work should explore MPC participation in real-time electricity markets and ancillary service programs. This includes quantifying the value of demand flexibility provided by SMCBs and designing MPC strategies that balance building-level costs with grid-level services such as ramping support or renewable energy integration.

Generalization Across Climate Zones and Building Types: The current study focused on mild California weather conditions. Validating performance across diverse climate zones (e.g., hot-humid, cold, mixed) and building types (e.g., retail, restaurants, schools) is essential for establishing the broader applicability of the hybrid MPC approach. Different climates may exhibit different disturbance patterns and control challenges.

Occupant-Centric Control: Incorporating occupant feedback and preferences into the MPC framework could improve comfort and acceptance. This might include learning individual thermal preferences, detecting real-time occupancy for demand-controlled ventilation, or providing occupants with control over price sensitivity parameters. However, given the stochastic nature of occupant behavior, the MPC should adapt its optimization criteria by detecting manual thermostat overrides in real-time and temporarily excluding overridden thermostats from the control loop [40].

7. Conclusion

This work presented a hybrid MPC framework that integrates first-principles building thermal modeling with data-driven disturbance estimation, validated through both simulation and experimental studies. The results demonstrate that the hybrid approach can significantly improve prediction accuracy and closed-loop control performance compared to purely model-based MPC when unmeasured disturbances are present.

In simulation studies, the Hybrid MPC achieved a 53.9% heating load shift during morning peak hours and a 38.4% cooling load shift during afternoon peak periods, while reducing peak demand by 16.0% for heating season and 4.4% for cooling season compared to the Baseline. Daily energy cost reductions reached 31.5% in heating season scenarios and 26.6% in cooling season scenarios, demonstrating performance that closely approaches or exceeds that of an ideal MPC with perfect

disturbance knowledge. Overall, the Hybrid MPC consistently outperforms the standard vanilla MPC and achieves performance comparable to the ideal MPC.

In experimental testing at the FLEXLAB facility, the Hybrid MPC reduced daily HVAC energy costs by 8.7% and peak building demand by 29.2% relative to the Baseline, demonstrating comparable benefits to the 11.6% cost savings and 23.2% peak reduction achieved by the ideal MPC. These results confirm that the hybrid modeling approach can deliver substantial load shifting, demand reduction, and cost savings in real-world operation without requiring direct measurement of internal heat gains.

However, the study also highlights key limitations. The Hybrid MPC's performance is sensitive to the quality and diversity of the training data, and stochastic unmeasured disturbances remain challenging to model accurately. In practice, extended data collection-on the order of one month or more-appears necessary for robust disturbance modeling. Furthermore, the stochastic nature of the disturbance motivates future work in stochastic and robust MPC formulations, which could explicitly incorporate uncertainty into the optimization problem.

Overall, while the Hybrid MPC approach offers a practical pathway to improving control performance under realistic operating conditions, its effectiveness depends on careful disturbance modeling, sufficient training data, and consideration of operational factors such as load profiles and price signals. Future research will focus on integrating stochastic MPC and adaptive learning strategies to further enhance robustness and adaptability in the presence of uncertainty. Ultimately, this study demonstrates that intelligent control can be democratized for small commercial buildings, unlocking a critical demand flexibility resource for sustainable urban energy systems.

CRediT authorship contribution statement

Sang Woo Ham: Writing – original draft, Visualization, Validation, Software, Methodology, Investigation, Formal analysis, Data curation; **Donghun Kim:** Supervision, Software, Project administration, Methodology, Funding acquisition, Conceptualization; **Lazlo Paul:** Writing – review & editing, Software, Data curation; **Armando Casillas:** Investigation, Data curation, Conceptualization; **Anand Krishnan Prakash:** Software, Data curation; **Richard Brown:** Supervision, Project administration, Funding acquisition, Conceptualization; **Marco Pritoni:** Writing – review & editing, Supervision, Project administration, Funding acquisition.

Data availability

The authors do not have permission to share data.

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

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) through the NextGen HVAC Innovation Challenge program and by California Energy Commission through grant EPC-19-013.

Appendix A. Multi-Zone Model Extension

The single-zone gray-box model (Eq. 5) can be extended to a multi-zone formulation that captures inter-zone thermal interactions. Fig. A.24 illustrates a three-zone thermal network, where each zone retains the same two-node structure (thermal mass $T_{w,i}$ and zone air $T_{za,i}$) as the single-zone model. Adjacent zones are thermally coupled through inter-zone resistances $R_{zz,ij}$, which represent heat transfer through shared walls or partitions between zones i and j. The state-space formulation follows the same structure as Eq. 5 but with an augmented state vector containing all zone states and off-diagonal coupling terms in the system matrix corresponding to inter-zone heat transfer. For zones without shared boundaries, the corresponding inter-zone resistance approaches infinity and the coupling terms vanish, recovering the independent single-zone formulation used in this study. The remaining hybrid modeling workflow (system identification, ID extraction, LSTM prediction) applies identically to each zone. Further details on multi-zone gray-box modeling can be found in [40].

