Honda Research Institute EU
Honda Research Institute Europe GmbH
https://www.honda-ri.de/

Hierarchical MPC for building energy management: Incorporating data-driven error compensation and mitigating information asymmetry

Jens Engel, Thomas Schmitt, Tobias Rodemann, Jürgen Adamy

2024

Preprint:

Copyright Elsevier 2024. This manuscript version is made available under the CC-BY-NC-ND 4.0 license http://creativecommons.org/licenses/by-nc-nd/4.0/

Applied Energy 372 (2024) 123780



Contents lists available at ScienceDirect

Applied Energy

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

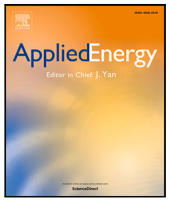

Hierarchical MPC for building energy management: Incorporating data-driven error compensation and mitigating information asymmetry

Jens Engel [a,b,*], Thomas Schmitt [a], Tobias Rodemann [a], Jürgen Adamy [b]

[a] Honda Research Institute Europe GmbH, Carl-Legien-Strasse 30, D-63073, Offenbach/Main, Germany
[b] Control Methods and Intelligent Systems Laboratory, Technical University of Darmstadt, Landgraf-Georg-Strasse 4, D-64283, Darmstadt, Germany

ARTICLE INFO

Keywords:
Energy management system
Model predictive control
Data-driven residual estimator
Disturbance prediction
Digital twin
Information asymmetry

ABSTRACT

The increasing adoption of renewable energy sources (RESs) in public power grids has led to a demand for more intelligent energy management systems (EMSs) in large-scale buildings. A common approach for controlling EMSs for buildings is Model Predictive Control (MPC). For large-scale buildings, hierarchical MPC schemes have been proposed, offering the advantage of scalability through problem decomposition into multiple layers. However, hierarchical schemes often suffer from information mismatch due to information asymmetry between layers, leading to suboptimal control performance. This issue is worsened by model errors inherent to the models underlying the MPC controllers. To address these challenges, we propose a hierarchical MPC approach, which includes data-driven error compensation. Additionally, to mitigate information mismatch, a one-iteration communication step is introduced between the hierarchical layers. The proposed approach comprises two layers: an aggregator layer that controls overall energy flows of the building, and a distributor layer that allocates thermal energy to individual temperature zones. The distributor may request additional thermal budget by providing the aggregator with an otherwise expected performance loss, which it can trade off accordingly. The approach is evaluated in a software-in-the-loop (SiL) simulation using a physics-based digital twin model of a multi-zone commercial building, showing notable improvements in overall control performance in comparison to a naive hierarchical baseline and similar performance to a monolithic baseline.

1. Introduction

The increasing influx of renewable energy sources (RESs) into the public power grid leads to an increasing demand for intelligent energy management systems (EMSs) for buildings. A promising control method for such EMSs is Model Predictive Control (MPC) [1,2]. For EMSs of large-scale buildings such as company buildings or factories, *hierarchical* MPC has attracted interest in the past years. In contrast to normal *monolithic* MPC, the overall control problem is decomposed into multiple layers in a hierarchical structure. A central upper layer acts as an *aggregator* that determines overall reference trajectories or budgets from a high-level perspective. The lower layer(s) act as *distributor(s)* that track or distribute the high-level plan from a more detailed low-level perspective. One advantage of the decomposition into a hierarchy is better scalability of the underlying optimization problems, e. g. if the building consists of many temperature zones [1,3]. Another is that different layers may run at different time scales, allowing for a decoupling of slow and fast dynamical processes or decisions [1,4]. Most often for building EMSs, a hierarchical structure is used to decompose the control of heating, ventilation, and air conditioning (HVAC) systems, but there are also other applications, e. g. the integration of electric vehicles (EVs) [5,6].

Various diverse approaches for hierarchical HVAC control schemes exist in literature. In some approaches, the aggregator only determines optimal temperature setpoints for a distributor, which can then freely decide on the actual thermal energy requirements [7,8]. In other approaches, the aggregator instead provides thermal budgets [9] and reference temperatures [3,10–12]. The decomposition of the system into different layers takes also various shapes: In [10], the authors propose a hierarchical control scheme for a 13-floor university building, consisting of 3 control layers: A “microgrid level”, managing the battery storage system, a “central HVAC level”, managing the central HVAC system, and a “zone level”, where individual temperature zones are controlled. Adjacent layers communicate to propose energy budgets to jointly achieve thermal comfort and price-optimal operation. In [3], a high-level controller aggregates the connected HVAC system of multiple buildings and determines temperature setpoints as well as an overall

* Corresponding author at: Honda Research Institute Europe GmbH, Carl-Legien-Strasse 30, D-63073, Offenbach/Main, Germany.
E-mail addresses: jens.engel@honda-ri.de (J. Engel), thomas.schmitt@honda-ri.de (T. Schmitt), tobias.rodemann@honda-ri.de (T. Rodemann), juergen.adamy@tu-darmstadt.de (J. Adamy).

https://doi.org/10.1016/j.apenergy.2024.123780
Received 4 March 2024; Received in revised form 24 May 2024; Accepted 21 June 2024
0306-2619/© 2024 The Authors. Published by Elsevier Ltd. This is an open access article under the CC BY-NC-ND license (http://creativecommons.org/licenses/by-nc-nd/4.0/).

cooling budget. Multiple low-level controllers, one for each building, track these setpoints and distribute the cooling budget. In [11], a high-level controller determines temperature setpoints and airflow budgets for virtual "meta zones", representing aggregates of individual zones, which are then tracked and distributed among individual zones by a lower-level controller. In [12], an aggregator determines hourly energy budgets for the whole building, while a distributor controls the individual temperature zones in near-real time while observing the allocated energy budget.

An inherent trait of hierarchical control approaches is the presence of information asymmetry. The process of aggregating information from the distributor level to the aggregator level introduces a loss of information, as differently composed information is available at different layers [4]. In the case of a building EMSs, where an aggregator combines multiple separate thermal zones into one larger aggregated zone, local details of the actual building, e. g. the inhomogeneous distribution of thermal zones or of thermal actuation capabilities, are overlooked. Hence, the aggregator may assign energy budgets that may be insufficient to achieve the anticipated control performance. To alleviate this information asymmetry, additional communication between layers can be introduced [4,10]. The additional communication overhead should be kept low, as to not outweigh the benefit of decomposing the control problem. In the context of HVAC control schemes, such communication usually involves upwards communication of energy budget requests by the distributors [6,12,13], or iterative consensus schemes [10,14,15].

Another challenge in MPC-based EMS control approaches is the availability of an accurate model. Multiple methodologies for constructing such a model exist, usually categorized as white-box modeling, gray-box modeling, and black-box modeling. White-box models, mostly developed using building energy performance simulation tools such as EnergyPlus or TRNSYS, can achieve high accuracy but are generally too intricate for direct implementation in the MPC's optimal control problem (OCP). Gray-box models, such as state space or Resistor–Capacitor (RC) models, exhibit lower accuracy but can be effectively employed in an OCP [16]. Both white- and gray-box modeling of buildings are highly complex and require building-specific expert knowledge, i. e. models cannot be easily transferred to other buildings [1,16,17]. Thus, data-driven black-box modeling has garnered increased interest [18–20], e.g., using Gaussian Processes (GPs), artificial neural networks (ANNs) or reinforcement learning or notably using physically consistent neural networks (PCNNs) [21,22]. PCNNs are hybrid models consisting of an affine part and an ANN, where consistency constraints on the affine part ensure physicality. While the primary advantage of these black-box models lies in the relatively low modeling effort, they demand a substantial amount of data, which poses challenges in data acquisition [19,20]. Simultaneously, the main strengths of white/gray-box modeling, namely the ability to incorporate specific domain knowledge, interpretability, and ease of adaptability of the resulting model, are lost. Therefore, a hybrid approach involving these modeling paradigms could be a promising way to enable employing building EMSs on a larger scale in the real world [20].

One such hybrid approach is to augment a gray-box model using a *residual estimator*. Instead of replacing an approximate gray-box model fully with a data-driven model or a more complex, and thus more expensive model, its model error is reduced by estimating the residual value using an additional data-driven regression model. In [23], the authors include long short-term memory (LSTM) and GP-based black-box residual estimators into an MPC for the control of a coupled tank system. The estimators can predict dynamical residual errors and are therefore integrated into the OCP. Since the estimators are highly nonlinear, they are (piece-wise) linearized at each iteration to simplify the OCP. In [24], GPs are used to learn the model error for an autonomous racing car. Training data is generated through simulation without the use of an MPC. The GPs are explicitly used in the OCP as part of the model dynamics without linearization. A similar approach is used in [25], where a physics-based model of a motorcycle is augmented using a GP residual estimator for the use in an MPC. The authors of [26] employ a similar approach to augment the model of a single-zone building using a GP-based estimator. While they show an enhanced prediction accuracy of the resulting model, they do not apply it in MPC.

Another common approach in building EMSs is to use residual estimators to directly predict heat disturbances acting on the building, caused through occupant behavior and ambient effects. These represent one of the largest error sources in buildings [1,27]. Integrating these estimators into the model of the MPC of the EMS allows for explicit error compensation. In [28], the authors predict thermal disturbances using linear and nonlinear regression models, based on cloud cover, solar irradiation and ambient temperature. In [29,30], the authors train a feed-forward ANN to estimate the heat disturbance in a single-zone building based on time information (and ambient temperature [29]). The authors of [30] also predict the temperature of the thermal mass of the zone envelope from past zone temperature measurements using an additional ANN. In [31], the authors use an XGBoost regressor [32] to estimate the model error as a time-variant heat disturbance for a multi-zone building, based on time information, ambient temperature and building load. A similar approach is proposed in [33], where the time-variant heat disturbance for a multi-zone building is predicted using an ANN, based on time information, irradiation, load and occupancy data.

The introduction of such explicit error compensation can greatly reduce the model error. Yet, it may also introduce another source of information asymmetry between layers: If heating and cooling influences appear simultaneously across different temperature zones, these influences cancel each other in the aggregated perspective. Thus, they are invisible to the aggregator, which therefore may derive wrong thermal budgets or suboptimal setpoints. In previous work, we have shown that this can have a significant influence on overall control performance [31].

In this work, we propose a hierarchical MPC approach for the control of a multi-zone building, which addresses these joint issues of information asymmetry and error compensation. The system is composed of two layers: An aggregator, which controls the overall energy flows considering aggregated thermal zones, and a distributor which allocates thermal energy to the individual zones of the building. The main contributions of our approach are:

- The inclusion of data-driven error compensation through heat influence prediction into both layers of the hierarchy.
- An aggregator–distributor scheme with one-iteration communication, in which the distributor can request additional budget while simultaneously providing the aggregator with an expected performance loss (EPL) associated with not granting the requested budget, based on which the aggregator can trade off the budget request.
- An estimation scheme of the proposed EPL based on predicted heat influences to avoid solving an additional optimization problem.
- A comparison of the proposed hierarchical approach against a monolithic baseline approach to show the efficacy of the proposed approach.
- The verification of the proposed approach in a digital twin simulation of a multi-zone commercial building based on real historical measurement data.

This paper is organized as follows: The building itself, the digital twin and the simplified gray-box models are described in Section 2. The hierarchical control approach with inter-layer communication, as well as a monolithic baseline approach is described in Section 3. The data-driven error compensation approach is explained in Section 4. The efficacy of the proposed approach is shown by digital twin simulation results in Section 5. Finally, a conclusion and outlook for future work is given in Section 6.

