Building Technologies & Urban Systems Division
Energy Technologies Area
Lawrence Berkeley National Laboratory

The First Field Application of a Low-cost MPC for Grid-interactive K-12 Schools:
Lessons-Learned and Savings Assessments

Sang woo Ham[1], Donghun Kim[1], Tanya Barham[2], Kent Ramseyer[3]

[1]Building Technology & Urban Systems Division, Lawrence Berkeley National Laboratory, [2]Community Energy Labs, [3]Newport-Mesa Unified School District

Energy Technologies Area
October 2023

10.1016/j.enbuild.2023.113351

This work was supported by the Assistant Secretary for Energy Efficiency and Renewable Energy, Building Technologies Office, of the US Department of Energy under Contract No. DE-AC02-05CH11231.

Disclaimer:

Highlights

The First Field Application of a Low-cost MPC for Grid-interactive K-12 Schools: Lessons-Learned and Savings Assessment

Sang woo Ham, Donghun Kim, Tanya Barham, Kent Ramseyer

- First MPC demonstration in K-12 school buildings for peak demand reduction and load shifting.
- Experimentally demonstrated a scalable MPC for multiple RTUs without hardware retrofit.
- Hierarchical MPC is improved in consideration of practical challenges such as user's override and poor data resolutions.
- Achieved 24% peak demand reduction and showed 16% of load shifting potential.
- Share lessons learned during the demonstration of practical challenges.

The First Field Application of a Low-cost MPC for Grid-interactive K-12 Schools: Lessons-Learned and Savings Assessment

Sang woo Ham[a], Donghun Kim[a], Tanya Barham[b], Kent Ramseyer[c]

[a]*Building Technology & Urban Systems Division, Lawrence Berkeley National Laboratory, Berkeley, CA, USA*
[b]*Community Energy Labs, Portland, OR, USA*
[c]*Newport-Mesa Unified School District, Costa Mesa, CA, USA*

Abstract

K-12 schools are the largest energy consumers in the public sector, with their HVAC energy consumption representing the largest portion of their total energy use. While transitioning these schools to grid-interactive HVAC system operation through advanced controls offers significant financial and environmental benefits, and model predictive control (MPC) has been identified as a promising solution to achieve that, very few MPCs are affordable and have been deployed in K-12 schools. This situation raises concerns about the unclear real-world benefits of MPC technology among facility managers and industries. To address this gap, this paper presents a low-cost MPC solution that requires minimal control infrastructure costs and a unique field demonstration at a K-12 school, conducted for both cooling and heating seasons. This work adopted a previously developed MPC and extended it for use in the school application. The MPC aims to coordinate multiple packaged units to eliminate unnecessary peaks and shift cooling or heating loads in response to grid signals based on load conditions, while maintaining thermostat temperatures within school-defined bounds. Throughout the field tests, the MPC achieved a 24% reduction in peak demand during the cooling season and shifted cooling or heating loads by up to 16% in response to the school's utility tariff, considering load conditions, while also allowing end-users to override thermostat setpoints. The paper also discusses the limitations of this study and future research directions for better performance of the MPC at K-12 schools.

Keywords: Model Predictive Control, K-12 school, load flexibility, load shifting, building decarbonization

Nomenclature

AHU: Air handling unit

API: Application programming interface

DX: Direct expansion air-conditioner

HVAC: Heating, ventilation, and air-conditioning

HMPC: Hierarchical MPC

LD: Lumped disturbance

LMPC: Lower-level MPC

MPC: Model predictive control

RTF: Runtime fraction

RTU: Rooftop unit

RMSE: Root mean square error

SYSID: System identification

TABS: Thermally activated building system

TES: Thermal energy storage

TOU: Time-of-Use utility tariff

UMPC: Upper-level MPC

VAV: Variable-air-volume system

($\mathbf{A}(\cdot)$, $\mathbf{B_u}(\cdot)$, $\mathbf{B_w}(\cdot)$, $\mathbf{C}(\cdot)$): A state space model structure that maps θ to building dynamics (i.e., G_{u} and G_{w})

($\mathbf{A_d}(\cdot)$, $\mathbf{B_{d},}_u(\cdot)$, $\mathbf{B_{d},}_w(\cdot)$, $\mathbf{C_d}(\cdot)$): A discretized state space model of ($\mathbf{A}(\cdot)$, $\mathbf{B}_{\mathrm{u}}(\cdot)$, $\mathbf{B}_{\mathrm{w}}(\cdot)$, $\mathbf{C}(\cdot)$)

$A_{\mathbf{win},i}$**:** Effective window area of ith zone windows [kW/m^2]

$C_{\mathbf{w},i}$**:** Thermal capacitance of wall mass of ith zone [kWh/K]

$C_{\mathbf{za},i}$**:** Thermal capacitance of zone air of ith zone [kWh/K]

$\mathcal{D}(k)$**:** Set of measured data from time step from timestep from 1 to k.

ER**:** Electricity cost rate [$/kWh]

e_i**:** Zero mean white noise of ith zone

($\mathcal{F}(\cdot)$, $\mathcal{G}(\cdot)$): A state space model structure that maps ρ to lumped disturbance dynamics (i.e., H)

f_i**:** Convective fraction of the incident solar radiation of ith zone windows [-]

$G_{\mathbf{u}}$**:** A dynamic system that maps $\mathbf{u}$ to $\mathbf{y}$

$G_{\mathbf{w}}$**:** A dynamic system that maps $\mathbf{w}$ to $\mathbf{y}$

$G_{\mathbf{g}}$**:** A dynamic system that maps $\dot{Q}_{g,1:n}$ to $\mathbf{y}$

H**:** Dynamics of lumped output disturbances

j**:** Prediction time step

m**:** Number of measured inputs

$P_{\mathbf{RTU},i}$**:** ith RTU power [kW]

$P_{\mathbf{bldg}}$**:** Sum of all building powers [kW]

$N_{\mathbf{p}}$**:** Prediction horizon

N**:** Number of data

n**:** Number of RTUs

$\dot{Q}_{\mathbf{g},i}$**:** Unmeasured heat gains of ith zone [kW]

$\dot{q}_{\mathbf{sol,win},i}$**:** Incident solar radiation per area of ith zone windows [kW/m^2]

$\dot{Q}_{\mathbf{hc},i}$**:** Rated heating($\dot{Q}_{\mathrm{h},i}$)/cooling($\dot{Q}_{\mathrm{c},i}$) capacity of ith zone RTU [kW]

($R_{\mathbf{zw},i}$, $R_{\mathbf{zo},i}$**):** Thermal resistances between temperature nodes of ith zone [K/kW]

($T_{\mathbf{l}}, T_{\mathbf{u}}$**):** Lower and upper temperature bounds [°C]

$T_{\mathbf{za},i}$**:** Air temperature of ith zone(i.e., classroom) [°C]

$T_{\mathbf{w},i}$**:** Wall thermal mass temperature of ith zone [°C]

$T_{\mathbf{oa}}$**:** Outdoor air temperature [°C]

t, k**:** Continuous time and discrete time

u: Vector of control inputs (i.e., RTU compressor stages, $[u_{\mathrm{h},1:n}, u_{\mathrm{c},1:n}]$)

$u_{\mathbf{h},i}$, $u_{\mathbf{c},i}$**:** Heating and cooling stages of ith RTU

$\hat{\mathbf{u}}(k)$**:** Vector of estimated runtime fraction of RTUs for a sampling time k

$\bar{\mathbf{u}}(k)$**:** Vector of runtime fraction of RTUs for a sampling time k)

w: Vector of measured disturbances (i.e., $[T_{\mathrm{oa}}, \dot{Q}_{\mathrm{sol,win},1:n}]$)

x: Vector of state variables (i.e., $[T_{\mathrm{w},1:n}, T_{\mathrm{za},1:n}]$)

$\hat{\mathbf{x}}(k|j)$**:** Vector of estimated(predicted) state variables at time k from the data at j

y: Vector of measured thermostat temperatures for all zones [°C]

($\Gamma_{\mathbf{l}}, \Gamma_{\mathbf{u}}$**):** Temperature violations from lower and upper temperature bounds

δ**:** An upper bound of instantaneous power

ε_i**:** One step ahead prediction error of ith zone

ϵ_y**:** Unmeasured noise resulted from rounding in an integer thermostat.

$\boldsymbol{\zeta}$**:** Vector of internal state of lumped output disturbances

θ**:** Physical parameters consisting of thermal resistances and capacitances, $[C_{\mathrm{w},1:n}, C_{\mathrm{za},1:n}, R_{\mathrm{zw},1:n}, R_{\mathrm{zo},1:n}, f_{1:n}, A_{\mathrm{win},1:n}, \dot{Q}_{\mathrm{h},1:n}, \dot{Q}_{\mathrm{c},1:n}]$

$\boldsymbol{\nu}$**:** Vector of lumped output disturbances [°C]

ρ**:** Parameters that constructs dynamics of lumped output disturbances, i.e. H

($\omega_{\mathbf{l}}, \omega_{\mathbf{u}}$**):** Weights on optimization variables for (Γ_{l}, Γ_{u})

ω_d**:** Weight on optimization variables for δ

1. Introduction

For K-12 schools (from kindergarten to the 12th grade), with the urgent call for climate change mitigation efforts [1], various energy-related activities such as building retrofit [2] and decarbonization [3, 4] are actively undergoing. A recent study [5] shows that K-12 schools account for 60 million metric tons of carbon emission [6] and have the potential to reduce carbon intensities from 1.66-40.00 kg/m^2-yr to 0.41-18.69kg/m^2-yr in 2050. Transitioning to dynamic HVAC operations in response to the grid status (e.g., adjusting HVAC setpoints to shift or curtail loads during high carbon-emitting or peak demand periods) is also an important and urgent research topic, since 1) there are, for example, about 100,000 K-12 schools of the U.S. [7], 2) they are the largest energy consumer in the public sector [8], 3) the HVAC energy accounts for 46% of the total energy consumption in K-12 schools [9], and thus 4) there could be substantial environmental and financial potentials of the grid-interactive school operations. In addition, unlike other buildings, the grid-interactive operation could provide unique educational benefits of improving equity and community resilience [4]. Furthermore, thanks to the consistent HVAC system configuration in a majority of K-12 schools (i.e., single- or multiple-staged RTUs controlled by their own thermostats), there is also a potential to have a universally applicable control solution for those buildings.

One of the most widely studied control approaches for dynamic HVAC operations is Model Predictive Control (MPC) [10], which could optimize the operation of an HVAC system with given constraints by utilizing mathematical models for buildings and disturbance forecasts (e.g., weather). By including price signals from the grid (e.g., Time-of-Use (TOU) rate, real-time price, and carbon emission), MPC could provide flexible load management for grid services [1] such as load shifting and peak demand management.

Although numerous studies of using MPC for building applications have been carried out and published over the last several decades, the majority of them are still based on pure simulation studies and only a limited number of papers reported MPC demonstrations at actual buildings. The lack of references to the MPC real-world applications in the literature raises significant concerns about unclear real-world benefits and technology readiness levels of the MPC technology from building owners and industries, which is clearly a barrier to the wide adoption of MPC.

Following is a review of the literature regarding MPC field demonstrations for commercial buildings. Siroky et al. [11] implemented MPC for the radiant heating system of an office building in a university for about 8 weeks. It showed 15-28% of HVAC energy saving (depending on the insulation level) by controlling the supply water temperature. Bengea et al. [12] applied MPC for a 5-zone building, conditioned by variable-air-volume (VAV) systems. By optimizing the temperatures of mixed-air, supply air, and zone supply air and flow rate of mixed-air, the MPC showed 20% and 70% of cost savings for transition and heating season (1 week each). West et al. [13] tested the performance of MPC for the two office buildings. Both offices were conditioned by air handling units (AHUs). For 51 and 10 days of trials, each building showed 19% and 32% of heating energy reduction by controlling setpoint or supply air temperatures. Li et al. [14] demonstrated MPC for an office building with VAV systems. It showed more than 20% of HVAC energy reduction by optimizing the supply air flow rate and temperatures of AHU and terminal boxes. Hilliard et al. [15] applied MPC for a district heating system in a university building. They showed 29% energy saving for 4 months by optimizing morning start time and zone setpoints. Commercially available MPCs were evaluated for several commercial buildings (i.e., office, hospital, and high school) [16]. During the 7-15 months of trials, they showed 0-9% of HVAC energy saving by optimizing the supply air setpoint and duct static pressure of AHUs. Kim et al. [17] developed a plug-and-play MPC for multiple rooftop units (RTUs) for small/medium commercial buildings and demonstrated it in a gymnasium building with 4 RTUs. The MPC is designed to reduce energy cost and ON/OFF cycling, resulting in 8% of energy reduction. Though peak demand reduction is not directly included in the objective function, it also showed 40% peak demand reductions by avoiding over-cooling via the coordination of multiple RTUs. Coninck and Helsen [18] implemented an MPC

for an office in Brussels, Belgium to minimize the energy cost of two heat pumps and one gas boiler. It reduces 34-40% of energy cost by increasing efficient operations (i.e., more heat pumps use and preheating with low-temperature water). Kim and Braun [19] implemented an MPC algorithm for multiple ON/OFF RTUs in a small retail store to reduce both energy and peak demand. For several months of trial, it reduced about 12% and 18% of energy and peak demand by coordinating operations of RTUs.

MPC has also been widely applied for buildings with thermally activated building systems (TABS) for efficient system operations to reduce energy costs. Sturzenegger et al. [20] implemented MPC for 7 months of both heating and cooling seasons to an office building with TABS for main heating and cooling and AHUs for ventilation. The MPC controlled TABS, AHUs, and blinds to minimize energy costs and maintain thermal comfort. For the savings estimation, a simulation study was utilized to evaluate the savings of MPC, which showed 17% of nonrenewable primary energy savings. Drgona et al. [21] conducted the field test of a cloud-based MPC on an office building with a ground source heat pump (GSHP) and TABS. With optimized operations of HVAC systems, they showed 53.5% energy saving on GSHP and 36.9% of thermal comfort improvements. Freund and Schmitz [22] applied MPC for large office buildings with TABS system during three heating months. By optimizing the heating system operations (i.e., heat pump, district heating, and domestic hot water), they achieved 30% of energy savings compared to the rule-based control. Morovat et al. [23] showed the load-shifting potential of various control scenarios on an electrically heated school with geothermal heat pumps, hydronic radiant floor, and thermal energy storage (slab). The calibrated model from the measured data was used for case studies. On-peak heating energy consumption is shifted to off-peak hours by pre-heating the thermal energy storage via radiant floor heating. By applying these controls, about 40-65% of energy from on-peak hours can be shifted to off-peak hours.

Cotrufo et al. [24] developed a black-box model-based MPC to reduce greenhouse gas emissions by restricting natural gas usage in the morning heating peak time through pre-heating using electric baseboard heaters during non-peak time. Blum et al. [25] implemented an MPC in a large office building to reduce HVAC energy. It reduced 40% of HVAC energy consumption by optimizing the control operation of direct expansion air-conditioners (DXs) and economizers. Zhang et al. [26] developed an MPC software called SolarPlus and implemented it in a retail store with a photovoltaic (PV) panel and a battery. The MPC is designed to handle various grid services, including energy bill reduction, real-time pricing, load shifting, load shedding, load tracking, demand limiting, and load shifting. The field tests showed the MPC could handle various grid services, resulting in 12% energy cost and 34% peak demand reductions. Kim et al. [27] developed a campus-scale MPC with thermal energy storage (TES). The MPC is designed to reduce redundant PV generation and electricity consumption during the high carbon-emitting time (i.e., decarbonization). The field demonstration reduced the excess PV generation by about 25%, greenhouse gas emission by 10%, and peak electricity demand by 10%. Merema et al. [28] implemented two types of MPCs for different timesteps for a large building during the heating season. The goal of those MPCs is to minimize energy use while guaranteeing indoor air quality and thermal comfort. The performances were compared to similar condition periods, and the savings were expected at 10%-40% for electricity usage and 21%-55% for thermal energy use.

Considering the limited MPC demonstration works, it is not surprising that there is no field demonstration study of MPC at K-12 schools. The lack of MPC field study at schools may be attributed to several points including the followings: 1) HVAC systems are the most outdated part of the facility in K-12 schools [2] and they often lack building management systems and financial resources. As a result, MPC strategies in the literature which typically require numerous sensors and complex data communication networking, are often unaffordable for many schools. 2) There are many stakeholders such as teachers, facility operators, energy managers, school boards, and utility companies in K-12 schools [29, 30]. When some of them are not familiar with new technology and the benefits are not clear, they may be reluctant to invest in or implement even

mature technology. 3) Due to concerns about the potential impact on students and teachers, school facility operators/managers maintain a conservative approach. As a result, testing state-of-the-art technologies that have not undergone comprehensive verification in real buildings is very difficult to be deployed at schools.

