Title

CEL: A Low-cost, Scalable Control Solution for Grid-Interactive Small and Medium Sized Commercial Buildings

Authors

Kim, Donghun

Ham, Sang Woo

Publication Date

2023-12-14

Peer reviewed

CEL: A Low-cost, Scalable Control Solution for Grid-Interactive Small and Medium Sized Commercial Buildings

Donghun Kim (PI)[a], Sang Woo Ham[a]

[a]*Building Technology & Urban Systems Division, Lawrence Berkeley National Laboratory, Berkeley, CA, USA*

Abstract

This project aims to develop and enhance a low-cost, highly scalable control solution for Small and Medium-Sized Commercial Buildings (SMCB), assess the business potential at multiple sites, and perform commercialization efforts in collaboration with Community Energy Labs (CEL). The technology can be applied to any buildings served by multiple units, with the benefits being greatest for open-spaced buildings, such as banks, retail stores, restaurants, and factories. This project aims to develop an affordable control solution for: 1) SMCB grid responsiveness, 2) reduction of GHG by changing unit operations, 3) greater reduction in utility costs, and 4) rapid adoption in the marketplace. The proposed technology will be built on a previously developed and demonstrated MPC solution. The minimal sensor requirement and less need of control expertise are the unique feature of the algorithm that leads to low capital and maintenance costs, and short installation and implementation time. These attributes contribute to low capital and maintenance costs, as well as a short installation and implementation time. However, these advantages come with a trade-off: increased difficulties and unreliability when applying traditional modeling and MPC control approaches due to limited information. This final report provides the field demonstration of the low-cost, highly scalable control solution on K-12 school buildings. This report begins the description of the school buildings and practical barriers of the application of MPC for real occupied buildings. Then, the details of the MPC solution to reduce peak demand and peak-time energy use. Finally, the results of the field demonstration are presented with lessons learned.

Nomenclature

API: Application programming interface

DX: Direct expansion air-conditioner

HVAC: Heating, ventilation, and air-conditioning

LD: Lumped disturbance

MPC: Model predictive control

RTF: Runtime fraction

RTU: Rooftop unit

SYSID: System identification

TABS: Thermally activated building systems

TES: Thermal energy storage

TOU: Time-of-Use utility tariff

($\mathbf{A}(\cdot)$, $\mathbf{B_u}(\cdot)$, $\mathbf{B_w}(\cdot)$, $\mathbf{C}(\cdot)$): 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]

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

($\mathcal{F}(\cdot)$, $\mathcal{G}(\cdot)$): 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

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

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($\dot{Q}_{\mathrm{c},i}$) of ith zone [kW]

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

($T_{\mathrm{l}}, T_{\mathbf{u}}$): 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]

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

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

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

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

$u_{\mathbf{h},i}$, $u_{\mathbf{c},i}$**:** 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)

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

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

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

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

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

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

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

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

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

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

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

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. 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 about 100,000 K-12 schools of the U.S. [5], 2) they are the largest energy consumer in the public sector [6], 3) the HVAC energy accounts for 46% of the total energy consumption in K-12 schools [7], 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, for example, 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) [8], which could optimize the operation of the 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 regulation.

Although numerous studies of using MPC for building applications have been carried out and published for the last decades, the majority of them are still simulation studies and only a limited number of papers that report MPC test results at actual buildings are available in the literature. The lack of references in the literature raises significant concerns of unclear real world benefits and technical readiness level of the MPC technology from building owners and industries, which is clearly a barrier to the wide-adoption of MPC.

Some MPC field works as follows. Kim et al., [9] 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 cyclings, 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 [10] 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 [11] 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. Some studies [12, 13] implemented an MPC in large office buildings with thermally activated building systems (TABS). By optimizing the control sequences of the complex HVAC system, they showed a significant amount of heating energy reductions. Cotrufo et al. [14] 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. [15] 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. Kun et al. [16] 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. [17] developed a campus-scale MPC with thermal energy storage (TES). The MPC is designed to reduce redundant PV generation and electricity consumption during

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%.

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 from several points including the followings: 1) HVAC systems are the most outdated part of the facility in K-12 schools [2], so the advanced control solution often requires major retrofit in HVAC hardware. 2) HVAC systems and thermostats are not usually complete one-way supervisory control systems. Specifically, the default setpoint schedules of thermostats are mainly managed by the school district's policy, but they can be overridden when feeling uncomfortable by end-users (i.e., teachers and students) like a thermostat control in a residential building [18]. 3) there are many stakeholders such as teachers, facility operators, energy managers, school boards and utility companies in K-12 schools [19, 20] and, when some of them are not familiar with the new technology and benefits, they are reluctant to invest in or implement the new technology. Nonetheless, to achieve the aforementioned environmental, financial and educational benefits of transitioning to the grid-interactive school operation, and to provide clear technical feasibility of the MPC technology, MPC field studies at schools are necessary.

This paper is to partially fill the gap. We present the first experimental study of an MPC at a K-12 school, and share the practical challenges and lessons-learned from the site 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 practical challenges in K-12 schools

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 was 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 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 have WiFi-enabled thermostats 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 3.3. The school is under the SCE-TOU-GS-2 grandfathered tariff: the grandfathered time-of-use (TOU) has the 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 [21]. 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 was collected in 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).

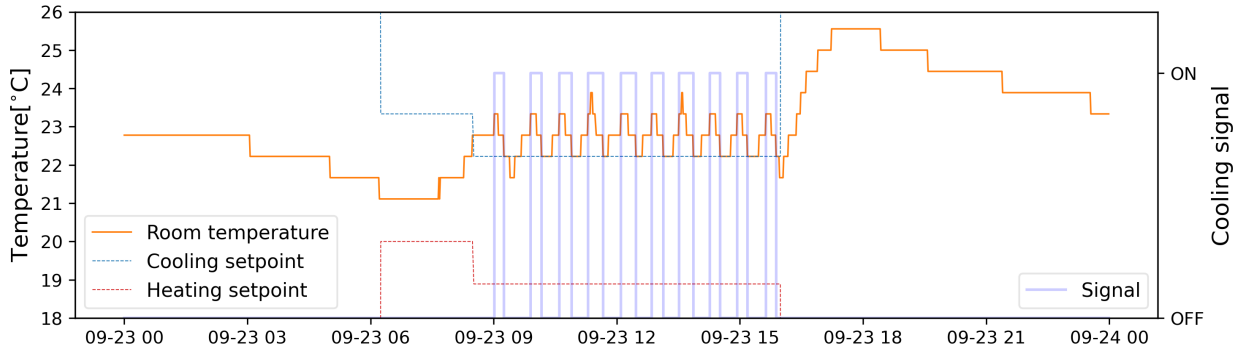

Figure 2: Typical RTU operation of one classroom for a day in the cooling season.

Figure 3 presents air temperatures (top) and cooling signals (middle) for the six classrooms, and electricity usages (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 middel 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 in 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 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. Practical 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

