Journal of Building Engineering 111 (2025) 113016

Contents lists available at ScienceDirect

Journal of Building Engineering

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

A commissioning-oriented approach to data-driven modeling of buildings with heat pumps for predictive demand response

Manuel Koch [a,b,*], Parantapa Sawant [b], Ralph Eismann [b], Colin N. Jones [a]

[a] EPFL, STI IGM LA3, Lausanne, 1015, Vaud, Switzerland
[b] FHNW, Hofackerstrasse 30, Muttenz, 4132, Basel-Land, Switzerland

ARTICLE INFO

Keywords:
Building energy systems
Model predictive control
Demand response
System identification

ABSTRACT

As the share of electricity generation from non-dispatchable sources like wind and photovoltaics grows, so does the need for demand response to stabilize the grid. Since the electricity consumption of heat pumps in buildings is both substantial and flexible, they offer a large potential in this regard. Looking beyond the well-established time-of-use schemes, we investigate the more challenging task of frequency control, in which a baseline consumption and flexibility band are pre-calculated for the following day, then executed based on instantaneous commands from the grid operator. Since most real-world building automation systems do not allow a direct manipulation of the heat pump power, we identify compatible, commissioning-oriented models for predictive control, treating thermostat setpoints as an input and heat pump power as an output. In a two-month simulation study, a piecewise-affine model structure shows significantly better prediction accuracy than a simpler linear model. However, the control performance with both models is similar, with a ratio of total flexibility offered to total energy consumed of 57.1 % and 59.0 %, and a normalized tracking error of 13.7 % and 12.2 %, respectively. We further provide estimates of the efficacy of local battery storage and aggregation over multiple buildings for improving the tracking accuracy, and find an exponential decay of the error as a function of battery size and aggregation volume.

1. Introduction

In the electric grid, production and consumption must be kept in a fine balance to prevent system failure. With some exceptions, the consumption can be treated as an uncontrollable but predictable value, which is usually lowest during the night, and peaks during the early evening hours [1]. Therefore, the balancing is mostly done on the production side. In the coming years, the challenge of this task is expected to increase because of two factors: The rapid growth of non-dispatchable electricity production from wind and photovoltaics [2] means that more active balancing must be done. At the same time, there are various reasons against building more of the types of power plant best suited for flexible production. Natural gas plants are meant to be phased out under many countries carbon emission reduction goals [3,4], hydropower is tightly constrained by geography, and the types of battery suitable for grid storage are either very expensive [5,6] or still in development [7].

Because of this, the demand response market is expected to grow in the coming years [8,9]. Demand response is an umbrella term for various methods to align consumption with production, ranging from fixed time-of-use tariffs incentivizing consumers to alter their consumption, to complex and dynamic electricity trading schemes between consumers, producers and the grid operator [10]. Heat pumps and electric chillers in buildings have been identified as electricity consumers with a high potential for demand response,

* Corresponding author at: EPFL, STI IGM LA3, Lausanne, 1015, Vaud, Switzerland.
E-mail address: manuelpascal.koch@epfl.ch (M. Koch).

https://doi.org/10.1016/j.jobe.2025.113016
Received 2 October 2024; Received in revised form 13 April 2025; Accepted 27 May 2025
Available online 11 June 2025

Nomenclature

a	AGC signal
I_{sol}	Solar radiation
j	Scenario
k	Discrete time step
M	Number of scenarios
N	Prediction horizon
P_{el}	Electric heat pump power
r	Slack variables
s	Slack variables
T_{amb}	Ambient temperature
T_{set}	Setpoint temperature
T_z	Zone temperature
u	Control variable
v	Disturbance scenario
w	Control variable
y	Observation variable
z	Observation variable
γ	Flexibility band (one sided)
AGC	Automatic generation control
ARX	Autoregressive model with exogenous inputs
HVAC	Heating, ventilation and air conditioning
MPC	Model predictive control
NMAE	Normalized mean absolute error
PI	Proportional–integral controller
PWARX	Piecewise-affine ARX

since buildings tend to have high thermal inertia, and the room temperature can be temporarily varied within a certain range, without causing unreasonable discomfort in the occupants.

To give an impression of the magnitude of this potential, we note that in Europe [11,12] and the United States [13–15], buildings consume around a third of the total primary energy, with roughly half of it used for heating and cooling, depending on the location and accounting method, resulting in around 15% of the total consumption being for heating and cooling of buildings. While cooling is almost exclusively provided by electric chillers [16,17], heating is still predominantly provided by fossil fuels [14,15,18]. However, in many countries, heat pumps are becoming the predominant method of space heating, electrifying this segment of consumption as well [19,20]. For the scope of this paper, we restrict our work to the heating case.

Demand response can be broadly classified into two groups [21]: In price-based programs, the utility provider charges higher tariffs during times of high demand or low supply, but it is left to the consumers to shift their consumption accordingly. Since the use of buildings for price-based programs has already been thoroughly studied [22–24], we focus on the more technically challenging incentive-based programs, in which the consumers agree to follow a direct request from the provider to increase or decrease their consumption. They are paid for providing this service, or penalized if they fail to do so. More specifically, we aim to participate in the secondary frequency control level of the Swiss national grid operator, detailed in Section 2.3.

While the use of buildings for frequency control is well established, an analysis of recent review papers on this topic has revealed a number of topics that merit further investigation, two of which we address in this paper: On the technical level, many studies assume direct control of a heat pump's compressor and thus electric power. However, real-world building automation systems usually do not allow direct control of the power. Instead, it can only be manipulated indirectly by shifting thermostat setpoints. The dynamics thereof, including response delays, minimum compressor speeds and operating times, are neglected in the corresponding studies [25,26]. Some studies aim to work around the challenge of indirect control by treating the actual temperature as equal to the thermostat setpoint [27] or by treating the heat pump power as a linear function of the thermostat setpoint [28], but these approaches equally neglect significant dynamics of the HVAC system. Research with a more socio-economic focus finds considerable skepticism among building owners regarding the financial sensibility of participating in a frequency control program, as well as concerns about granting a third party control over their building equipment [21,29].

Before outlining the contribution of this study, we briefly introduce the control method used: Since participating in secondary frequency control requires us to predict our intended power consumption ahead of time, we use model predictive control (MPC). As the name suggests, this method uses a model of the controlled plant to predict future control actions and system states. However, obtaining this model can be a costly process, especially for buildings, which are all unique, thus require a custom model [30].

In this study, we address the aforementioned technical and socio-economic gaps in the existing literature, as well as the challenge of obtaining the building model for the MPC, by introducing a "commissioning-oriented" and data-driven approach to modeling buildings with heat pumps, based on piecewise-affine input–output models, detailed in Section 2.2. The model treats thermostat

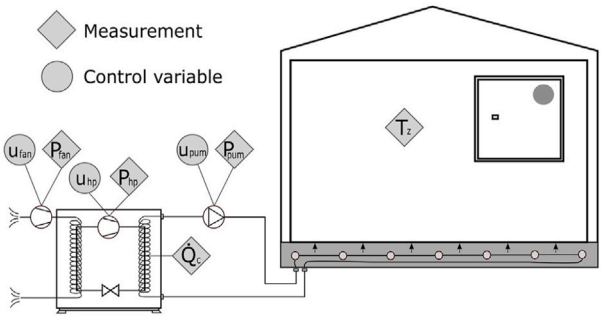

Fig. 1. Schematic of building model [33].

setpoints as an input and the resulting power consumption of the heat pump as an output, capturing the aforementioned nonlinear dynamics of the heat pump, and making it compatible with real-world building automation systems. Using a data-driven approach that only requires a few days of training data lowers the commissioning costs [28,31], addressing the financial concerns. Lastly, our approach allows an owner or occupant to maintain substantial control over the building, by setting a permissible upper and lower indoor temperature limit, to which any load shifting for frequency control is constrained. With the sum of these properties, we provide a template framework to facilitate the harnessing of heat pumps in buildings for frequency control.

2. Methodology

In this section, we further detail the proposed model structure and evaluate it for secondary frequency control with a simulated building. The evaluation is split into two parts: In the first part, we assess the open-loop prediction accuracy and compare it to a simpler linear ARX model. In the second part, we close the loop and deploy both models to predictively control the building with the objective of maximum demand response. The results are presented separately in Section 3, before being jointly discussed and concluded in Sections 4 and 5.