Nonetheless, to achieve the aforementioned environmental, financial, and educational benefits of transitioning to the grid-interactive school operation, it is critical to have a low-cost MPC solution that can overcome financial scarcity, engage an industry to resolve liability issues and build trust among multiple stakeholders, and conduct MPC field studies that showcase compelling savings.

This paper is to partially fill the gap for the realization of the grid-interactive K-12 schools. We present a low-cost grid-interactive MPC solution and field test results conducted over a month for each cooling and heating season at a K-12 school, and share the practical challenges and lessons learned from the site implementation. The main contributions of this paper are:

- A unique implementation and real-word assessment of an MPC in a K-12 school, aimed at peak demand reduction and load shifting, without requiring hardware retrofitting
- Proposing low-cost MPC algorithms considering practical challenges
- Sharing lessons learned from the site MPC implementation

Section 2 describes the demonstration site and baseline HVAC operation data during typical days. Several practical challenges of applying the MPC are presented through data analysis. Section 3 presents technical details of the MPC from modeling to an MPC algorithm. In addition, some algorithm improvements to handle some practical challenges are described. The experimental details follow in Section 4, and Section 5 shows the experimental results regarding peak demand reduction and load shifting.

2. Building description and technical challenges of MPC in school buildings

2.1. Building description

The demonstration site is an elementary school located in California, U.S., and the MPC software and hardware discussed in the later sections were deployed to two pilot buildings in the school as indicated in Figure 1. Each building is composed of three classrooms, and each classroom is conditioned by its own thermostat and a rooftop unit (RTU).

The summary of building information, including HVAC, HVAC schedules, and utility tariff, is listed in Table 1. Buildings and RTUs are old, but WiFi-enabled thermostats were installed. The RTU consists of a direct expansion (DX) unit for cooling and natural gas for heating. The default thermostat schedule is programmed by the school district policy, but an end-user (e.g., teachers) can override it by changing the setpoint or occupied mode. The thermostat provides its operation data (zone air temperature, setpoint, compressor, and fan stage signals) through the thermostat vendor's cloud application programming interface (API), and it is stored in our cloud database (Section 4.1). It is important to mention that the data including thermostat temperature has an integer resolution. The MPC design issues associated with the override and integer resolution are discussed in Section 3.5. The school is under the SCE-TOU-GS-2 grandfathered tariff: the grandfathered time-of-use (TOU) has an on-peak price period of 12:00-18:00, although a typical TOU in CA has on-peak on 16:00-21:00 due to increased solar generation [31]. For the performance evaluation, we installed RTU-level power meters for Building 5 and a building-level power meter for Building 7 (see Figure 1). Thermostat and power meter data were collected at the 1-minute interval.

Figure 1: Demonstration buildings in an elementary school in CA.

2.2. Typical HVAC operation and data analytics in the cooling season

Figure 2 shows a typical cooling scenario of a classroom. The HVAC occupied schedule for the school is 6:00-16:00, but there are fan-only periods of 6:00-6:15 and 16:00-18:00 due to the ventilation policy. The default cooling setpoint is 23.3°C (74°F), but the teacher manually lowered the setpoint to 22.2°C (72°F) near 8:00 when students arrived. During the daytime, the room air temperature fluctuated in 22.2-23.3°C (72-74°F) as the RTU cycled with the thermostat deadband of 1.1°C (2°F). Therefore, the classroom was maintained at around 22.8°C (73°F) despite the setpoint of 22.2°C (72°F).

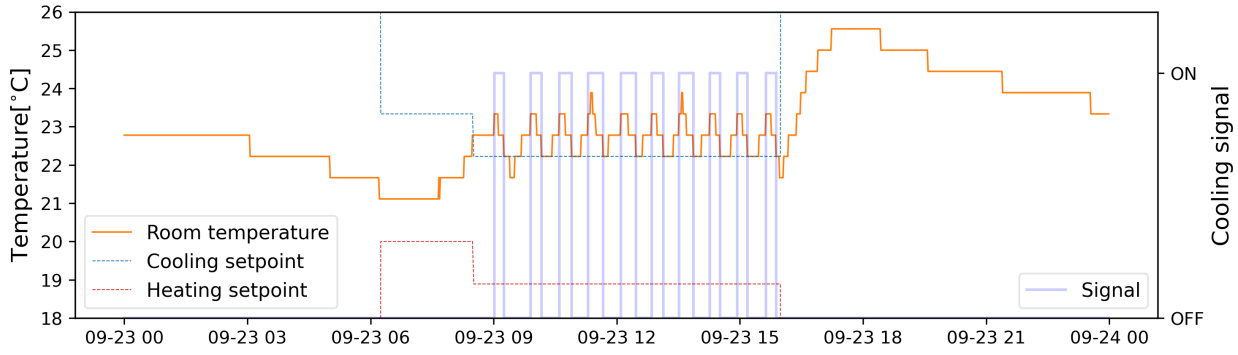

Figure 2: Typical RTU operation of one classroom for a day in the cooling season.

Figure 3 presents air temperatures (top) and cooling signals (middle) for the six classrooms, and electricity usage (bottom) for the two buildings for a typical cooling day. Though the default heating and cooling setpoints were 20.0°C (68°F) and 23.3°C (74°F), respectively, only one classroom followed the default setpoint (i.e., temperature cycles in 23.3-24.4°C). Most classrooms lowered setpoints to 22.2°C (i.e., temperature cycled in 22.2-23.3°C). One classroom had a setpoint at 20.6°C (69°F), although a setpoint

Table 1: Building information, HVAC, schedule, and tariff summary.

	Description
Building information	Building: Two buildings composed of 6 classrooms. • Classroom: Each classroom for 20 occupants (900 ft^2 or 84 m^2), built 1960. • Exterior wall: concrete exterior walls with R15 insulation. • Glazing: Double pane windows on the east or west side (WWR[1]: 0.18).
HVAC	• York ZYG05L2B RTU AC/Gas heating with a Wifi-enabled thermostat[2] for each classroom. • RTU rated power: 3.3kW in Bldg.5 and 2.3kW in Bldg.7.
Default schedule	Schedule and setpoints • Occupied 1: 6:00-6:15, 12.8-29.4°C (55-85°F) with Fan On. • Occupied 2: 6:15-16:00, 20.0-23.3°C (68-74°F) with Fan On. • Occupied 3: 16:00-18:00, 12.8-29.4°C (55-85°F) with Fan On. • Unoccupied: all except mentioned above, 7.2-35.0°C (45-95°F) with Fan Auto. Override: end-user can override the default schedule setpoints[3]
Utility tariff	Time-of-Use (SCE-TOU-GS-2[4]). • 00:00-08:00: Off-peak (0.05673 $/kWh). • 08:00-12:00 and 18:00-24:00: Mid-peak (0.10289 $/kWh). • 12:00-18:00: On-peak (0.20833 $/kWh). Demand charge 17.57 $/kW[5].

[1] WWR: window wall ratio.
[2] Venstar ColorTouch Model T8850 (integer resolution).
[3] Maximum heating and minimum cooling override setpoint is 22.2°C (72°F).
[4] SCE-TOU-GS-2: Southern California Edison Time-of-Use General service (grandfathered).
[5] Maximum demand is the measured maximum average kW during any 15-minute metered interval.

below 22.2°C was prohibited and locked according to the school district's policy: there was a software issue in this thermostat of the classroom which allowed violating the setpoint lower limit. In addition, one classroom's schedule ended early and was set to the unoccupied mode near 14:00, resulting in high temperatures of over 26°C.

Each RTU's cooling signal and the sum of all those signals are presented in the middle of Figure 3. All RTUs are single-stage units and operate independently so that, as indicated with the thick blue line, 5-6 RTUs could turn ON simultaneously. The impact of this simultaneous operation on power consumption can be found at the bottom of Figure 3. Since the peak demand charge is based on the electricity usage in a 15-min window (Table. 1), the 15-minute moving average of the total power is visualized in the bold-red line as well. The total building power spiked when the RTUs operated simultaneously and reached 16.48kW. Considering the non-RTU powers, estimated at approximately 5kW in 6:00-8:00, there is a high potential to reduce peak power demand when the RTUs' ON/OFF operations are coordinated and distributed.

2.3. Technical challenges: override and thermostat resolution

MPC is usually designed to minimize energy-related terms (e.g., energy cost) while maintaining room temperatures within a comfort band. The comfort band is typically defined by the weekly setpoint schedule (Table 1). However, when an end-user (e.g., a teacher) overrides the setpoint schedule (as shown in Figure

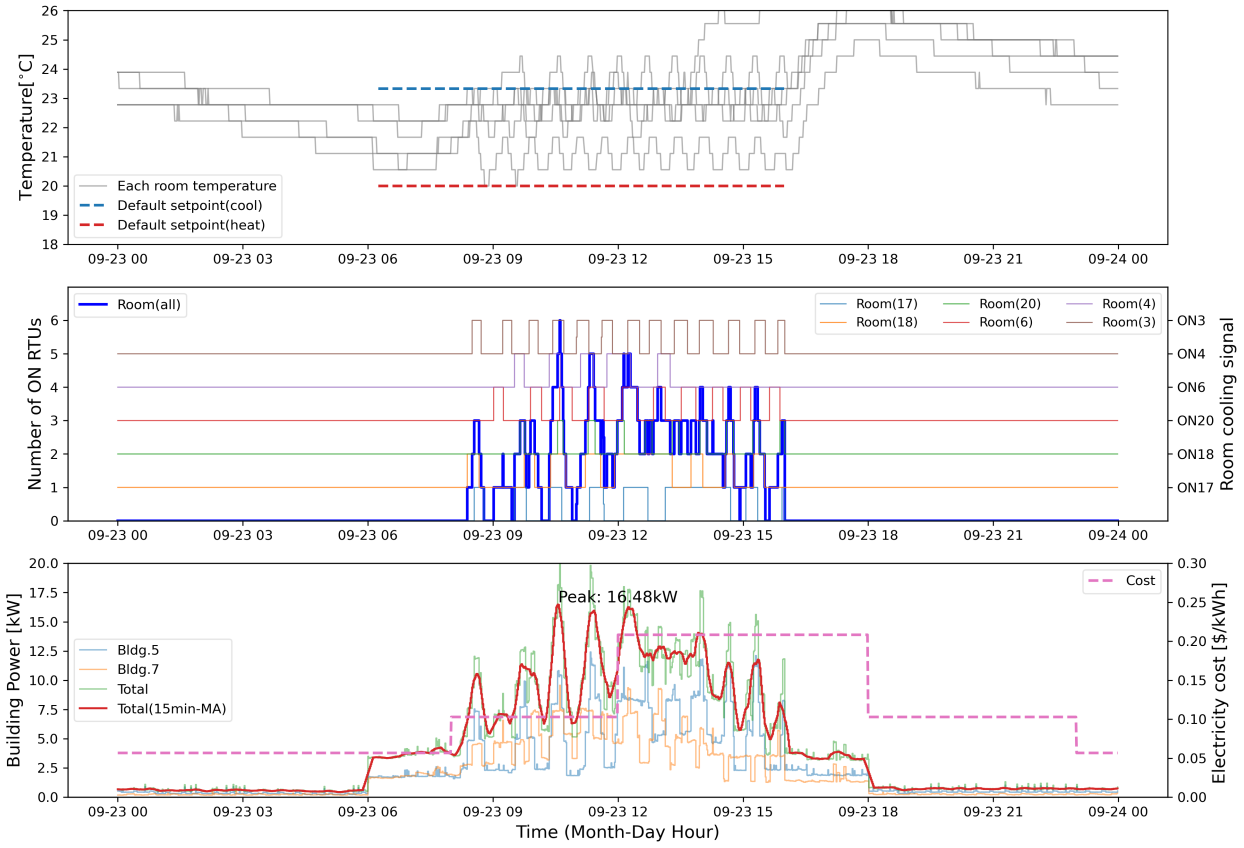

Figure 3: Air temperatures (top), cooling signals (mid), and the power(electricity) use (bottom) of six classrooms in the two buildings for a typical cooling day.

2, the comfort band constraint in MPC should be updated to reflect the personal preference. Otherwise, the end user's action will be overridden by the MPC again, resulting in conflicts with control actions and loss of credibility between the technology and the end user [32, 33]. Therefore, it is important to understand the end users' thermostat usage behaviors and design how to respond to the users' override actions.

Figure 4 shows the percentage of schedule overridden hours during RTU operation hours by days and classrooms (sorted by daily mean outdoor air temperature). To explain the override behaviors, the data is sorted by outdoor air temperature from left to right. One can notice that all classrooms except the conference room (Room 20) showed significant levels of override actions and that there is no apparent correlation between override and outdoor air temperature. Therefore, it can be interpreted as the override behavior is often stochastic (e.g., habits on arrival/leaving time, opening conditions of windows, doors, and blinds, etc.). We regularly surveyed teachers during the tests and discovered that the override actions were predominantly habitual. For example, the teacher lowered the setpoint near 8:00 (student arrival time) even though the room temperature was still at the setpoint (Figure 4). The same occurred right after lunchtime. Although it is not visualized in this paper, this habitual override behavior is consistently observed throughout most classrooms.

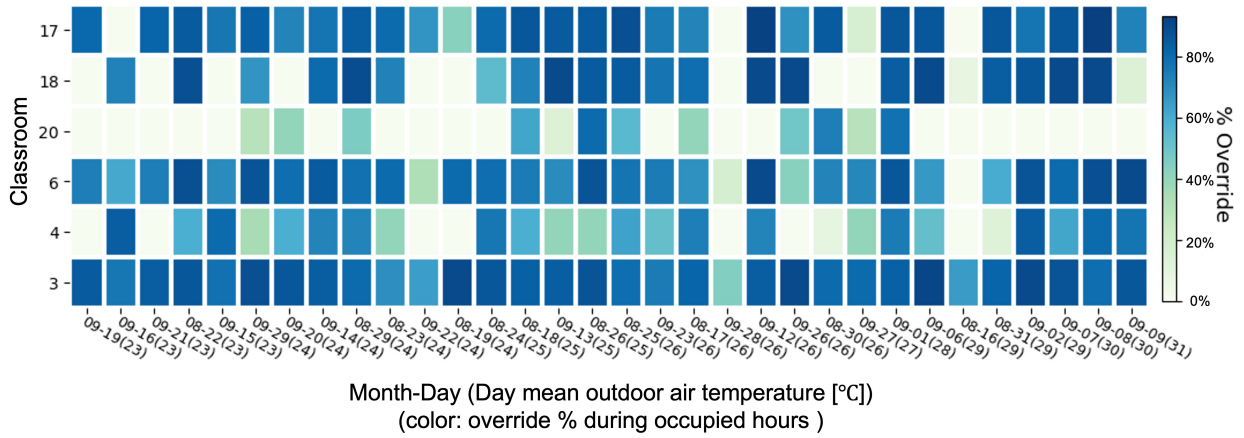

Figure 4: Percentage of schedule overridden hours during RTU operation hours by days and classrooms (sorted by daily mean outdoor air temperature).

It is important to mention that, in a typical programmable thermostat, there are two types of override actions on the end user's side: Setpoint override and Schedule override. They are designed to temporarily change the scheduled setpoint or mode (i.e., occupied or unoccupied). Figure 5 shows a snapshot of the thermostat dashboard on a default schedule day with various override behaviors. Room 05 is at a default schedule (i.e., cooling setpoint at 74°F in occupied hours). When the end user changes the setpoint by clicking the warmer or cooler buttons, it is Setpoint override (Room 06). Once this action happens, the Setpoint override typically stays until the next scheduled mode change unless the user sets a timer for the Setpoint override. Another override action is the Schedule override. This is originally designed to change the current schedule mode (i.e., from an occupied to an unoccupied mode or vice versa). Room 04 was at Schedule override, so it was supposed to follow unoccupied setpoints (e.g., 95°F). However, the end user changed setpoints as well (i.e., both overrides happened at the same time), so it had the setpoint of 72°F like Room 06. In this case, Rooms 04 and 06 have the same cooling setpoints, but Room 04's fan operation follows the unoccupied fan setting (see Table 1). Although the override behaviors at the school were habitual for most of the time, an MPC algorithm should incorporate the user's override action in the decision-making process since there might be the case when the user actually feels uncomfortable or when there is an actual schedule change (e.g., temporal absence).