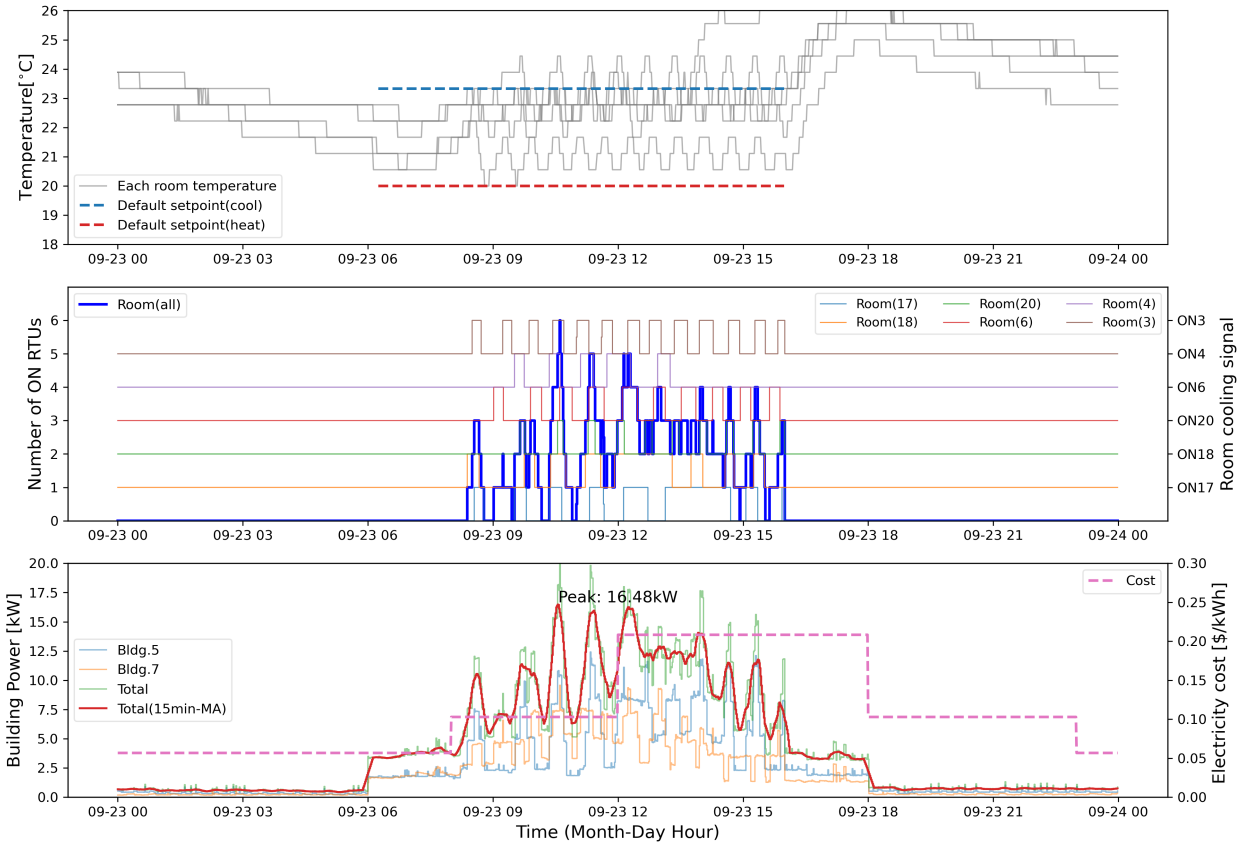

Figure 3: Air temperatures (top), cooling signals (mid), and the power(electricity) use (bottom) of six classrooms in the two buildings for a typical cooling day.

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 [18, 22]. 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 habitual. For example, the teacher lowered the setpoint near 8:00 (student arrival time) though the room temperature was still at the setpoint (Figure 4). Although it is not visualized in this paper, this habitual override behavior is consistently observed throughout most classrooms.

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

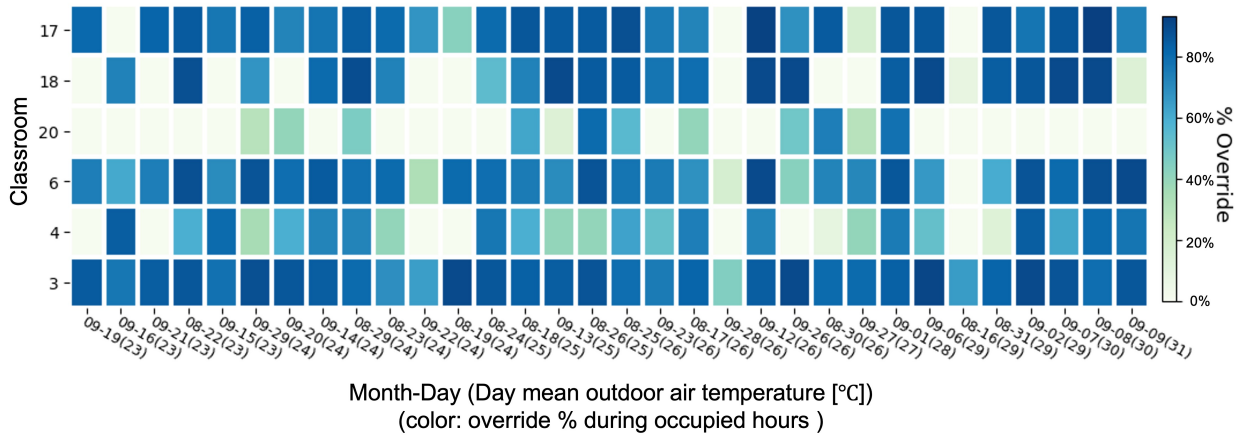

Figure 4: Percentage of schedule overridden hours during RTU operation hours by days and classrooms (sorted by daily mean outdoor air temperature).

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 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 a 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 for a simple feedback controller which is to track a setpoint. It also challenges in developing a good building thermal model because the transient response is poorly captured.

3. MPC design

In this study, the MPC algorithm introduced in [23] was applied and modified. The MPC has a hierarchical control structure as shown in Figure 7. The upper-level MPC (UMPC) generates the desired setpoint

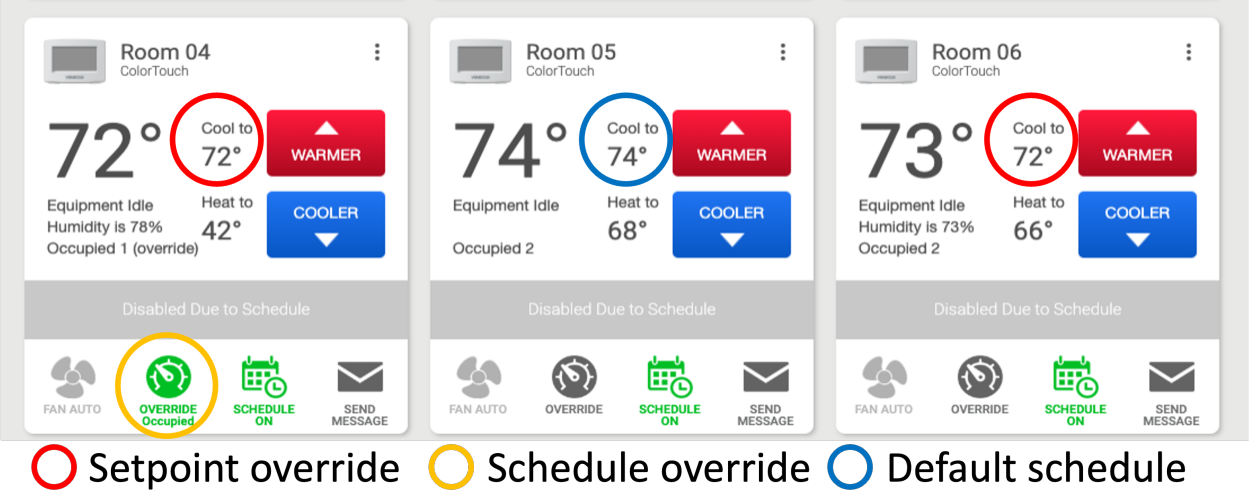

Figure 5: Snapshot of thermostat dashboard for a default schedule day with various override behaviors.