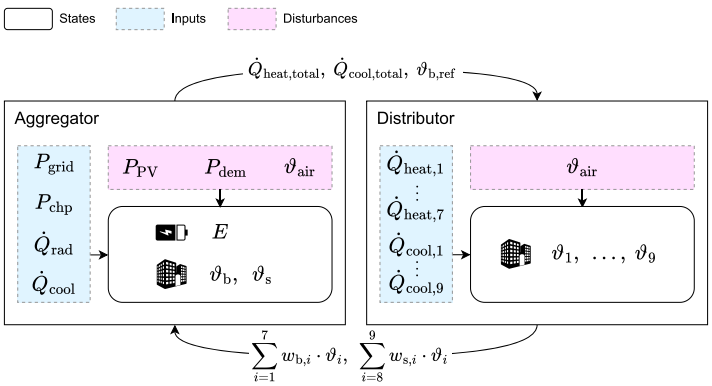

Fig. 1. Structure of the proposed hierarchical control approach. The aggregator plans operation of electrical and aggregated thermal components, the distributor distributes the thermal budgets between individual temperature zones.

2. Building models

The building considered in this work is a medium-sized company building situated in Offenbach, Germany. Its footprint of approx. 13, 000 m^2 is partitioned into 9 distinct temperature zones, including offices, halls, workshops and an emissions lab. Apart from a connection to the public power grid, a gas-fired combined heat and power plant (CHP) for co-production of electricity and heat with 199 kW_{el}, and a photovoltaic (PV) plant with 750 kWp serve as energy sources which supply an average load demand of approx. 250 kW. The building's HVAC system is supplied with heat from both the CHP and gas-fired heating boilers. Cooling power is supplied by a central cooling system consisting of three water-chillers. Ventilation and air conditioning is provided through various electrical ventilation systems. Apart from that, a second-life battery with 98 kWh capacity is available as electrical energy storage [34].

The model used in the control approach is split into two layers: An *aggregator*, representing the overall, aggregated perspective of the entire building, and a *distributor*, controlling the thermal behavior of the 9 individual temperature zones. This structure is illustrated in Fig. 1. In the following, the models of these two layers are described. Furthermore, a brief description of the digital twin model used for evaluation purposes is given.

2.1. Aggregator model

The aggregator represents the overall building, including both its thermal and electrical components. We model the aggregator using a simplified state space model representing the most important entities. Of the total 9 temperature zones, 7 are aggregated as a single 'building zone' with the state ϑ_{b} representing the average temperature of the zones, calculated as a weighted average by means of relative thermal capacity. The remaining 2 zones refer to server rooms and are aggregated with an average 'server zone' temperature ϑ_{s}. Together with the stationary battery's stored energy E in kWh, these temperatures form the state vector x_{agg}. As input u_{agg} to the system, we consider the grid power P_{grid}, the (electrical) CHP power P_{chp}, the gas heating power $\dot{Q}_{\text{rad}}$, and the HVAC cooling power $\dot{Q}_{\text{cool}}$.

The PV power P_{PV}, the building's electrical load demand P_{dem}, the ambient air temperature ϑ_{air}, (constant) losses to the ground $\dot{Q}_{\text{other,b}}$, and (constant) internal heatings $\dot{Q}_{\text{other,s}}$ are considered as disturbances d_{agg}. The unit of all electrical and thermal powers is kW, and that of all temperatures is °C. With this, the time-continuous state space description of the model is given by

$$\begin{bmatrix} \dot{E}(t) \\ \dot{\vartheta}_{\text{b}}(t) \\ \dot{\vartheta}_{\text{s}}(t) \end{bmatrix} = \begin{bmatrix} 0 & 0 & 0 \\ 0 & -\frac{H_{\text{air,b}}-\beta_{\text{bs}}}{C_{\text{th,b}}} & \frac{\beta_{\text{bs}}}{C_{\text{th,b}}} \\ 0 & \frac{\beta_{\text{bs}}}{C_{\text{th,s}}} & -\frac{H_{\text{air,s}}-\beta_{\text{bs}}}{C_{\text{th,s}}} \end{bmatrix} \begin{bmatrix} E(t) \\ \vartheta_{\text{b}}(t) \\ \vartheta_{\text{s}}(t) \end{bmatrix}$$

Table 1
Numerical values of the building parameters. Note that $C_{\text{th,b}} = \sum_{i=1}^{7} C_{\text{th},i}$, $C_{\text{th,s}} = C_{\text{th,8}} + C_{\text{th,9}}$, $H_{\text{air,b}} = \sum_{i=1}^{7} H_{\text{air},i}$, $H_{\text{air,s}} = H_{\text{air,8}} + H_{\text{air,9}}$, $\beta_{\text{bs}} = \beta_{29} + \beta_{58} + \beta_{68}$, $\beta_{ij} = \beta_{ji}$, and all other β_{ij} not listed below are zero, e. g. $\beta_{12} = 0$.

	in kWh/K		in kW/K		in W/K
$C_{\text{th,1}}$	289.40	$H_{\text{air,1}}$	3.69	β_{29}	48.40
$C_{\text{th,2}}$	595.01	$H_{\text{air,2}}$	9.82	β_{34}	345.60
$C_{\text{th,3}}$	268.58	$H_{\text{air,3}}$	3.65	β_{56}	1100.48
$C_{\text{th,4}}$	129.96	$H_{\text{air,4}}$	2.79	β_{58}	23.40
$C_{\text{th,5}}$	413.82	$H_{\text{air,5}}$	4.79	β_{68}	8.00
$C_{\text{th,6}}$	413.82	$H_{\text{air,6}}$	6.19		
$C_{\text{th,7}}$	124.67	$H_{\text{air,7}}$	3.19		
$C_{\text{th,8}}$	3.01	$H_{\text{air,8}}$	0.03		
$C_{\text{th,9}}$	6.02	$H_{\text{air,9}}$	0.04		

$$+ \begin{bmatrix} 1 & 1 & 0 & \frac{1}{\epsilon_{\text{c}}} & \frac{1}{\epsilon_{\text{c}}} \\ 0 & \frac{1}{c_{\text{cur}} \cdot C_{\text{th,b}}} & \frac{1}{C_{\text{th,b}}} & \frac{1}{C_{\text{th,b}}} & 0 \\ 0 & 0 & 0 & 0 & \frac{1}{C_{\text{th,s}}} \end{bmatrix} \begin{bmatrix} P_{\text{grid}}(t) \\ P_{\text{chp}}(t) \\ \dot{Q}_{\text{rad}}(t) \\ \dot{Q}_{\text{cool,b}}(t) \\ \dot{Q}_{\text{cool,s}}(t) \end{bmatrix}$$

$$+ \begin{bmatrix} 1 & 1 & 0 & 0 & 0 \\ 0 & 0 & \frac{H_{\text{air,b}}}{C_{\text{th,b}}} & \frac{1}{C_{\text{th,b}}} & 0 \\ 0 & 0 & \frac{H_{\text{air,s}}}{C_{\text{th,s}}} & 0 & \frac{1}{C_{\text{th,s}}} \end{bmatrix} \begin{bmatrix} P_{\text{PV}}(t) \\ P_{\text{dem}}(t) \\ \vartheta_{\text{air}}(t) \\ \dot{Q}_{\text{other,b}}(t) \\ \dot{Q}_{\text{other,s}}(t) \end{bmatrix}, \tag{1}$$

where $C_{\text{th,b}}$ and $C_{\text{th,s}}$ are the thermal capacities of the building and server zone, respectively, in $\frac{\text{kWh}}{\text{K}}$, $H_{\text{air,b}}$ and $H_{\text{air,s}}$ are the heat transfer coefficients to the ambient air of the building and the server zone, respectively, in $\frac{\text{kW}}{\text{K}}$, β_{bs} is the heat transfer coefficient between the two zones in $\frac{\text{kW}}{\text{K}}$. Numerical values are given in Table 1. c_{cur} is the ratio of the CHP's electrical to thermal power and ϵ_{c} is the energy efficiency ratio of the cooling machines. Their values are given in Table 3. The values assumed for the constant heat losses and gains $\dot{Q}_{\text{other,b}}$ and $\dot{Q}_{\text{other,s}}$ are given in Table 2. The losses to the ground making up $\dot{Q}_{\text{other,b}}$ were derived by calculating the expected heat flow through the ground covering faces, assuming a temperature difference of 10 K between the zones and the ground, which approximates the yearly average temperature difference. The internal heat gains $\dot{Q}_{\text{other,s}}$ were derived from historical average electrical power demand in the server zones, to approximate the idle load of the server zones.

Discretizing the time-continuous state-space model yields

$$\begin{aligned} x_{\text{agg}}(k+1) = {} & A_{\text{agg}}(T_{\text{s}})x_{\text{agg}}(k) + B_{\text{agg}}(T_{\text{s}})u_{\text{agg}}(k) \\ & + S_{\text{agg}}(T_{\text{s}})d_{\text{agg}}(k) + \epsilon_{\text{agg}}(k) \end{aligned} \tag{2}$$

with $\epsilon_{\text{agg}}(k) = \begin{bmatrix} 0 & \epsilon_{\text{b}}(k) & \epsilon_{\text{s}}(k) \end{bmatrix}^{\intercal}$ being the residual model error of the gray-box model, as initially motivated. We consider the residual

Table 2
Static heat losses to the ground (Zone 1-7) and internal heat gains (Zones 8 and 9) $\dot{Q}_{\text{other},i}$ assumed to act on the individual temperature zones. Note that $\dot{Q}_{\text{other,b}} = \sum_{i=1}^{7} \dot{Q}_{\text{other},i}$ and $\dot{Q}_{\text{other,s}} = \sum_{i=8}^{9} \dot{Q}_{\text{other},i}$.

	kW		kW		kW
$\dot{Q}_{\text{other},1}$	−10.58	$\dot{Q}_{\text{other},4}$	0	$\dot{Q}_{\text{other},7}$	−4.56
$\dot{Q}_{\text{other},2}$	−13.10	$\dot{Q}_{\text{other},5}$	−16.64	$\dot{Q}_{\text{other},8}$	16.87
$\dot{Q}_{\text{other},3}$	−10.80	$\dot{Q}_{\text{other},6}$	0	$\dot{Q}_{\text{other},9}$	51.71
$\dot{Q}_{\text{other,b}}$	−58.58	$\dot{Q}_{\text{other,s}}$	68.58		

model error ϵ_{agg} as an additive, time-variant disturbance, as this is a good approximation of the main sources of error in buildings, namely unknown internal heat flows caused by occupant behavior, which are independent of state and inputs, as well as heat flows to and from the environment, which are mainly dependent on external influences (e. g. ambient temperature, solar irradiation), assuming the building stays close to a certain temperature setpoint. Note that we can respect the model error $\epsilon_{\text{agg}}(k)$ only in the discretized form since it is estimated from discretely sampled data points. For more details on the modeling itself, the reader is referred to [35]. The estimation of $\epsilon_{\text{agg}}(k)$ is described and discussed in Section 4.