2.1. Building simulation and controller implementation

We use BOPTEST,[1] version 0.2.0, to run the "SingleZoneResidentialHydronicHeatPump" model from the IDEAS Modelica library,[2] shown in Fig. 1. This library is built on first principles, representing thermal masses like floors, walls and furniture as chains of discretized nodes to account for the dynamics of heat flowing through each mass. The building is a 192 m^2 single-family house, modeled as one zone, with a 15 kW air-source heat pump and a radiant slab floor heating system. The compressor of the heat pump can run at speeds corresponding to an electric power of around 1.1 kW to 3.5 kW. The occupancy is five people from 19:00 to 7:00. The thermal insulation is rather poor and the energy consumption thus high by modern standards. During sunny days, the building is prone to overheating from solar gains. Temperate weather data from Brussels, Belgium is used, along with Boptest's native weather forecasting function. The zone temperature is controlled by a thermostat with a PI controller. We choose a time step of 15 min for the system identification and MPC, which is a common choice for buildings [32]. All optimization problems are implemented using CVXPY,[3] version 1.3, and solved using Gurobi,[4] version 10.0., except for the ARX-based MPC in Section 2.6, which is solved using OSQP,[5] which we found to execute faster than Gurobi for this particular problem. The LinearSVC command from scikit-learn[6] is used to classify the operational data for the PWARX. Computations are executed on an 8-core *Intel Core i9-11900K* CPU.

2.2. Model structure and system identification

To motivate the chosen model structure, we first analyze the situation in the building simulation model used for this study, further detailed in Section 2.1: The building is actively heated, but not cooled, and its heat pump has a minimum power of around one third of its maximum power, which is typical and stems from the hydrodynamics of the lubrication of the compressor. This makes the relationship between temperature setpoint and heating power not just nonlinear, but also discontinuous, shown in a simplified sketch in Fig. 2. The figure also shows a linear and a piecewise-affine fit to the heat pump (HP) power curve. Unsurprisingly, the latter yields a much more accurate representation, especially on the left side, which represents a setpoint below the current zone

[1] https://www.energy.gov/eere/buildings/boptest-building-operations-testing-framework.
[2] https://github.com/open-ideas/IDEAS.
[3] https://www.cvxpy.org/.
[4] https://www.gurobi.com/.
[5] Included in CVXPY installation. https://osqp.org/.
[6] https://scikit-learn.org/stable/modules/generated/sklearn.svm.LinearSVC.html.

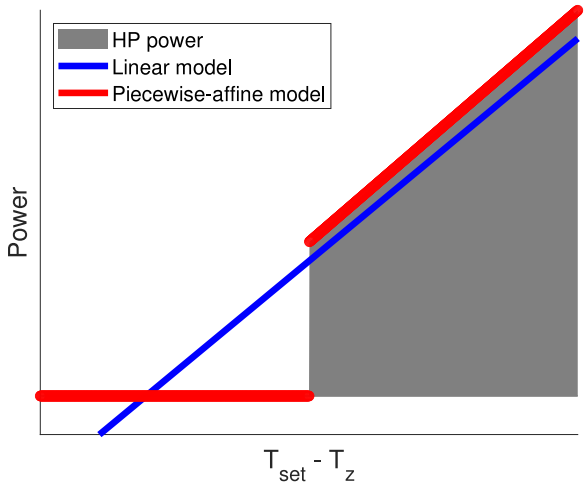

Fig. 2. Simplified illustration of a linear and a piecewise-affine fit to the nonlinear and discontinuous heat pump power as a function of the setpoint and zone temperature.

temperature. We thus select a piecewise-affine autoregressive model with exogenous inputs (PWARX) with two pieces to map the zone temperature setpoint to the electric power of the heat pump. This model structure also has the advantage that it can be trained with just a few days or weeks of data, as opposed to more general nonlinear modeling techniques, such as neural networks, Gaussian processes or random forests, which typically require months or years of training data [34,35].

To provide a quick overview of the most relevant literature, [36] described the idea of hybrid models for the identification of building thermodynamics as virtually unexplored in 2020, and presented an algorithm to fit a PWARX model with an unknown number of sub-models, followed-up by [37]. They cluster the data points by iteratively reducing the number of local models until a convergence criterion is reached. Refs. [38,39] describe an algorithm for identifying a piecewise linear (or nonlinear) model with dynamic partition specifically for input–output stable bilinear systems, briefly demonstrated on a building simulation. While bilinear dynamics are common in buildings, this does not match our needs. Refs. [40–42] use a decision tree within a manually reduced partition space to build a Takagi–Sugeno fuzzy model based on local ARX models. A number of publications identify piecewise-affine models with a predetermined partition of the data [43,44]. However, none of these studies cover the prediction of heating power as a function of thermostat setpoints.

With the stage being set, we now introduce the specific procedures used to identify three different models, of which only two are used for the open-loop system identification part of this study, in which mathematical models of the system are extracted from measured operational data. For this part, a data set is generated by randomly switching the setpoint temperature of the building between 20 °C and 24 °C every two hours. The simulation starts at the beginning of the year, but the first two days are discarded because of initialization effects. The first one of the two models is the third-order PWARX (1) with the inputs $u = \begin{pmatrix} T_{set} & T_{amb} & I_{sol} \end{pmatrix}^\top$, and the outputs $y = \begin{pmatrix} T_z & P_{el} \end{pmatrix}^\top$, corresponding to the "inverse model" in Fig. 3. The state space is partitioned by a linear support vector classifier with the parameter vectors α and β, based on whether the heat pump is on or off at the next time step.

$$y_{k+1} = \begin{cases} A_1 \tilde{y}_k + B_1 \tilde{u}_k, & \text{if } \alpha \begin{bmatrix} \tilde{y}_k^\top & \tilde{u}_k^\top \end{bmatrix}^\top \leq \beta \\ A_2 \tilde{y}_k + B_2 \tilde{u}_k, & \text{otherwise} \end{cases} \tag{1}$$

With the shorthand notation $\tilde{x}_k = \begin{pmatrix} x_k & x_{k-1} & x_{k-2} \end{pmatrix}^\top$. The second model is an equivalent linear ARX, which serves as a reference method. From January 13 to July 19, both models are identified daily at midnight, based on the last 10 days of data. Subsequently, a 24 h prediction (i.e. 96 steps) of the zone temperature and heat pump power is performed every 15 min for that day until the new model for the next day is identified. The moving data windows are visualized in the top bar of Fig. 4.

Similarly, the closed-loop frequency control part of the study also uses models identified at midnight and applied during the following day. As is common in system identification, the closed-loop data is not as well-excited as the open-loop data, which tends to result in less accurate models. We thus heuristically increased the number of days used for system identification from 10 to 14, aiming to compensate for this loss in data quality with increased data volume. The day before these 14 days is used for validation, corresponding to the bottom bar in Fig. 4, and further detailed in Section 2.4. In addition to the two "inverse models", a "physics model" is identified from the same data. This model is used for the top-level of the frequency control scheme, detailed in Section 2.5. Since it does not directly interact with the building automation system, there is no need for the inversion of the physical inputs and outputs. It is thus a conventional, third-order, linear ARX model. The frequency control simulation runs for 60 days, from January 17 to March 18.

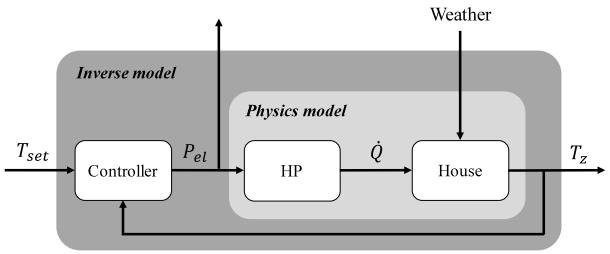

Fig. 3. Inputs and outputs of the physics model for the day-ahead planner, which is always an ARX, and the inverse model for the low-level MPC, which is either an ARX or PWARX. We note that the gray boxes show the scope of the models, not an embedding of one inside of the other.

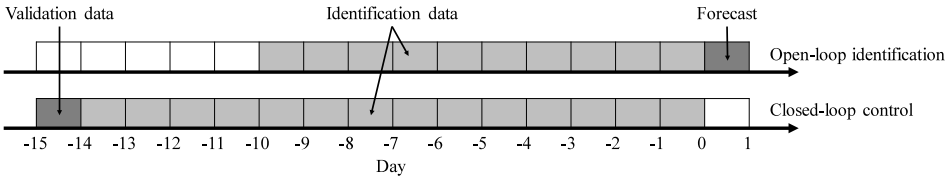

Fig. 4. Moving time windows for open-loop identification and forecast, and closed-loop identification and validation. Executed every midnight. Zero is the current time.