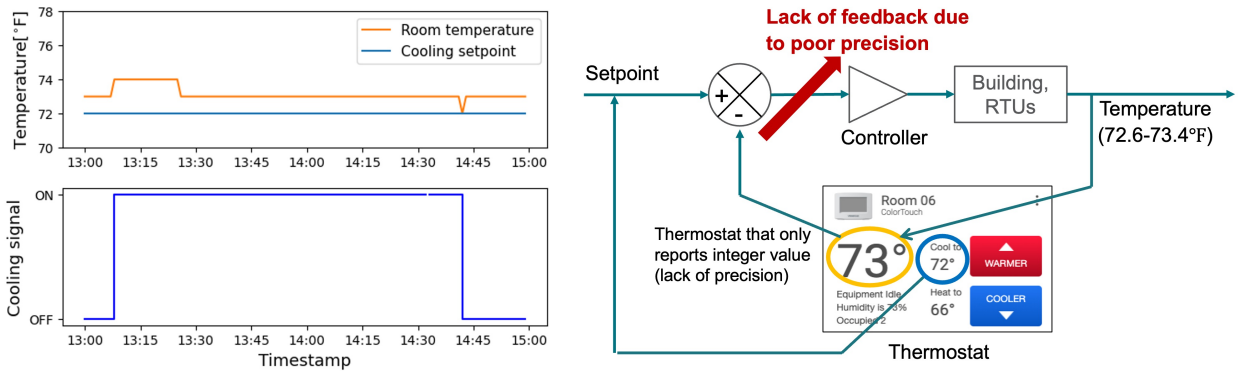

(a) Temperature/cooling response for 13:00-15:00

(b) Issues in integer thermostat in feedback control

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

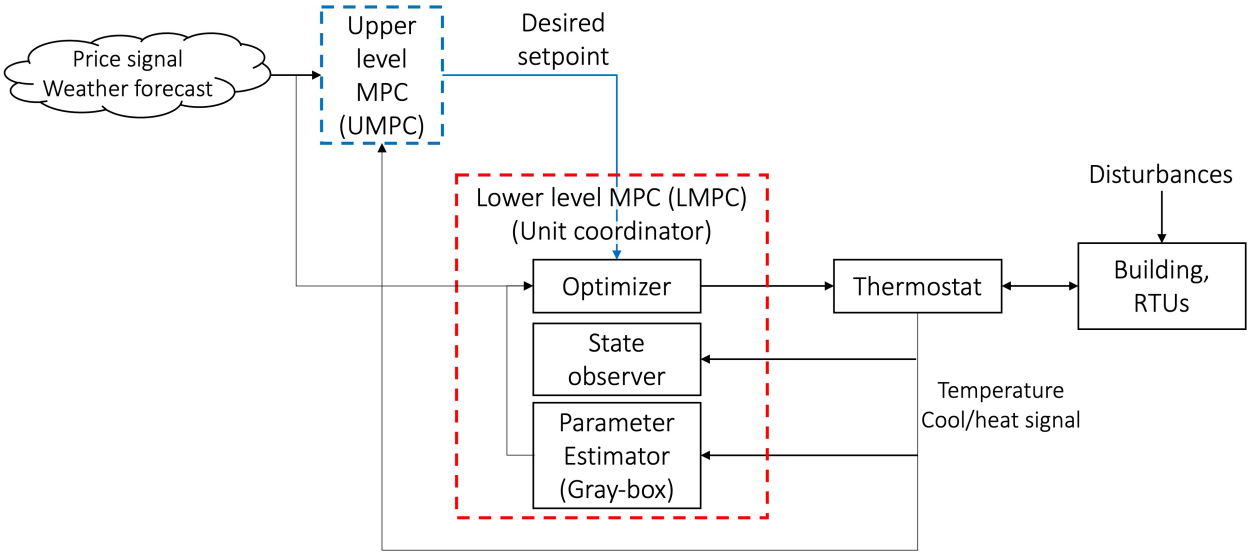

Figure 7: Schematic diagram of a hierarchical MPC for optimal load shifting and peak demand reduction

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). In this study, we briefly introduce the conceptual framework of the building thermal model and hierarchical MPC. More details can be found in our previous work [23].

3.1. 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 structure simplicity but robust prediction performance by holding physical principles.

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[24] 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 [24, 25], 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 [24, 25, 23], 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 [26, 24] 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 [27] which only looks at simulation errors and thus ignores unmeasured disturbances can be found in [24, 25].

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

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

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

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

We investigated both single-zone and multi-zone thermal network model structures (as shown in Figure 8 (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[25] 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

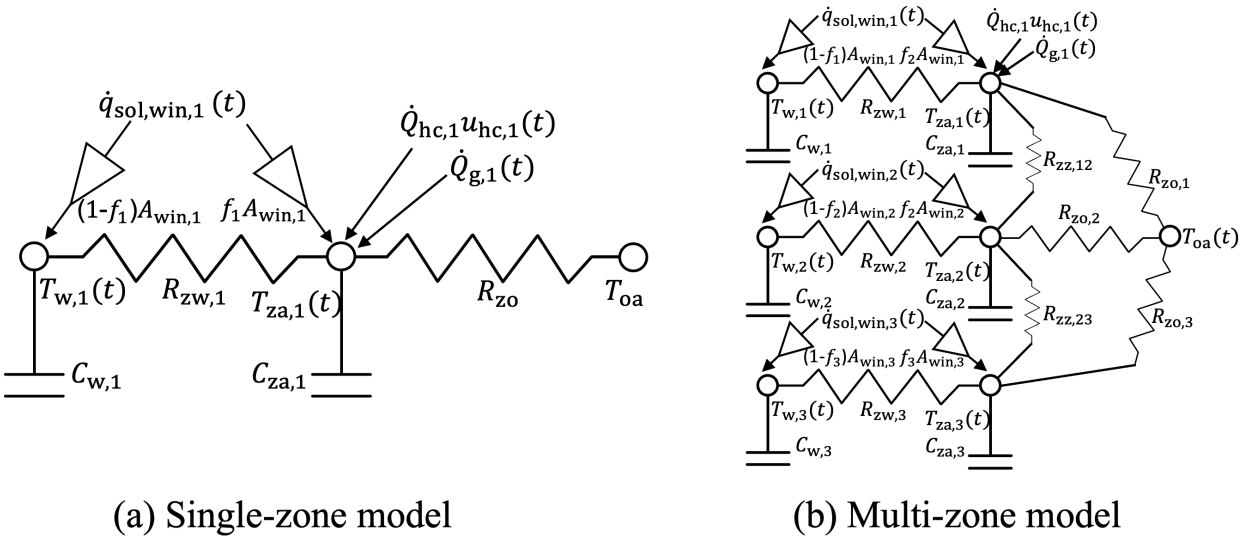

Figure 8: Two types of gray-box model structures.

Table 2: Rules to generate optimization boundary.

Required metadata input	• $A_{\text{fl},i}$: Floor area of ith classroom. • $A_{\text{mw},i}$,$tilt_{\text{mw},i}$,$Az_{\text{mw},i}$: 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)]. • $\dot{Q}^*_{\text{rated,hc}}$ or $\dot{Q}_{\text{rated,hc}}$: Area normalized or nominal rated heat/cool rate [kW/m^2 or kW]. • L: Max. and min. values of scale variable [-]
	Rules for [minimum, maximum] values
Variables	• $C_{\text{w},i}$: $[C_{\text{w,min}}A_{\text{fl},i}th_{\text{min}}$, $C_{\text{w,max}}A_{\text{fl},i}th_{\text{max}}]$ • $C_{\text{za},i}$: $[C_{\text{za,min}}A_{\text{fl},i}H_{\text{min}}$, $C_{\text{za,max}}A_{\text{fl},i}H_{\text{max}}]$ • $R_{\text{zw},i}$: $[1/(4(\sqrt{A_{\text{fl},i}})H_{\text{max}}U_{\text{max}})$, $1/(4(\sqrt{A_{\text{fl},i}})H_{\text{min}}U_{\text{min}})]$ • $R_{\text{zo},i}$: [1e-4, $R_{\text{zw,max},i}/20$] • f_i: [1e-9,1] • $A_{\text{win},i}$: $[L_{\text{min}}A_{\text{fl},i}$, $L_{\text{max}}A_{\text{fl},i}]$ • $Q_{\text{h},i}$: $[L_{\text{min}}\dot{Q}^*_{\text{rated,h}}A_{\text{fl},i}$, $L_{\text{max}}\dot{Q}^*_{\text{rated,h}}A_{\text{fl},i}]$ or $[L_{\text{min}}\dot{Q}_{\text{rated,h}}$, $L_{\text{max}}\dot{Q}_{\text{rated,h}}]$ • $Q_{\text{c},i}$: $[L_{\text{min}}\dot{Q}^*_{\text{rated,c}}A_{\text{fl},i}$, $L_{\text{max}}\dot{Q}^*_{\text{rated,c}}A_{\text{fl},i}]$ or $[L_{\text{min}}\dot{Q}_{\text{rated,c}}$, $L_{\text{max}}\dot{Q}_{\text{rated,c}}]$ • ρ: [-0.999,0.999]

problem for the LD is solved with a multi-start method using the `least_squares` function of Scipy optimizer [28]. The first 12-hour data is used for initializing the initial states using a Kalman filter [29]. 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 $[1e-4,\ 3e-3]$, [2.0, 5.0], [0.05, 1.0], [5, $3e3$] and $[1e-5,\ 5]$ (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 ($\dot{Q}^{*}_{\mathrm{rated,hc}} = 0.1262\mathrm{kW/m^2}$) [30]. 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.