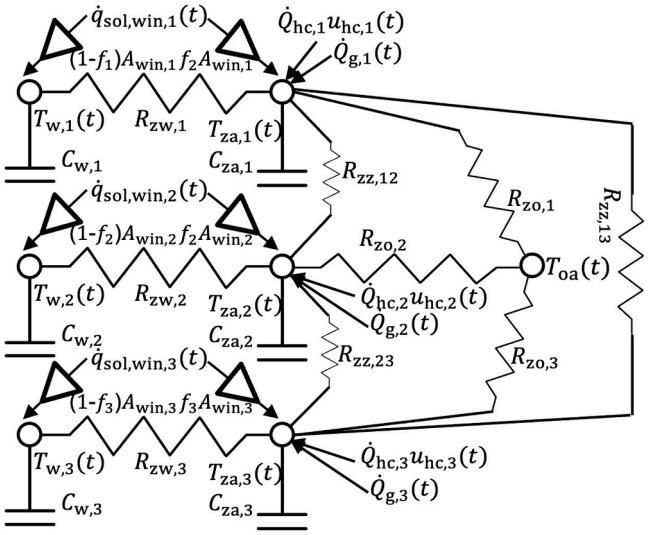

Fig. A.24. Multi-zone thermal network model showing inter-zone thermal interactions through shared wall resistances ($R_{zz,ij}$).

Appendix B. Sensitivity Analysis

A systematic sensitivity analysis is conducted over a discrete grid of the peak demand weight $\omega_d \in \{0, 0.1, 1, 10, 25, 50\}$ and the demand target fraction $\delta \in \{0.55, 0.70, 0.85, 1.00\}$ for all three MPC variants (MPC_{ideal}, MPC_{hybrid}, $MPC_{vanilla}$) under both cooling and heating seasons. Five performance metrics are evaluated: peak-time energy saving, energy cost saving, peak demand saving, comfort violation saving, and compressor cycling ratio saving, each expressed as a percentage relative to the baseline.

B.1. Trade-off analysis

Fig. B.25 presents pairwise trade-off plots for three metric pairs-cost saving vs. peak-time saving, peak demand saving vs. peak-time saving, and peak demand saving vs. cost saving-across all (ω_d, δ) combinations. Marker color distinguishes the MPC variant, marker shape encodes δ, and marker size encodes ω_d; circled points indicate the parameter combination that maximizes the sum of the two plotted metrics for each MPC variant. Cost saving and peak-time saving are largely proportional with respect to ω_d for both seasons; however, the heating season exhibits a dramatic drop in cost saving when ω_d exceeds approximately 25, due to excessive pre-heating that shifts energy into off-peak periods but increases total consumption. When $\omega_d = 0$ and $\delta = 1.0$, the optimizer focuses exclusively on cost minimization without demand constraints, yielding the highest cost and peak-time savings. Peak demand saving is inversely related to both cost saving and peak-time saving, confirming