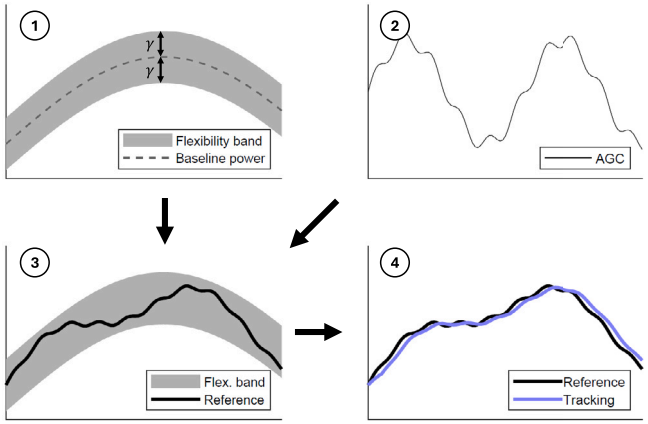

Fig. 5. Baseline power, flexibility band γ and AGC signal are merged into a reference power to be tracked with the heat pump.

2.3. Secondary frequency control framework

The Swiss national electric grid operator uses three distinct levels of frequency control: Primary frequency control, which is activated within seconds; secondary frequency control, which is activated within minutes; and tertiary frequency control, which is activated should an imbalance last longer than 15 min [10]. For this study, we chose the secondary level as the best-suited based on a trade-off between control frequency and thermal storage capacity. Tracking the faster primary level would require more frequent switching of the heat pump. This would increase the strain on its mechanical components, potentially shortening its service life. On the other hand, covering the longer modulation requests of the slower tertiary level with the given thermal storage capacity of a building would require lower modulation power, resulting in lower payments.

In this section, we detail the secondary frequency control framework, which closely follows the work of [45]. A baseline power consumption and a flexibility band γ are calculated ahead of time at midnight for the following day, and communicated to the grid operator, corresponding to the first plot of the simplified visualization in Fig. 5. We note that the flexibility band is symmetrically applied to both sides of the baseline and constant over each day. Within this band, the power consumption is then modulated based on the instantaneous automatic generation control (AGC) signal from the grid operator, which has a range of $[-1, 1]$, corresponding to the minimum and maximum power of the pre-calculated flexibility band (Plots 2 and 3). The resulting reference trajectory must then be tracked with a contractually specified level of accuracy (Plot 4). Since the grid operator rewards the amount of flexibility provided, we seek to maximize the width of the flexibility band, occasionally at the expense of optimal energy efficiency.

From [45], we have one year of AGC data with a 1 min resolution. This dataset was artificially generated to match the statistical properties of a sample of real data from the Swiss national grid operator. Figs. B.19 and B.20 show a two-day sample of the signal,

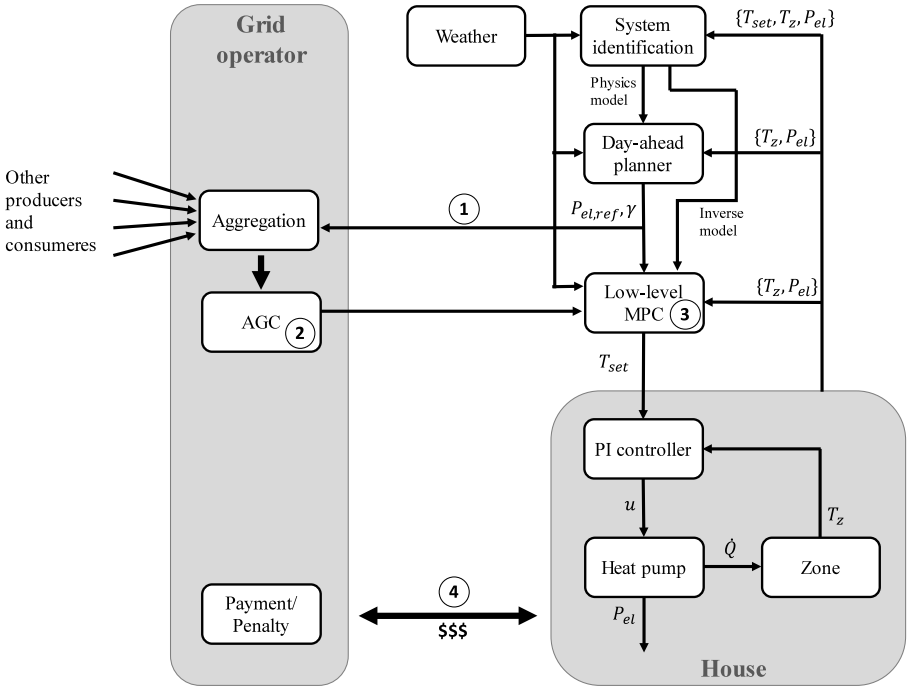

Fig. 6. Bi-level control structure for frequency control. Numbers correspond to Fig. 5.

and a frequency decomposition, which is mostly concentrated into a range of multiple minutes to hours. These are the frequencies that HVAC systems are well-suited to cover, considering the typical thermal time constants of buildings [46,47], especially radiant slab systems like the one we use [28,48]. To match the time step of the MPC, the original AGC signal is resampled from 1 min to 15 min. We note that doing so has the side effect of attenuating some peaks in the signal.

2.4. Closed-loop frequency control overview

With the situation on the building-side and the grid-side laid out, we now give an overview of the control structure, before further detailing its two main levels in Sections 2.5 and 2.6. Because of the uncertainty in the AGC signal and the temperature constraints in a building, a bi-level predictive control structure is used, with a day-ahead planning level, and a subordinate intra-day execution level with a shorter horizon, shown in Fig. 6. Once an initial physics and inverse model are identified, the day-ahead planner and low-level MPC commence operation. The day-ahead planner calculates a baseline consumption and a flexibility band, which is then combined with a prediction of the AGC signal to obtain a power tracking reference for the low-level MPC, as described before. The low-level MPC then outputs optimized thermostat setpoints that are executed by the heat pump controller.

In reality, the output of the day-ahead planner would be communicated to the grid operator, who determines and returns the AGC signal based on the bids submitted by all the participants in the frequency control scheme. However, since the power of our heat pump is negligible compared to the total contracted power for secondary frequency control [49], we neglect the impact of our bid on the overall market, instead using the aforementioned AGC dataset as is. Furthermore, we use an ideal 6 h prediction of the AGC signal, rather than implementing a data-driven forecasting method.

Fig. 7 shows the flow of identification, validation and back-up control. The physics model for the day-ahead planner and the inverse model for the low-level MPC are identified and validated simultaneously once per day. The secondary frequency control is only enabled if both models meet their respective accuracy thresholds. If one of the models fails this test, an identification mode with pseudo-random binary switching temperature setpoints is activated for that day, which provides well-excited data for the next system identification. However, even if the accuracy tests are passed, the output of the day-ahead planner may still be discarded, if the realizable flexibility band for that day is below 250 W. In this case, the building goes into a default mode for that day, with a constant setpoint of 22 °C and a nightly set-back to 20.5 °C from 20:00 in the evening to 6:00 in the morning. The straightforward implementation of such a back-up scheme is a secondary benefit of the inverse-modeling approach studied in this paper, which does not alter or go around the pre-existing, stable heat pump controller. Overriding its setpoints is trivial.