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

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

3.2. 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).

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

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 (3.1) given the data $\mathcal{D}_k = \{y(k-1), y(k-2), \cdots, u(k+j-1), u(k+j-2), \cdots\}$. ($T^{\mathrm{L}}_{\mathrm{l},i}$, $T^{\mathrm{L}}_{\mathrm{u},i}$) 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 $\Gamma_{\text{l}}, \Gamma_{\text{u}} (\in \mathbb{R}^+)$ 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 [11] 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., $0.7 \times \sum_{i=1}^{n} P_{\text{RTU},i}$, 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 are 16.5kW (δ=11.6kW, see Figure 1), 0.3°C of temperature violation is equivalent to the use of one more RTU.

While LMPC coordinates unit stages, UMPC is supposed to provide optimal setpoints that minimize energy cost ($ER \times P_{\text{RTU}}$) 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:

$$\min \quad \sum_{j=1}^{N_p^U} \sum_{i=1}^{p} ER(k+j-1) P_{\text{RTU},i}(k+j-1) \bar{u}_{\text{hc},i}(k+j-1) \tag{10}$$

$$+\omega_{\text{d}}\delta + \omega_{\text{l}}\Gamma_{\text{l}} + \omega_{\text{u}}\Gamma_{\text{u}} \tag{11}$$

$$s.t. \quad T_{\text{l},i}^{\text{U}} - \Gamma_{\text{l}} \leq E(\bar{y}_i(k+j)|\mathcal{D}_k) \leq T_{\text{u},i}^{\text{U}} + \Gamma_{\text{u}} (\forall i \in \{1, \cdots, n\})$$

$$\sum_{i=1}^{n} P_{\text{RTU},i}(k+j-1) \bar{u}_{\text{hc},i}(k+j-1) \leq \delta$$

$$0 \leq \bar{u}_{\text{hc},i}(k+j-1) \leq 1 \quad (\forall j \in \{1, \cdots, N_{\text{p}}^{\text{U}}\}), \tag{12}$$

where $\bar{u}_{\text{hc},i}$ and $\bar{y}_i$ are the moving averaged RTU stages (normalized 0-1 scale, i.e., RTF) and thermostat temperatures, respectively, for UMPC timestep (i.e., 30 minutes in this study).

While the upper and lower temperature bounds of UMPC (i.e., $T_{\text{l},i}^{\text{U}}$, $T_{\text{u},i}^{\text{U}}$) are a comfortable temperature range specified by users, the optimal input trajectories of UMPC are converted to the desired setpoint trajectories through Eq. 13, which are used for temperature boundaries for LMPC ($T_{\text{l},i}^{\text{L}}$, $T_{\text{u},i}^{\text{L}}$.

$$T_{\text{u},i}^{\text{L}}(k) = \begin{cases} \bar{y}_i^*(k), & \text{if } \bar{u}_{\text{c},i}^*(k) > 0 \\ T_{\text{u},i}^{\text{U}}(k), & \text{if } \bar{u}_{\text{c},i}^*(k) = 0 \end{cases} \tag{13}$$

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

where $\bar{u}_{\text{hc}i}^*(k)$ 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 bound of the LMPC.

LMPC and UMPC are linear and mixed integer linear programming problems and are solved by using `pyglpk` [31]. 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. 13. 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.

3.3. MPC improvement 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 idea of reflecting the user's override actions is to 1) detect which unit(s) is overridden, 2) predict the stage status for the corresponding unit until the next sampling time, and 3) impose the unit 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 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 9, 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 ok for the LMPC due to the short prediction horizon and sampling time.

Another remedy for the school applications is to overcome 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):

$$\begin{aligned} \mathbf{y} &= \sigma(\mathbf{z}_{\text{true}}) \\ \mathbf{z}_{\text{true}} &= \mathbf{y} + \boldsymbol{\epsilon}_y \end{aligned} \tag{14}$$

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:

$$\boldsymbol{\varepsilon}(k) = \mathbf{z}_{\text{true}}(k) - \boldsymbol{\epsilon}_y(k) - \hat{\mathbf{y}}(k|k-1) \tag{15}$$

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

$$\begin{aligned}
\hat{\mathbf{y}}(k|k-1) &= \mathbf{C}_{\mathrm{d}}\hat{\mathbf{x}}(k|k-1) + \hat{\boldsymbol{\zeta}}(k|k-1) \\
&= \mathbf{C}_{\mathrm{d}}\hat{\mathbf{x}}(k|k-1) + \mathcal{F}\hat{\boldsymbol{\zeta}}(k-1|k-2) + \mathcal{G}\boldsymbol{\varepsilon}(k-1) \\
&= \mathbf{C}_{\mathrm{d}}\hat{\mathbf{x}}(k|k-1) + \mathcal{F}\hat{\boldsymbol{\zeta}}(k-1|k-2) + \\
&\quad \mathcal{G}\cdot(\mathbf{z}_{\mathrm{true}}(k-1) - \boldsymbol{\epsilon}_y(k-1) - \hat{\mathbf{y}}(k-1|k-2)) \\
&= \hat{\mathbf{y}}_{\mathrm{true}}(k|k-1) + \sum_{i=2}^{k} \mathcal{G}\boldsymbol{\epsilon}_y(i-1)
\end{aligned} \tag{16}$$

The true innovation ($\epsilon_{\mathrm{true}}(k) = \mathbf{z}_{\mathrm{true}}(k) - \hat{\mathbf{y}}_{\mathrm{true}}(k|k-1)$) 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. 14, 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. 17) is applied to estimate the innovation.

$$\begin{aligned}
\mathbb{E}[\boldsymbol{\varepsilon}_{\mathrm{true}}(k)] &= \mathbb{E}[\mathbf{z}_{\mathrm{true}}(k)] - \mathbb{E}[\hat{\mathbf{y}}_{\mathrm{true}}(k|k-1)] \\
&= \mathbb{E}[\mathbf{y}(k) + \boldsymbol{\epsilon}(k)] - \mathbb{E}\left[\hat{\mathbf{y}}(k|k-1) - \sum_{i=2}^{k-1}\mathcal{G}\boldsymbol{\epsilon}_y(i-1)\right] \\
&= \sigma(\mathbf{y}(k)) - \sigma(\hat{\mathbf{y}}(k|k-1))
\end{aligned} \tag{17}$$

4. Experiment description

4.1. Control system architecture

Figure 9 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 (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 was 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) [32]. The solar radiation is collected by recording the current solar forecast data. Solar forecast is estimated from NOAA's cloud cover forecast [33] and pvlib's [34] 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 Controls service.