The thermostat data resolution is an important quantity that affects the performance of MPC. Some thermostats provide only integer-valued data (i.e., no decimal digits), although their internal sensors have several decimal digits. Practically, this is not a problem for monitoring purposes, but it could cause performance degradation for a feedback controller relying on the thermostats. In Figure 6 (a), the room temperature, cooling setpoint, and cooling signal of one classroom are visualized for 2 hours in a day. Due to the poor resolution, the obtained temperature is an integer number (73°F) for most time and does not respond to the cooling action (although the actual temperature would respond within say 72.5-73.5°F). The lack of the system response gives a challenge in designing a feedback controller: Figure 6 (b) exemplifies the difficulty of a simple feedback controller which is to track a setpoint. It also challenges developing a good building thermal model because the transient response is poorly captured.

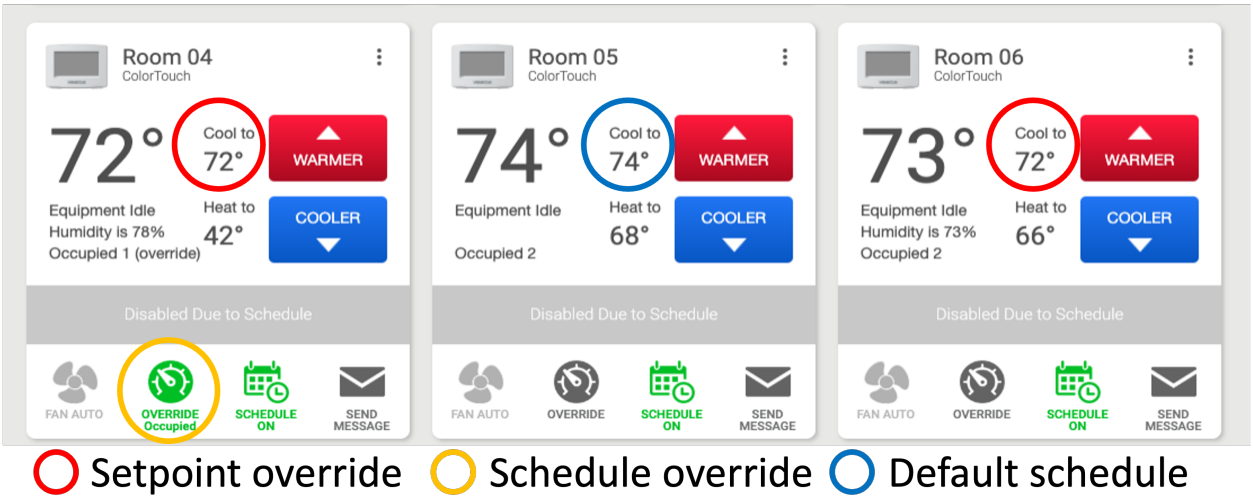

Figure 5: Snapshot of thermostat dashboard for a default schedule day with various override behaviors.

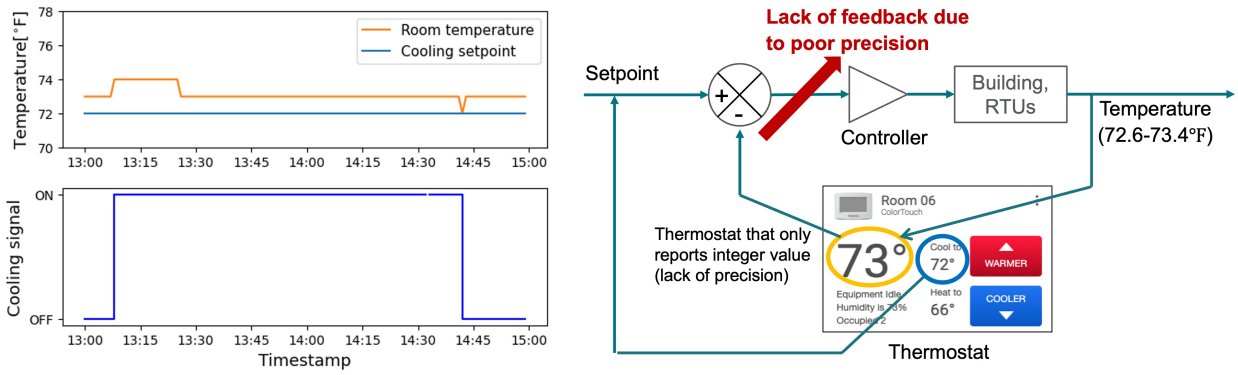

(a) Temperature/cooling response for 13:00-15:00

(b) Issues in integer thermostat in feedback control

Figure 6: Operation data and issues of an integer thermostat in feedback control.

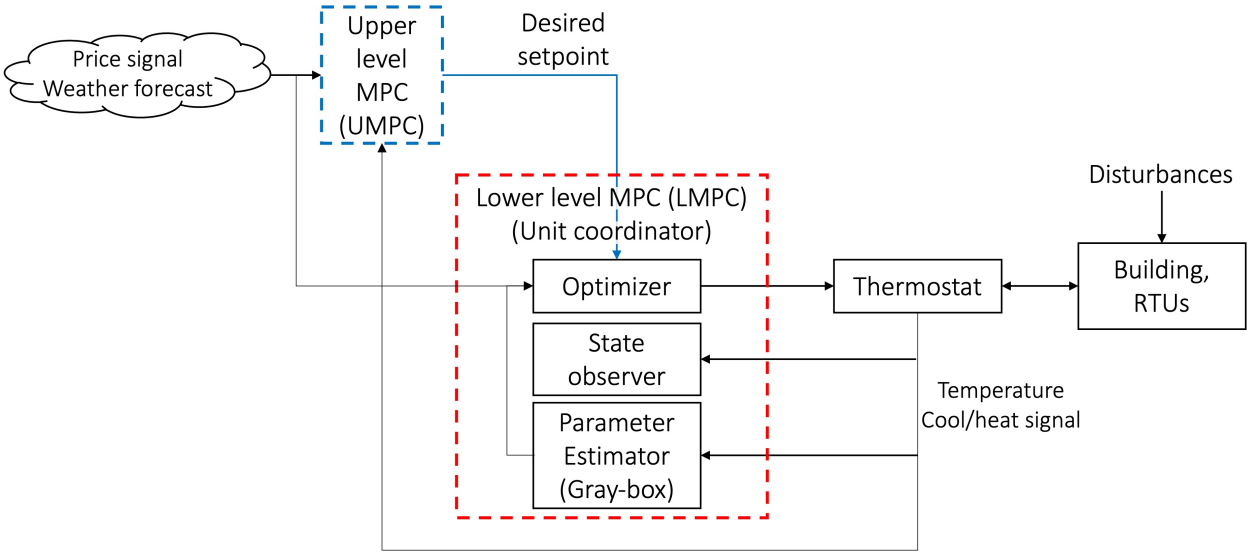

Figure 7: Schematic diagram of a hierarchical MPC for optimal load shifting and peak demand reduction

3. MPC design

3.1. Overview

In this study, the MPC algorithm introduced in [34] was applied and modified. The MPC only requires commercially viable web-enabled thermostats as on-site control hardware. The thermostat data, such as temperatures and compressor stages, is transmitted to the thermostat vendors' database via a standard Wi-Fi router and modem. A cloud MPC server solves optimization problems based on the thermostat signals retrieved from the database and weather forecast information available from public weather stations. It then sends the setpoint back to each thermostat. Due to the absence of the need for additional sensors or communication networks in schools, the MPC can offer a cost-effective solution by minimizing control infrastructure requirements. Moreover, for schools that have already installed smart thermostats, there is no installation cost involved.

The MPC has a hierarchical control structure as shown in Figure 7. The upper-level MPC (UMPC) generates the desired setpoint profile targeting to minimize the energy cost through load shifting by considering future weather and utility price signals for a long-term time horizon (e.g., one day). On the other hand, the lower-level MPC (LMPC) works as a slave controller which aims to track the desired setpoint while optimally coordinating multiple RTUs to reduce unnecessary simultaneous operations by looking at a short-term prediction horizon (e.g., 15 minutes) as depicted in Figure 8.

In this section, we discuss the rationale behind the hierarchical MPC structure and overview the implementation procedure in Section 3.2. The frameworks of adopted modeling approach and MPC are described in Section 3.3 and Section 3.4, respectively. Section 3.5 discusses modifications and customizations of the MPC aimed at addressing a certain technical difficulty discussed in Section 2.3.

3.2. Hierarchical MPC for unit staging and load sifting

The control algorithm structure follows the standard hierarchical control approach [35], which is commonly adopted in various systems to solve complex control design problems. These problems often involve

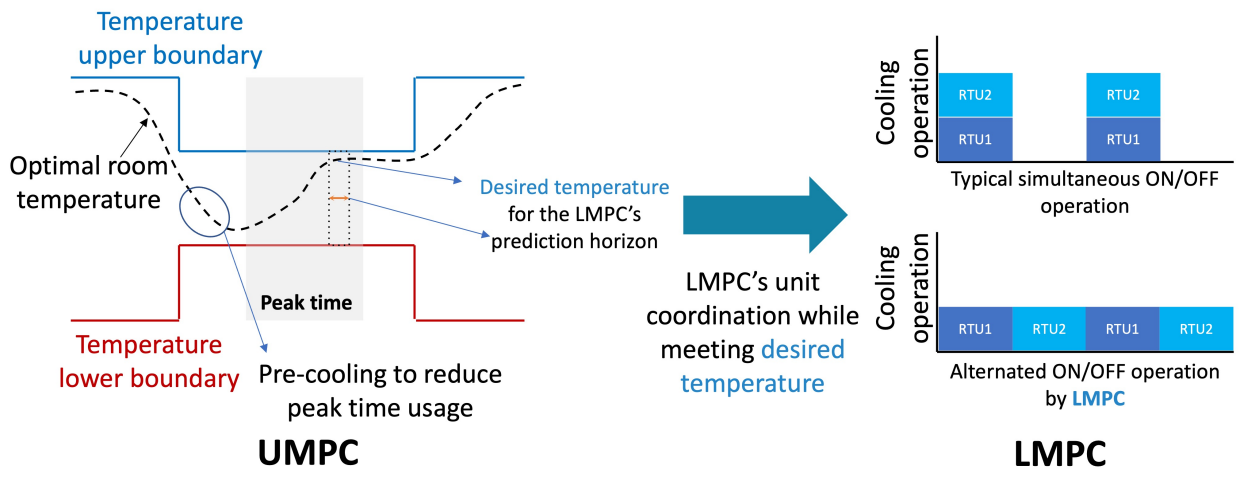

Figure 8: Conceptual diagram of the MPC for load shifting (UMPC) and peak demand reduction(LMPC).

different timescales and multiple control objectives that are difficult to address with a centralized control approach. The key advantages of hierarchical control include modularity, flexibility, and scalability, achieved by independently designing the control for each layer and by separating control tasks into distinct layers based on timescale separation. The hierarchical MPC [10, 36] belongs to the hierarchical control framework and, as such, inherits the same benefits. In our case, according to the hierarchical MPC structure, our control design tasks were split into the load shifting problem (a longer-term control problem) and the multiple unit staging problem to eliminate unnecessary peaks (a shorter-term control problem) to facilitate design, management, and updates.

Figure 9 depicts the implementation process of the proposed MPC. The initial step involved gathering building metadata, such as floor area, along with thermostat data including room temperatures and RTU operating signals, for system identification. This step corresponds to Section 3.3. The Modeling step encompasses system identification and model validation, which correspond to Section 3.3 and Section 5.1, respectively. Subsequently, an assessment plan was developed, as described in Section 4.2, and MPC and Baseline controls were implemented accordingly.

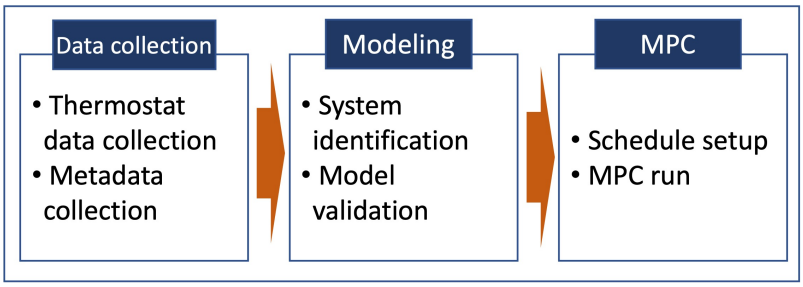

Figure 9: Implementation process of MPC.

3.3. Building system model and system identification

A gray-box model structure is one of the most common methods in the building control field to predict zone air temperature with given control input profiles and measured disturbances such as weather due to several advantages including the physical meaningfulness of each parameter, ease of using a priori knowledge, and guaranteed conservation of energy.

One of the practical challenges of system identification (SYSID) for buildings is that unmeasured heat gains (e.g., occupant gain, lighting/plug loads, in/exfiltration, and incident solar irradiances on various surfaces) can easily result in a poor model [37] regardless of a selection of model structures (e.g., gray-box or black-box model structures from the classical ARX/ARMAX to modern neural networks) and identification methods (e.g., the Prediction Error Method and the Subspace Method). Since, for K-12 schools, installing additional sensors for those unmeasured heat gains could be prohibited due to the cost increment, it is critical to reduce the negative effect of unmeasured disturbances by other means. An identification algorithm developed for this purpose, namely the lumped disturbance (LD) approach [37, 38], is thus adopted for this study because it showed more robust performance in system identification under significant unmeasured disturbances such as occupant heat gains. While the algorithm details can be found in the previous papers [37, 38, 34], we brief and present the practical application procedure (from an experimental design to a model validation strategy) of the LD approach for the school buildings.

A discretized thermal network model has the following form:

$$\mathbf{y}(k) = G_{\mathrm{u}} \circ \mathbf{u}(k) + G_{\mathrm{w}} \circ \mathbf{w}(k) + G_{\mathrm{g}} \circ \dot{Q}_{\mathrm{g},1:n}(k) \tag{1}$$

where z^{-1} is the backward time shift operator such that $z^{-1}x(k) = x(k-1)$ for a sequence of x.

In the LD approach, the unmeasured heat gain is defined as a lumped disturbance term ($\nu(k) := G_{\mathrm{g}}\dot{Q}_{\mathrm{g},1:n}(k)$), and it can be modeled as a filtered process of white noise $G_{\mathrm{g}}\dot{Q}_{\mathrm{g},1:n}(k) = \nu(k) = H(z)e(k)$. Then, this model can be written as a state-space model:

$$\begin{aligned} \mathbf{x}(k+1) &= \mathbf{A}_{\mathrm{d}}(\theta)\mathbf{x}(k) + \mathbf{B}_{\mathrm{d,u}}(\theta)\mathbf{u}(k) + \mathbf{B}_{\mathrm{d,w}}(\theta)\mathbf{w}(k) \\ \mathbf{y}(k) &= \mathbf{C}_{\mathrm{d}}(\theta)\mathbf{x}(k) + \boldsymbol{\nu}(k) \\ \boldsymbol{\zeta}(k+1) &= \mathcal{F}(\rho)\boldsymbol{\zeta}(k) + \mathcal{G}(\rho)\boldsymbol{\epsilon}(k) \\ \boldsymbol{\nu}(k) &= \boldsymbol{\zeta}(k) + \boldsymbol{\varepsilon}(k). \end{aligned} \tag{2}$$

Parameters (i.e., θ^* and ρ^*) were estimated by using the Prediction Error Method [39, 37] which minimizes the square sum of one-step prediction errors (ε) (Eq. 3) with measurements of $\mathbf{u}, \mathbf{w}, \mathbf{y}$. The calculation follows the three steps: (1) innovation (Eq. 4), (2) filtering (Eq. 5), and (3) prediction (Eq. 6). The mathematical details and performance of the LD approach compared to a traditional SYSID [40] which only looks at simulation errors and thus ignores unmeasured disturbances can be found in [37, 38].

$$\theta^*, \rho^* = \arg\min_{\theta} \sum_{k=1}^{N} \left(\boldsymbol{\varepsilon}(k;\theta)\right)^2 \tag{3}$$

$$\begin{aligned} \boldsymbol{\varepsilon}(k;\theta) &= \mathbf{y}(k) - \hat{\mathbf{y}}(k|k-1) \\ \hat{\mathbf{y}}(k|k-1) &= \mathbf{C}_{\mathrm{d}}\hat{\mathbf{x}}(k|k-1;\theta) + \hat{\boldsymbol{\zeta}}(k|k-1;\theta) \end{aligned} \tag{4}$$