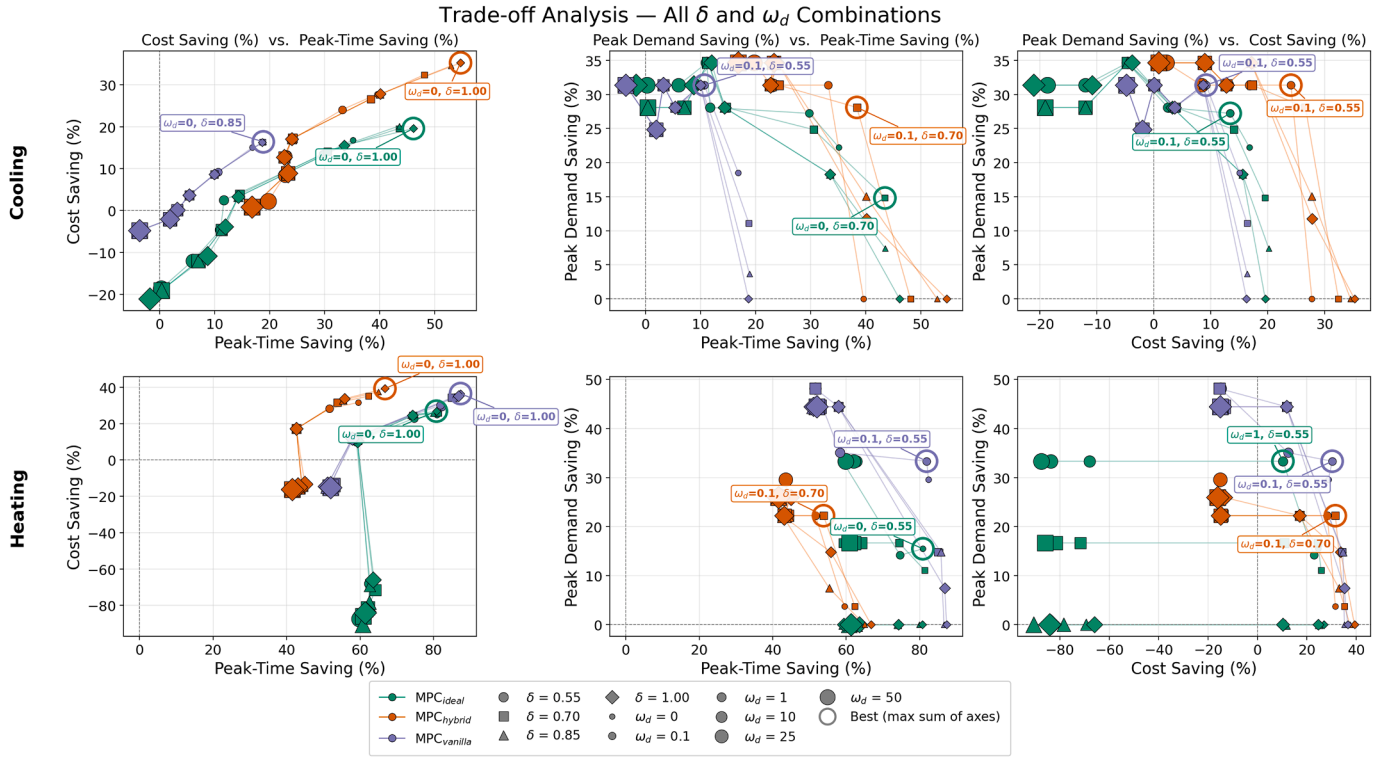

Fig. B.25. Trade-off analysis across all ω_d and δ combinations. Columns show pairwise trade-offs between cost saving vs. peak-time saving (left), peak demand saving vs. peak-time saving (center), and peak demand saving vs. cost saving (right); rows correspond to cooling (top) and heating (bottom) seasons. Marker color indicates the MPC variant, marker shape encodes δ, and marker size encodes ω_d (larger markers correspond to higher demand penalty weights). Circled points indicate the parameter combination that maximizes the sum of the two plotted metrics for each MPC variant. Thin lines connect points sharing the same MPC variant and δ across increasing ω_d.

the fundamental trade-off between load-shifting and demand-limiting objectives.

B.2. Sensitivity to δ at fixed $\omega_d = 0.1$

Fig. B.26 isolates the effect of δ while holding $\omega_d = 0.1$. Peak-time saving and cost saving remain largely insensitive to δ across all MPC variants in both seasons, whereas peak demand saving decreases as δ increases (i.e., the demand constraint is relaxed). The most restrictive setting, $\delta = 0.55$, achieves the highest peak demand saving but slightly reduces cost and peak-time savings. The setting $\delta = 0.70$ is therefore selected as a balanced choice that maintains strong demand reduction without noticeably penalizing economic performance.

B.3. Sensitivity to ω_d at fixed $\delta = 0.7$

Fig. B.27 shows the impact of ω_d with δ fixed at 0.7. Peak-time saving and cost saving decrease monotonically with increasing ω_d, while peak demand saving increases. The transition is sharpest between $\omega_d = 0.1$ and $\omega_d = 1$; beyond $\omega_d = 10$, additional demand reduction is marginal while cost penalties grow substantially, particularly in the heating season. The value $\omega_d = 0.1$ is identified as a balanced operating point that retains most of the cost-optimal performance while providing meaningful peak demand reduction.

B.4. Heatmap overview

Figs. B.28 and B.29 present heatmaps of all five metrics across the full (ω_d, δ) grid for the cooling and heating seasons, respectively. All MPC variants achieve positive comfort violation savings across the entire parameter space, confirming that thermal comfort is consistently maintained. The cycling ratio saving column reveals a notable distinction: $\text{MPC}_{\text{ideal}}$, which tracks target schedules more precisely, tends to increase compressor cycling (negative cycling saving), whereas $\text{MPC}_{\text{hybrid}}$ achieves comparable cost and peak-time performance with lower cycling ratios owing to the smoothing effect of its less precise predictions. $\text{MPC}_{\text{vanilla}}$ exhibits the least sensitivity to both ω_d and δ but delivers lower overall savings. Performance differences between $\text{MPC}_{\text{ideal}}$ and $\text{MPC}_{\text{hybrid}}$ vary across seasons, with the hybrid variant showing competitive or superior results in most cooling-season configurations.