[1] https://communityenergylabs.com/

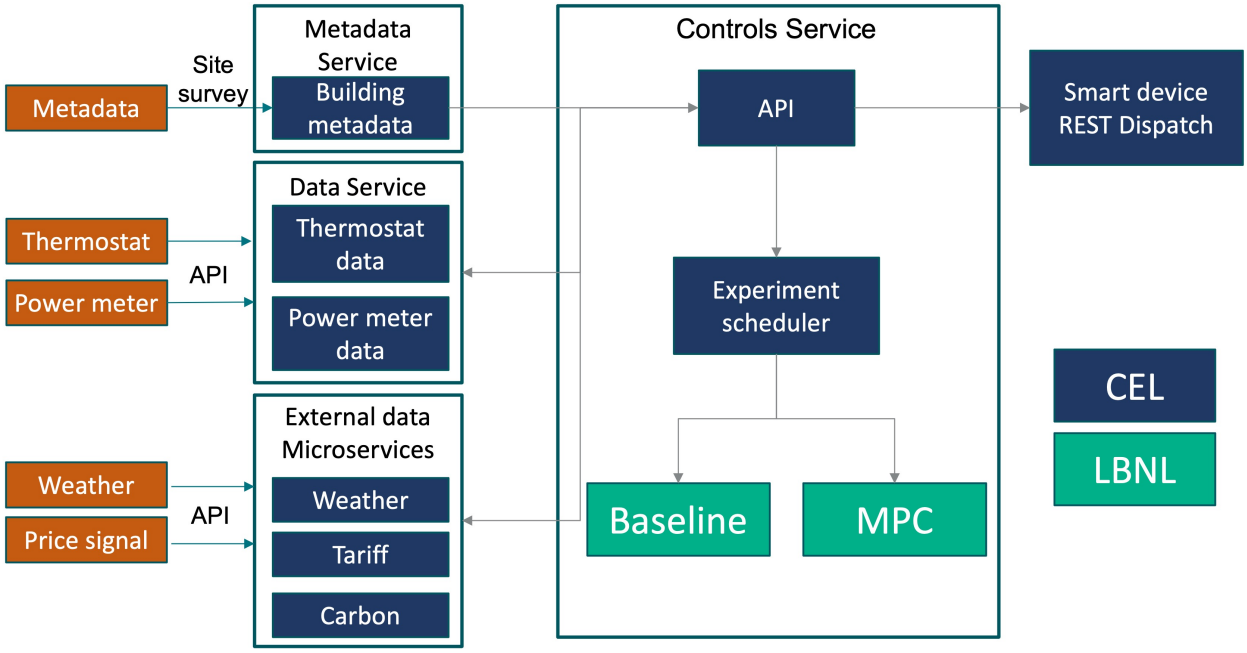

Figure 9: Schematic diagram of data collection and MPC service.

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 [35] will use "blocking" to ensure that the same number of days are assigned to each control for each block period (see Figure 10 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 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

Figure 11 shows the comparison between the measured cooling RTF and predicted one calculated by Eq. 8 for a classroom. Overall, the model can capture the timing of cooling operation (i.e., zigzag shape) though the magnitude is smaller than the measurement. This could be attributed from 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

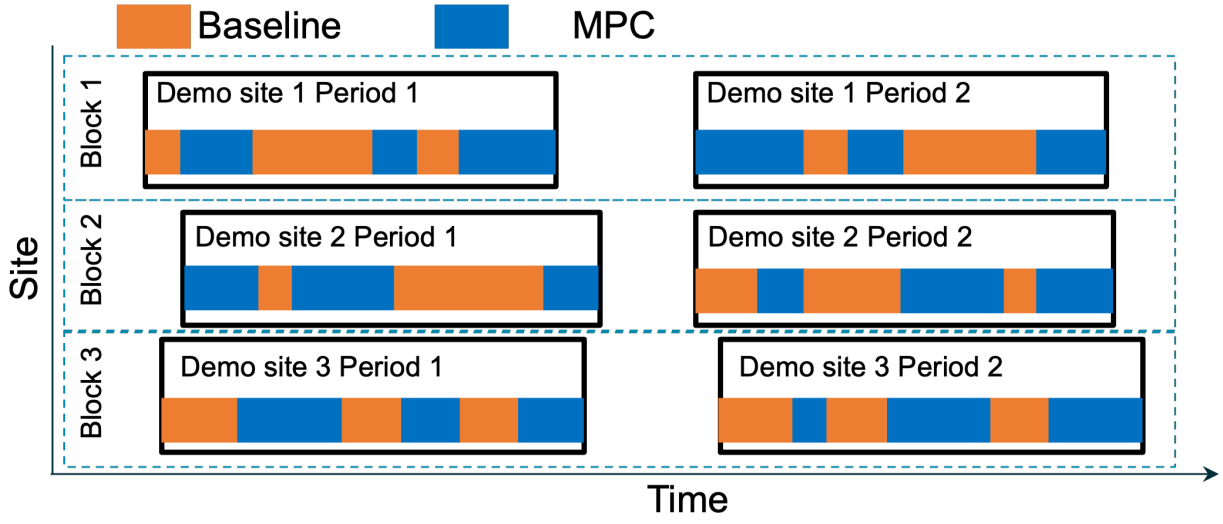

Figure 10: Conceptual diagram of the randomized block schedule.

cooling RTF included all unmeasured disturbance terms such as infiltration. To summarize, the model is well identified through the LD approach and can capture the thermal characteristics of the building, but it shows some discrepancy due to unmeasured heat gains.

5.2. Operational profiles

Figure 12 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 2.1 and 2.2). Due to the decentralized and independent thermostat control scheme, 5 to 6 RTUs operated simultaneously several times, resulting in the peak power of 18.16kW.

On the other hand, Figure 13 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 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-min 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 12 with a reduced peak power of 13.98kW.

5.3. Peak demand reduction

Figure 14 shows the daily peak demand for cooling days with respect to daily averaged outdoor air temperature. Remind 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 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

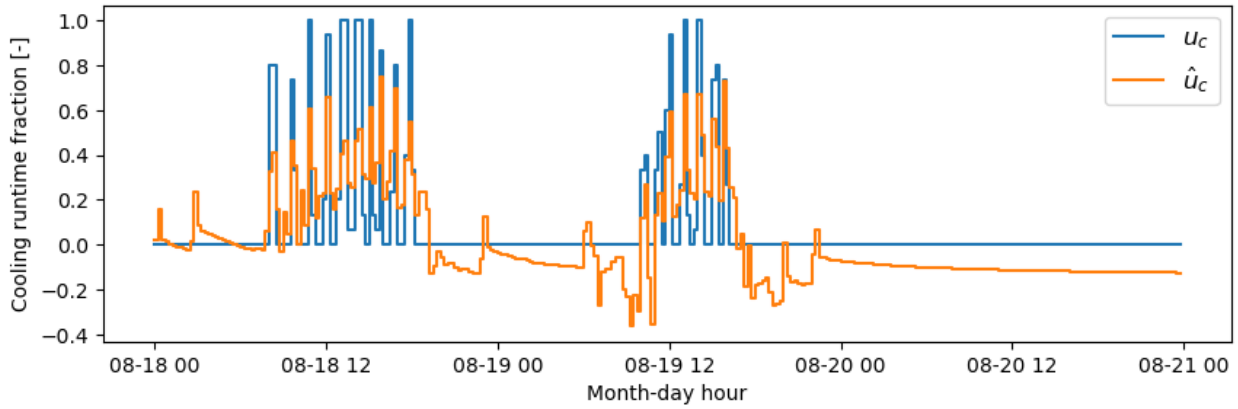

Figure 11: Comparison between the required ($\hat{u}_c$) and measured (u_c) amount of cooling runtime fraction in one classroom.