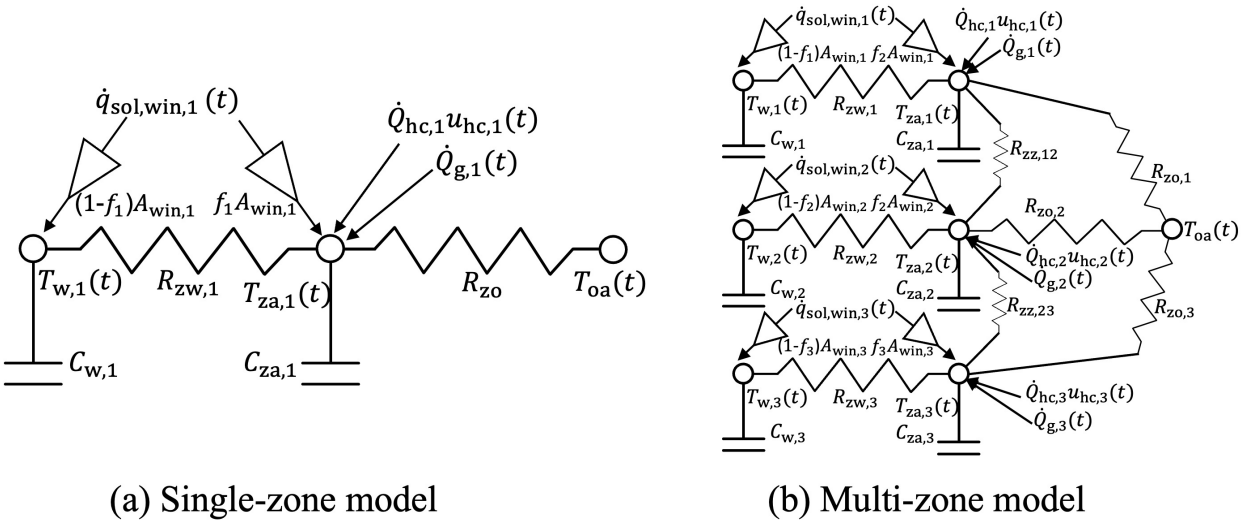

Figure 10: Two types of gray-box model structures.

$$\begin{bmatrix}\hat{\mathbf{x}}(k|k;\theta)\\ \hat{\boldsymbol{\zeta}}(k|k;\theta)\end{bmatrix} = \begin{bmatrix}\hat{\mathbf{x}}(k|k-1;\theta)\\ \hat{\boldsymbol{\zeta}}(k|k-1;\theta)\end{bmatrix} + \begin{bmatrix}\mathbf{0}\\ \mathcal{G}(\rho)\end{bmatrix}\left[\boldsymbol{\varepsilon}(k)\right] \tag{5}$$

$$\begin{aligned}\hat{\mathbf{x}}(k+1|k;\theta) &= \mathbf{A}_{\mathrm{d}}(\theta)\hat{\mathbf{x}}(k|k;\theta) + \mathbf{B}_{\mathrm{d,u}}(\theta)\mathbf{u}(k) + \mathbf{B}_{\mathrm{d,w}}(\theta)\mathbf{w}(k)\\ \hat{\zeta}(k+1|k;\theta) &= \mathcal{F}(\theta)\hat{\boldsymbol{\zeta}}(k|k;\theta)\end{aligned} \tag{6}$$

We investigated both single-zone and multi-zone thermal network model structures (as shown in Figure 10 (a) and (b), respectively), and chose the single-zone 2R2C model structure because there was no significant performance improvement when using the multi-zone model.

Although the LD approach considers unmeasured disturbances explicitly, the accuracy of the resulting model highly depends on the quality of the training dataset[38] like all other identification approaches. As an effort for better data quality, we designed the experiments as follows: For one week during vacation, the cooling setpoints were perturbed according to a pseudo-binary random signal (PRBS) with a 2-hour time-scale and the 4th order to turn ON/OFF RTUs. Once the data was collected, the non-linear optimization problem for the LD is solved with a multi-start method using the `least_squares` function of Scipy optimizer [41]. The first 12-hour data is used for initializing the initial states using a Kalman filter [42]. The boundaries of model parameters (i.e., θ) are generated by the following rules (Table 2).

The required metadata input needs to be specified from site inspection or common knowledge. For example, $A_{\mathrm{fl},i}$ and major window-related values are obtained from building drawings or satellite images. The minimum and maximum values of U, H, th, C_{w}, and C_{za} were set to 1×10^{-4}, 3×10^{-3}], [2.0, 5.0], [0.05, 1.0], [5, 3×10^{3}] and [1×10^{-5}, 5] (with appropriate units in the Table 2), which are generally applicable ranges for a typical building. The rated heating or cooling rates were not available because the RTU is very aged. Therefore, the empirical HVAC design value for a classroom building was used for ($\dot{Q}^{*}_{\mathrm{rated,hc}} = 0.1262\mathrm{kW/m^2}$)

Table 2: Rules to generate optimization boundary.

Required metadata input	• $A_{\mathrm{fl},i}$: Floor area of ith classroom. • $A_{\mathrm{mw},i}$,$tilt_{\mathrm{mw},i}$,$Az_{\mathrm{mw},i}$: Area [m^2], tilt [°], and azimuth [-] of ith major window. • U: Max. and min. values of U-value [kW/(m^2·K)]. • H, th: Max. and min. values of zone height (H) and thermal mass thickness (th) [m]. • C^*_{w}, C^*_{z}: Max. and min. values of volume normalized capacitance of wall/zone [kJ/(m^3·K)]. • $\dot{Q}^*_{\mathrm{rated,hc}}$ or $\dot{Q}_{\mathrm{rated,hc}}$: Area normalized or nominal rated heat/cool rate [kW/m^2 or kW]. • L: Max. and min. values of scale variable [-]
	Parameters and rules to generate minimum and maximum bounds for the nonlinear least squares optimization
Parameters	• $C_{\mathrm{w},i}$: [$C_{\mathrm{w,min}}A_{\mathrm{fl},i}th_{\mathrm{min}}$, $C_{\mathrm{w,max}}A_{\mathrm{fl},i}th_{\mathrm{max}}$] • $C_{\mathrm{za},i}$: [$C_{\mathrm{za,min}}A_{\mathrm{fl},i}H_{\mathrm{min}}$, $C_{\mathrm{za,max}}A_{\mathrm{fl},i}H_{\mathrm{max}}$] • $R_{\mathrm{zw},i}$: [$1/(4(\sqrt{A_{\mathrm{fl},i}})H_{\mathrm{max}}U_{\mathrm{max}})$, $1/(4(\sqrt{A_{\mathrm{fl},i}})H_{\mathrm{min}}U_{\mathrm{min}})$] • $R_{\mathrm{zo},i}$: [1e-4, $R_{\mathrm{zw,max},i}/20$] • f_i: [1e-9,1] • $A_{\mathrm{win},i}$: [$L_{\mathrm{min}}A_{\mathrm{fl},i}$, $L_{\mathrm{max}}A_{\mathrm{fl},i}$] • $Q_{\mathrm{h},i}$: [$L_{\mathrm{min}}\dot{Q}^*_{\mathrm{rated,h}}A_{\mathrm{fl},i}$, $L_{\mathrm{max}}\dot{Q}^*_{\mathrm{rated,h}}A_{\mathrm{fl},i}$] or [$L_{\mathrm{min}}\dot{Q}_{\mathrm{rated,h}}$, $L_{\mathrm{max}}\dot{Q}_{\mathrm{rated,h}}$] • $Q_{\mathrm{c},i}$: [$L_{\mathrm{min}}\dot{Q}^*_{\mathrm{rated,c}}A_{\mathrm{fl},i}$, $L_{\mathrm{max}}\dot{Q}^*_{\mathrm{rated,c}}A_{\mathrm{fl},i}$] or [$L_{\mathrm{min}}\dot{Q}_{\mathrm{rated,c}}$, $L_{\mathrm{max}}\dot{Q}_{\mathrm{rated,c}}$] • ρ_i: [-0.999,0.999]

[43]. Finally, the relatively wide range of scale variables (L_{min} : 0.1, and L_{min} : 3.0) is used for generating the boundaries.

On the other hand, there were several missing data periods due to unstable WiFi. So, the data was split into multiple datasets by omitting the missing periods, and the prediction errors of each dataset were aggregated for the objective function in the optimization. The collected data was moving averaged by a 15-minute window. This is the sampling time of SYSID and was chosen to smooth the ON/OFF heating and cooling operations because of the time lag and performance degradation during the start-up time.

For evaluating model accuracy, in addition to the typical cross-validation strategy that directly compares output predictions (in our case, the thermostat temperatures) with measurements of a validation dataset, we also compared the predicted heating/cooling loads with measurements. This is because it is important for our MPC to accurately estimate the required amount of heating or cooling rate (more precisely, runtime fraction (RTF) of each RTU stage for our case, i.e., $u_{\mathrm{h}}(k)$ and $u_{\mathrm{c}}(k)$). For doing this, the state-space model in Eq. 2 was rearranged to Eq. 7 by neglecting the error noise. Then, the required heating and cooling RTF of RTUs ($\hat{\mathbf{u}}(k)$) and the next time states ($\hat{\mathbf{x}}(k+1)$) can be calculated through Eq. 8. Here, ($\mathbf{C}_{\mathrm{d}}\mathbf{B}_{\mathrm{d,u}}(\theta)$) is not invertible, so the pseudo-inverse was utilized.

$$\begin{aligned}\mathbf{y}(k+1) &= \mathbf{C}_{\mathrm{d}}(\theta)\mathbf{x}(k+1)\\ &= \mathbf{C}_{\mathrm{d}}(\theta)\left(\mathbf{A}_{\mathrm{d}}\mathbf{x}(k)+\mathbf{B}_{\mathrm{d,u}}(\theta)\mathbf{u}(k)+\mathbf{B}_{\mathrm{d,w}}(\theta)\mathbf{w}(k)\right)\end{aligned} \tag{7}$$

$$\begin{aligned}\hat{\mathbf{u}}(k) &= \left(\mathbf{C}_{\mathrm{d}}\mathbf{B}_{\mathrm{d,u}}(\theta)\right)^{\dagger}\left[\mathbf{y}(k+1)-\mathbf{C}_{\mathrm{d}}(\theta)\left(\mathbf{A}_{\mathrm{d}}\hat{\mathbf{x}}(k)+\mathbf{B}_{\mathrm{d,w}}(\theta)\mathbf{w}(k)\right)\right]\\ \hat{\mathbf{x}}(k+1) &= \mathbf{A}_{\mathrm{d}}\hat{\mathbf{x}}(k)+\mathbf{B}_{\mathrm{d,u}}(\theta)\hat{\mathbf{u}}(k)+\mathbf{B}_{\mathrm{d,w}}(\theta)\mathbf{w}(k)\end{aligned} \tag{8}$$

3.4. Hierarchical MPC

The control problem of LMPC at a current time step k can be written as Eq. 9. LMPC is mainly designed to coordinate the operation of RTUs by looking at relatively short prediction horizons (e.g., 15 minutes to 1 hour).

$$
\begin{aligned}
\min \quad & \sum_{j=1}^{N_{\mathrm{p}}^{\mathrm{L}}} \sum_{i=1}^{n} P_{\mathrm{RTU},i} u_{\mathrm{hc},i}(k+j-1) + \omega_{\mathrm{d}}\delta + \omega_{\mathrm{l}}\Gamma_{\mathrm{l},i} + \omega_{\mathrm{u}}\Gamma_{\mathrm{u},i} \\
s.t. \quad & T_{\mathrm{l},i}^{\mathrm{L}} - \Gamma_{\mathrm{l},i} \leq \mathbb{E}(y_i(k+j)|\mathcal{D}_k) \leq T_{\mathrm{u},i}^{\mathrm{L}} + \Gamma_{\mathrm{u},i} \; (\forall i \in \{1, \cdots, n\}) \\
& \sum_{i=1}^{n} P_{\mathrm{RTU},i} u_{\mathrm{hc},i}(k+j-1) \leq \delta \quad (\forall j \in \{1, \cdots, N_{\mathrm{p}}^{\mathrm{L}}\}),
\end{aligned}
\tag{9}
$$

where j is the time step for control changes from the current measured value of k. $u_{\mathrm{hc},i}(k+j)$ and $y_i(k+j)$ are RTU unit stage (heat or cool) and thermostat temperature of the i^{th} room at the $k+j^{\mathrm{th}}$ timestep. $\mathbb{E}(\mathbf{y}(k+j)|\mathcal{D}_k)$ is the optimal j-step temperature prediction from the building model (3.3) given the data $\mathcal{D}_k = \{y(k-1), y(k-2), \cdots, u(k+j-1), u(k+j-2), \cdots\}$. $(T_{\mathrm{l},i}^{\mathrm{L}}, T_{\mathrm{u},i}^{\mathrm{L}})$ are the desired setpoints obtained from UMPC for i^{th} RTU. $\omega_l, \omega_u (\in \mathbb{R}^+)$ and $\omega_d (\in \mathbb{R}^+)$ are weights on variables of $\Gamma_{\mathrm{l}}, \Gamma_{\mathrm{u}} (\in \mathbb{R}^+)$ and $\delta (\in \mathbb{R}^+)$. Γ_{l} and Γ_{u} can be viewed as comfort violations. The optimizer finds the optimal sequence of RTUs over the prediction horizon and $\delta, \Gamma_l, \Gamma_u$, which forms a mixed integer linear programming problem.

In the last constraint, δ is the upper bound of the electric demand for each time prediction horizon, and it guarantees the performance of peak demand reduction in LMPC. Please refer to [19] for details and long-term field experiment results.

In this study, the upper bound of δ was set to 70% of the summation of RTU power (i.e., $0.7 \times \sum_{i=1}^{n} P_{\mathrm{RTU},i}$, which acts as the target peak. ω_{d}, ω_{l}, and ω_{u} are set to 10, 1000, and 1000, respectively. Considering the total power of RTUs is 16.5kW (δ=11.6kW, see Figure 1), 0.3°C of temperature violation is equivalent to the use of one more RTU. The timestep for LMPC was set at 5 minutes due to the typical anti-cycling time of RTU. This anti-cycling time represents the minimum interval during which the compressor should remain off before it can be restarted, and it is commonly set to 5 minutes.

While LMPC coordinates unit stages, UMPC is supposed to provide optimal setpoints that minimize energy cost ($ER \times P_{\mathrm{RTU}}$) by looking at relatively long prediction horizons (e.g., more than 6 hours). The control problem of the UMPC at the current time k can be written as:

$$
\begin{aligned}
\min \quad & \sum_{j=1}^{N_p^U} \sum_{i=1}^{n} ER(k+j-1) P_{\mathrm{RTU},i}(k+j-1) \bar{u}_{\mathrm{hc},i}(k+j-1) \\
& + \omega_{\mathrm{d}}\delta + \omega_{\mathrm{l}}\Gamma_{\mathrm{l},i} + \omega_{\mathrm{u}}\Gamma_{\mathrm{u},i} \\
s.t. \quad & T_{\mathrm{l},i}^{\mathrm{U}} - \Gamma_{\mathrm{l},i} \leq \mathbb{E}(\bar{y}_i(k+j)|\mathcal{D}_k) \leq T_{\mathrm{u},i}^{\mathrm{U}} + \Gamma_{\mathrm{u},i} (\forall i \in \{1, \cdots, n\}) \\
& \sum_{i=1}^{n} P_{\mathrm{RTU},i}(k+j-1) \bar{u}_{\mathrm{hc},i}(k+j-1) \leq \delta \\
& 0 \leq \bar{u}_{\mathrm{hc},i}(k+j-1) \leq 1 \quad (\forall j \in \{1, \cdots, N_{\mathrm{p}}^{\mathrm{U}}\}),
\end{aligned}
\tag{10}
$$

where $\bar{u}_{\mathrm{hc},i}$ and $\bar{y}_i$ are the moving averaged RTU stages (normalized 0-1 scale, i.e., RTF) and thermostat temperatures, respectively, for a timestep (i.e., 30 minutes in this study). The timestep for UMPC was determined to be approximately five times that of LMPC in order to achieve the desired time scale separation. This factor is commonly utilized in cascade or hierarchical controls, e.g., [10].

While the upper and lower temperature bounds of UMPC (i.e., $T^{\mathrm{U}}_{\mathrm{l},i}$, $T^{\mathrm{U}}_{\mathrm{u},i}$) are a comfortable temperature range specified by users, the optimal input trajectories of UMPC are converted to the desired setpoint trajectories through Eq. 11, which are used for temperature boundaries for LMPC ($T^{\mathrm{L}}_{\mathrm{l},i}$, $T^{\mathrm{L}}_{\mathrm{u},i}$.

$$
\begin{aligned}
T^{\mathrm{L}}_{\mathrm{u},i}(k) &= \begin{cases} \bar{y}^*_i(k), & \text{if } \bar{u}^*_{\mathrm{c},i}(k) > 0 \\ T^{\mathrm{U}}_{\mathrm{u},i}(k), & \text{if } \bar{u}^*_{\mathrm{c},i}(k) = 0 \end{cases} \\
T^{\mathrm{L}}_{\mathrm{l},i}(k) &= \begin{cases} \bar{y}^*_i(k), & \text{if } \bar{u}^*_{\mathrm{h},i}(k) > 0 \\ T^{\mathrm{U}}_{\mathrm{l},i}(k), & \text{if } \bar{u}^*_{\mathrm{h},i}(k) = 0 \end{cases}
\end{aligned}
\tag{11}
$$

where $\bar{u}^*_{\mathrm{hc}i}(k)$ indicates the optimal heating or cooling RTF for the averaging window for the i^{th} unit at the k^{th} timestep of the UMPC and $\bar{y}^*_i(k)$ is the corresponding desired temperature. When the required RTU operations are zero, the boundaries of UMPC are used for LMPC. Otherwise, the desired temperature profiles are used for the upper or lower bounds of the LMPC for cooling or heating operation, respectively.