2.2. Distributor model

The distributor represents the 9 individual temperature zones from a thermal perspective, neglecting the electrical part of the building, which is covered by the aggregator. The temperature ϑ_i of zone i is described by

$$\begin{aligned}\dot{\vartheta}_i(t) = &\frac{1}{C_{\text{th},i}}\left(\dot{Q}_{\text{heat},i}(t) + \dot{Q}_{\text{cool},i}(t) + \dot{Q}_{\text{other},i}(t)\right)\\ &- \sum_{j\neq i}\frac{\beta_{ij}}{C_{\text{th},i}}\left(\vartheta_i(t) - \vartheta_j(t)\right)\\ &- \frac{H_{\text{air},i}}{C_{\text{th},i}}\left(\vartheta_i(t) - \vartheta_{\text{air}}(t)\right)\end{aligned} \tag{3}$$

with $C_{\text{th},i}$ being the thermal capacity of zone i in $\frac{\text{kWh}}{\text{K}}$, β_{ij} the heat transfer coefficient between zones i and j in $\frac{\text{kW}}{\text{K}}$, $H_{\text{air},i}$ the heat transfer coefficient between zone i and the outside air in $\frac{\text{kW}}{\text{K}}$, and $\dot{Q}_{\text{heat},i}$ and $\dot{Q}_{\text{cool},i}$ in kW the heating and cooling powers allocated to zone i. Note that $\dot{Q}_{\text{heat},i}(k) = 0\ \forall i = 8, 9$, as the server zones are not actively heated. Analogous to the aggregator, $\dot{Q}_{\text{other},i}$ are uncontrollable disturbances, which are assumed constant and represent heat losses to the ground (for the building zones 1–7) or internal heat gains (for the server zones 8 and 9). The values we assume for these heat flows are listed in Table 2. Analogously to the aggregator, we again consider a residual model error $\epsilon_i(k)$ in each temperature zone.

With the 9 zone temperatures ϑ_i making up the state vector x_{dis}, $\dot{Q}_{\text{heat},i}$ and $\dot{Q}_{\text{cool},i}$ the input vector u_{dis}, the ambient temperature ϑ_{air} and $\dot{Q}_{\text{other},i}$ the disturbance vector d_{dis}, Eq. (3) can be expressed in the form of a continuous state space model analogous to Eq. (1). Discretizing the model with a sample time T_{s} yields the discrete state-space model of the distributor,

$$\begin{aligned}x_{\text{dis}}(k+1) = &\ A_{\text{dis}}(T_{\text{s}})x_{\text{dis}}(k) + B_{\text{dis}}(T_{\text{s}})u_{\text{dis}}(k)\\ &+ S_{\text{dis}}(T_{\text{s}})d_{\text{dis}}(k) + \epsilon_{\text{dis}}(k),\end{aligned} \tag{4}$$

with residual model errors $\epsilon_{\text{dis}}(k) = \begin{bmatrix}\epsilon_1(k) & \dots & \epsilon_9(k)\end{bmatrix}^{\intercal}$, where A_{dis} is the system matrix, B_{dis} the input matrix and S_{dis} the disturbance matrix. Numerical values are given in Table 1. For brevity, the reader is referred to [35, Chapter 2] for more details on the state space model.

2.3. Digital twin

As a surrogate of the real facility, we employ a Modelica-based digital twin model implemented in SimulationX [36,37]. This digital twin is a physical white-box model, encompassing 9 different temperature zones. It covers their connections and the heat losses to both the ambient air and the ground. Moreover, it takes into account internal heat gains from electrical usage and building occupancy. Additionally, it features a detailed HVAC system. This model deviates from the state-space model, presented earlier, in its ability to capture the nonlinear dynamics and efficiencies of the zones, HVAC, and CHP systems. For example, the digital twin model incorporates the CHP's minimum power output of 50 percent and the HVAC's efficiency that varies with ambient temperature. The heat capacities of the zones are identical between the digital twin and the previously presented state space models. The digital twin model has a more detailed representation of the boundaries of each thermal zone, the values for the heat transfer coefficients $H_{\text{air},i}$ and β_{ij} in Table 1 were derived as total values covering all these respective boundary components. The model uses historic measurement data for ambient temperature, solar irradiation and PV power, as well as the electrical power demand (in each zone).

3. Control approach

In this section, we present our proposed two-layer hierarchical control approach. For each of the two layers, the aggregator and distributor, a control problem is formulated. The control problem of the aggregator decides on aggregated electrical and thermal operation of the building. The distributor's problem solves the optimal distribution of the aggregated thermal budgets. Due to the aggregation of the thermal part of the building in the aggregator, information asymmetry is present in the hierarchy. The aggregator is only partially aware of heat disturbances acting on individual zones, leading to suboptimal thermal budgets being allocated. To solve this, we introduce a so called *negotiation* between the layers, where the distributor can propose an adjusted budget. In the following, we first describe the individual control problems. Then, we explain how they interact through the negotiation. Finally, we describe an equivalent monolithic controller, against which we benchmark the proposed hierarchical approach in Section 5. All parameters of the equations presented in this section are listed in Table 3.

3.1. Aggregator

The aggregator plans the operation of the overall building. The goal is to regulate the building temperatures while minimizing the monetary costs. This is expressed as a weighted sum of multiple cost functions.

Objectives. The building temperatures are controlled by two cost functions. The first one keeps the temperature of the building zone comfortable by minimizing the deviation from 22 °C, i. e.

$$J_{\text{comf,agg}}(k) = \sum_{n=1}^{N_{\text{p}}}\left(\vartheta_{\text{b}}(n|k) - 22\,^{\circ}\text{C}\right)^2. \tag{5}$$

The notation $a(n|k)$ refers to the value of $a(k+n)$ as predicted at time step k, and will be used from here on. Secondly, the server zone temperature is to be kept within a safe range by

$$\begin{aligned}J_{\text{s,agg}}(k) = &\sum_{n=1}^{N_{\text{p}}}\max\left(15\,^{\circ}\text{C} - \vartheta_{\text{s}}(n|k),\ 0\right)\\ &+ \max\left(\vartheta_{\text{s}}(n|k) - 21\,^{\circ}\text{C},\ 0\right).\end{aligned} \tag{6}$$

The monetary costs consist of the costs for gas used in the CHP and for heating, i. e.

$$\ell_{\text{mon,chp}}(k) = c_{\text{gas}} \cdot \frac{1}{\eta_{\text{chp}}}\left(1 + \frac{1}{c_{\text{cur}}}\right) \cdot T_{\text{s}} \cdot P_{\text{chp}}(k), \tag{7}$$

$$\ell_{\text{mon,heat}}(k) = c_{\text{gas}} \cdot \frac{1}{\eta_{\text{rad}}} \cdot T_{\text{s}} \cdot \dot{Q}_{\text{rad}}(k), \tag{8}$$

as well as the electricity costs for buying or selling energy to/from the public grid,

$$\ell_{\text{mon,grid}}(k) = \left(c_{\text{buy}}^{\text{grid}} \cdot \max\left(0, P_{\text{grid}}(k)\right) - c_{\text{sell}}^{\text{grid}} \cdot \max\left(0, -P_{\text{grid}}(k)\right) \right) \cdot T_{\text{s}}, \tag{9}$$

which are summed to the stage costs

$$\ell_{\text{mon}}(k) = \ell_{\text{mon,grid}}(k) + \ell_{\text{mon,chp}}(k) + \ell_{\text{mon,heat}}(k). \tag{10}$$

Furthermore, we have to consider costs on the peak power demand of the calendar year through

$$J_{\text{peak}}(k) = c_{\text{peak}}^{\text{grid}} \cdot \max(0, \max_{n} \left(P_{\text{grid}}(n|k) - P_{\text{grid,peak}}(k)\right)), \tag{11}$$

where $P_{\text{grid,peak}}(k)$ is the peak incurred until time step k. The total monetary cost function is then given by

$$J_{\text{mon}}(k) = J_{\text{peak}}(k) + \sum_{n=0}^{N_{\text{p}}-1} \ell_{\text{mon}}(n|k). \tag{12}$$

Note that while the max terms in the previous equations would in general lead to a nonlinear optimization problem, it is possible to reformulate them into linear expressions using epigraph reformulations [38].

Constraints. The input constraints of the aggregator are given by

$$-P_{\text{grid,max}} \le P_{\text{grid}}(k) \le P_{\text{grid,max}}, \tag{13a}$$
$$0 \le P_{\text{chp}}(k) \le P_{\text{chp,max}}, \tag{13b}$$
$$0 \le \dot{Q}_{\text{rad}}(k) \le \dot{Q}_{\text{rad,max}}, \tag{13c}$$
$$\dot{Q}_{\text{cool,b,max}} \le \dot{Q}_{\text{cool,b}}(k) \le 0, \tag{13d}$$
$$\dot{Q}_{\text{cool,s,max}} \le \dot{Q}_{\text{cool,s}}(k) \le 0. \tag{13e}$$

The state constraints are given by

$$0.15 \cdot C_{\text{bat}} \le E(k) \le 0.85 \cdot C_{\text{bat}}, \tag{14a}$$
$$P_{\text{charge,max}} \le \frac{E(k+1) - E(k)}{T_{\text{s}}} \le P_{\text{charge,max}}. \tag{14b}$$

Optimal control problem. Together, the aggregator's OCP is described by

$$\min_{\boldsymbol{u}_{\text{agg}}} \; w_{\text{comf}} \cdot J_{\text{comf,agg}}(k) + w_{\text{s,agg}} \cdot J_{\text{s,agg}}(k) + w_{\text{mon}} \cdot J_{\text{mon}}(k), \tag{15a}$$
$$\text{s.t. } (2), (13) \; \forall n = 0, \dots, N_{\text{p}} - 1, \tag{15b}$$
$$(14) \; \forall n = 1, \dots, N_{\text{p}}, \tag{15c}$$

with $\boldsymbol{u}_{\text{agg}} = \left(u_{\text{agg}}(0|k), \dots, u_{\text{agg}}(N_{\text{p}}-1|k)\right)$ being the sequence of control inputs. We consider a prediction horizon of $N_{\text{p}} = 48$ steps of $T_{\text{s}} = 0.5\,\text{h}$ each, i. e. 1 day in total. The time step notation (k) and $(k+1)$ in (2), (13) and (14) are to be read as $(n|k)$ and $(n+1|k)$, respectively.

For the weights we chose $w_{\text{comf}} = 0.8$, $w_{\text{s,agg}} = 1$, $w_{\text{mon}} = 0.2$ as a reasonable compromise [39]. Alternatively, multi-objective optimization could be used to dynamically determine these weights [38,40], since the aggregator's OCP is always solvable quickly enough due to the hierarchization. Note that in the actual implementation, additional slack variables are necessary for the epigraph reformulations of the max-terms.

3.2. Distributor

The distributor manages the optimal distribution of the thermal budgets of the HVAC system across the individual 9 temperature zones.

Table 3
Parameters of the considered building EMS.