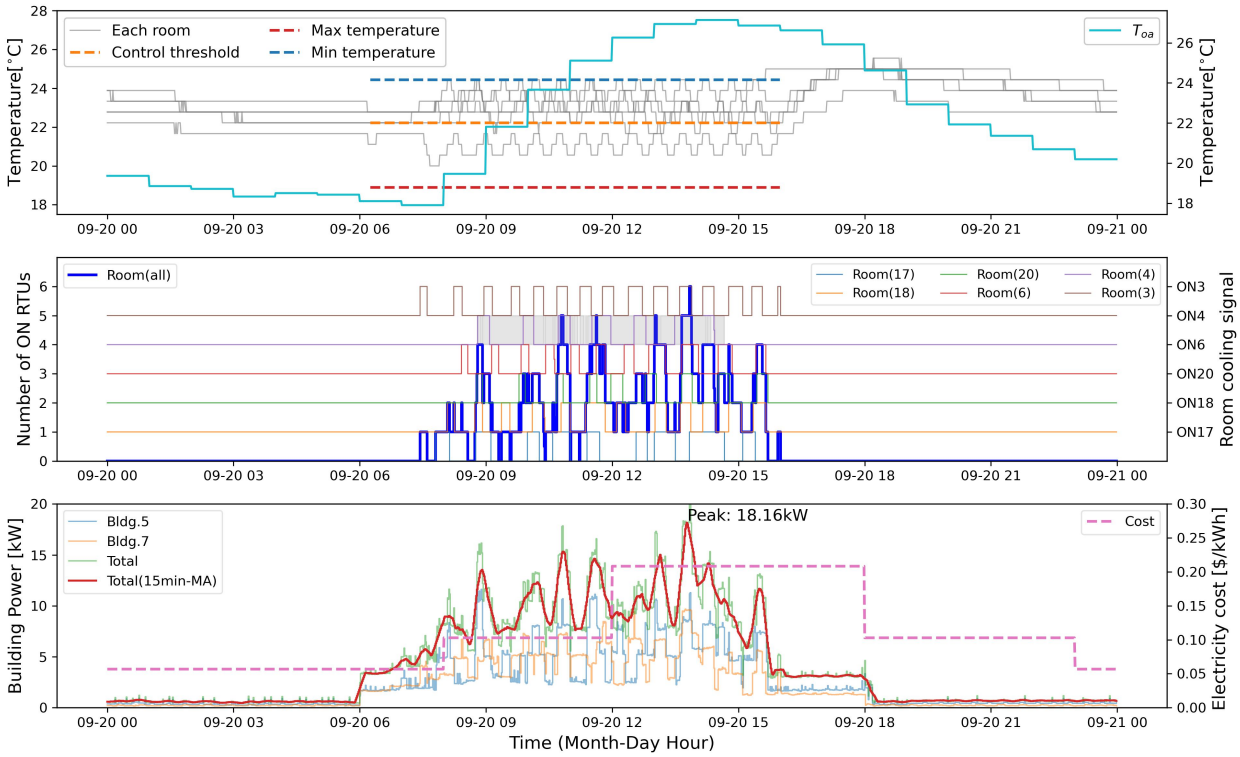

Figure 12: Summary of RTU operation in a default schedule day (Baseline); top: room temperatures with min/max comfort boundary (thermostat deadband adjusted setpoints) and outdoor air temperature, middle: RTU signals and override period (grayed area), bottom: building power and electricity price signal.

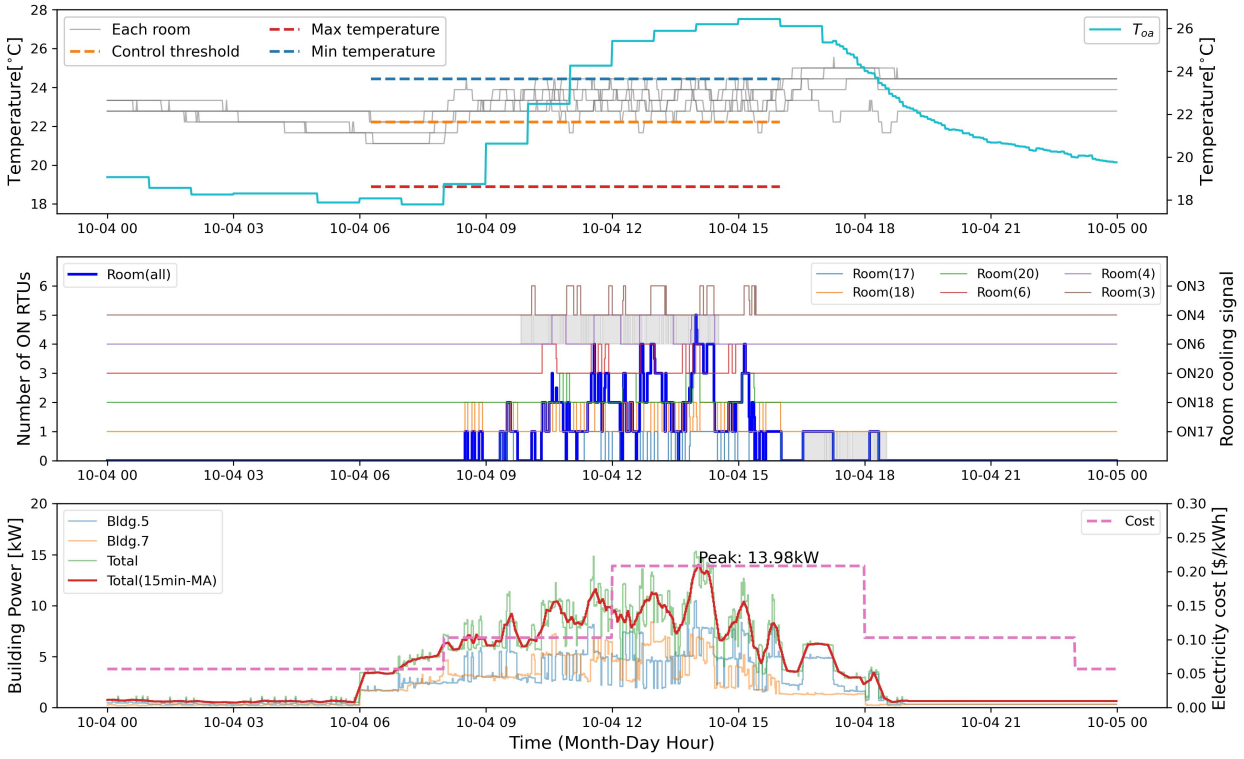

Figure 13: Summary of RTU operation in an MPC day (MPC); top: room temperatures with min/max comfort boundary (thermostat deadband is included) and outdoor air temperature, middle: RTU signals and override period (grayed area), bottom: building power and electricity price signal.

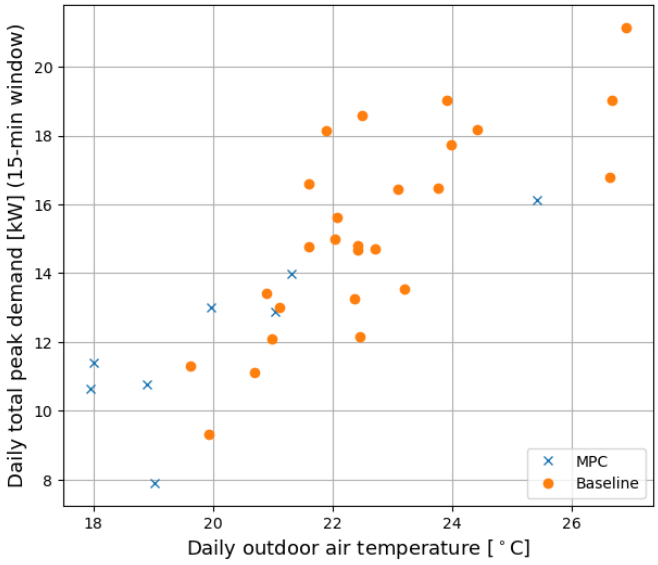

Figure 14: Peak demand vs. daily outdoor air temperature for the whole data collecting period (Sep-Oct, 2022).

outdoor temperatures. Thus, for demand savings estimation, the mean temperature line was divided into two groups (below and above 22^oC) and compared the maximum peaks for each group. The indices of $B1$ and $M1$ in Table 3 represent the maximum daily peaks for the Baseline and MPC respectively for the milder period. Likewise $B2$ and $M2$ indicate those for the hotter period. The MPC reduced the total peak and HVAC peak powers around 25% and 30% respectively for both conditions. For the hotter condition, there is only one data point for the MPC (due to the change of the season once the MPC test initiated and the randomized test schedule) and thus the estimation might be highly inaccurate. However, a closer data analysis indicates that this data point is the worst-case scenario for the MPC and therefore the values are not overestimated. Figure 15 shows the responses for the date. Note that four of six classrooms were in the override mode, so that the MPC could only control two classrooms (see Section 2.3 for MPC operation under override mode). Even with this limited controllability, MPC finds a way to reduce the peak demands (see the thick red-line in the bottom figure) by looking at other unit operations while maintaining all temperatures within the prescribed bounds.