LMPC and UMPC are linear and mixed integer linear programming problems and are solved by using `pyglpk` [44]. While UMPC is solved every 30 minutes, LMPC is solved every 5 minutes. UMPC is solved first, and then LMPC is solved based on the obtained comfort boundaries via Eq. 11. This boundary is used until the next UMPC execution. Once the LMPC provides optimal input trajectories, it is decided to turn on which RTUs based on the next time input trajectories. Technically, we send high (78°F) or low (72°F) setpoints to turn off or to the RTU in a cooling scenario.

3.5. MPC improvements for schools

The MPC discussed in the preceding section needs modifications for K-12 school applications to handle the challenges discussed in Section 2.3.

The main idea of reflecting the user's override actions is to 1) detect which RTU(s) is overridden, 2) predict the stage status for the corresponding RTU until the next sampling time, and 3) impose the RTU stage to the LMPC constraint. More precise descriptions are as follows. First, the data is labeled as `override` in the database when users make override actions. The Schedule overrides are easily detected because the current schedule mode data is usually available through the thermostat API. Since MPC only changes setpoints, it is marked as `override` whenever there is a change in the schedule mode. Setpoint override can be detected when we read the real-time data. As shown in Figure 11, the thermostat data is recorded in the database every minute and whenever MPC sends the optimal setpoints. When we read data from thermostat API, if setpoints differ from previous values and deviate from the default schedule, it is Setpoint override and labeled as `override` unless the changes are made by the MPC. Second, when running MPC, it first checks the most recent data to see if there are any RTUs currently in `override`. When running LMPC, the current cooling or heating ON/OFF signals of RTUs in `override` are assumed to be held for the prediction horizon. In other words, the MPC only calculates the optimal ON/OFF signal profiles for non-overridden RTUs by assuming the overridden RTUs' current operation continues for the prediction horizon. Although this method is not precisely correct (e.g., what if the identified unit switches ON from OFF for the prediction horizon), it worked in the desired manner for the LMPC due to the short prediction horizon and sampling time.

Another remedy for the school applications corresponds to the integer-valued temperature measurement (Figure 6). From the MPC perspective, the one-step-ahead prediction error (i.e., innovation, $\boldsymbol{\varepsilon}$) is used for estimating the current states $\hat{\mathbf{x}}(k|k)$ (Eq. 5). However, with the thermostat rounding behavior, the true measurement $\mathbf{z}_{\text{true}}$ ($\in \mathbb{R}^n$) is rounded to the ones place as $\mathbf{y}$ ($\in \mathbb{N}^n$), which can be viewed as an uncertain observation (i.e., hidden variable):

$$\begin{aligned} \mathbf{y} &= \sigma(\mathbf{z}_{\text{true}}) \\ \mathbf{z}_{\text{true}} &= \mathbf{y} + \boldsymbol{\epsilon}_y \end{aligned} \tag{12}$$

where $\epsilon_{y,i} \sim \text{Uniform}(-0.5, 0.5)$ and σ is a rounding function to the decimal point.

Therefore, the innovation in Eq. 4 can be rewritten as:

$$\boldsymbol{\varepsilon}(k) = \mathbf{z}_{\text{true}}(k) - \boldsymbol{\epsilon}_y(k) - \hat{\mathbf{y}}(k|k-1) \tag{13}$$

The one-step-ahead prediction can be rewritten as:

$$\begin{aligned} \hat{\mathbf{y}}(k|k-1) &= \mathbf{C}_{\text{d}}\hat{\mathbf{x}}(k|k-1) + \hat{\boldsymbol{\zeta}}(k|k-1) \\ &= \mathbf{C}_{\text{d}}\hat{\mathbf{x}}(k|k-1) + \mathcal{F}\hat{\boldsymbol{\zeta}}(k-1|k-2) + \mathcal{G}\boldsymbol{\varepsilon}(k-1) \\ &= \mathbf{C}_{\text{d}}\hat{\mathbf{x}}(k|k-1) + \mathcal{F}\hat{\boldsymbol{\zeta}}(k-1|k-2) + \\ &\quad \mathcal{G} \cdot (\mathbf{z}_{\text{true}}(k-1) - \boldsymbol{\epsilon}_y(k-1) - \hat{\mathbf{y}}(k-1|k-2)) \\ &= \hat{\mathbf{y}}_{\text{true}}(k|k-1) + \sum_{i=2}^{k} \mathcal{G}\boldsymbol{\epsilon}_y(i-1) \end{aligned} \tag{14}$$

The true innovation ($\epsilon_{\text{true}}(k) = \mathbf{z}_{\text{true}}(k) - \hat{\mathbf{y}}_{\text{true}}(k|k-1)$) is also a random variable due to $\epsilon_y(k) - \sum_{i=2}^{k} \mathcal{G}\boldsymbol{\epsilon}_y(i-1)$. Thus, it can be estimated by taking expectations in MPC. From Eq. 12, the expectation can be obtained by taking the rounding function to each value. Thus, whenever the current state is estimated in MPC (i.e., state observer in 7), the rounding function (Eq. 15) is applied to estimate the innovation.

$$\begin{aligned} \mathbb{E}[\boldsymbol{\varepsilon}_{\text{true}}(k)] &= \mathbb{E}[\mathbf{z}_{\text{true}}(k)] - \mathbb{E}[\hat{\mathbf{y}}_{\text{true}}(k|k-1)] \\ &= \mathbb{E}[\mathbf{y}(k) + \boldsymbol{\epsilon}(k)] - \mathbb{E}\left[\hat{\mathbf{y}}(k|k-1) - \sum_{i=2}^{k-1} \mathcal{G}\boldsymbol{\epsilon}_y(i-1)\right] \\ &= \sigma(\mathbf{y}(k)) - \sigma(\hat{\mathbf{y}}(k|k-1)) \end{aligned} \tag{15}$$

4. Experiment description

4.1. Control system architecture

Figure 11 shows the overall schematic diagram of the data collection and MPC service. Our industry partner[1] collects all data into their serverless cloud service and provides it to MPC via three service API

[1] https://communityenergylabs.com/

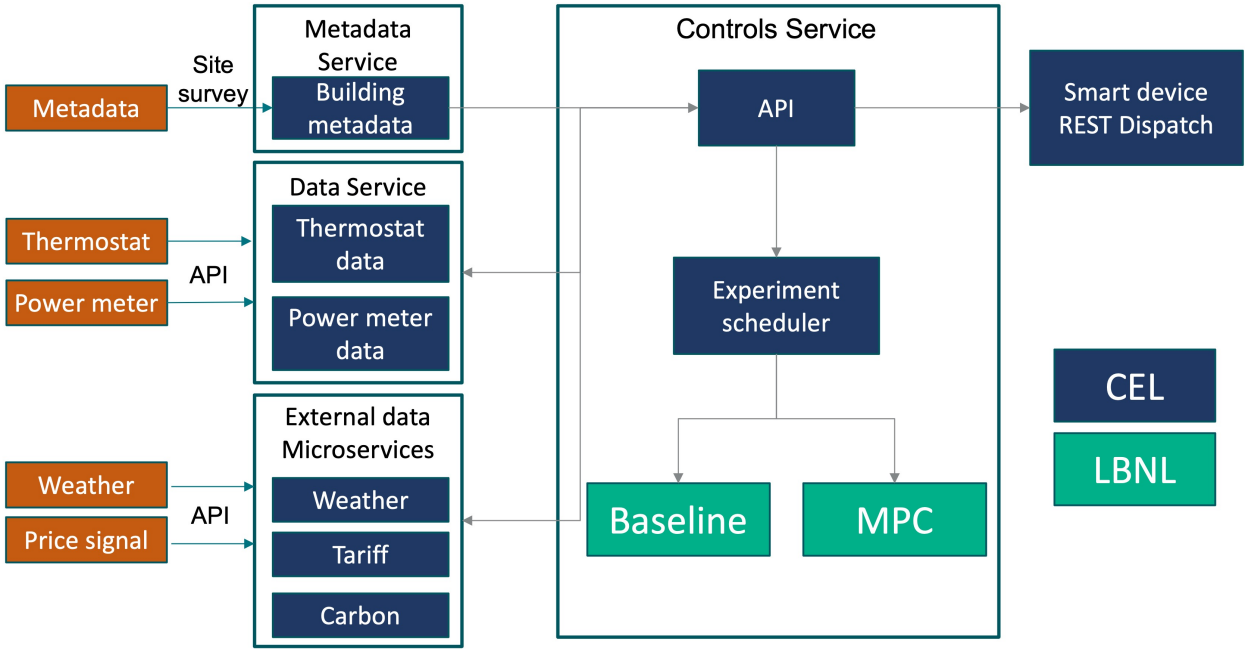

Figure 11: Schematic diagram of data collection and MPC service.

(Metadata service, Data service, and External data microservices).

Metadata was obtained through site survey and stored in the `json` format. This was mainly used for SYSID (Table 2) and experiment status control. The thermostat and power meter data were recorded every 1-min interval. Because some thermostat vendors only provide instantaneous values for the current operating signal, it is necessary to reduce data record intervals so as not to lose the operation data. The current outdoor air temperature is recorded every 5-min interval, obtained from National Oceanic and Atmospheric Administration (NOAA) [45]. The solar radiation is collected by recording the current solar forecast data. The solar forecast is estimated from NOAA's cloud cover forecast [46] and pvlib's [47] function to calculate global horizontal, direct normal, and diffuse horizontal irradiance from cloud cover. Utility tariff (i.e., SCE's TOU rate) is obtained from the utility company website.

Based on our experimental schedule (Section 4.2), the Controls service allows MPC to send setpoints. MPC is currently running on our server and executed when the Controls service lets the MPC server know the day is MPC day. Once the MPC calculation is done, the obtained setpoints are sent to each thermostat via the Controls service.

4.2. Experimental schedule

Instead of a typical measurement and verification (M&V) approach that compares "pre-installation" to "post-installation", we choose the following randomized block scheme: for each day, either the Baseline or MPC will be randomly selected and implemented. The randomization process [48] will use "blocking" to ensure that the same number of days are assigned to each control for each block period (see Figure 12 for the conceptual diagram). Currently, we have only one demonstration site, so only one block schedule is generated and assigned. Though school starts end of Aug 2022, the experiment started Oct 2022 due to

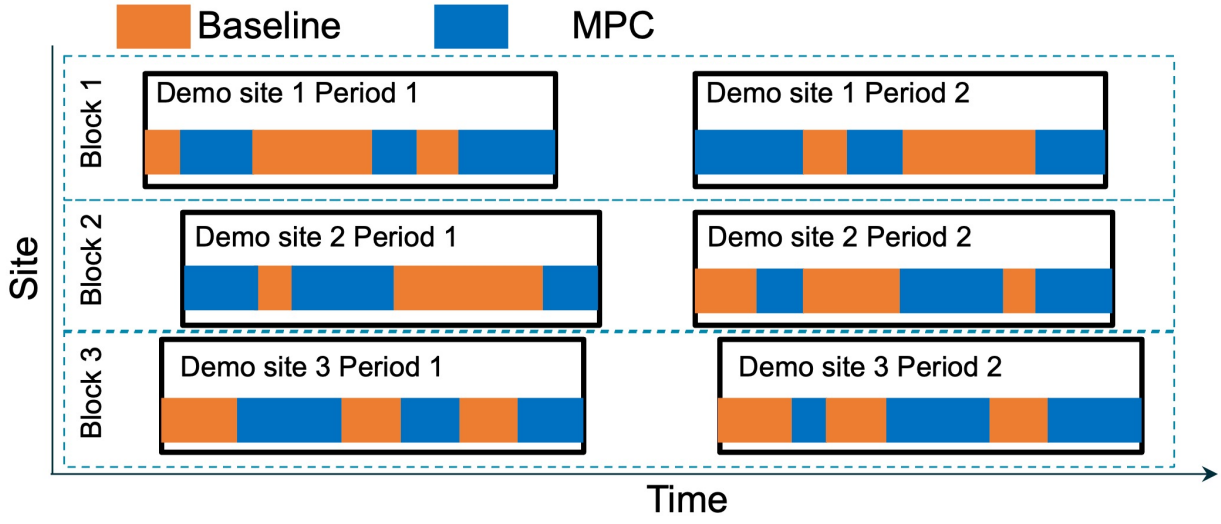

Figure 12: Conceptual diagram of the randomized block schedule.

several delays (e.g., installation of sensors, heat wave, and Data and Controls service troubleshooting). Since the weather was getting colder, the MPC experiment is mainly done in Oct 2022, but the Sep 2022 data is also included for data analysis due to the cold weather in Oct 2022.

5. Results

5.1. Building model validation

When handling the optimal load shifting problem, it makes more sense to compare the predicted thermal loads with the actual loads instead of solely comparing indoor temperatures. Figure 13 shows the comparison between the measured cooling RTF and the predicted one calculated by Eq. 8 for a classroom. Overall, the model can capture the timing of cooling operation (i.e., zigzag shape). However, the magnitude is smaller than the measurement, which results in 0.14 (training data) and 0.23 (test data) of RMSEs. This could be attributed to the unmeasured occupant heat gains from students. For example, 08-18 was a typical classroom day, and there were substantial unmeasured body heat gains from students during the daytime. Therefore, the underestimated RTFs can be viewed as the summation of true RTFs and unmeasured body heat gains. Similarly, RTFs were also underestimated in 08-19, but the magnitude was smaller because it was a less crowded classroom day (Friday). On the other hand, there were some discrepancies during the non-cooling time. This is mainly because the estimated cooling RTF included all unmeasured disturbance terms such as infiltration.

To understand the effect of the unmeasured disturbances, the 24-hour ahead temperature predictions on training and test data are visualized in Figure 14. On the training data, the thermal response of the room temperature was relatively well captured though it is not perfect due to small unmeasured disturbances such as infiltration or device stand-by heat gains. However, during the typical occupied days (test data in Figure 14 (b)), the temperature prediction diverges more from the measured values after the occupied times due to unmeasured internal heat gains. The occurrence of prediction errors is inevitable in the low-cost MPC

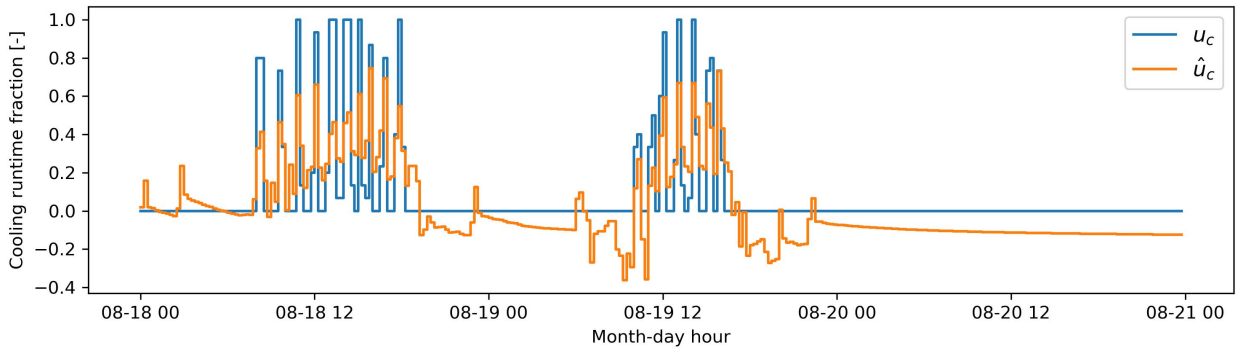

Figure 13: Comparison between the required ($\hat{u}_c$) and measured (u_c) amount of cooling runtime fraction in one classroom (test data).

infrastructure for two primary reasons: (1) Achieving accurate predictions, even with a perfect thermal building model, requires knowledge of future disturbances that are not directly measured in the control architecture, and (2) there always be model errors when significant unmeasured disturbances present during the training data collection period and the input data is correlated, regardless of the system identification methods or model structures employed (for a comprehensive and analytical discussion, please refer to [38]). To reduce this concern during the control phase, we utilized the off-set free MPC approach that was developed to reject the model uncertainty and unmeasured disturbance in the control field. Like the lumped disturbance modeling approach, it also augments a disturbance model to a physical system model, and track a underlying dynamic state of the unmeasured disturbances by the state-observer (Figure 7).