2.5. Day-ahead planning

Planning the necessary margins of safety to be able to react to the AGC signal without violating the temperature constraints of the building makes this a robust MPC problem, albeit an unusual one. Rather than guaranteeing constraint satisfaction under

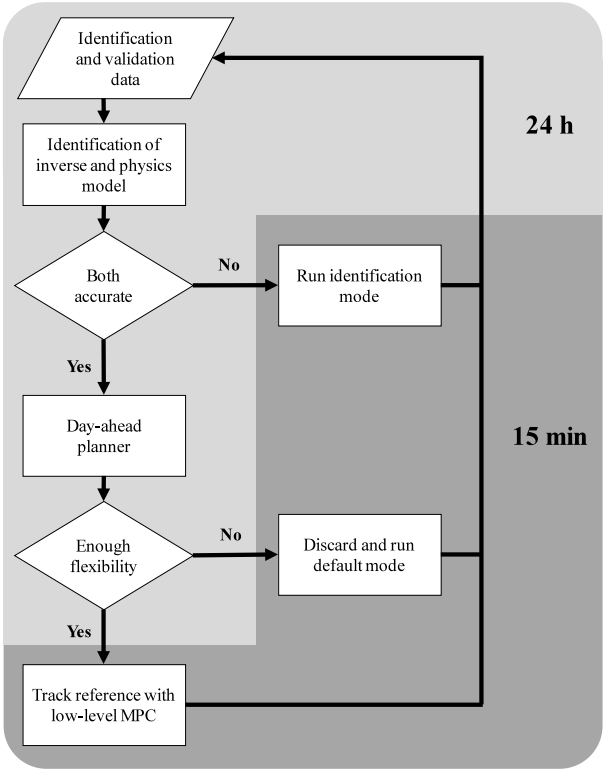

Fig. 7. Back-up control structure with time steps.

a disturbance, we aim to guarantee constraint satisfaction while tracking an uncertain reference signal. Conventional robust MPC schemes tend to be very conservative, as well as assuming a uniform noise distribution [50]. We therefore use the scenario-approach to robust MPC, which is less conservative and does not assume uniform noise [51]. Instead, the optimization is performed over a number of measured disturbance trajectories (the "scenarios"), to find an input sequence that satisfies the constraints for all of them. This method has been applied to buildings in several publications [52–54]. There is theoretical work on how many scenarios to use, but it is highly conservative. Ref. [54] reports good performance with four scenarios, capturing the uncertainties of the weather forecast and internal gains, in a building control study, when the theory suggested to use 4799 scenarios. For our study, we compiled a collection of 10 scenarios that is well-diversified and representative of the whole AGC data set, shown in Fig. B.21.

Since the output of the day-ahead planning is used indirectly as a setpoint for the low-level controller, the conventional way of modeling with power as input and temperature as output can be used, corresponding to the physics model in Fig. 3. To avoid confusion with the inverse model, we introduce new nomenclature with the input $w = P_{el}$ and the output $z = T_z$. We thus obtain the third-order ARX model (2), denoted by $\mathcal{A}$ and $\mathcal{B}$, with the time step k, the scenario j, the AGC disturbance-scenarios. v and the flexibility band γ. Note that we use the symbol v for the pre-sampled AGC scenarios and a for the actual or predicted values.

$$z_{k+1}^{j} = \mathcal{A}\tilde{z}_{k}^{j} + \mathcal{B}\begin{bmatrix} \tilde{w}_k + \gamma\tilde{v}_k^j \\ \tilde{T}_{amb,k} \\ \tilde{I}_{sol,k} \end{bmatrix} = f(\tilde{z}_k^j,\, \tilde{w}_k,\, \gamma,\, \tilde{v}_k^j) \tag{2}$$

The optimization problem (3) has a horizon of 36 h, i.e. $N = 144$. This covers the 24 h of the current day, the 6 h horizon of the low-level controller at the end of the day, plus an additional 6 h to prevent turnpike behaviors, i.e. sharp changes at the end of the optimized trajectory that would not be there if the optimization horizon were longer, at 30 h. As indicated above, the number of scenarios is $M = 10$. The cost function (3a) seeks to maximize the energy flexibility, while confining the reference trajectory to the range of power that the heat pump can actually deliver (3c), (3d), and respecting the zone temperature constraints (3e). For this purpose, we introduce the slack variables r and s. The hierarchy of these goals is enforced by the weights of 1 for the flexibility, 100 for the minimum heat pump power and 10^5 for the zone temperature (i.e. comfort) constraints. We choose a quadratic cost for the zone temperature, because it is not very sensitive to small violations, but the 1-norm for the minimum heat pump power, because its binary nature makes it much more sensitive to small violations. If the control command for the heat pump's compressor speed drops below its minimum threshold, it switches off. The optimization problem is completed by the dynamics (2) in (3b), and the constraints on the slack variable r (3f) and the flexibility band (3g). The minimum and maximum heat pump powers are $P_{HP,min} = 1.2\,\text{kW}$ and $P_{HP,max} = 3.5\,\text{kW}$. The former is chosen slightly above the 1.1 kW estimate in Section 2.1 to provide a small

safety margin. The temperature constraints are $T_{z,min} = 20$ °C and $T_{z,max} = 24$ °C. We note that all inputs, outputs and models are normalized to a mean of 0 and a standard deviation of 1. For legibility, this is not explicitly reflected in the equations.

$$\max_{w,z,\gamma,r,s} \gamma - 100\|r\|_1 - 10^5 \sum_{j=1}^{M} \|s^j\|_2^2 \tag{3a}$$

$s.t.$

$$z_{k+1}^j = f(\tilde{z}_k^j, \tilde{w}_k, \gamma, \tilde{v}_k^j) \tag{3b}$$

$w_{[-2:-1]}$, $\tilde{z}_0$ and v_k^j known

$$\gamma \leq w_k \leq P_{HP,max} - \gamma \tag{3c}$$
$$w_k - \gamma + r_k \geq P_{HP,min} \tag{3d}$$
$$T_{z,min} \leq z_k^j + s_k^j \leq T_{z,max} \tag{3e}$$
$$r_k \geq 0 \tag{3f}$$
$$\gamma \geq 0 \tag{3g}$$

$\forall k \in \{0, \dots, N-1\}$
$\forall j \in \{1, \dots, M\}$

2.6. Low-level controller

From the day-ahead planner, we get an optimal power trajectory P_{el}^*. Combined with the flexibility band γ and the predicted AGC signal a, we obtain the tracking reference $P_{el,k}^{ref} = P_{el,k}^* + \gamma a_k$ for the low-level MPC (4). A short horizon of 6 h (i.e. $N = 24$) is chosen for the low-level MPC to prevent excessive computation times with the nonlinear PWARX model. For the sake of comparison, the same horizon is used with the linear ARX model. The cost function (4a) minimizes the tracking error while respecting the soft output constraints (4d). The MPC is completed by the input constraints (4c) and the dynamics constraints (4b). If the ARX model is used, the dynamics constraints are $g(\tilde{y}_k, \tilde{u}_k) = A\tilde{y}_k + B\tilde{u}_k$. If the PWARX model is used, Eq. (4b) is replaced by Eq. (1), implemented using a big-M formulation, creating a mixed-integer problem.

$$\min_{u,y,\delta} \|P_{el} - P_{el}^{ref}\|_2^2 + 10^5 \|s\|_2^2 \tag{4a}$$

$s.t.$

$$y_{k+1} = g(\tilde{y}_k, \tilde{u}_k) \tag{4b}$$

$u_{[-2:-1]}$ and $\tilde{y}_0$ known

$$T_{z,min} \leq u_k \leq T_{z,max} \tag{4c}$$
$$\begin{pmatrix} T_{z,min} \\ 0 \end{pmatrix} \leq y_k + s_k \leq \begin{pmatrix} T_{z,max} \\ P_{HP,max} \end{pmatrix} \tag{4d}$$

$\forall k \in \{0, \dots, N-1\}$

3. Results

3.1. Open-loop system identification

Figs. 8 and 9 show the mean and 0.25/0.75 percentile of the absolute error of the prediction of the zone temperature and electric power over a horizon of one day, based on the entire open-loop data set. While distributions are commonly characterized by either the mean and standard deviation, or the median with the 0.25 and 0.75 percentile, we use the mean with the 0.25 and 0.75 percentile. The distributions of the absolute errors are non-Gaussian and, in case of the PWARX, zero-heavy, thus the standard deviation and median are both poor characterizations.. All four error trajectories start at zero, since they are initialized with the most recent measurements. The plots begin at 0.25 h, which is one time step into the future, and the shortest time-scale considered in our frequency control framework, as described in Sections 2.1 to 2.3. A few hours into the future, the trajectories converge toward

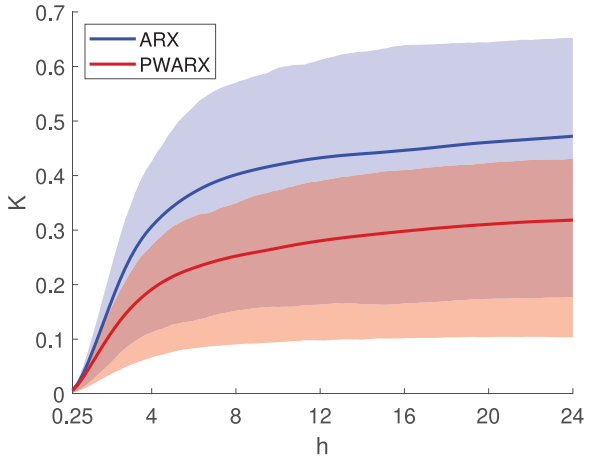

Fig. 8. Comparison of the mean absolute error and 0.25/0.75 percentile of the zone temperature over the prediction horizon.