Parameter	Description	Value
T_{s}	Sample rate	0.5 h
N_{p}	Prediction horizon	48 ≙ 24 h
$c_{\text{peak}}^{\text{grid}}$	Peak cost factor	101 €/kW
$c_{\text{buy}}^{\text{grid}}$	Electricity cost	0.13 €/kWh
$c_{\text{sell}}^{\text{grid}}$	Feed-in Tariff	0.07 €/kWh
c_{gas}	Gas costs	0.045 €/kWh
ϵ_{c}	Energy efficiency ratio of the cooling machines	1.78
c_{cur}	Power-to-heat ratio of the CHP	0.677
η_{chp}	Efficiency of the CHP	0.892
η_{rad}	Efficiency of the gas radiator	0.97
$P_{\text{grid,max}}$	Maximum grid power	1000 kW
$P_{\text{chp,max}}$	Maximum CHP power	199 kW
$\dot{Q}_{\text{rad,max}}$	Maximum gas heating power	1500 kW
$\dot{Q}_{\text{cool,b,max}}$	Maximum building zone cooling power	−1353 kW
$\dot{Q}_{\text{cool,s,max}}$	Maximum server zone cooling power	−197 kW
C_{bat}	Second-life battery capacity	98 kWh
$P_{\text{charge,max}}$	Maximum battery charge power	32.9 kW
$P_{\text{grid,peak}}(0)$	Initial peak grid demand	474.6 kW

Objectives. The first goal of the distributor is to keep the occupant zones 1–7 at a comfortable temperature. This is ensured by minimizing the weighted deviation of the individual zone temperatures from the reference temperature $\vartheta_{\text{b}}(n|k)$ proposed by the aggregator, which is encoded by

$$J_{\text{comf,dis}}(k) = \sum_{n=1}^{N_{\text{p}}} \sum_{i=1}^{7} w_{\text{b},i} \cdot \left(\vartheta_i(n|k) - \vartheta_{\text{b}}(n|k)\right)^2, \tag{16}$$

where $w_{\text{b},i}$ are the relative weights of the temperature zones, i. e.

$$w_{\text{b},i} = \frac{C_{\text{th},i}}{\sum_{j=1}^{7} C_{\text{th},j}} \; \forall i = 1, \dots, 7. \tag{17}$$

The second goal of the distributor is, analogously to the objective in the aggregator, to keep the temperatures of server zones 8 and 9 within a safe range, i. e.

$$J_{\text{s,dis}}(k) = \sum_{n=1}^{N_{\text{p}}} \sum_{i=8}^{9} w_{\text{s},i} \left(\max\left(15\,^{\circ}\text{C} - \vartheta_i(n|k),\ 0\right) + \max\left(\vartheta_i(n|k) - 21\,^{\circ}\text{C},\ 0\right) \right), \tag{18}$$

where $w_{\text{s},i}$ are the relative weights of the server zones, i. e.

$$w_{\text{s},i} = \frac{C_{\text{th},i}}{\sum_{j=8}^{9} C_{\text{th},j}} \; \forall i = 8,\ 9. \tag{19}$$

We further consider slack variables to adjust the budget in the distributor during the negotiation, namely $z_{\text{cool,b}}$, $z_{\text{cool,s}}$, and z_{heat}. When these slack variables are active, we want to avoid that they take arbitrarily large values. Therefore, we consider an additional cost function

$$J_{\text{slack}}(k) = 10^{-4} \cdot \sum_{n=0}^{N_{\text{p}}-1} z_{\text{heat}}(n|k) - z_{\text{cool,b}}(n|k) - z_{\text{cool,s}}(n|k). \tag{20}$$

Constraints. The inputs of the distributor are subject to constraints stemming from the structure of the building's HVAC system,

$$0 \le \sum_{i=1}^{7} \dot{Q}_{\text{heat},i}(k) \le \dot{Q}_{\text{heat,max}} \tag{21a}$$
$$0 \le \dot{Q}_{\text{heat},i}(k) \le 894\,\text{kW} \quad \forall i \in \{1, \dots, 7\}, \tag{21b}$$
$$0 \le \dot{Q}_{\text{heat},i}(k) \le 0 \quad \forall i \in \{8, 9\}, \tag{21c}$$

$$\dot{Q}_{\text{cool,b,max}} \leq \sum_{i=1}^{7} \dot{Q}_{\text{cool},i}(k) \leq 0, \tag{21d}$$

$$-800\,\text{kW} \leq \dot{Q}_{\text{cool},i}(k) \leq 0 \quad \forall i \in \{1,2,3,4,7\}, \tag{21e}$$

$$-330\,\text{kW} \leq \dot{Q}_{\text{cool},i}(k) \leq 0 \quad \forall i \in \{5,6\}, \tag{21f}$$

$$-53\,\text{kW} \leq \dot{Q}_{\text{cool},8}(k) \leq 0, \tag{21g}$$

$$-144\,\text{kW} \leq \dot{Q}_{\text{cool},9}(k) \leq 0, \tag{21h}$$

where $\dot{Q}_{\text{heat,max}} = \dot{Q}_{\text{rad,max}} + {}^{1}/_{c_{\text{cur}}} P_{\text{chp,max}}$. Furthermore, the distributor shall follow the heat budget allocated by the aggregator, hence the constraints

$$\sum_{i=1}^{7} \dot{Q}_{\text{heat},i}(k) = \dot{Q}_{\text{heat}}(k) + z_{\text{heat}}(k), \tag{22a}$$

$$\sum_{i=1}^{7} \dot{Q}_{\text{cool},i}(k) = \dot{Q}_{\text{cool,b}}(k) + z_{\text{cool,b}}(k), \tag{22b}$$

$$\sum_{i=8}^{9} \dot{Q}_{\text{cool},i}(k) = \dot{Q}_{\text{cool,s}}(k) + z_{\text{cool,s}}(k) \tag{22c}$$

with

$$\dot{Q}_{\text{heat}}(k) = \dot{Q}_{\text{rad}}(k) + \frac{1}{c_{\text{cur}}} P_{\text{chp}}(k) \tag{23}$$

apply. The aforementioned slack variables can therefore be used to adjust the thermal budgets. These slack variables are constrained by

$$-M_{\text{slack}} \leq z_{\text{cool,s}}(k) \leq 0, \tag{24a}$$

$$-M_{\text{slack}} \leq z_{\text{cool,b}}(k) \leq 0, \tag{24b}$$

$$0 \leq z_{\text{heat}}(k) \leq M_{\text{slack}}. \tag{24c}$$

During the negotiation, the slack variables are activated and deactivated by varying the value of the parameter M_{slack}.

Optimal control problem. The distributor's OCP is hence described by

$$\min_{\boldsymbol{u}_{\text{dis}}} \quad J_{\text{comf,dis}}(k) + J_{\text{s,dis}}(k) + J_{\text{slack}}(k) \tag{25a}$$

$$\text{s.t.} \quad (4), (21), (22), (24)\ \forall n = 0, \dots, N_{\text{p}} - 1 \tag{25b}$$

with $\boldsymbol{u}_{\text{dis}} = \left(u_{\text{dis}}(0|k), \dots, u_{\text{dis}}(N_{\text{p}}-1|k)\right)$ being the sequence of control inputs, and the same prediction horizon as in the aggregator. Again, the time step notation (k) and $(k+1)$ in (4), (21)–(24) are to be read as $(n|k)$ and $(n+1|k)$, respectively.

3.3. Negotiation

The negotiation process has 4 steps. This is illustrated in Fig. 2. In the first step, the aggregator first proposes an initial plan. To do this, the OCP (15) is solved without any adjustments. In the second step, the distributor receives this plan and solves its OCP as defined in (25). In this step, the distributor shall be able to freely adjust the aggregator's budget $\boldsymbol{u}_{\text{agg},0}$, therefore the budget slack variables are activated by setting $M_{\text{slack}} = 2000\,\text{kW}$. Then, the distributor communicates the predicted value of the slack variables $z_{\text{cool,b}}$, $z_{\text{cool,s}}$ and z_{heat}, as well as the EPL I_{comf} caused by the influence of individual residual errors on temperature tracking performance to the aggregator. In the third step, the aggregator receives these values and solves a modified OCP, where it can trade off the requested budget against the expected influence of the individual residual errors on the zone temperature trajectories. This is motivated by the assumption that the mismatch in information that leads to insufficient budget allocation is caused directly by the residual model errors.

To perform this tradeoff, an additional decision variable $p(n|k)$ is added to the aggregator's OCP, together with the constraints

$$\dot{Q}_{\text{cool,b}}(n|k) = \dot{Q}_{\text{cool,b},0}(n|k) + p(n|k) \cdot z_{\text{cool,b}}(n|k) \tag{26a}$$

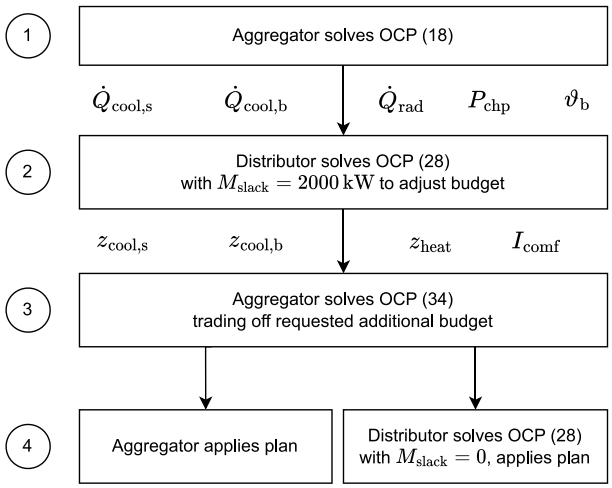

Fig. 2. Diagram of the negotiation process between the aggregator and distributor layers in the proposed approach. The negotiation allows the distributor to propose an adjusted thermal budget to account for the information mismatch between the layers.

$$\dot{Q}_{\text{cool,s}}(n|k) = \dot{Q}_{\text{cool,s},0}(n|k) + z_{\text{cool,s}}(n|k) \tag{26b}$$

$$\dot{Q}_{\text{heat}}(n|k) = \dot{Q}_{\text{heat},0}(n|k) + p(n|k) \cdot z_{\text{heat}}(n|k) \tag{26c}$$

which apply $\forall n = 0, \dots, N_{\text{p}} - 1$. The index 0 on the right hand side of the constraints refers to the trajectories planned in the first step of the negotiation. The variable $p(n|k)$ controls what fraction of the budget requested by the distributor is actually allotted, so

$$0 \leq p(n|k) \leq 1. \tag{27}$$

Note that the server zones' thermal budget request is always granted, as it is used to avoid soft-constraint violations.

Increasing the thermal budgets increases monetary costs. Therefore, this additional cost is traded off against the gain in tracking performance in the distributor. This is achieved with the compensation cost function

$$J_{\text{comp}}(k) = \sum_{n=0}^{N_{\text{p}}-1} (1 - p(n|k)) \cdot \sum_{j=n+1}^{N_{\text{p}}} I_{\text{comf}}(n,j). \tag{28}$$

Here, the EPL $I_{\text{comf}}(n,j)$ effectively approximates the weighted deviation of the individual zones from the reference trajectory $\vartheta_{\text{b},0}$ expected if the error ϵ_i is not compensated once it occurs, i. e.

$$I_{\text{comf}}(n,j) = \sum_{i=1}^{7} w_{\text{b},i} \cdot \Big((\vartheta_i(j|k) + \tilde{\epsilon}_{\text{res},i}(n,j) - \vartheta_{\text{b},0}(j|k))^2 - (\vartheta_i(j|k) - \vartheta_{\text{b},0}(j|k))^2 \Big). \tag{29}$$

Here, $\tilde{\epsilon}_{\text{res}}(n,j)$ is the residual influence of $\tilde{\epsilon}_{\text{dis}}(n|k)$ as expected at $(n+j+1|k)$, assuming that an uncompensated error decays with the passive dynamics of the building, i. e.

$$\tilde{\epsilon}_{\text{res}}(n,j) = A_{\text{dis}}^{j-n-1} \tilde{\epsilon}_{\text{dis}}(n|k). \tag{30}$$

The motivation for this compensation function is that an uncompensated error $\epsilon_i(n|k)$ will propagate along the horizon and add onto other errors until it eventually decays, hence an uncompensated error will also create costs after its initial occurrence. We approximate this by assuming that errors are independent and the temperature influences add and can be compensated linearly. This avoids multiplications of the decision variables $p(n|k)$ with themselves, which would make the OCP highly nonlinear. An alternative to estimating the EPL in this way would be to resolve the OCP without budget adjustment, and deriving the EPL from the difference in tracking performance. By estimating the EPL based on the heat influence, we avoid this additional computational effort. In summary, in this third step, the aggregator solves the

modified OCP

$$\min_{\boldsymbol{u}_{\text{agg}}} \quad w_{\text{comf}} \cdot J_{\text{comf,agg}}(k) + w_{\text{s,agg}} \cdot J_{\text{s,agg}}(k) + w_{\text{mon}} \cdot J_{\text{mon}}(k) + w_{\text{comp}} \cdot J_{\text{comp}}(k) \tag{31a}$$

$$\text{s.t.} \quad (2), (13), (26), (27)\ \forall n = 0 \dots N_{\text{p}} - 1, \tag{31b}$$

$$(14)\ \forall n = 1 \dots N_{\text{p}}. \tag{31c}$$

To heat or cool the entire building by 1 K, the heating or cooling must run at full power for approximately 1.25 h. Additional budget needs are caused by positive and negative heat flows canceling in the overall error, meaning that one half of the building encounters a heating disturbance, the other half of the building a cooling disturbance. Hence, 1 K in weighted absolute deviation occurs when half the building would need to be cooled fully, while the other half would need to be heated by the same amount, meaning a total heat power of

$$\dot{Q}_{\text{max}} = 0.5 \cdot \dot{Q}_{\text{cool,max}} + 0.5 \cdot \dot{Q}_{\text{heat,max}} \tag{32}$$

must be applied for approx. 2 time steps to compensate the disturbance. Thus, we make the assumption that $(1\,\text{K})^2$ of weighted squared deviation from the reference trajectory is equivalent to the maximum possible costs incurred by increasing the budget accordingly, i. e.