5.2. Operational profiles

Figure 15 shows the result summary for a sample day for the baseline RTU operation. In the top figure, all rooms were well-controlled within the comfort band. One thermostat was below the control threshold due to the thermostat's malfunction on the threshold lock feature (see Sections 2.2 and 2.3). Due to the decentralized and independent thermostat control scheme, 5 to 6 RTUs operated simultaneously several times, resulting in a peak power of 18.16kW.

On the other hand, Figure 16 shows the result summary for a sample day for the MPC. For a fair comparison, the date was carefully selected considering the outdoor air temperature profile and school schedule. In the top figure, all rooms were well-controlled within the comfort band (including the thermostat deadband). One room was in override mode during the occupied period, and another room was in override mode after the occupied period. The thick red line at the bottom figure, representing the 15-minute averaged total power consumption for the two buildings clearly shows that the MPC coordinates and distributes the unit operations to avoid unnecessary peaks. This resulted in a smoothed power profile compared to the Baseline in Figure 15 with a reduced peak power of 13.98kW.

Figure 17 shows the comparison of UMPC's desired setpoints and room temperatures for three classrooms during a cooling day (the overridden classrooms are excluded) to depict how well LMPC adheres to UMPC solutions. We visualize the UMPC's setpoints during the occupied times because the control actions were the same as the default setpoint schedule during the unoccupied times. Due to the thermostat's integer resolution, there were some variations. Nonetheless, the room temperature followed the UMPC's desired setpoint through LMPC's ON/OFF actions.

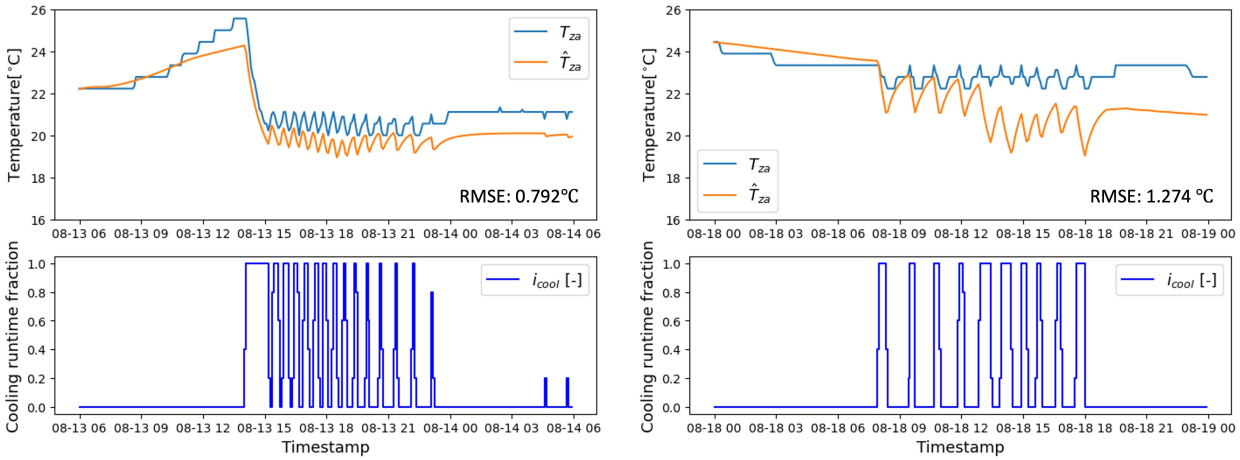

(a) 24-hour temperature prediction (Train data)

(b) 24-hour temperature prediction (Test data)

Figure 14: 24-hour future thermostat-temperature predictions of a classroom using training and test data.

5.3. Peak demand reduction during a cooling season

Figure 18 shows the daily peak demand for cooling days with respect to daily averaged outdoor air temperature. Remember that the demand charge is not based on peaks over a day but on maximum values over a longer time period (i.e., the billing period). Therefore, comparisons between daily peaks for the MPC and Baseline at a certain mean daily outdoor temperature are not very meaningful for estimating demand reduction. It would be more meaningful to compare the two maximum peaks over a range of the mean outdoor temperatures. Thus, for demand savings estimation, the mean outdoor temperature line (the x-axis) was divided into two groups (below and above 23°C, marked as (a) Mild days and (b) Hot days in Figure 18) and compared the maximum peaks for each group. The indices of ($B1$, $M1$) and ($B2$, $M2$) in Table 3 represent the maximum daily peaks for the Baseline and MPC respectively for two different conditions in the milder days. Likewise, $B3$ and $M3$ indicate those for the hot DAYS. The MPC reduced the total peak and HVAC peak powers by around 22-24% and 29-32%, respectively, for both conditions. For the hot DAYS, there is only one data point for the MPC (because the season changed once the MPC test was initiated, and the randomized test schedule prevented MPC from running in hot conditions) and thus the estimation might be highly inaccurate. However, closer data analysis indicates that this data point is the worst-case scenario. Figure 19 shows the corresponding profiles to that MPC day (M3). Note that four of six classrooms were in the override mode so the MPC was only able to control two classrooms (see Section 2.3 for MPC operation under override mode). Even with this limited controllability, the MPC found a way to reduce the peak demand (by looking at other unit statuses) while maintaining all temperatures within the prescribed bounds. Therefore, the M3 data point can be viewed as a worst-case scenario.

5.4. Temperature profiles during a cooling season

It is critical to make a comparison of the indoor temperature distributions between Baseline and MPC days (Figure 20) to ensure there was no compromise in thermal comfort. The temperature distributions of MPC days are in a similar range (22-24°C) with Baseline days during the occupied time (i.e., 08:00-16:00). Therefore, the MPC could successfully control the indoor temperature during the trials. One should note

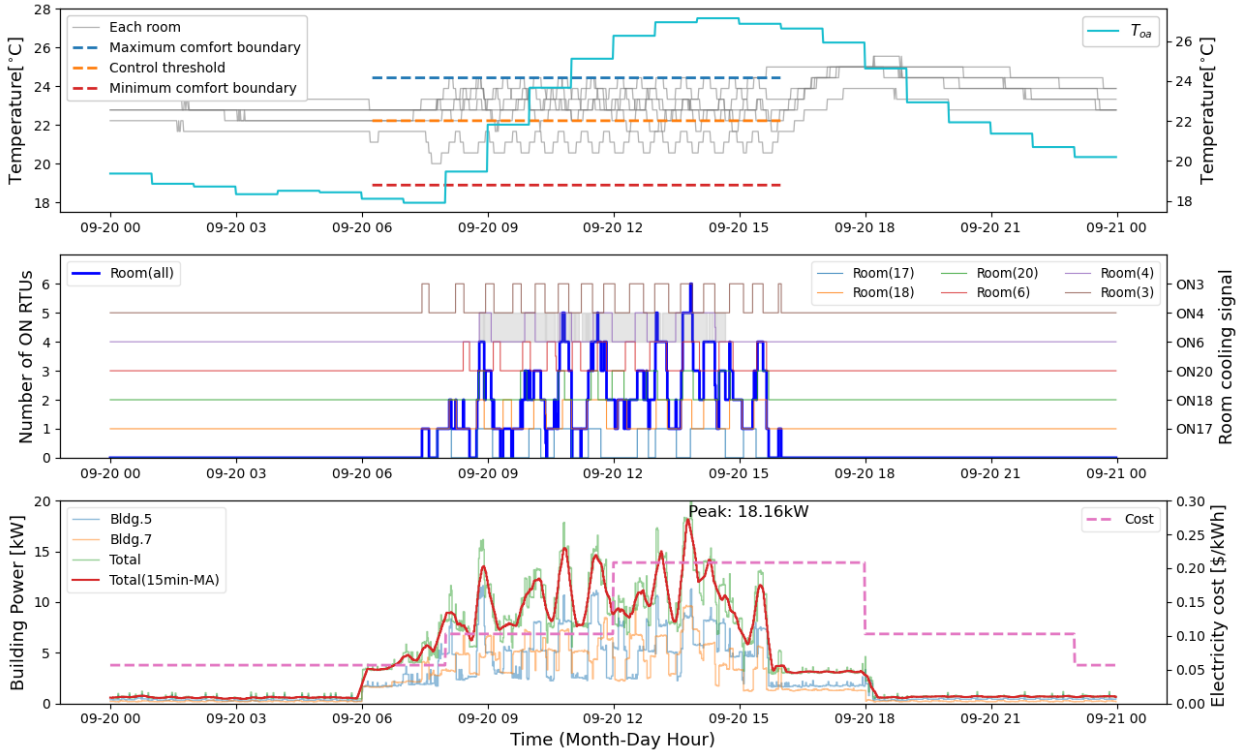

Figure 15: Summary of RTU operation in a default schedule day (Baseline); top: each room temperatures with min/max comfort boundary (thermostat deadband adjusted setpoints) and control threshold (i.e., minimum adjustable cooling setpoint, see Table 1), and outdoor air temperature, middle: RTU signals and override period (grayed area), bottom: building power and electricity price signal.

Table 3: Peak demand reduction summary (the parenthesis represents a reference Baseline data point for calculating a relative reduction)

Index	Case	Mean T_{oa}	Total Peak Power	HVAC Peak Power	% Total Peak Reduction	% HVAC Peak Reduction
B1	Baseline	21.9°C	18.1kW	13.8kW	-	-
B2	Baseline	18.3°C	14.5kW	10.5kW	-	-
B3	Baseline	26.9°C	21.1kW	17.3kW	-	-
M1	MPC	21.3°C	14.0kW	9.4kW	23% (B1)	32% (B1)
M2	MPC	18.0°C	11.4kW	7.5kW	22% (B2)	29% (B2)
M3	MPC	25.4°C	16.1kW	12.1kW	24% (B3)	30% (B3)

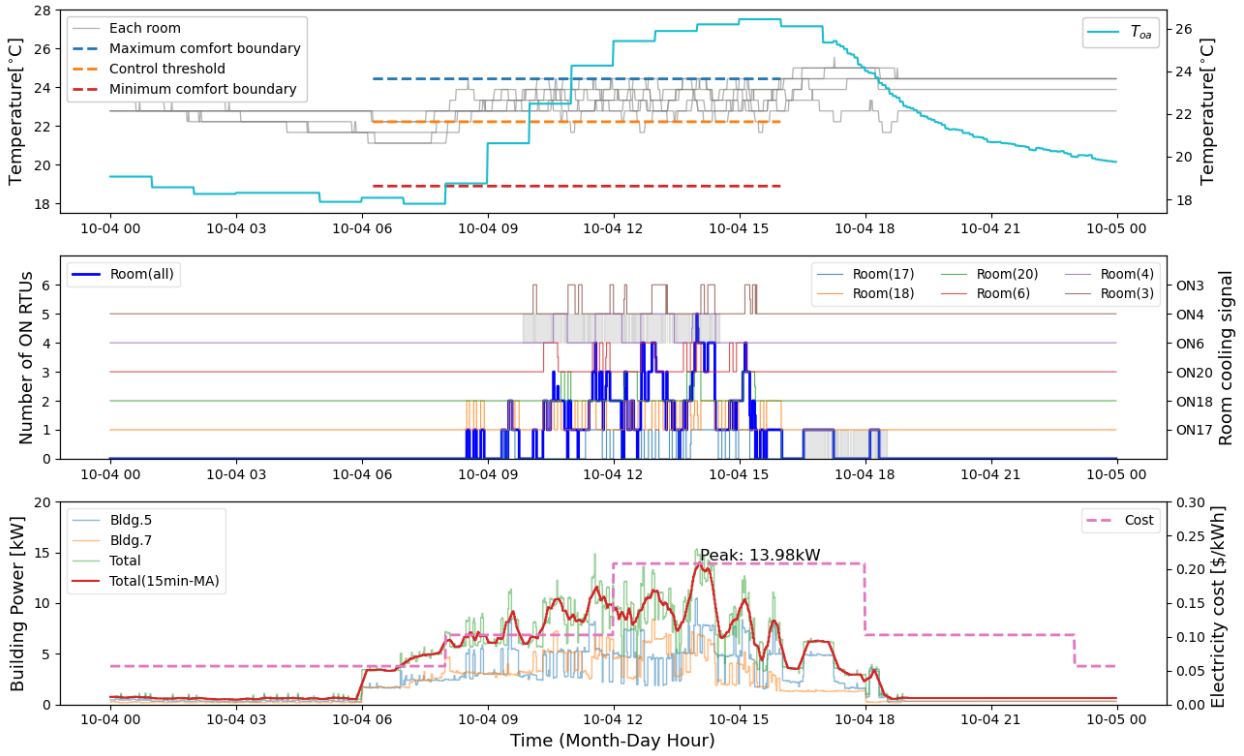

Figure 16: Summary of RTU operation in an MPC day (MPC); top: each room temperatures with min/max comfort boundary (thermostat deadband adjusted setpoints) and control threshold (i.e., minimum adjustable cooling setpoint, see Table 1) and outdoor air temperature, middle: RTU signals and override period (grayed area), bottom: building power and electricity price signal.

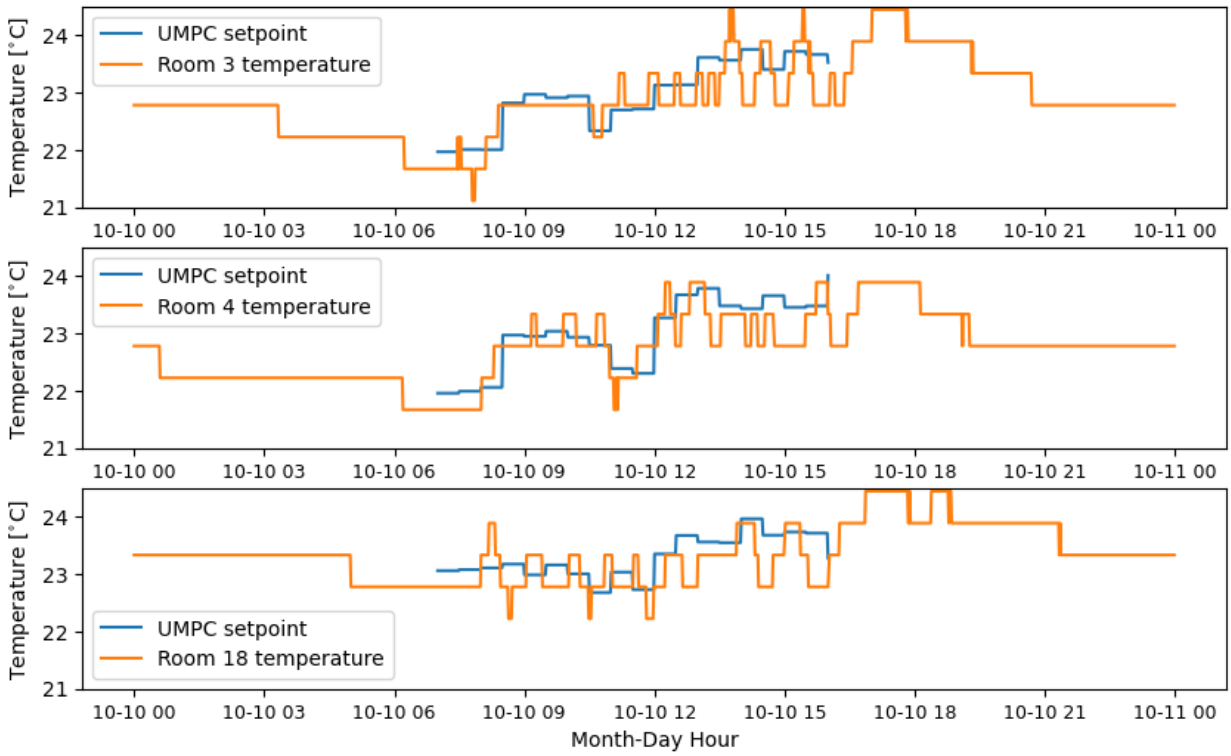

Figure 17: Comparison of UMPC setpoints (occupied times) and room temperatures for three classrooms during a cooling day (the overridden three classrooms are excluded).