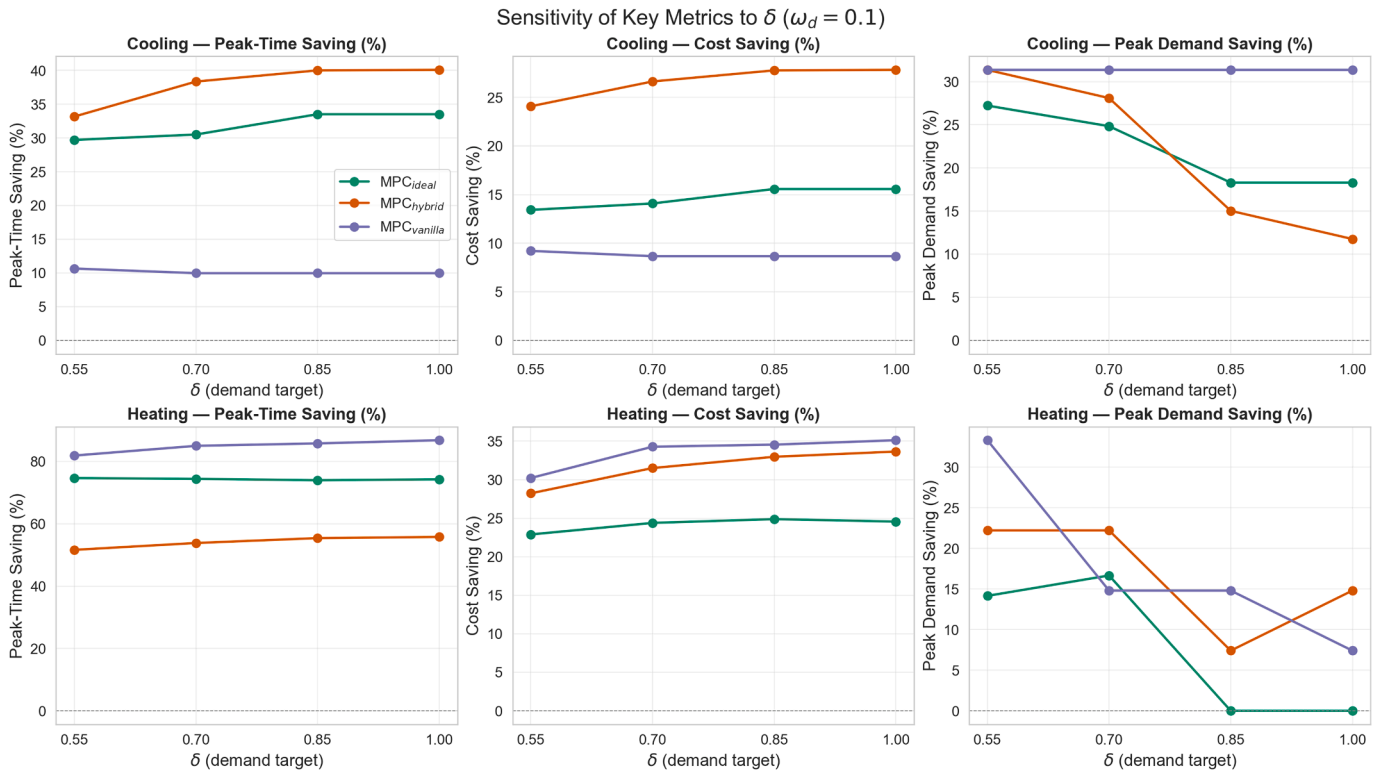

Fig. B.26. Sensitivity of key performance metrics to the demand target fraction δ at fixed $\omega_{\mathrm{d}} = 0.1$. Columns correspond to peak-time saving (left), cost saving (center), and peak demand saving (right); rows correspond to cooling (top) and heating (bottom) seasons. Peak-time and cost savings are largely insensitive to δ, whereas peak demand saving decreases as δ increases. The setting $\delta = 0.70$ offers a balanced trade-off.