$$(1\,\text{K})^2 \mathrel{\hat{=}} T_{\text{s}} \cdot \left(c_{\text{buy}}^{\text{grid}} \cdot \dot{Q}_{\text{cool,max}} + c_{\text{gas}} \cdot \dot{Q}_{\text{heat,max}}\right). \tag{33}$$

We use this to calibrate the weight w_{comp} conservatively, i. e.

$$w_{\text{comp}} = w_{\text{mon}} \frac{T_{\text{s}}}{(1\,\text{K})^2} \left(c_{\text{buy}}^{\text{grid}} \dot{Q}_{\text{cool,max}} + c_{\text{gas}} \dot{Q}_{\text{heat,max}}\right). \tag{34}$$

In general, the choice of this weight is arbitrary and should be made such that a good compromise between budget increase and tracking error is achieved.

The new plan resulting from this step is final and applied by the aggregator and distributor, i. e. the distributor distributes the newly allocated budgets by resolving its OCP with $M_{\text{slack}} = 0$. In the next time step, the described negotiation process is repeated.

3.4. Monolithic baseline approach

The monolithic control approach solves a single OCP to derive all control inputs of both the aggregator and distributor, by optimizing a joint objective function, considering monetary costs and the thermal objectives for the individual zones. For this approach, we combine the distributor's state-space model (4) with a representation of the stationary battery

$$E(k+1) = E(k) + P_{\text{bat}}(k) \cdot T_{\text{s}}, \tag{35}$$

with the balance equation

$$P_{\text{bat}}(k) = P_{\text{grid}}(k) + P_{\text{chp}}(k) + \frac{1}{\varepsilon_{\text{c}}} \dot{Q}_{\text{cool,b}}(k) + \frac{1}{\varepsilon_{\text{c}}} \dot{Q}_{\text{cool,s}}(k) + P_{\text{dem}}(k) + P_{\text{PV}}(k). \tag{36}$$

Furthermore, aggregator and distributor inputs are now directly coupled through

$$P_{\text{chp}}(k) = c_{\text{cur}} \cdot \dot{Q}_{\text{chp}}(k), \tag{37a}$$

$$\dot{Q}_{\text{heat}}(k) = \dot{Q}_{\text{rad}}(k) + \dot{Q}_{\text{chp}}(k), \tag{37b}$$

$$\dot{Q}_{\text{heat}}(k) = \sum_{i=1}^{i=7} \dot{Q}_{\text{heat},i}(k), \tag{37c}$$

$$\dot{Q}_{\text{cool,b}}(k) = \sum_{i=1}^{i=7} \dot{Q}_{\text{cool},i}(k), \tag{37d}$$

$$\dot{Q}_{\text{cool,s}}(k) = \sum_{i=8}^{i=9} \dot{Q}_{\text{cool},i}(k). \tag{37e}$$

Objectives. The monolithic approach should behave equivalently to the hierarchical approach. Therefore, it needs to consider the monetary costs $J_{\text{comf,agg}}$, as well as the combination of the thermal objectives of the aggregator and distributor. The hierarchical approach combines two thermal objectives for the building zones 1–7 through the hierarchy: Tracking the reference temperature across the entire building using $J_{\text{comf,agg}}$ and minimizing the spread between individual zones using $J_{\text{comf,dis}}$. Both objectives can be transferred to the monolithic approach by replacing $\vartheta_{\text{b}}(k)$ in Eq. (5) and (16) with the weighted average of the individual zones, i. e.

$$\vartheta_{\text{b}}(k) = \sum_{i=1}^{7} w_{\text{b},i} \cdot \vartheta_i(k). \tag{38}$$

For controlling the server zones 8 and 9, the distributor objective $J_{\text{s,dis}}$ can be used directly.

Optimal control problem. Combining the objectives with the modified system description, the OCP of the monolithic approach is given by

$$\min_{\boldsymbol{u}_{\text{agg}}, \boldsymbol{u}_{\text{dis}}} \quad w_{\text{comf}} \cdot J_{\text{comf,agg}}(k) + w_{\text{comp}} \cdot J_{\text{comf,dis}}(k) + w_{\text{s,agg}} \cdot J_{\text{s,dis}}(k) + w_{\text{mon}} \cdot J_{\text{mon}}(k), \tag{39a}$$

$$\text{s.t.} \quad (4), (13), (21),\ \forall n = 0, \dots, N_{\text{p}} - 1, \tag{39b}$$

$$(35)\text{–}(37),\ \forall n = 0, \dots, N_{\text{p}} - 1, \tag{39c}$$

$$(14)\ \forall n = 1, \dots, N_{\text{p}}, \tag{39d}$$

with the same weights as in the hierarchical scheme.

4. Data-driven error compensation

As described in Section 2, we consider a residual model error $\epsilon(k)$ in both aggregator and distributor. In the hierarchical control approach presented in this paper, we want to compensate this unknown residual error. Therefore, we want to design an estimator $\tilde{\epsilon}(k)$ of this residual error $\epsilon(k)$, such that $\tilde{\epsilon}(k) \approx \epsilon(k)$. For this, we follow the approach developed in [31]. Integrating the estimator into the control approach allows for active compensation of the model error and thus improved control performance. As motivated in Section 2.1, we assume the time-variant model error to represent unknown exogenous influences. While these influences are unknown and difficult to model explicitly, they are by nature strongly correlated with observable and predictable data. Therefore, we use machine learning regression models as estimators. We train a regressor that estimates the residual error of the distributor $\tilde{\epsilon}_{\text{dis}}(k) = \left[\tilde{\epsilon}_1(k) \quad \dots \quad \tilde{\epsilon}_9(k)\right]^{\intercal}$. The estimates in the aggregator are the weighted sum of the individual zone estimates, which is analogous to ϑ_{b} and ϑ_{s} being the weighted averages of the individual zone temperatures ϑ_i, i. e.

$$\tilde{\epsilon}_{\text{b}}(k) = \frac{\sum_{i=1}^{7} C_{\text{th},i} \cdot \tilde{\epsilon}_i(k)}{\sum_{i=1}^{7} C_{\text{th},i}}, \tag{40a}$$

$$\tilde{\epsilon}_{\text{s}}(k) = \frac{\sum_{i=8}^{9} C_{\text{th},i} \cdot \tilde{\epsilon}_i(k)}{\sum_{i=8}^{9} C_{\text{th},i}}. \tag{40b}$$

The estimator predicts the model error one time step at a time, i. e. N_{p} separate predictions will be made to calculate $\tilde{\epsilon}_i(n|k)$ over the horizon $n = 0, \dots, N_{\text{p}} - 1$ at each time step k.

4.1. Feature selection

The first step to training a regression model is feature selection. The target variable (i. e. labels) of the regression model are the observed model errors. These can be calculated as the difference between measured and predicted state, i. e.

$$\epsilon(k) = x(k+1) - x(1|k). \tag{41}$$

Generally, this difference will also include prediction errors of the disturbances and control inputs: Both the disturbance $d(0|k)$ predicted and input $u(0|k)$ derived for the current time step k by the MPC will differ from the actual values $d(k)$ and $u(k)$ realized through the building (or digital twin). This may be unwanted, as prediction errors of the inputs are generally dynamic in nature, i. e.

$$u(k) - u(0|k) = f(u(0|k), x(k)) \tag{42}$$

which is a type of error we do not aim to remedy with the proposed approach. To solve this, a static time-variant error $\epsilon(k)$ can be retrieved by recalculating $x(1|k)$ using the state space model and measurements $u(k)$ and known disturbances $d(k)$.

When selecting the features of the regressors, i. e. their input values, we want to consider easy (re)trainability and employability. This means that only such features should be used, which are readily available as historical measurements offline and are also both measurable and predictable online. Hence, we consider all (non-constant) disturbances of the aggregator as good candidate features, since by definition, they are both measurable and predictable in the case of the proposed EMS. Additionally, all states of the aggregator and distributor, as well as all inputs of both layers are measurable in our setup.

We want to estimate errors which are caused by internal heat flows and ambient influences. The internal heat flows are mainly caused through occupant behavior, which is time dependent, seasonal and correlated with the electrical consumption inside the building, which is also a direct cause of internal heating. Therefore, the total building load $P_{\text{dem}}(k)$, the time of day $\text{ToD}(k) : \mathbb{N} \to [0, 24)$ in h and the day of the year $\text{DoY}(k) : \mathbb{N} \to \{1, \dots, 365\}$ are chosen as features. Ambient influences are directly dependent on the ambient temperature $\vartheta_{\text{air}}(k)$, yielding another feature, but are also time-dependent (e. g. warm/cold air drawn in by the HVAC system). Furthermore, ambient influences may propagate through the building through heat diffusion, hence we additionally consider past values of the ambient temperature $\vartheta_{\text{air}}(k-1), \vartheta_{\text{air}}(k-2)$. The full set of features we consider is thus $P_{\text{dem}}(k), \vartheta_{\text{air}}(k), \vartheta_{\text{air}}(k-1), \vartheta_{\text{air}}(k-2), \text{ToD}(k), \text{DoY}(k)$.

4.2. Generating training data

We used the digital twin model described in Section 2.3 to generate the training data for training the estimators. We simulated the full calendar year 2021 by applying the control approach presented in Section 3 to the digital twin in a software-in-the-loop (SiL) setup. Here, no compensation was applied, i. e. $\tilde{\epsilon}_i(k) = 0$. Note, that since all $\tilde{\epsilon}_i(k) = 0$, this is equivalent to omitting the negotiation in the controller and having the distributor always follow the proposed budget of the aggregator. The MPC controller is implemented in Python. The digital twin model is exported as a functional mock-up unit (FMU) and connected to the controller using the `fmpy` package. At each time step k, the controller receives the updated state from the digital twin model, determines the control input $u(0|k)$ and applies it to the digital twin. For the SiL simulation, real measurement data collected for the year 2021 at the Honda R&D facility in Offenbach, Germany was used in both the digital twin and the MPC. In this study, we want to showcase the general method of estimating the intrinsic model error and how this relates to control performance. To ease interpretability of the results and to exclude other sources of uncertainty, the MPC uses perfect predictions for PV power, the building's load and ambient air temperature. To evaluate the impact of actual predictions on the facility discussed in this paper, readers are directed to [41,42]. We collected both the states and inputs predicted by the MPC as well as the resulting states and inputs in the digital twin. From these we calculated the model error $\epsilon(k)$ for training as described initially. In a real world setting, one could use historical measurement data that can be mapped to the states, inputs and disturbances of the state-space models, to perform an a posteriori simulation using the state-space model, to calculate or approximate the model error.