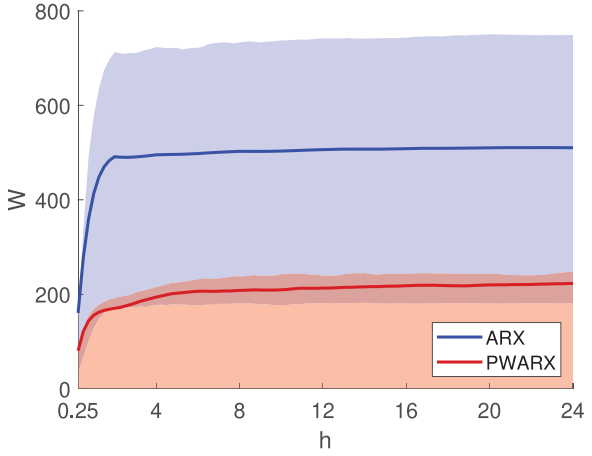

Fig. 9. Comparison of the mean absolute error and 0.25/0.75 percentile of the electric power over the prediction horizon.

Table 1
Mean of the mean absolute errors and 0.25/0.75 percentiles over the prediction horizon.

	ARX		PWARX	
	T_z	P_{el}	T_z	P_{el}
Unit	K	W	K	W
0.75 percentile	0.54	721	0.35	230
Mean	0.39	495	0.25	206
0.25 percentile	0.14	175	0.09	0

different values, similar to what has been observed in other studies on system identification for buildings [55,56]. The trajectories for the electric power converge much faster than the ones for the zone temperature, since the corresponding dynamics have much shorter time constants. The PWARX significantly outperforms the ARX across the horizon in both cases. We note that the mean and 0.75 percentile for the PWARX in Fig. 9 are very close, while the 0.25 percentile is zero, because all the correctly predicted zero-values (i.e. the HP is off) create a very zero-heavy distribution. The mean values over the horizon are summarized in Table 1.

Since we consider the prediction of the electric power to be more crucial for our application, we analyze it further: The distribution of the raw (i.e. non-absolute) one-step prediction error of the electric power in Fig. 10 shows a roughly Gaussian form for the ARX. For the PWARX, the distribution is significantly less Gaussian, with the aforementioned zero-heavy distribution clearly visible. Furthermore, there is a noticeable increase around an error of $-1\,\mathrm{kW}$, which stems from incorrect predictions that the heat pump will be off at the next time step. The absence of an equivalent increase around an error of $+1\,\mathrm{kW}$ reveals that the model is more likely to make a false-off prediction than a false-on. The presence of such a bias in the state-space partition is not surprising, considering we are fitting a linear classifier to data that is not guaranteed to be linear.

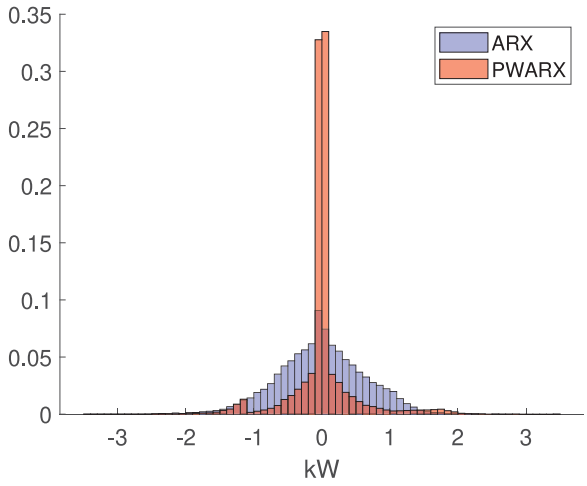

Fig. 10. Distribution of one-step prediction error of the electric power.

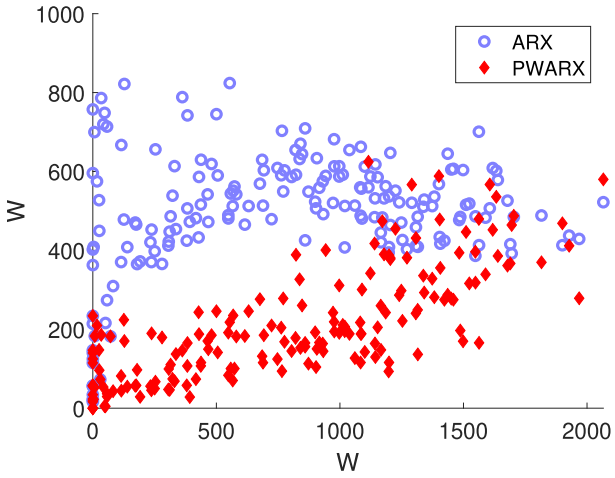

Fig. 11. Mean absolute prediction error of the electric power depending on the actual mean electric power during the prediction window.

Fig. 11 shows the mean absolute error of the prediction of the electric power, averaged over the horizon, depending on the mean of the actual heat pump power during that prediction window. A low power on the x-axis corresponds to warm weather, a high power to cold weather. We see that the PWARX has a significant advantage over the ARX for warm days, but not for cold days. This is unsurprising, since there is more heating during cold days, creating a stronger coupling between the setpoint temperature and the heating power. During warm days, ambient heat, solar and internal gains will frequently lift the zone temperature above the setpoint. This decoupling between the input and output variables is captured by the piecewise-affine modeling.

3.2. Closed-loop frequency control

The key performance indicators in Table 2 show very similar results for the frequency control with the ARX-based and the PWARX-based low-level MPCs, which stands in contrast to the significantly better prediction performance of the PWARX in Section 3.1. This can be explained by the findings from Fig. 11, since most of the frequency control action of the building happens during colder days. With both methods, the building participates in frequency control on around two-thirds of the simulation days, with most of the other days spent in default mode, when the flexibility band is below the threshold. Only one day with the ARX and two days with the PWARX are spent in identification mode, which means that the closed-loop system identification works quite

Table 2
Key performance indicators for ARX-based and PWARX-based low-level control. The symbols are defined in Section 3.2.

		Unit	ARX	PWARX
Mode	Active	Days	40	38
	Identification	Days	1	2
	Default	Days	19	20
Energy	E_{tot}	kWh	2114	2095
	γ_{mean}	W	634	639
	γ_{total}	kWh	624	598
	r_{flex}	%	59.0	57.1
Error	$NMAE$	%	12.2	13.7

reliably. The mean flexibility band γ_{mean} for days with active frequency control is similar with both methods, and the difference in the total flexibility offered during the simulation period γ_{total} stems from the number of active days, rather than the average flexibility. The ratios of total flexibility offered (two-sided) to total energy consumed $r_{flex} = 2\,\gamma_{total}/E_{tot}$ are also similar with 59.0% and 57.1%.

Grid operators contract providers of frequency control with mandatory tracking accuracy requirements, which are normalized to the magnitude of the flexibility band that the provider commits to. Correspondingly, we define the tracking error as the normalized mean absolute error (5), with the number of samples n. The ARX-based method slightly outperforms the PWARX-based method in this regard, with a NMAE of 12.2%, compared to 13.7%.

$$NMAE = \| P_{el} - P_{el}^{ref} \|_1 / (\gamma\, n) \tag{5}$$

While the key performance indicators are very similar, there is a stark difference in the computational costs. While the ARX-based MPC reliably solves in a fraction of a second, the PWARX-based MPC usually takes around one second, but occasionally several minutes on an 8-core CPU. The violations of the zone temperature constraints during active frequency control are negligible, which was expected due to the setpoint-based control method. Detailed time-series plots of the simulations are shown in Appendix A.

Fig. 12 compares the performance of the ARX-based method on the left, and the PWARX-based method on the right. The top plots provide insights into the quality of the identified models for the day-ahead planner and the low-level MPC, by showing the mean absolute errors of the zone temperature predictions based on the validation day, as introduced in Section 2.2. It can be seen that both models rarely exceed the permissible threshold. The middle plots show the same measure for the electric power prediction, which only happens on the lower control level. The bottom plots show the daily flexibility band γ, which falls short of the minimum threshold on around a quarter of the all days.

3.3. Estimating the impact of aggregation and battery storage

A grid operator may require a higher level of tracking accuracy than what is shown in our results. This may be achieved by adding battery storage to the building system, and using it to compensate the tracking error of the heat pump at a higher frequency, or by aggregating many buildings into a single node interacting with the grid. Aggregation allows to average out individual tracking errors, while also addressing the unwillingness of grid operators to engage with a large number of small consumers individually. The electric power of HVAC systems is typically between a few kW for single houses, and a few MW for commercial buildings, while power plants commonly fall into the GW range. A great variety of aggregation methods for buildings are described in the existing literature [57–59]. Refs. [57,58] report an exponential decrease of the prediction uncertainty with an increasing number of buildings.