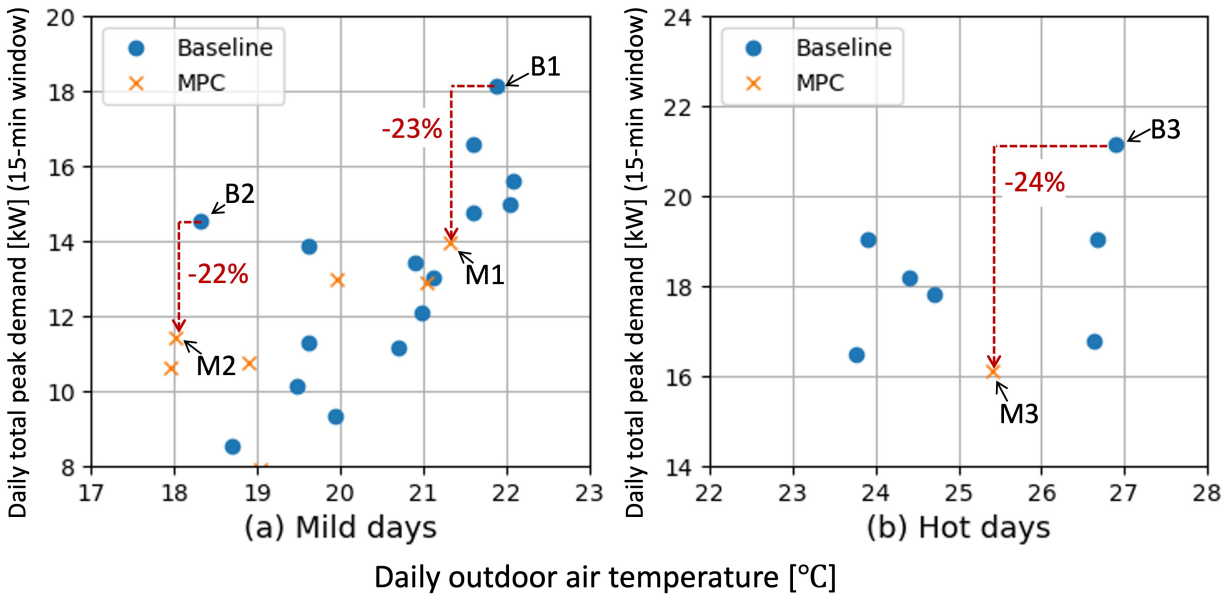

Figure 18: Peak demand vs daily averaged outdoor air temperature for the whole data collecting period (Sep-Oct, 2022). For the estimation of demand cost reduction, the data was split into two parts: below (Mild) and above (Hot) 23^{o}C of the average outdoor temperature.

that Baseline had lower and tighter temperature distributions during the peak cooling times (i.e., 12:00-16:00). The reason is that MPC tried to maintain higher temperatures within the comfort boundaries to reduce unnecessary cooling operations.

5.5. Load shifting potential during a cooling season

Figure 21 shows the comparison between the zone-wise averaged room temperatures and load profiles for Baseline and MPC days with the utility tariff to show the load shifting potential. A 1-hour moving average was applied for clarification purposes. The two days were selected to have similar outdoor temperatures: the mean temperatures for the Baseline and MPC days are 21.3^{o}C and 20.8^{o}C, respectively.

On the Baseline day, the power demand suddenly increased near 13:00. This is because the students had lunchtime at noon and came back to the classrooms near 13:00, i.e., increasing cooling loads. Note also that the average temperature (blue) decreases gradually near lunchtime, and reaches and maintains the coolest temperature of about 23^{o}C. This implies a manual setpoint change that could be due to the habitual override action (as mentioned in Section 2.3) or a warmer thermal sensation. This change also contributes to the peak power near 13:00.

The comparison between the powers in the lower sub-figure clearly shows that the MPC shifts the load in response to the price signal: the MPC used more energy before the electricity price jumps (12:00) and less energy after that. This precooling effort can also be confirmed in the temperature plot: the MPC's zonal averaged temperature was around 1^{o}C lower than that of the Baseline control. The MPC couldn't precool further because the lower temperature bound during the occupied time was set to around 22.2^{o}C (Table 1). The shifted HVAC electrical loads for 12:00-15:00 were calculated to be 16%.

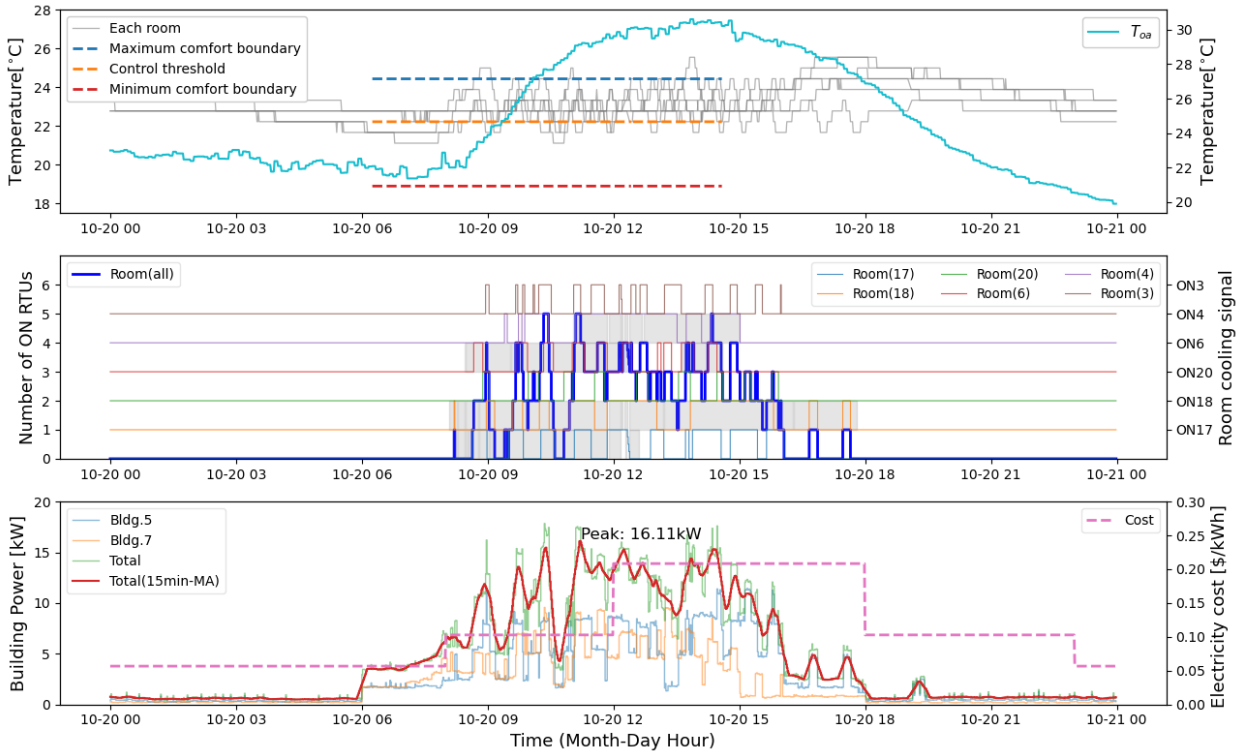

Figure 19: Summary of RTU operation in a peak cooling day (MPC); top: each room temperatures with min/max comfort boundary (thermostat deadband adjusted setpoints) and control threshold (i.e., minimum adjustable cooling setpoint, see Table 1) and outdoor air temperature, middle: RTU signals and override period (grayed area), bottom: building power and electricity price signal.

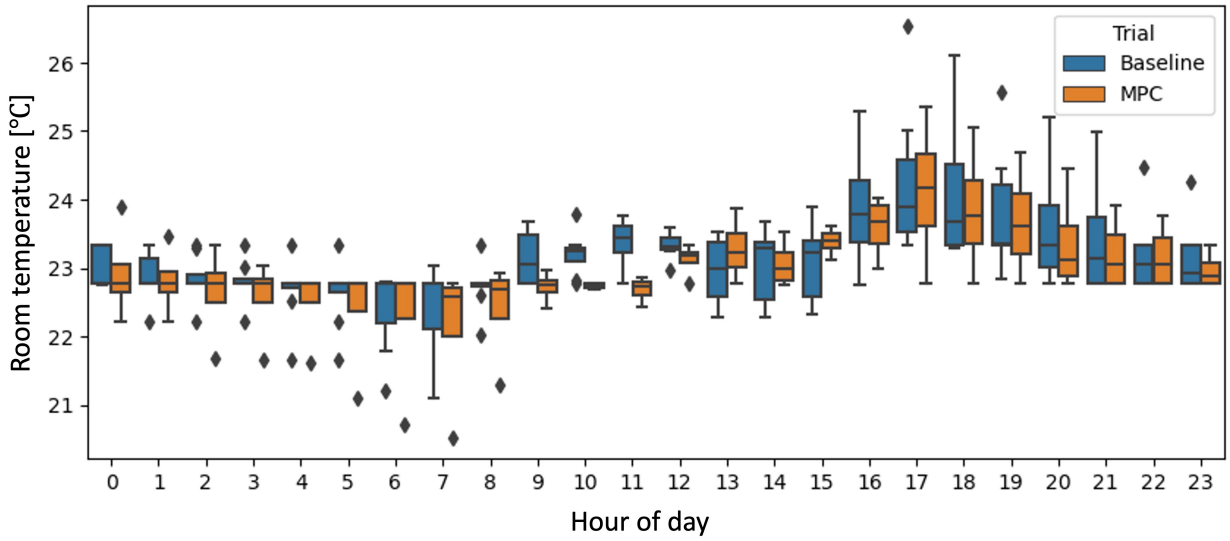

Figure 20: Comparison of the indoor temperature profiles between Baseline and MPC days in a classroom.

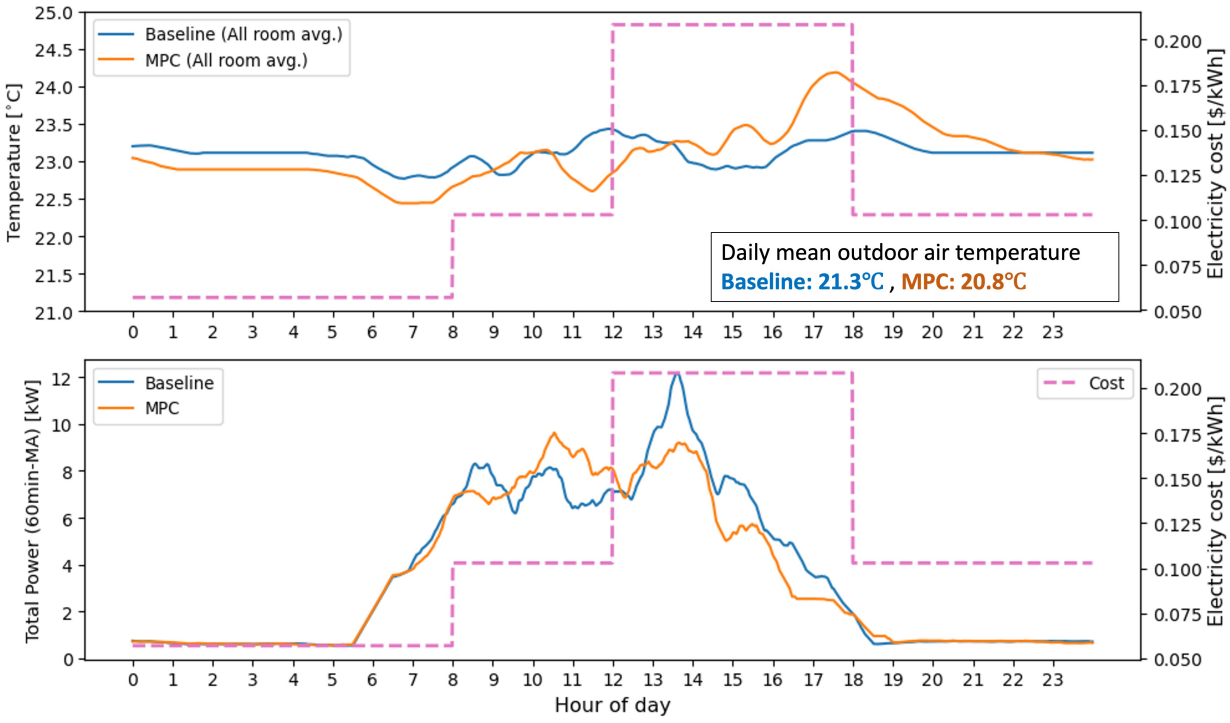

Figure 21: Comparison between the temperature and load profiles (zone-wise averaged) for a Baseline and an MPC day with the utility tariff (dashed pink line)

5.6. Peak demand reduction load shifting potential during a heating season

In this school, the RTUs have gas heating, so there are no benefits of MPC for grid services. However, due to increased momentum and pressure towards the electrification for building decarbonizations, it is of interest to test the MPC assuming that the RTUs are heat pumps. To do this, we retrained the building models using heating operation data. However, since the building envelope characteristics such as Rs and Cs remain unchanged, only the parameters of ρ_i and $Q_{h,i}$ in Table 2, which represent the unmeasured disturbance model and rated heating capacity for the i^{th} RTU, were retrained while keeping the others fixed at the values trained using cooling data. In addition, the thermostat's heating operating signals were used as a proxy for the heating energy consumption. Figure 22 shows the result comparison of heating operations between Baseline and MPC for heating days. For the Baseline, all RTUs were running simultaneously to heat up the space at 6:00. After some heating operations, the heating load (estimated by the RTF) decreases quickly as the internal and external heat increased. However, for the MPC, only 3-4 RTUs were running simultaneously with pre-heating and the RTU coordination. One should note that the preheating started at 5:00 AM. It could start earlier, but we enforced the MPC to start controlling thermostats after 5:00 AM because the fan operation during the early pre-dawn hours could bother the residential neighbors. Nevertheless, the MPC can successfully reduce peak demand with some preheating while showing a similar temperature control to the Baseline.

6. Further discussion, limitations and future works

The experimental test results have shown that K-12 schools can serve as valuable resources for providing load flexibility to the grid, and the proposed MPC can be one of the enabling technologies. However, several limitations still exist, particularly in the experimental design, modeling approach, and applicable HVAC types. Further work is required to address these limitations and enhance the performance of the MPC system for K-12 schools.

Experimental data: The data for the MPC days in Hot days (refer to Figure 19 and Section 5.3) is insufficient. While we anticipate that savings can be achieved, as discussed in Section 5.3, it is not clear enough to definitively conclude the demand reduction value of 24% for the Hot days. Further testing is necessary in the future to obtain a more comprehensive and reliable assessment of the savings.

Modeling: While the minimal infrastructure requirement of the proposed MPC may be economically appealing, it presents significant technical challenges in both modeling and control. To address these issues, this work adopted the lumped disturbance system identification algorithm and the offset-free MPC method. These approaches are specifically designed to mitigate the impact of unmeasured disturbances during the modeling and control phases, respectively. They both incorporate explicit models of unknown disturbances and are supported by strong theoretical foundations. However, the approaches still exhibit limited performance, particularly for the UMPC, which needs a high-quality model in a low-frequency domain (or long-term time scale). As depicted in Figures 13 and 14, the model's ability to forecast temperatures or loads is not highly accurate. The prediction errors are unavoidable within the low-cost MPC infrastructure due to two main reasons: (1) Accurate predictions, even with a perfect thermal building model, necessitate knowledge of future disturbances but are not measured in the control architecture, and (2) asymptotic model errors can arise from training data whenever dominant unmeasured disturbances present in the data collection period and input data exhibits correlations among themselves, irrespective of the system identification methods or model structures employed (we refer to [38] for the analytical discussion).