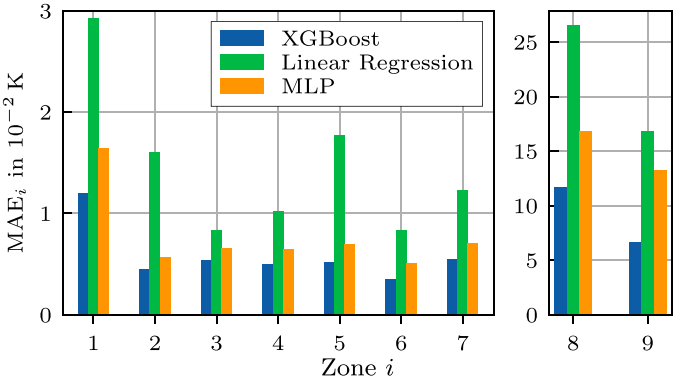

Fig. 3. Mean absolute error (MAE) of the three candidate regressor models across the 9 individual temperature zones on the test set.

4.3. Regressor selection

We evaluate three different candidate regression models for the error estimation:

- XGBoost (eXtreme Gradient Boosting),
- multiple linear regression,
- and a multilayer perceptron (MLP) with 2 ReLU layers with 128 neurons each.

XGBoost is an open-source software library that implements gradient boosted decision trees efficiently and effectively [32]. Gradient tree boosting is a technique for regression and classification tasks, which makes use of ensembles of decision trees. We use the same general features for all candidates, but apply a cyclical transformation of the time features $\text{ToD}(k)$ and $\text{DoY}(k)$ for the linear regression and MLP. To train the estimators, we applied a 75/25 train-test split on the training data, and used scikit-learn [43] and PyTorch to train the regressors. Fig. 3 shows the mean absolute error (MAE) in each temperature zone, i. e.

$$\text{MAE}_i = \frac{1}{N} \sum_{k=0}^{N-1} \left| \epsilon_i(k) - \tilde{\epsilon}_i(k) \right|, \tag{43}$$

on the test set for all three candidate models. Out of the candidates, XGBoost exhibits the best performance across all zones.

5. Simulation results

To evaluate our proposed approach, we apply it in a SiL simulation of the year 2021 as described in Section 4.2. Fig. 4 shows the building load demand, available PV power and ambient temperature over the course of the evaluated year. The OCPs where implemented in Python using CVXPY [44]. We conducted all simulations on an Intel Xeon E3-1280 v6 CPU using Gurobi [45] as the solver. As a first evaluation, we compare the proposed hierarchical approach against a naive hierarchical controller without information mitigation in Section 5.1. To confirm the efficacy of our approach, we then benchmark it against the monolithic baseline approach in Section 5.2. Since our proposed approach relies on the data-driven residual estimation, we evaluate how estimation performance and control performance relate in Section 5.3.

5.1. Comparison against hierarchical baseline

We compare our approach to two comparison scenarios:

1. **Baseline**: a reduced hierarchical control approach, without error compensation, i. e. $\tilde{\epsilon}(k) = 0$, and without negotiation. This is equivalent to only solving the aggregator's OCP (15) and then having the distributor distribute the unadjusted plan by solving (25) with $M_{\text{slack}} = 0$, i. e. executing only step 1 and 4 of the negotiation process shown in Fig. 2.

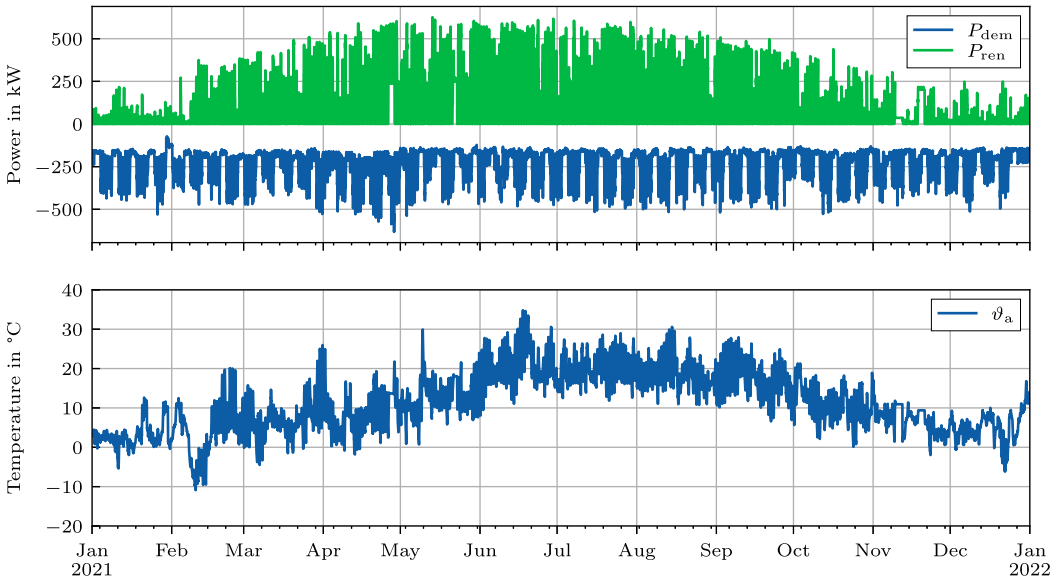

Fig. 4. Available PV power P_{PV}, building load demand P_{dem} and ambient temperature ϑ_{air} during the year 2021.

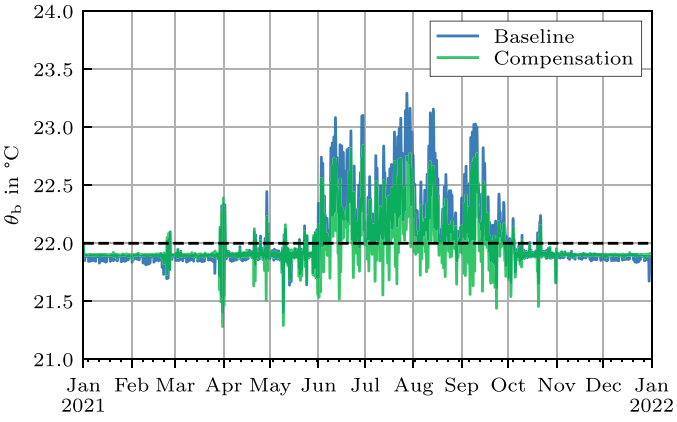

Fig. 5. Comparison of building zone temperature ϑ_{b} between a simulation without (baseline, in blue) and with error compensation (compensation, in green), but both without negotiation between the layers.

2. **Compensation**: the same as baseline, but with active error compensation, i. e. $\tilde{e}(k) \neq 0$.

Fig. 5 presents the resulting building zone temperature ϑ_{b} for the evaluated scenarios, showing that active error compensation improves overall temperature control performance, as ϑ_{b} is significantly closer to the desired 22 °C during summer in the compensation scenario. Figs. 6(a) and 6(b) show the individual zone temperatures for the baseline and compensation scenarios, respectively. Comparing the resulting control performance shows only minor improvement in the compensation scenario compared to baseline. Fig. 7 shows the mean absolute residual error (MARE) for each zone in the two scenarios, i. e.

$$\text{MARE}_i = \frac{1}{N} \sum_{k=0}^{N-1} |\epsilon_i(k)|, \tag{44}$$

where $\epsilon_i(k)$ is the left-over model error recorded during the simulations. This shows that the model error is indeed significantly reduced in all zones when error compensation is active. This suggests that the low control performance regarding temperature tracking is not caused by insufficient error compensation, but instead by the information mismatch present between the layers of the hierarchy, as motivated in Section 3. This is confirmed when looking at the control performance achieved by our proposed approach with negotiation shown in Fig. 6(c). Here, the temperature in all zones stays close to 22 °C and zone temperatures no longer diverge significantly. This is further illustrated in Fig. 8, which shows the root mean square error (RMSE) of the deviation from the reference temperature ϑ_{b} set by the aggregator in zones 1–7, i. e.

$$\text{RMSE}_i = \sqrt{\frac{1}{N} \sum_{k=1}^{N} \left(\vartheta_i(k) - \vartheta_{\text{b}}(k)\right)^2}. \tag{45}$$

In the baseline case, all zones strongly deviate from the reference temperature, while in the compensation scenario the tracking error is only very slightly reduced in zones 1, 2, 3, 5, and 6, and even slightly increased in zones 4 and 7. With both active compensation and the proposed negotiation however, the tracking error is reduced significantly in all zones.

Fig. 9 exemplarily shows the behavior of the proposed approach with negotiation on a representative period of three days in May 2021. In this period, high fluctuations of the ambient temperature occur, rising from approx. 7 °C to 30 °C during one day, which is a common occurrence in late spring in Germany.

Table 4 shows the monetary costs, the weighted RMSE (WRMSE) of the tracking error, calculated as

$$\text{WRMSE} = \sum_{i=1}^{7} w_{\text{b},i} \cdot \text{RMSE}_i, \tag{46}$$

and the overall comfort performance

$$\text{RMSE}_{\text{comf}} = \sqrt{\frac{1}{N} \sum_{k=1}^{N} \left(\sum_{i=1}^{7} w_{\text{b},i} \cdot \vartheta_i(k) - 22\,°\text{C} \right)^2} \tag{47}$$

for the scenarios evaluated with the hierarchical approach as well as the monolithic baseline. Note that for the monolithic approach, $\vartheta_{\text{b}}(k) = \sum_{i=1}^{7} w_{\text{b},i} \cdot \vartheta_i(k)$ in (45). In terms of overall comfort performance, the initial findings are confirmed, i. e. compensation without negotiation already enables better tracking of the overall building temperature. With both active compensation and the proposed negotiation, the overall comfort performance is again improved. Compared to baseline, the tracking error could only be reduced by 3.9% with active compensation but without negotiation. The proposed approach with negotiation however incurs only 4.2% higher monetary costs, while reducing the tracking error by 98.2%.

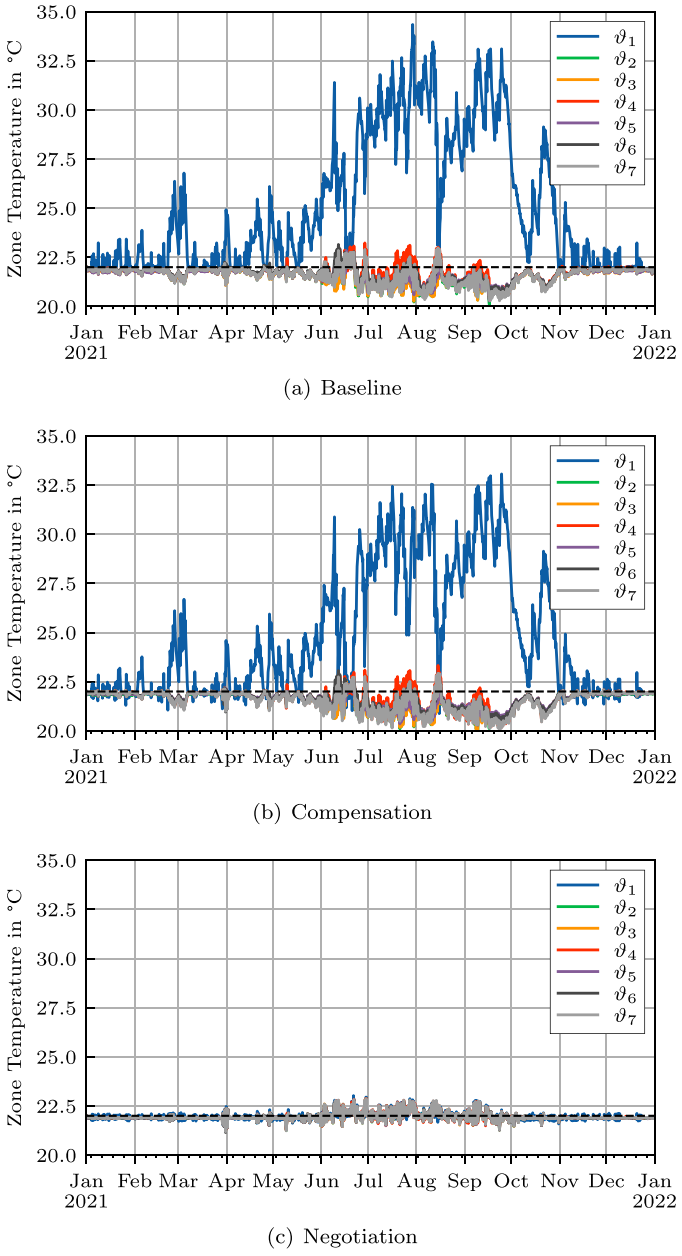

Fig. 6. Temperatures θ_i of the individual building temperature zones of zones 1–7, (a) without error compensation and without negotiation, (b) with error compensation and without negotiation, (c) with both error compensation and negotiation.