To estimate the impact of a battery, its stand-by power consumption, and the offset of the tracking error over the whole simulation are compensated. Such a measure would have to be implemented in reality to prevent the battery from converging toward its empty state, where it can only provide error compensation in one direction. We assume a charging and discharging efficiency of $\eta = 95\%$ each and a maximum charging and discharging power of $P_{bat,max} = 500\,\mathrm{W/kWh}$, proportional to the battery capacity. To simulate an aggregation over multiple houses, the tracking error is randomly shifted, before being averaged. This implies the assumption that the tracking errors of multiple houses are not correlated in time over the simulation period. While we cannot guarantee the applicability of this assumption for a given aggregation district, we consider it to be an acceptable simplification for the purpose of this estimation. The results in this section are based on the simulation with the ARX-based low-level MPC. A separate analysis with the PWARX is omitted, since the results in Section 3.2 are found to be very similar.

Fig. 13 shows an estimate of the reduction of the NMAE as a function of the battery capacity in each house, and the number of aggregated houses, based on the simulation results with the ARX-based low-level MPC. We use normalized errors for this figure because of the requirements established in Section 3.2. Because of the offset removal, the baseline NMAE for one house with no battery at 12.1% is slightly lower than the unmodified value of 12.2% in Table 2. Following the x- and y-axes of Fig. 13, we see an exponential decrease of the tracking error of almost one order of magnitude in both directions. Aggregating over 25 houses reduces the NMAE about as much as installing a 6 kWh battery in each house, i.e. a total of 150 kWh of capacity. Considering the costs of battery storage, aggregation can therefore provide substantial financial savings for the same improvement in tracking accuracy.

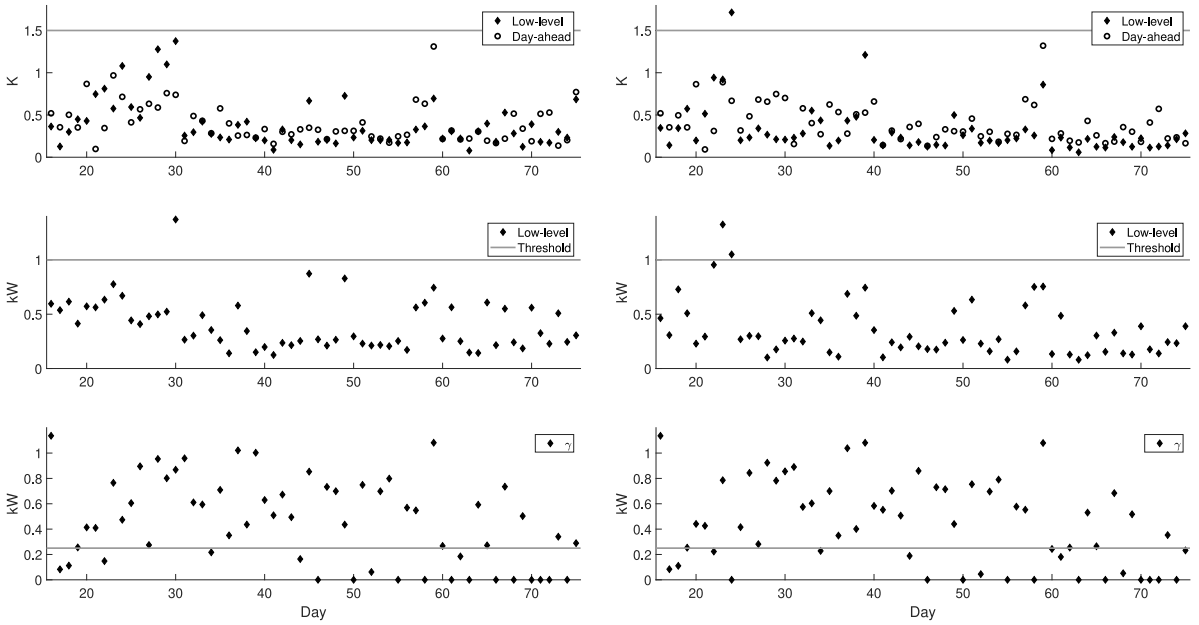

Fig. 12. System identification and flexibility bands with the ARX-based (left) and PWARX-based (right) methods. Top plots: Daily mean absolute error of the zone temperature prediction from the model validation for the day-ahead (physics) and the low-level (inverse) ARX models, compared to the maximum threshold. Middle plots: The equivalent value for the heat pump power prediction. Only applies to the inverse model. Bottom plots: Daily flexibility bands γ for the frequency control, compared to the minimum threshold.

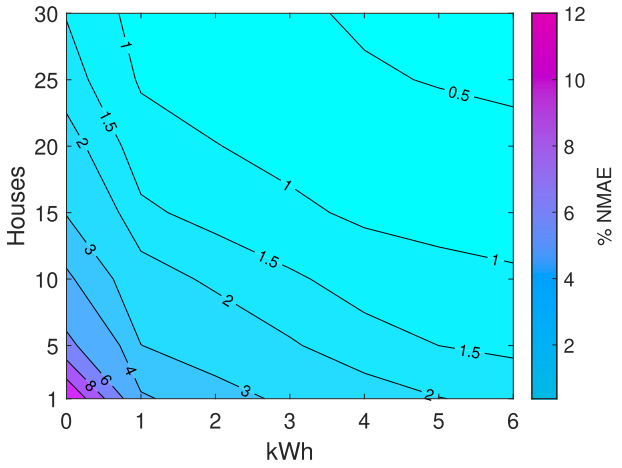

Fig. 13. NMAE as a function of the battery capacity in each house, and the number of aggregated houses.

Fig. 14 shows the reduction of the NMAE, along with its 0.25/0.75 percentile, and the normalized maximum absolute tracking error, as a function of aggregation only, but up to 10^4 houses. It requires an averaging over 10^3 houses to reduce the NMAE to around 1%, and 10^4 houses to reduce the normalized maximum absolute error to just below 2%. Such levels of accuracy are required by some grid operators. We note that the highly uneven distribution of the errors results in the mean being marginally higher than the 0.75 percentile at one house.

4. Discussion

While the simulation model is fairly detailed, there are still simplifications, compared to a real building. The model only has one thermal zone. The building occupants are modeled as a repetitive time-series of internal gains, without any of the hard-to-predict,

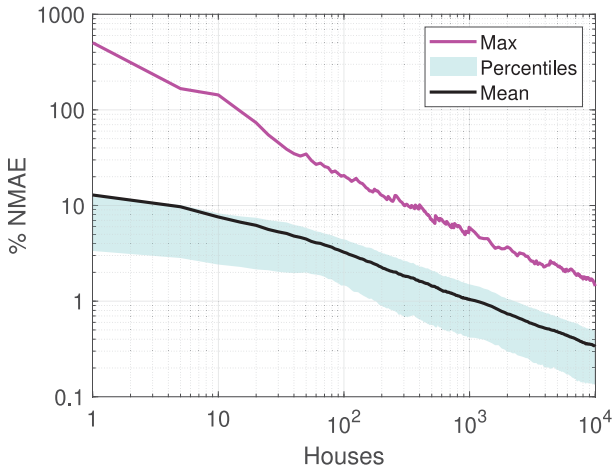

Fig. 14. NMAE, 0.25/0.75 percentile, and maximum error as a function of the number of aggregated houses.

random behaviors real humans may display, like opening windows or placing an unintended source of heat near the thermostat. There are no ageing effects or changes of use that could alter the thermodynamic properties of the building. We also use idealized forecasting functions for the weather and AGC signal with no uncertainties. Estimating the impact of these simplifications on the results is no trivial task. Moreover, another building with a different shape, construction style, local climate, HVAC system or utilization could yield different results. We thus recommend a continuation of this line of research with both different simulation models and real buildings.

Furthermore, different identification and control methods may yet result in superior control performance or faster computation. Acquiring well-excited identification data in an occupied building while keeping occupant discomfort at a minimum can be a challenge. Possible avenues of improvement include measuring the current occupancy to use unoccupied times for more aggressive excitation [28], or the acquisition of a much larger data set, in combination with various data mining techniques [60] to select high-quality data for identification. Since long-term prediction accuracy is desirable for day-ahead planning, model optimization techniques minimizing the corresponding error should also be considered [61].