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	26.9°C	21.1kW	17.3kW	-	-
M1	MPC	21.3°C	14.0kW	9.4kW	23% (B1)	32% (B1)
M2	MPC	25.4°C	16.1kW	12.1kW	24% (B2)	30% (B2)

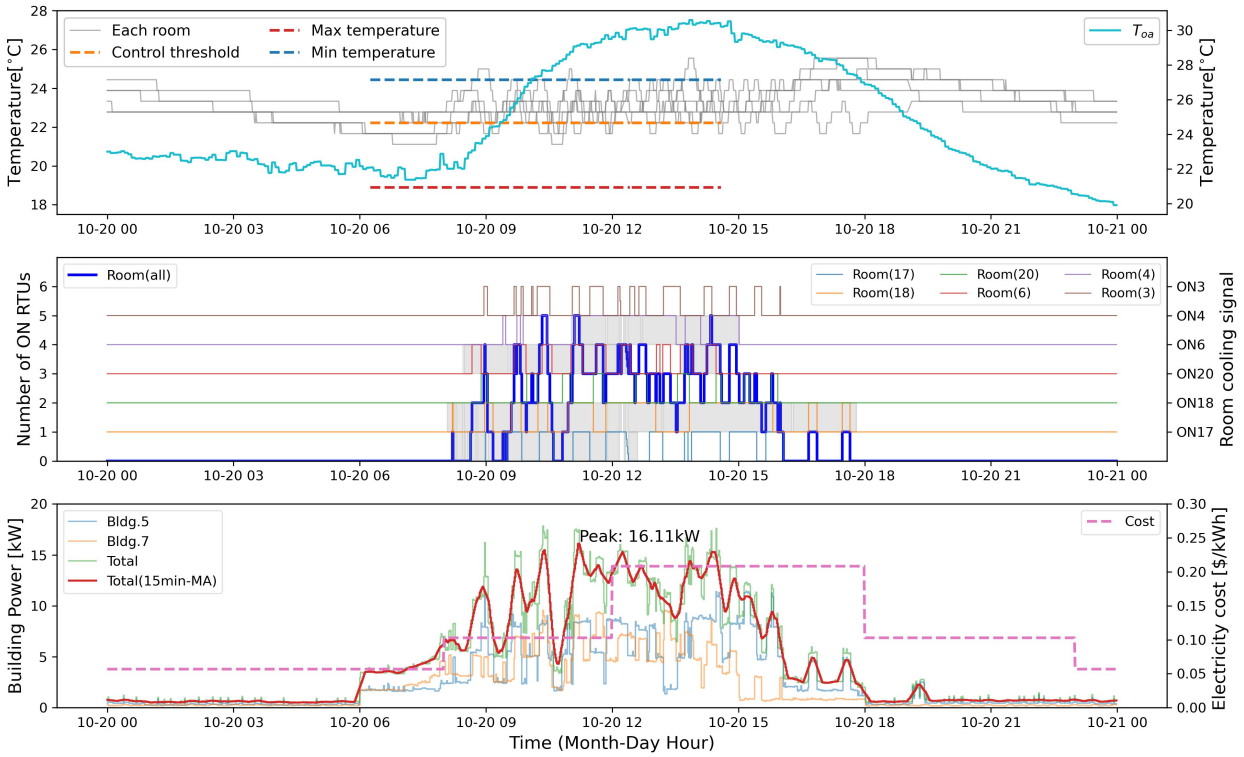

Figure 15: Summary of RTU operation in a peak cooling day (MPC); top: room temperatures with min/max comfort boundary (thermostat deadband adjusted setpoints) and outdoor air temperature, middle: RTU signals and override period (grayed area), bottom: building power and electricity price signal.

5.4. Temperature profiles

It is critical to make a comparison of the indoor temperature distributions between Baseline and MPC days (Figure 16) 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 that Baseline had lower and tighter temperature distributions during the peak cooling times (i.e., 12:00-16:00). The reason is MPC tried to maintain higher temperature within the comfort boundaries to reduce unnecessary cooling operations.

5.5. Load shifting potential

Figure 17 shows the comparison between the load profiles for a Baseline and MPC day with the utility tariff to show the load shifting potential. A 1-hour moving average was applied for a clarification purpose. 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, resulting in high cooling power demand. On the contrary, the MPC used more power for cooling before noon to avoid the high power demand during the high cost time. As a result, it showed smoother and less power use during high cost time and reduced 16% of load during 12:00-15:00.

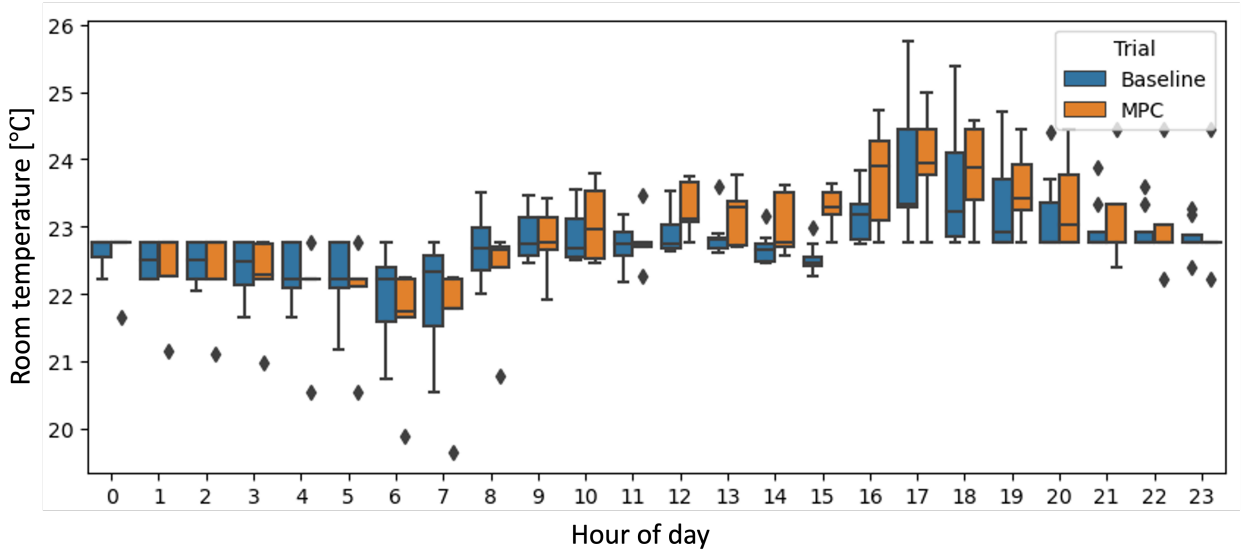

Figure 16: Comparison of the indoor temperature profiles between Baseline and MPC days in a classroom.