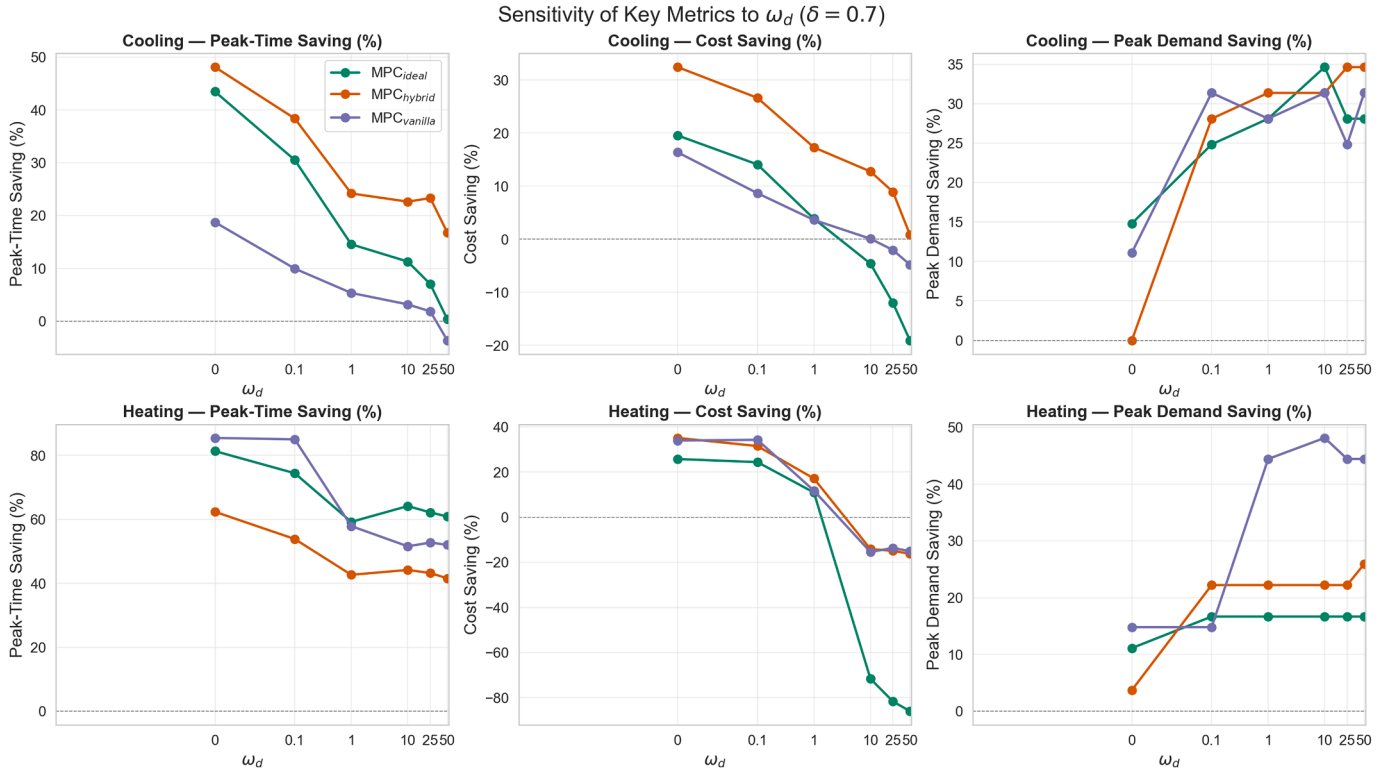

Fig. B.27. Sensitivity of key performance metrics to the demand penalty weight ω_{d} at fixed $\delta = 0.7$. Columns correspond to peak-time saving (left), cost saving (center), and peak demand saving (right); rows correspond to cooling (top) and heating (bottom) seasons. Peak-time and cost savings decrease monotonically with increasing ω_{d}, while peak demand saving increases. The value $\omega_{\mathrm{d}} = 0.1$ is identified as a balanced operating point.

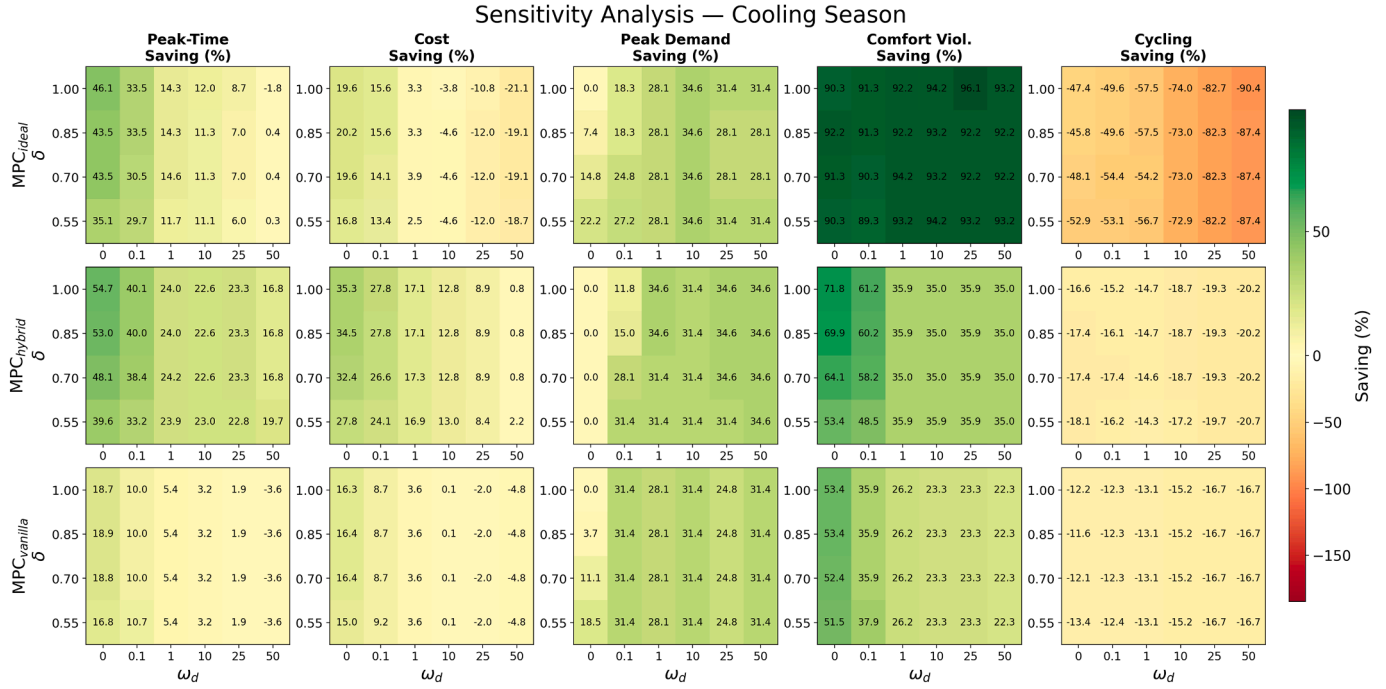

Fig. B.28. Sensitivity heatmaps for the cooling season across the full (ω_d, δ) grid. Rows correspond to MPC_{ideal} (top), MPC_{hybrid} (middle), and $MPC_{vanilla}$ (bottom); columns correspond to peak-time saving, cost saving, peak demand saving, comfort violation saving, and cycling ratio saving. Cell values indicate percentage savings relative to the baseline. A single diverging colormap (red–yellow–green) is used across all panels, centered at zero.