Table 4
Energy costs,weighted root mean square error (WRMSE) of the temperature tracking error between individual zone temperatures and reference temperature, and overall comfort performance $\text{RMSE}_{\text{comf}}$ for all evaluated scenarios and approaches.

Simulation	Energy costs in Euro	WRMSE in K	$\text{RMSE}_{\text{comf}}$ in K
Baseline	379 941	1.125	0.274
Compensation	380 576	1.081	0.176
Negotiation	395 841	0.020	0.171
Monolithic	394 287	0.016	0.160

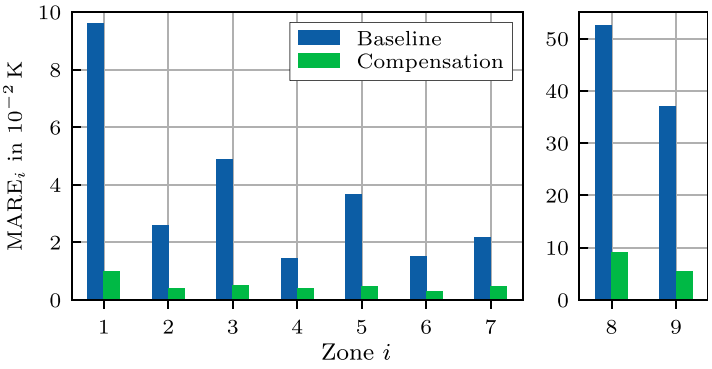

Fig. 7. Mean absolute residual error (MARE) without error compensation (Baseline) and with error compensation (Compensation) in each temperature zone.

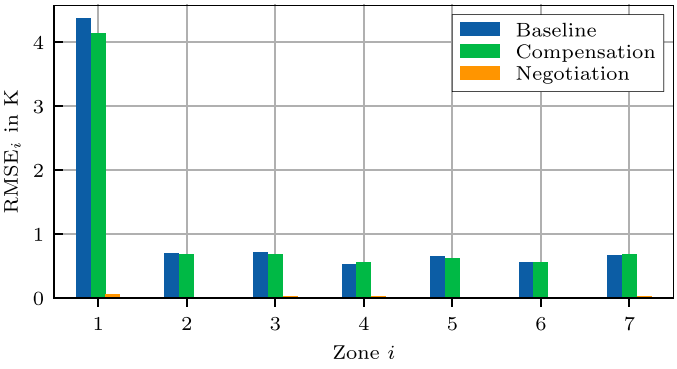

Fig. 8. Root mean square error (RMSE) of the tracking error between the reference temperature ϑ_b and individual building zone temperatures ϑ_i in zones 1–7, without error compensation and without negotiation (Baseline), with error compensation and without negotiation (Compensation), and with both error compensation and negotiation (Negotiation).

5.2. Comparison against monolithic baseline

To further benchmark our proposed approach, we simulated the monolithic baseline approach in the identical setting with activate error compensation. Table 4 shows the results of the simulations in terms of energy costs and control performance. When compared to the monolithic approach, our proposed hierarchical approach achieves very similar performance. The monolithic approach incurs 0.4% lower costs than the hierarchical approach, and achieves better, but very similar tracking error reduction of 98.6% against the naive baseline. In terms of overall comfort performance, the proposed hierarchical approach again performs similar to the monolithic baseline. These results indicate that the proposed approach is able to successfully mitigate the information mismatch.

5.3. Impact of estimation performance on control performance

To evaluate how the control performance of the proposed control approach is influenced by the performance of the data-driven residual estimation, we performed the initially described simulations with all three estimators proposed in Section 4.3, namely XGBoost, a multiple linear regression model and an MLP. Fig. 10 shows the MARE in each zone for all three estimators. The results are comparable to those during training (see Fig. 3): XGBoost performs best, the MLP exhibits very similar performance to XGBoost, and the linear model performs worst, but still reasonably well. Fig. 11 shows the RMSE of the deviation from the reference temperature in each zone for each of the three regressors.

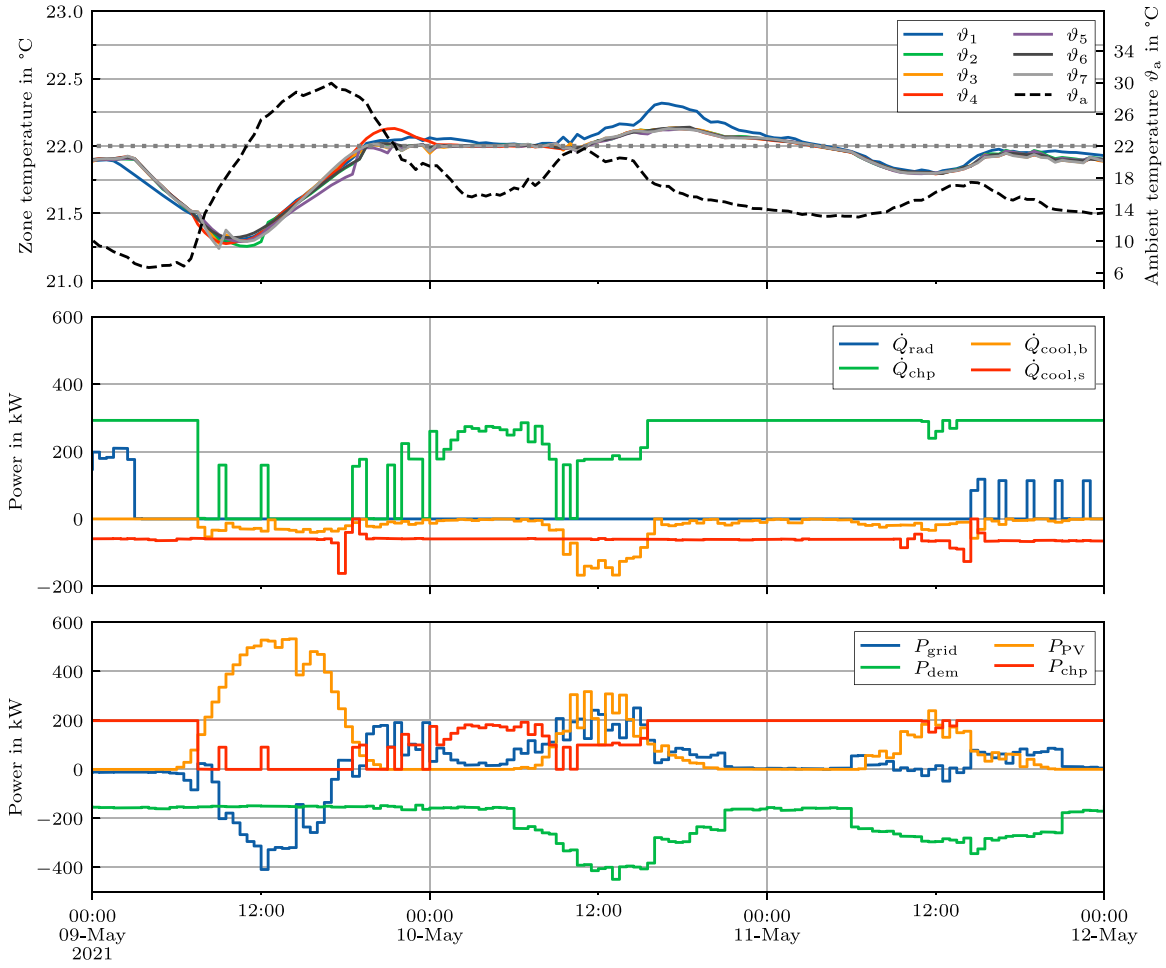

Fig. 9. Building zone temperatures, ambient temperature, thermal and electrical control inputs during 3 representative days in May as realized by the proposed hierarchical control approach. In anticipation of the high ambient temperatures in the afternoon of the first day, the controller lets the building cool down in the first half of the day, applying only minimal thermal power to compensate for internal heat gains, keeping the temperature uniform across zones. During the afternoon of the second day, the controller relaxes the temperature tracking in zone 1 to account for the receding availability of PV power and to make use of the decreasing ambient temperatures to passively cool the zone during the evening and night.

The RMSE of the tracking error mostly follows the trends in MARE performance, but is less pronounced. XGBoost and the MLP exhibit almost equal tracking performance, and the linear model is competitive in zones 3, 4 and 7. Since the tracking error is an objective of the control approach, i. e. it is weighed against monetary costs. Hence, the tracking error is partly a deliberate choice of the controller, meaning that under certain conditions, a better estimations performance may not lead to a reduced tracking error. Table 5 shows the WRMSE of the tracking error and the overall comfort performance with all three estimators. In terms of these metrics, XGBoost and the MLP produce almost equal control performance, with XGBoost slightly outperforming XGBoost in terms of tracking error, and the MLP very slightly outperforming XGBoost in terms of overall comfort performance. With the linear model, which exhibits significantly worse estimation performance, worst performance in both metrics is achieved, which still significantly outperforms the naive hierarchical control approach. Overall, the results indicate that the proposed approach is sensitive to estimation performance, with a better estimation performance leading to better control performance. However, even with reduced estimation performance, good control performance can be achieved.

Table 5
WRMSE of the temperature tracking error between individual zone temperatures and reference temperature and overall comfort performance $RMSE_{comf}$ with the proposed control approach for all evaluated regression models.

Simulation	WRMSE in 10^{-2} K	$RMSE_{comf}$ in 10^{-2} K
XGBoost	1.975	17.072
Linear Regression	2.686	17.448
MLP	2.003	16.968

6. Conclusion & outlook

We have shown that our proposed hierarchical control approach can significantly improve control performance, by successfully compensating the model error and accounting for the information mismatch between the layers of the hierarchy. Control performance in terms of temperature tracking performance between aggregator and distributor is greatly increased (98.2% improvement), while incurring only slightly increased monetary costs (4.2%). Furthermore, it achieves only slightly decreased performance compared to a monolithic baseline approach.

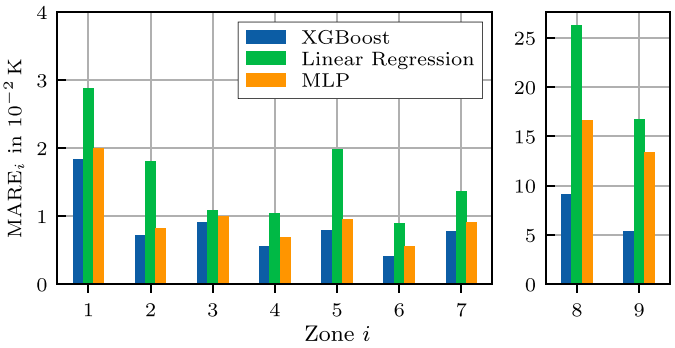

Fig. 10. Mean absolute residual error (MARE) of the evaluated regressor models, XGBoost, multiple linear regression and the multilayer perceptron (MLP) in each temperature zone.