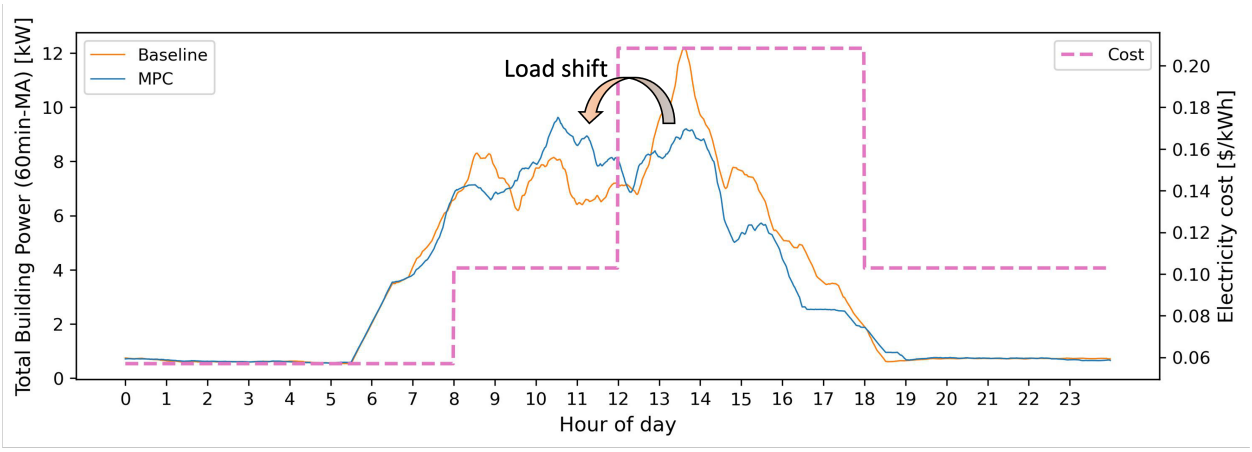

Figure 17: Comparison between the load profiles in a Baseline and an MPC day with the utility tarrif

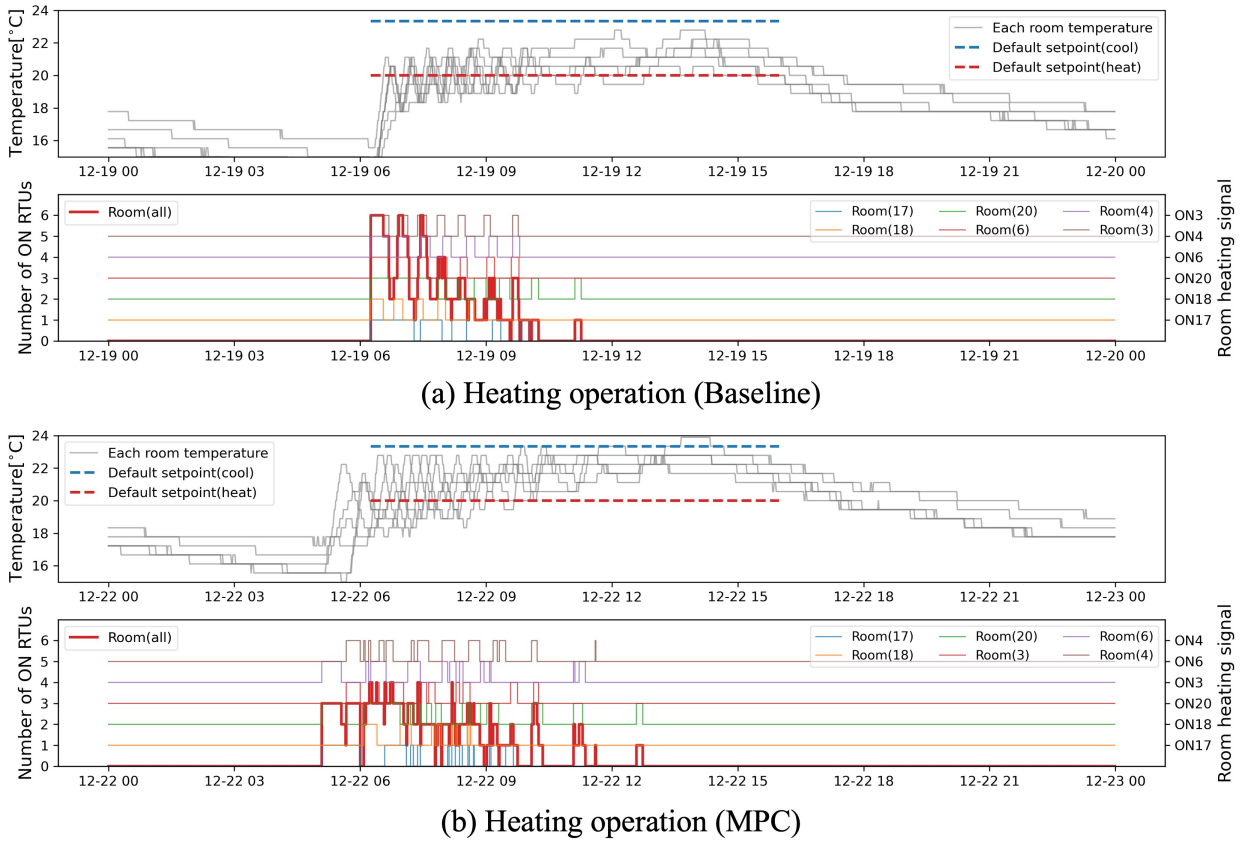

Figure 18: Comparison of heating operations between Baseline and MPC on a heating day.

5.6. Peak demand reduction potential during the heating season

In this school, the RTUs have gas heating, so there are no benefits of MPC for grid services. However, due to increasing momentum and pressure towards the electrification for building decarbonizations, it is of interest to test the MPC assuming that the RTUs are heat pumps. Figure 18 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 increases. 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 of the Baseline.

6. Conclusions

In this final report, 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 , but very few advanced controls for load flexibility has 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 potentials of providing grid services. The proposed MPC is a hierarchical MPC where the upper-level MPC is responsible for load shifting and the lower-level MPC is responsible for eliminating unnecessary peak and/or curtailing power. The hierarchical MPC has been modified to take accounts for practical challenges of user's override actions and poor data resolutions. The MPC was deployed to six pilot classrooms during a cooling season, and the results are compared with baseline data. The MPC reduced 24% of total peak power and 30% of HVAC peak power, and shifted 16% of cooling load for an on-peak price period to a low price period. This MPC solution will enable the realization of a low-cost and scalable MPC solution for applications in SMCBs.

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. None, 1784302, ark:/13030/qt78k303s5 (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] 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).

[6] 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).

[7] 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).

[8] 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, Annu. Rev. Control 50 (2020) 190–232.

[9] 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 Build. 107 (2015) 279–293.

[10] R. De Coninck, L. Helsen, Practical implementation and evaluation of model predictive control for an office building in brussels, Energy Build. 111 (2016) 290–298.

[11] 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 Build. 178 (2018) 49–60.

[12] S. Freund, G. Schmitz, Implementation of model predictive control in a large-sized, low-energy office building, Build. Environ. 197 (2021) 107830.

[13] J. Drgoňa, D. Picard, L. Helsen, Cloud-based implementation of white-box model predictive control for a GEOTABS office building: A field test demonstration, J. Process Control 88 (2020) 63–77.

[14] 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 Build. 206 (2020) 109563.

[15] D. Blum, Z. Wang, C. Weyandt, D. Kim, M. Wetter, T. Hong, others, Field demonstration and implementation analysis of model predictive control in an office HVAC system, Appl. Energy (2022).

[16] 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.

[17] D. Kim, Z. Wang, J. Brugger, D. Blum, M. Wetter, T. Hong, others, Site demonstration and performance evaluation of MPC for a large chiller plant with TES for renewable energy integration and grid decarbonization, Appl. Energy (2022).

[18] 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.

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

[20] 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).

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

[22] K. Buchanan, R. Russo, B. Anderson, The question of energy reduction: The problem(s) with feedback, Energy Policy 77 (2015) 89–96.

[23] 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 Build. (2022).

[24] 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, Build. Environ. 107 (2016) 169–180.

[25] 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 Build. 167 (2018) 359–369.

[26] 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.

[27] J. Braun, N. Chaturvedi, An inverse Gray-Box model for transient building load prediction, HVAC&R Res. 8 (1) (2002) 73–99.

[28] 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, Nat. Methods 17 (3) (2020) 261–272.

[29] S. Rouchier, M. J. Jiménez, S. Castaño, Sequential monte carlo for on-line parameter estimation of a lumped building energy model, Energy Build. 187 (2019) 86–94.

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

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

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

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

[34] W. F. Holmgren, C. W. Hansen, M. A. Mikofski, Pvlib python: A python package for modeling solar energy systems, J. Open Source Softw. 3 (29) (2018) 884.

[35] R. N. Edmondson, Multi-level block designs for comparative experiments, JABES 25 (4) (2020) 500–522.