We have planned two approaches for addressing the challenges in the MPC implementation for K-12 schools. (1) Fortunately, K-12 schools possess unique characteristics that make their class and occupancy schedules relatively predictable with relatively small effort. Additionally, the occupancy load often represents

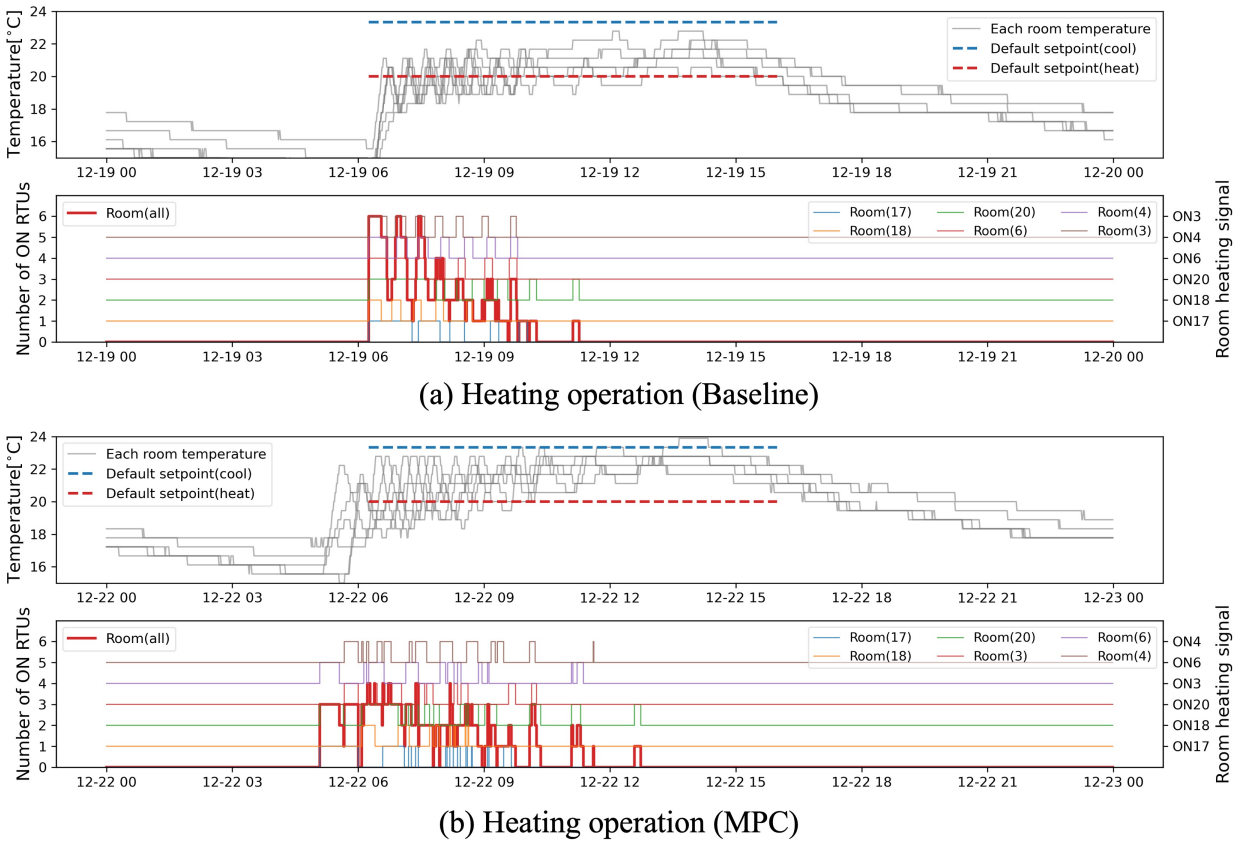

Figure 22: Comparison of heating operations between Baseline and MPC on a heating day.

a significant heat load in most school buildings. Although there are scalability concerns related to human involvement required to accurately track and update the school schedule for each classroom, and related to the use of design values for internal heat gain estimation which may introduce another uncertainty, we think it is worthwhile to explore including these factors in future endeavors. (2) We are developing a modeling approach, namely the hybrid modeling approach [49], to enhance prediction accuracy by modeling the prediction errors. That new approach has not been adopted in this initial implementation, since schools are not test facilities. However, as the research progresses using simulations and laboratories, we will incorporate the hybrid modeling approach to achieve improved long-term predictions.

Applicable building types: Since the proposed MPC is specifically designed to coordinate multiple packaged units, it is not applicable to buildings that are conditioned by central HVAC systems, such as an air handling unit (AHU) with variable air volume (VAV) boxes or a variable refrigerant flow (VRF) system. These centralized systems are commonly found in medium to large-sized commercial buildings. In such cases, the LMPC component of the MPC should be removed, and the UMPC component should be modified to incorporate a new constraint and a new power mapping strategy. These modifications would depend on the characteristics of the centralized system and the trajectories of the sum of zonal loads and new controllable devices. It is worth noting that newly constructed schools often employ central HVAC systems. Therefore, extending this MPC approach to accommodate various centralized systems is a topic for future research.

Comfort assessment: We adopted an industrial approach to assessing comfort by evaluating controls based on thermostat temperatures and by addressing teachers' comfort complaints if there are because those are the primary concern of facility operators and other measurements such as mean radiant temperature and air velocity for precise thermal comfort evaluation are unavailable in real buildings. However, we acknowledge the need for a precise assessment of thermal comfort. We will conduct a comfort survey using a comfort survey tool (e.g., https://cbe.berkeley.edu/resources/occupant-survey/) in the future.

7. Conclusions

In this work, an MPC was developed for flexible operation of K-12 school buildings and the performance was tested at K-12 school buildings. The demonstration is unique since very few advanced controls for load flexibility have been demonstrated in the K-12 market, despite the fact that there are more than 100,000 K-12 schools in the U.S. and that K-12 schools have tremendous potential to provide grid services. The MPC was deployed to six pilot classrooms during both heating and cooling seasons, and the results are compared with the existing control. Remarkable achievements of the MPC include a reduction of 24% in total peak power and a significant 30% decrease in HVAC peak power. Moreover, it shifted 16% of the load from high-priced on-peak periods to more affordable low-price periods. These results underscore the MPC's potential to harness the substantial demand response resources available in K-12 schools, simultaneously reducing utility costs and maintaining comfort. The experimental test results demonstrate that K-12 schools can serve as valuable resources for offering load flexibility to the grid. The proposed MPC can play a significant role as one of the enabling technologies. Nevertheless, limitations still remain, particularly in long-term predictions, due to unmeasured disturbances and theoretical model errors inherent to the low-cost MPC infrastructure. Further work is required to address these limitations and enhance the performance of the MPC system for K-12 schools.

8. Acknowledgement

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, through

the CRADA project entitled "A Low-cost, Scalable Control Solution for Grid-Interactive Small and Medium Sized Commercial Buildings".

The authors would like to thank Lance Bidnick, Patrick Bullock, and maintenance and operations staff from the Newport Mesa Unified School District for their strong support for this site demonstration work. Special thanks to Jennifer Worrall, Tracy Harton, Malakai Johnson and Daniel Quintana from the Community Energy Labs for collaboration and constructive feedback, and to Erika Gupta and Brian Walker from the Building Technologies Office of the U.S. Department of Energy for their support on the project.

References

[1] A. Satchwell, M. Piette, A. Khandekar, J. Granderson, N. Frick, R. Hledik, A. Faruqui, L. Lam, S. Ross, J. Cohen, K. Wang, D. Urigwe, D. Delurey, M. Neukomm, D. Nemtzow, A national roadmap for Grid-Interactive efficient buildings, Tech. rep., Lawrence Berkeley National Lab.(LBNL), Berkeley, CA (United States) (May 2021).

[2] U.S. Government Accountability Office, K-12 education: School districts frequently identified multiple building systems needing updates or replacement, Tech. rep., United States Government Accountability Office, Washington, D.C. (2020).

[3] New Buildings Institute, Decarbonization roadmap guide for school building decision makers, Tech. rep., New Buildings Institute, Portland, OR (2022).

[4] New Buildings Institute, Why K-12 should feature in america's national climate strategy, Tech. rep., New Buildings Institute, Portland, OR (2021).

[5] Y. Yang, Y. Lou, C. Payne, Y. Ye, W. Zuo, Long-term carbon intensity reduction potential of K-12 school buildings in the united states, Energy and Buildings 282 (2023) 112802. doi:10.1016/j.enbuild.2023.112802.

[6] Generation, Brighter future: A study on solar in U.S. K-12 schools, Tech. rep., Generation180 (2022).

[7] National Center for Education Statistics, Digest of education statistics - table 105.50. number of educational institutions, by level and control of institution: 2009-10 through 2019-20, https://nces.ed.gov/programs/digest/d21/tables/dt21_105.50.asp, accessed: 2022-12-9 (2022).

[8] U.S. Energy Information Administration (EIA), Commercial buildings energy consumption survey (CBECS) data: Table PBA3. sum of major fuel consumption totals and gross energy intensities by building activity subcategories, 2012, Tech. rep., U.S. Department of Energy, Washington, DC (2012).

[9] U.S. Energy Information Administration (EIA), Commercial buildings energy consumption survey (CBECS) data: Table B38-B46, Tech. rep., U.S. Department of Energy, Washington, DC (2012).

[10] 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, Annual reviews in control 50 (2020) 190–232. doi:10.1016/j.arcontrol.2020.09.001.

[11] J. Široký, F. Oldewurtel, J. Cigler, S. Prívara, Experimental analysis of model predictive control for an energy efficient building heating system, Applied energy 88 (9) (2011) 3079–3087. doi:10.1016/j.apenergy.2011.03.009.

[12] S. C. Bengea, A. D. Kelman, F. Borrelli, R. Taylor, S. Narayanan, Implementation of model predictive control for an HVAC system in a mid-size commercial building, HVAC&R Research 20 (1) (2014) 121–135. doi:10.1080/10789669.2013.834781.

[13] S. R. West, J. K. Ward, J. Wall, Trial results from a model predictive control and optimisation system for commercial building HVAC, Energy and Buildings 72 (2014) 271–279. doi:10.1016/j.enbuild.2013.12.037.

[14] P. Li, D. Vrabie, D. Li, S. C. Bengea, S. Mijanovic, Z. D. O'Neill, Simulation and experimental demonstration of model predictive control in a building HVAC system, Science and Technology for the Built Environment 21 (6) (2015) 721–732. doi:10.1080/23744731.2015.1061888.

[15] T. Hilliard, L. Swan, Z. Qin, Experimental implementation of whole building MPC with zone based thermal comfort adjustments, Building and environment 125 (2017) 326–338. doi:10.1016/j.buildenv.2017.09.003.

[16] J. Granderson, G. Lin, R. Singla, S. Fernandes, S. Touzani, Field evaluation of performance of HVAC optimization system in commercial buildings, Energy and Buildings 173 (2018) 577–586. doi:10.1016/j.enbuild.2018.05.048.

[17] D. Kim, J. E. Braun, J. Cai, D. L. Fugate, Development and experimental demonstration of a plug-and-play multiple RTU coordination control algorithm for small/medium commercial buildings, Energy and Buildings 107 (2015) 279–293. doi:10.1016/j.enbuild.2015.08.025.

[18] R. De Coninck, L. Helsen, Practical implementation and evaluation of model predictive control for an office building in brussels, Energy and Buildings 111 (2016) 290–298. doi:10.1016/j.enbuild.2015.11.014.

[19] 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, Energy and Buildings 178 (2018) 49–60. doi:10.1016/j.enbuild.2018.08.019.

[20] D. Sturzenegger, D. Gyalistras, M. Morari, R. S. Smith, Model predictive climate control of a swiss office building: Implementation, results, and Cost–Benefit analysis, IEEE Transactions on Control Systems Technology 24 (1) (2016) 1–12. doi:10.1109/TCST.2015.2415411.

[21] J. Drgoňa, D. Picard, L. Helsen, Cloud-based implementation of white-box model predictive control for a GEOTABS office building: A field test demonstration, Journal of process control 88 (2020) 63–77. doi:10.1016/j.jprocont.2020.02.007.

[22] S. Freund, G. Schmitz, Implementation of model predictive control in a large-sized, low-energy office building, Building and environment 197 (2021) 107830. doi:10.1016/j.buildenv.2021.107830.

[23] N. Morovat, A. K. Athienitis, J. A. Candanedo, B. Delcroix, Model-Based control strategies to enhance energy flexibility in electrically heated school buildings, Buildings 12 (5) (2022) 581. doi:10.3390/buildings12050581.

[24] N. Cotrufo, E. Saloux, J. M. Hardy, J. A. Candanedo, R. Platon, A practical artificial intelligence-based approach for predictive control in commercial and institutional buildings, Energy and Buildings 206 (2020) 109563. doi:10.1016/j.enbuild.2019.109563.

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

[26] K. Zhang, A. Prakash, L. Paul, D. Blum, P. Alstone, J. Zoellick, R. Brown, M. Pritoni, Model predictive control for demand flexibility: Real-world operation of a commercial building with photovoltaic and battery systems, Advances in Applied Energy 7 (100099) (2022) 100099. doi:10.1016/j.adapen.2022.100099.

[27] D. Kim, Z. Wang, J. Brugger, D. Blum, M. Wetter, T. Hong, M. A. Piette, Site demonstration and performance evaluation of MPC for a large chiller plant with TES for renewable energy integration and grid decarbonization, Applied energy 321 (2022) 119343.

[28] B. Merema, D. Saelens, H. Breesch, Demonstration of an MPC framework for all-air systems in non-residential buildings, Building and environment 217 (2022) 109053. doi:10.1016/j.buildenv.2022.109053.

[29] New Buildings Institute, Getting to zero: Zero energy schools stakeholder engagement and messaging, Tech. rep., New Buildings Institute, Portland, OR (2017).

[30] P. A. Torcellini, T.-K. L. Trenbath, N. Allen, M. McIntyre, A guide to zero energy and zero energy ready K-12 schools, Tech. Rep. NREL/TP-5500-72847, National Renewable Energy Lab. (NREL), Golden, CO (United States) (Aug. 2019). doi:10.2172/1557416.

[31] California Independent System Operator (CAISO), What the duck curve tells us about managing a green grid, CAISO, Folsom, CA (2016).

[32] R. Yang, M. W. Newman, Learning from a learning thermostat: lessons for intelligent systems for the home, in: Proceedings of the 2013 ACM international joint conference on Pervasive and ubiquitous computing, UbiComp '13, Association for Computing Machinery, New York, NY, USA, 2013, pp. 93–102. doi:10.1145/2493432.2493489.

[33] K. Buchanan, R. Russo, B. Anderson, The question of energy reduction: The problem(s) with feedback, Energy policy 77 (2015) 89–96. doi:10.1016/j.enpol.2014.12.008.

[34] D. Kim, J. E. Braun, MPC solution for optimal load shifting for buildings with ON/OFF staged packaged units: Experimental demonstration, and lessons learned, Energy and Buildings 266 (2022) 112118.

[35] S. Skogestad, I. Postlethwaite, Multivariable Feedback Control: Analysis and Design, John Wiley & Sons, 2005.

[36] R. Scattolini, P. Colaneri, Hierarchical model predictive control, in: 2007 46th IEEE Conference on Decision and Control, ieeexplore.ieee.org, 2007, pp. 4803–4808. doi:10.1109/CDC.2007.4434079.

[37] 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, Building and environment 107 (2016) 169–180. doi:10.1016/j.buildenv.2016.07.007.

[38] 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, Energy and Buildings 167 (2018) 359–369. doi:10.1016/j.enbuild.2017.12.007.

[39] L. Ljung, System identification, in: A. Procházka, J. Uhlíř, P. W. J. Rayner, N. G. Kingsbury (Eds.), Signal Analysis and Prediction, Birkhäuser Boston, Boston, MA, 1998, pp. 163–173. doi:10.1007/978-1-4612-1768-8_11.

[40] J. Braun, N. Chaturvedi, An inverse Gray-Box model for transient building load prediction, HVAC&R Research 8 (1) (2002) 73–99. doi:10.1080/10789669.2002.10391290.

[41] P. Virtanen, R. Gommers, T. E. Oliphant, M. Haberland, T. Reddy, D. Cournapeau, E. Burovski, P. Peterson, W. Weckesser, J. Bright, S. J. van der Walt, M. Brett, J. Wilson, K. J. Millman, N. Mayorov, A. R. J. Nelson, E. Jones, R. Kern, E. Larson, C. J. Carey, İ. Polat, Y. Feng, E. W. Moore, J. VanderPlas, D. Laxalde, J. Perktold, R. Cimrman, I. Henriksen, E. A. Quintero, C. R. Harris, A. M. Archibald, A. H. Ribeiro, F. Pedregosa, P. van Mulbregt, SciPy 1.0 Contributors, SciPy 1.0: fundamental algorithms for scientific computing in python, Nature methods 17 (3) (2020) 261–272. doi:10.1038/s41592-019-0686-2.

[42] S. Rouchier, M. J. Jiménez, S. Castaño, Sequential monte carlo for on-line parameter estimation of a lumped building energy model, Energy and Buildings 187 (2019) 86–94. doi:10.1016/j.enbuild.2019.01.045.

[43] A. A. Bell, HVAC: equations, data, and rules of thumb, McGraw-Hill, New York, 2000.

[44] B. Boyle, pyglpk: Updated fork of t. finley's PyGLPK module (2014).

[45] National Oceanic and Atmospheric Administration (NOAA), XML feeds of current weather conditions, https://w1.weather.gov/xml/current_obs/, accessed: 2022-12-21 (2022).

[46] National Oceanic and Atmospheric Administration (NOAA), Hourly tabular forecast, https://www.weather.gov/wrh/wxtable, accessed: 2022-12-20 (2022).

[47] W. F. Holmgren, C. W. Hansen, M. A. Mikofski, Pvlib python: A python package for modeling solar energy systems, Journal of open source software 3 (29) (2018) 884. doi:10.21105/joss.00884.

[48] R. N. Edmondson, Multi-level block designs for comparative experiments, Journal of Agricultural, Biological and Environmental Statistics 25 (4) (2020) 500–522. doi:10.1007/s13253-020-00416-0.

[49] S. W. Ham, D. Kim, Hybrid modeling approach for better identification of building thermal network model and improved prediction, 2022, p. Paper 420.