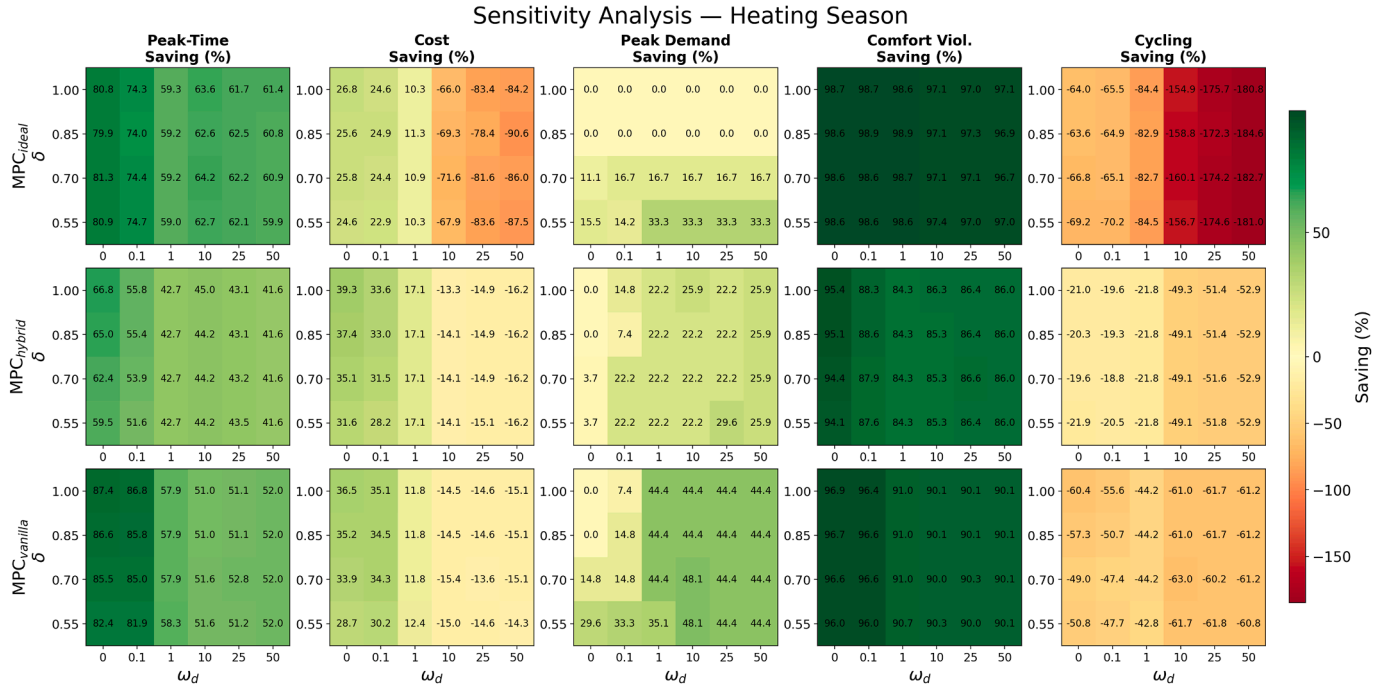

Fig. B.29. Sensitivity heatmaps for the heating season across the full (ω_d, δ) grid. Layout and encoding are identical to Fig. B.28. MPC_{ideal} follows target schedules most precisely, yielding the highest comfort violation savings but increased cycling ratios. MPC_{hybrid} achieves comparable cost and peak-time performance with lower cycling owing to the smoothing effect of its predictions.

References

[1] NERC, (2024). long-term reliability assessment, https://www.nerc.com/pa/RAPA/ra/Reliability

[2] NERC, Urgent need for resources over 10-year horizon as electricity demand growth accelerates, 2024 LTRA finds, https://www.nerc.com/news/Headlines

[3] U.S. EIA, U.S. energy consumption by source and sector, 2024 https://www.eia.gov/energyexplained/us-energy-facts/images/ consumption-by-source-and-sector.pdf.

[4] M. Neukomm, V. Nubbe, R. Fares, Grid-interactive efficient buildings technical report series: Overview of research challenges and gaps https://doi.org/10.2172/1577966.

[5] J.E. Braun, Reducing energy costs and peak electrical demand through optimal control of building thermal storage 96, 876–888. https://eta-publications.lbl.gov/sites/default/files/dbc-paper90.pdf.