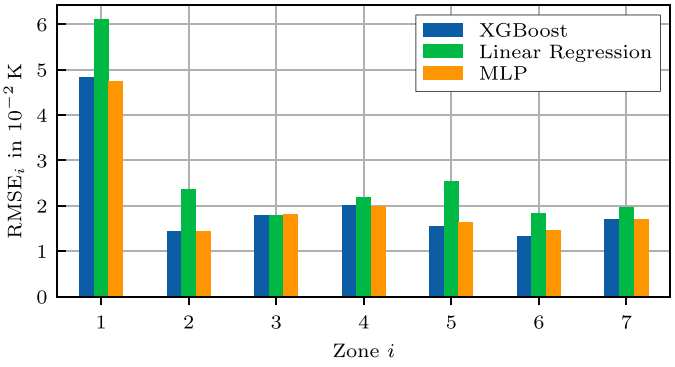

Fig. 11. Root mean square error (RMSE) of the tracking error between the reference temperature ϑ_b and individual building zone temperatures ϑ_i in zones 1–7, with the proposed control approach for all evaluated regressor models, XGBoost, multiple linear regression and the multilayer perceptron (MLP).

For this paper, we have used data from a full calendar year for training the error estimators. We have shown that the proposed approach relies on reasonably good estimation performance to achieve good control performance. Hence, further investigation could be conducted to understand how much and which training data is necessary for effective (re-)training of compensators in a real life setting.

Furthermore, the proposed error compensation approach can only account for additive disturbances, which already greatly improves performance. Extending the approach to cover dynamical errors may be worth investigating, as this may further improve control performance. However, such an extension is challenging, as a dynamical error estimator must enter the OCP directly, hence certain structural restrictions must be accounted for. Also, generating and selecting training data to sufficiently capture a building's behavior provides another challenge.

CRediT authorship contribution statement

Jens Engel: Writing – original draft, Methodology, Investigation, Data curation, Conceptualization. **Thomas Schmitt:** Writing – review & editing, Methodology, Conceptualization. **Tobias Rodemann:** Writing – review & editing, Supervision, Project administration. **Jürgen Adamy:** Writing – review & editing, Supervision.

Declaration of competing interest

The authors declare the following financial interests/personal relationships which may be considered as potential competing interests: Jens Engel and Thomas Schmitt have a patent pending to Honda Motor Co., Ltd. If there are other authors, they declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Data availability

The authors do not have permission to share data.

References

[1] Drgoňa J, Arroyo J, Figueroa IC, Blum D, Arendt K, Kim D, et al. All you need to know about model predictive control for buildings. Annu Rev Control 2020.
[2] Péan T, Costa-Castelló R, Salom J. Price and carbon-based energy flexibility of residential heating and cooling loads using model predictive control. Sustainable Cities Soc 2019;50:101579.
[3] Rawlings JB, Patel NR, Risbeck MJ, Maravelias CT, Wenzel MJ, Turney RD. Economic MPC and real-time decision making with application to large-scale HVAC energy systems. Comput Chem Eng 2018;114:89–98.
[4] Scattolini R. Architectures for distributed and hierarchical model predictive control – a review. J Process Control 2009;19(5):723–31.
[5] Kennel F, Görges D, Liu S. Energy management for smart grids with electric vehicles based on hierarchical MPC. IEEE Trans Ind Inf 2013;9(3):1528–37.
[6] Engel J, Schmitt T, Rodemann T, Adamy J. Hierarchical economic model predictive control approach for a building energy management system with scenario-driven EV charging. IEEE Trans Smart Grid 2022;3082–93.
[7] Lin D, Zhang L, Xia X. Model predictive control of a venlo-type greenhouse system considering electrical energy, water and carbon dioxide consumption. Appl Energy 2021;298:117163.
[8] Kim D, Braun JE. Hierarchical model predictive control approach for optimal demand response for small/medium-sized commercial buildings. In: 2018 annual American control conference (ACC). IEEE; 2018.
[9] Mtibaa F, Nguyen K-K, Dermardiros V, McDonald S, Venne J-S, Cheriet M. Hierarchical multi-agent control framework for energy efficiency and carbon emission reduction in multi-zone buildings. J Build Eng 2023;68:106204.
[10] Vasak M, Banjac A, Hure N, Novak H, Marusic D, Lesic V. Modular hierarchical model predictive control for coordinated and holistic energy management of buildings. IEEE Trans Energy Convers 2021;36(4):2670–82.
[11] Raman NS, Chaturvedi RU, Guo Z, Barooah P. Model predictive control-based hierarchical control of a multi-zone commercial HVAC system. ASME J Eng Sustain Build Cities 2021;2(2).
[12] Khatibi M, Rahnama S, Vogler-Finck P, Bendtsen JD, Afshari A. Towards designing an aggregator to activate the energy flexibility of multi-zone buildings using a hierarchical model-based scheme. Appl Energy 2023;333:120562.
[13] Hure N, Martincevic A, Vasak M. Model predictive control of building HVAC system employing zone thermal energy requests. In: 2019 22nd international conference on process control (PC19). IEEE; 2019.
[14] Lankeshwara G, Sharma R, Yan R, Saha TK. A hierarchical control scheme for residential air-conditioning loads to provide real-time market services under uncertainties. Energy 2022;250:123796.
[15] Wang Z, Zhao Y, Zhang C, Ma P, Liu X. A general multi agent-based distributed framework for optimal control of building HVAC systems. J Build Eng 2022;52:104498.
[16] Prívara S, Cigler J, Váňa Z, Oldewurtel F, Sagerschnig C, Žáčeková E. Building modeling as a crucial part for building predictive control. Energy Build 2013;56:8–22.
[17] Tian W, Heo Y, de Wilde P, Li Z, Yan D, Park CS, et al. A review of uncertainty analysis in building energy assessment. Renew Sustain Energy Rev 2018;93:285–301.
[18] Wang Z, Chen Y. Data-driven modeling of building thermal dynamics: Methodology and state of the art. Energy Build 2019;203:109405.
[19] Maddalena ET, Lian Y, Jones CN. Data-driven methods for building control — A review and promising future directions. Control Eng Pract 2020;95:104211.
[20] Halhoul Merabet G, Essaaidi M, Ben Haddou M, Qolomany B, Qadir J, Anan M, et al. Intelligent building control systems for thermal comfort and energy-efficiency: A systematic review of artificial intelligence-assisted techniques. Renew Sustain Energy Rev 2021;144:110969, URL https://www.sciencedirect.com/science/article/pii/S1364032121002616.
[21] Di Natale L, Svetozarevic B, Heer P, Jones C. Physically consistent neural networks for building thermal modeling: Theory and analysis. Appl Energy 2022;325:119806.
[22] Di Natale L, Svetozarevic B, Heer P, Jones C. Towards scalable physically consistent neural networks: An application to data-driven multi-zone thermal building models. Appl Energy 2023;340:121071.
[23] Scheurenberg D, Stemmler S, Abel D. Evaluation of data enhanced model predictive control for a coupled tank system. In: 2023 IEEE conference on control technology and applications. CCTA, IEEE; 2023, p. 79–84.
[24] Jain A, O'Kelly M, Chaudhari P, Morari M. BayesRace: Learning to race autonomously using prior experience. 2020, arXiv preprint arXiv:2005.04755.
[25] Picotti E, Bianchin F, Bruschetta M. Real-time learning-based nonlinear model predictive control of a virtual motorcycle employing grey-box modeling through Gaussian processes. Control Eng Pract 2024;144(Germany):105837.
[26] Gray FM, Schmidt M. A hybrid approach to thermal building modelling using a combination of Gaussian processes and grey-box models. Energy Build 2018;165:56–63.

[27] Oldewurtel F, Sturzenegger D, Morari M. Importance of occupancy information for building climate control. Appl Energy 2013;101:521–32.
[28] Thilker CA, Madsen H, Jørgensen JB. Advanced forecasting and disturbance modelling for model predictive control of smart energy systems. Appl Energy 2021;292(Denmark):116889.
[29] Ellis MJ. Machine learning enhanced grey-box modeling for building thermal modeling. In: 2021 American control conference (ACC). IEEE; 2021.
[30] Kumar P, Rawlings JB, Wenzel MJ, Risbeck MJ. Grey-box model and neural network disturbance predictor identification for economic MPC in building energy systems. Energy Build 2023.
[31] Schmitt T, Engel J, Rodemann T. Regression-based model error compensation for hierarchical MPC building energy management system. In: 2023 IEEE conference on control technology and applications. CCTA, IEEE; 2023, URL https://arxiv.org/abs/2306.09080.
[32] Chen T, Guestrin C. XGBoost: A scalable tree boosting system. In: Proceedings of the 22nd ACM SIGKDD international conference on knowledge discovery and data mining. ACM; 2016.
[33] Liang W, Li H, Zhan S, Chong A, Hong T. Energy flexibility quantification of a tropical net-zero office building using physically consistent neural network-based model predictive control. Adv Appl Energy 2024;100167.
[34] Stadie M, Rodemann T, Burger A, Jomrich F, Limmer S, Rebhan S, et al. V2B vehicle to building charging manager. In: EVTeC: 5th international electric vehicle technology conference 2021. Japanese Society of Automotive Engineers; 2021.
[35] Schmitt T. Multi-objective building energy management optimization with model predictive control (Ph.D. thesis), Darmstadt: Technische Universität Darmstadt; 2022, XVIII, 148 Seiten, URL http://tuprints.ulb.tu-darmstadt.de/22344/.
[36] Unger R, Mikoleit B, Schwan T, Bäker B, Kehrer C, Rodemann T. Green building - modeling renewable building energy systems with eMobility using modelica. In: Proceedings of modelica 2012 conference. Munich, Germany: Modelica Association; 2012.
[37] Castellani A, Schmitt S, Squartini S. Real-world anomaly detection by using digital twin systems and weakly supervised learning. IEEE Trans Ind Inf 2021;17(7):4733–42.
[38] Schmitt T, Rodemann T, Adamy J. Multi-objective model predictive control for microgrids. at - Automatisierungstechnik 2020;68(8):687–702, URL https://www.honda-ri.de/pubs/pdf/4361.pdf.
[39] Schmitt T, Engel J, Rodemann T, Adamy J. Application of Pareto optimization in an economic model predictive controlled microgrid. In: 28th mediterranean conference on control and automation. MED, IEEE; 2020, p. 868–74, URL https://www.honda-ri.de/pubs/pdf/4341.pdf.
[40] Schmitt T, Hoffmann M, Rodemann T, Adamy J. Incorporating human preferences in decision making for dynamic multi-objective optimization in model predictive control. Inventions 2022;7(3).
[41] Wang C, Bäck T, Hoos HH, Baratchi M, Limmer S, Olhofer M. Automated machine learning for short-term electric load forecasting. In: 2019 IEEE symposium series on computational intelligence. SSCI, 2019, p. 314–21.
[42] Schmitt T, Rodemann T, Adamy J. The cost of photovoltaic forecasting errors in microgrid control with peak pricing. Energies 2021;14(9). URL https://www.mdpi.com/1996-1073/14/9/2569.
[43] Pedregosa F, Varoquaux G, Gramfort A, Michel V, Thirion B, Grisel O, et al. Scikit-learn: Machine learning in python. J Mach Learn Res 2011;12:2825–30.
[44] Diamond S, Boyd S. CVXPY: A python-embedded modeling language for convex optimization. J Mach Learn Res 2016;17(83):1–5.
[45] Gurobi Optimization, LLC. Gurobi optimizer reference manual. 2023, URL https://www.gurobi.com.