In the spirit of the commissioning-oriented approach, we would like to minimize the required computational power for the controllers. Ideally, a small single-board computer should be able to solve the optimizations in real-time. However, the scenario-approach and the PWARX make it quite computationally expensive. We already made two compromises in this regard: Using a linear classifier to partition the state space for the PWARX model, when the data is not guaranteed to be linear, and using a short horizon for the low-level controller, when a longer horizon might yield better performance. Without the restrictions on training data and computational power we imposed for this study, it might be possible to improve the control performance by using shorter time steps, longer horizons or more complex nonlinear models. In large commercial buildings, an adequate logging of operational data may already be in place, and the costs of a high-performance CPU would be less significant, compared to the energy costs. In smaller buildings, these conditions may become more common in the future, as smart home devices are increasingly popular. This would also enable moving the data storage and computation to a cloud service, which would reduce the hardware costs at the building level. Lastly, we will point out that there are still ongoing advancements in the efficiency of optimization algorithms [62], which we find particularly interesting for small-scale building applications for the aforementioned reasons.

5. Conclusions

We compare the performance of two control-oriented model structures for predicting zone temperature and heat pump power as a function of thermostat setpoint in a building simulation. In open-loop prediction, the PWARX models are significantly more accurate than the ARX models. The difference is most pronounced on warmer days with less heating, but shrinks significantly on colder days with high heating power. However, the superior prediction accuracy does not translate to superior closed-loop frequency control performance. This is likely due to the fact that the PWARX model has its biggest advantage in warm weather conditions, when the potential for frequency control is smaller. Nevertheless, substantial energy flexibility for secondary frequency control is provided with either method, albeit not at the level of tracking accuracy currently required by the Swiss national grid operator. Despite its challenges, we believe that this topic merits further investigation, for the reasons introduced at the beginning.

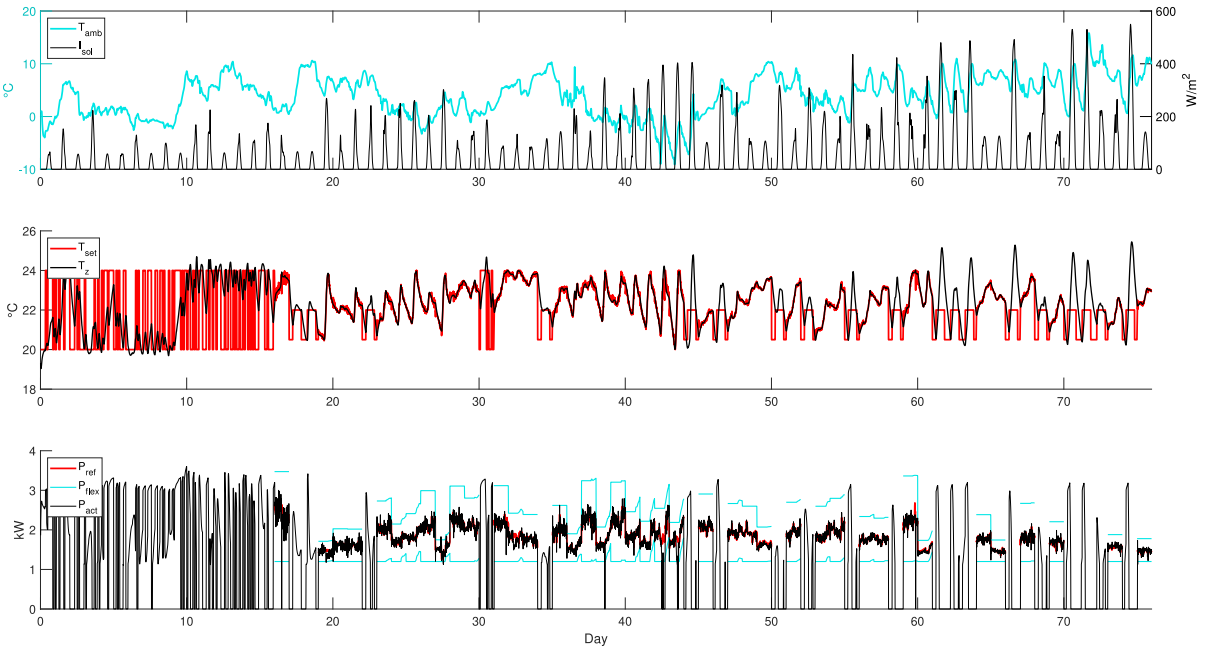

Fig. A.15. Full time-series of frequency control simulation with ARX-based low-level MPC, including initial system identification.

CRediT authorship contribution statement

Manuel Koch: Writing – original draft, Visualization, Software, Methodology, Formal analysis, Conceptualization. **Parantapa Sawant:** Writing – review & editing, Supervision, Methodology, Project administration. **Ralph Eismann:** Funding acquisition, Project administration, Supervision, Writing – review & editing, Methodology. **Colin N. Jones:** Writing – review & editing, Supervision, Project administration, Funding acquisition, Conceptualization, Methodology.

Funding

This work has received no external funding.

Declaration of competing interest

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

Appendix A. Time-series plots of frequency control

Figs. A.15 and A.16 show the complete time-series for the initial system identification phase and the subsequent frequency control phase for the simulations with the ARX-based and PWARX-based low-level MPCs, including the weather conditions in the top plots, the zone temperature and setpoint in the middle plot, and the frequency control in the bottom plot. The weather is typical for winter in Central Europe, being cold and dark, but without any extreme conditions. The temperature setpoint is tracked well, except for sharp changes, because of the high inertia of the radiant slab. The days on which the control system does not participate in the frequency control due to an insufficient flexibility band, but goes into default mode, mostly occur during warmer weather, often in an alternating pattern with days with active frequency control. This is unsurprising, since the best days for frequency control are when the average daily power of the heat pump is centered between $P_{HP,min}$ and $P_{HP,max}$, which is around two-thirds of the maximum power. The alternating default days allow the building to cool down, which increases the required heating power for the next day. The extracts of the same time window with both methods in Figs. A.17 and A.18 show just how similarly they perform at times. We note that the weather conditions, internal gains and AGC signal are identical for both figures, contributing to the similar results.

Appendix B. Supplementary plots

See Figs. B.19–B.21.

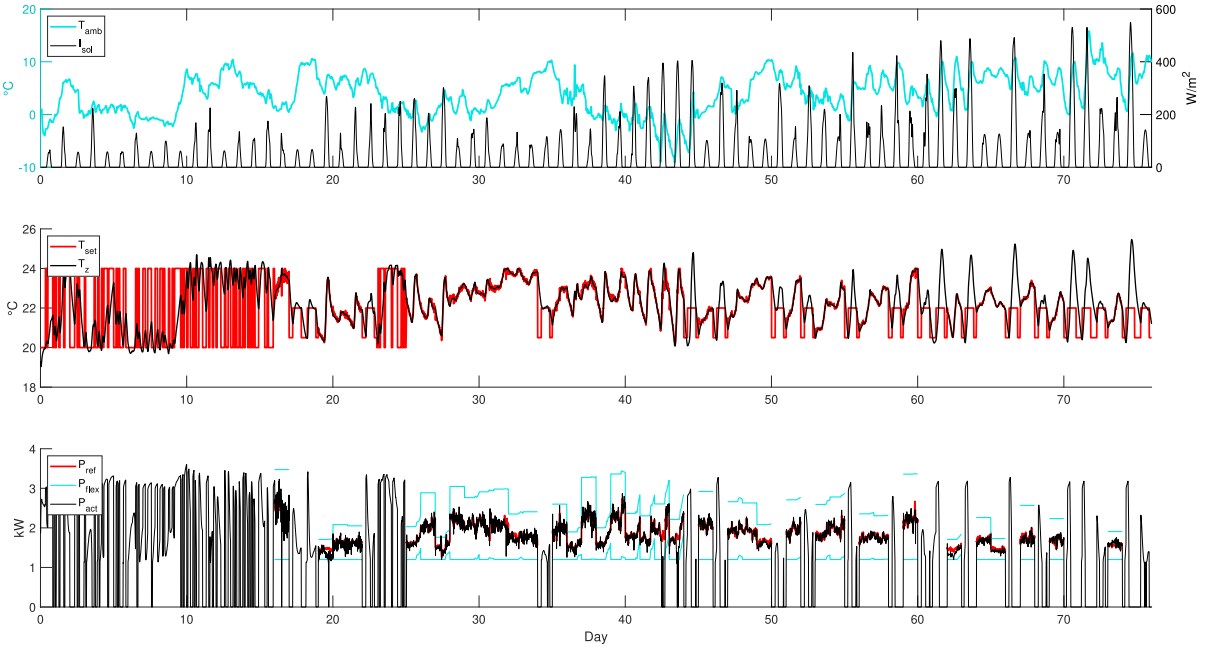

Fig. A.16. Full time-series of frequency control simulation with PWARX-based low-level MPC, including initial system identification.