[6] J.A. Candanedo, V.R. Dehkordi, M. Stylianou, Model-based predictive control of an ice storage device in a building cooling system 111, 1032–1045. https://doi.org/10.1016/j.apenergy.2013.05.081

[7] F. Oldewurtel, A. Ulbig, M. Morari, G. Andersson, Building control and storage management with dynamic tariffs for shaping demand response, in: 2011 2nd IEEE PES International Conference and Exhibition on Innovative Smart Grid Technologies, IEEE, pp. 1–8. https://doi.org/10.1109/isgteurope.2011.6162694

[8] G. Ceusters, R.C. Rodríguez, A.B. García, R. Franke, G. Deconinck, L. Helsen, A. Nowé, M. Messagie, L.R. Camargo, Model-predictive control and reinforcement learning in multi-energy system case studies 303, 117634. https://doi.org/10.1016/j.apenergy.2021.117634

[9] D. Patteeuw, G.P. Henze, L. Helsen, Comparison of load shifting incentives for low-energy buildings with heat pumps to attain grid flexibility benefits 167, 80–92. https://www.sciencedirect.com/science/article/pii/S0306261916300162.

[10] F.A. Qureshi, C.N. Jones, Hierarchical control of building HVAC system for ancillary services provision 169, 216–227. https://doi.org/10.1016/j.enbuild.2018.03.004

[11] M. Avci, M. Erkoc, A. Rahmani, S. Asfour, Model predictive HVAC load control in buildings using real-time electricity pricing 60, 199–209. https://doi.org/10.1016/j.enbuild.2013.01.008

[12] G. Bianchini, M. Casini, D. Pepe, A. Vicino, G.G. Zanvettor, An integrated MPC approach for demand-response heating and energy storage operation in smart buildings, in: 2017 IEEE 56th Annual Conference on Decision and Control (CDC), IEEE, pp. 3865–3870. https://doi.org/10.1109/cdc.2017.8264228

[13] F. Oldewurtel, A. Ulbig, A. Parisio, G. Andersson, M. Morari, Reducing peak electricity demand in building climate control using real-time pricing and model predictive control, in: 49th IEEE Conference on Decision and Control (CDC), IEEE, pp. 1927–1932. https://doi.org/10.1109/cdc.2010.5717458

[14] T. Hilliard, M. Kavgic, L. Swan, Model predictive control for commercial buildings: trends and opportunities 10, 172–190. https://doi.org/10.1080/17512549.2015.1079240

[15] M. Killian, M. Kozek, Ten questions concerning model predictive control for energy efficient buildings 105, 403–412. https://doi.org/10.1016/j.buildenv.2016.05.034

[16] J. Drgoňa, J. Arroyo, I. Cupeiro Figueroa, D. Blum, K. Arendt, D. Kim, E.P. Ollé, J. Oravec, M. Wetter, D.L. Vrabie, L. Helsen, All you need to know about model predictive control for buildings 50, 190–232. https://doi.org/10.1016/j.arcontrol.2020.09.001

[17] U.S. Energy Information Administration (EIA), Commercial buildings energy consumption survey (CBECS) data.

[18] Y. Li, Z. O'Neill, L. Zhang, J. Chen, P. Im, J. DeGraw, Grey-box modeling and application for building energy simulations - A critical review 146, 111174. https://doi.org/10.1016/j.rser.2021.111174

[19] A.J. Khabbazi, E.N. Pergantis, L.D. Reyes Premer, P. Papageorgiou, A.H. Lee, J.E. Braun, G.P. Henze, K.J. Kircher, Lessons learned from field demonstrations of model predictive control and reinforcement learning for residential and commercial HVAC: A review 399, 126459. https://doi.org/10.1016/j.apenergy.2025.126459

[20] S. Katipamula, R.M. Underhill, J.K. Goddard, others, Small-and medium-sized commercial building monitoring and controls needs: A scoping study. https://www.osti.gov/biblio/1063081.

[21] K. Fenaughty, E. Martin, D. Parker, B. Nigusse, An integrated control method for supplemental minisplit heat pumps in existing homes https://doi.org/10.2172/1814927.

[22] L. Paul, P.F. De Andrade, S. Ham, M. Pritoni, R. Brown, J. Feng, Open Building Operating System: an Open-Source Grid Responsive Control Platform for Buildings, in: 2023 ASHRAE Annual Conference at Tampa, FL,

[23] R. Ford, M. Pritoni, A. Sanguinetti, B. Karlin, Categories and functionality of smart home technology for energy management 123, 543–554. https://doi.org/10.1016/j.buildenv.2017.07.020

[24] T.S. Pedersen, P. Andersen, K.M. Nielsen, Central control of heat pumps for smart grid purposes tested on single family houses, in: 2013 10th IEEE INTERNATIONAL CONFERENCE ON NETWORKING, SENSING AND CONTROL (ICNSC), IEEE, pp. 118–123. https://doi.org/10.1109/icnsc.2013.6548722

[25] A. Afram, F. Janabi-Sharifi, Supervisory model predictive controller (MPC) for residential HVAC systems: Implementation and experimentation on archetype sustainable house in Toronto 154, 268–282. https://doi.org/10.1016/j.enbuild.2017.08.060

[26] S. Thorsteinsson, A.A.S. Kalaee, P. Vogler-Finck, H.L. Stærmose, I. Katic, J.D. Bendtsen, Long-term experimental study of price responsive predictive control in a real occupied single-family house with heat pump 347, 121398. https://doi.org/10.1016/j.apenergy.2023.121398

[27] Y. Jiang, A. Ejenakevwe, J. Wang, C.Y. Tang, L. Song, Development, Implementation, and Impact Analysis of Model Predictive Control-Based Optimal Precooling Using Smart Home Thermostats 113790. https://doi.org/10.1016/j.enbuild.2023.113790

[28] F. Bünning, B. Huber, P. Heer, A. Aboudonia, J. Lygeros, Experimental demonstration of data predictive control for energy optimization and thermal comfort in buildings 211, 109792. https://doi.org/10.1016/j.enbuild.2020.109792

[29] D. Kim, J.E. Braun, Development, implementation and performance of a model predictive controller for packaged air conditioners in small and medium-sized commercial building applications 178, 49–60. https://doi.org/10.1016/j.enbuild.2018.08.019

[30] D. Kim, J.E. Braun, MPC solution for optimal load shifting for buildings with ON/OFF staged packaged units: Experimental demonstration, and lessons learned 266, 112118. https://doi.org/10.1016/j.enbuild.2022.112118

[31] J. Braun, N. Chaturvedi, An Inverse Gray-Box Model for Transient Building Load Prediction 8, 73–99. https://doi.org/10.1080/10789669.2002.10391290

[32] D. Kim, J. Cai, K.B. Ariyur, J.E. Braun, System identification for building thermal systems under the presence of unmeasured disturbances in closed loop operation: Lumped disturbance modeling approach 107, 169–180. https://doi.org/10.1016/j.buildenv.2016.07.007

[33] S.W. Ham, D. Kim, Hybrid Modeling Approach For Better Identification Of Building Thermal Network Model And Improved Prediction Paper 420. https://docs.lib.purdue.edu/ihpbc/420.

[34] S.W. Ham, D. Kim, Machine learning-enhanced hybrid modeling approach for better identification of a building thermal network model and improved prediction xxx, 117285. https://doi.org/10.1016/j.enbuild.2026.117285

[35] M. Blonsky, K. McKenna, J. Maguire, T. Vincent, Home energy management under realistic and uncertain conditions: A comparison of heuristic, deterministic, and stochastic control methods 325, 119770. https://doi.org/10.1016/j.apenergy.2022.119770

[36] B. Dong, Z. Li, S.M.M. Rahman, R. Vega, A hybrid model approach for forecasting future residential electricity consumption 117, 341–351. https://doi.org/10.1016/j.enbuild.2015.09.033

[37] G. Dhaliwal, B. Gunay, I. Beausoleil-Morrison, Parametric analysis of model predictive control for residential HVAC systems 351, 116662. https://doi.org/10.1016/j.enbuild.2025.116662

[38] S.W. Ham, L. Paul, D. Kim, M. Pritoni, R. Brown, J. Feng, Decarbonization of heat pump dual fuel systems using a practical model predictive control: Field demonstration in a small commercial building 361, 122935. https://doi.org/10.1016/j.apenergy.2024.122935

[39] F. Massa Gray, M. Schmidt, A hybrid approach to thermal building modelling using a combination of Gaussian processes and grey-box models 165, 56–63. https://doi.org/10.1016/j.enbuild.2018.01.039

[40] S.W. Ham, D. Kim, T. Barham, K. Ramseyer, The first field application of a low-cost MPC for grid-interactive K-12 schools: Lessons-learned and savings assessment 296, 113351. https://doi.org/10.1016/j.enbuild.2023.113351

[41] D. Kim, J. Cai, J.E. Braun, K.B. Ariyur, System identification for building thermal systems under the presence of unmeasured disturbances in closed loop operation: Theoretical analysis and application 167, 359–369. https://doi.org/10.1016/j.enbuild.2017.12.007

[42] Lawrence Berkeley National Laboratory, FLEXLAB the world's most advanced integrated building and grid technologies testbed. https://flexlab.lbl.gov/.

[43] B.F. Gerke, M. Stuebs, S. Murthy, A. Khandekar, P. Cappers, R.E. Brown, M.A. Piette, Potential bill impacts of dynamic electricity pricing on California utility customers. https://escholarship.org/uc/item/6fr1v1z9.

[44] S. Katipamula, J. Haack, G. Hernandez, B. Akyol, J. Hagerman, VOLTTRON: An Open-Source Software Platform of the Future 4 (n.d.), 15–22. https://doi.org/10.1109/MELE.2016.2614178

[45] L. Paul, F. De Andrade Pereira, A.K. Prakash, S.W. Ham, J.D. Feng, R. Brown, M. Pritoni, Open building operating system: a grid-responsive semantics-driven control platform for buildings 1–18. https://doi.org/10.1080/23744731.2024.2444819