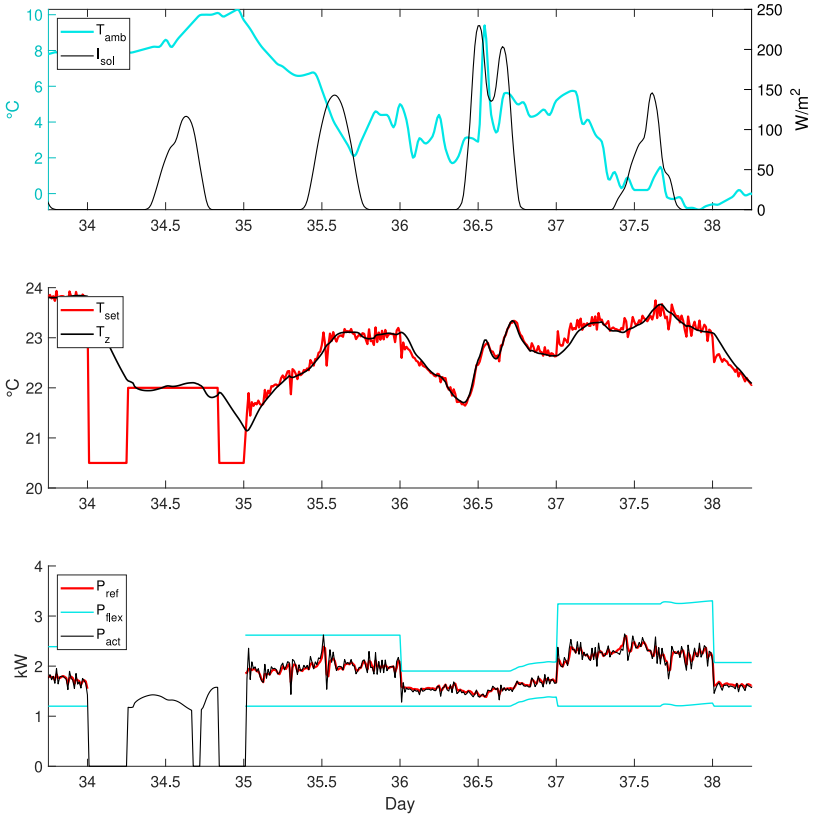

Fig. A.17. Extract of time-series with ARX-based low-level MPC.

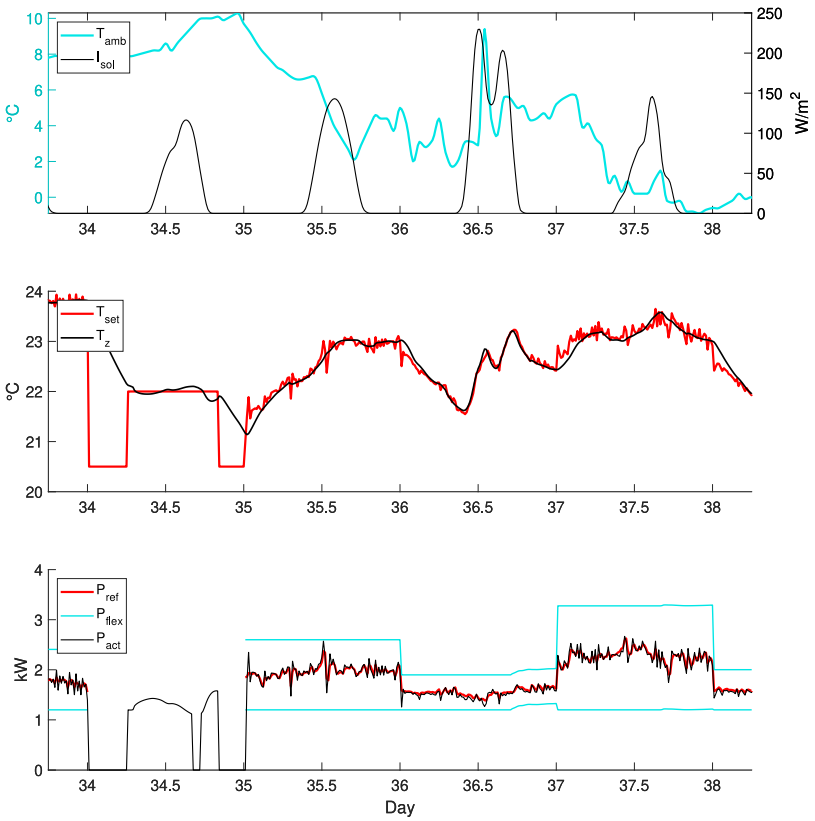

Fig. A.18. Extract of time-series with PWARX-based low-level MPC.

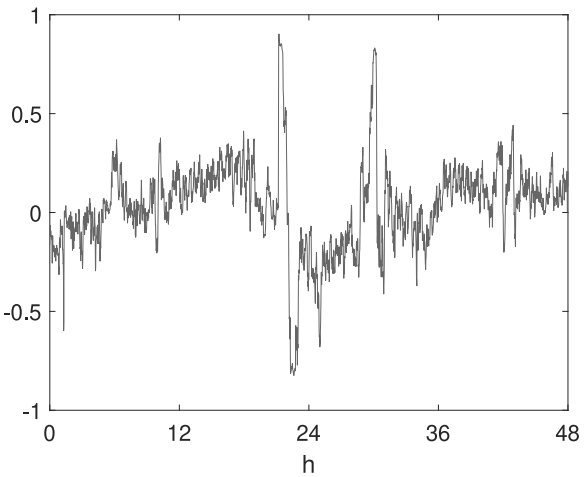

Fig. B.19. Two-day sample of AGC signal.

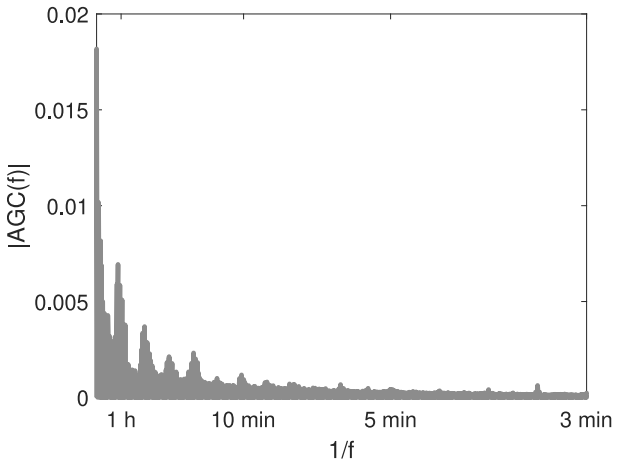

Fig. B.20. Spectral decomposition of AGC signal.

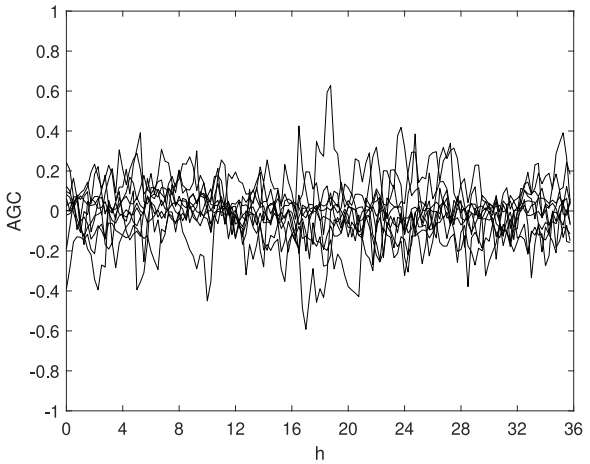

Fig. B.21. The ten AGC scenarios for the day-ahead planner.

Data availability

We would have to obtain permission before sharing the AGC signal dataset. The building model and weather data are publicly available.

References

[1] L.P.C. Do, K.-H. Lin, P. Molnár, Electricity consumption modelling: A case of Germany, Econ. Model. 55 (2016) 92–101, http://dx.doi.org/10.1016/j.econmod.2016.02.010.

[2] Our World in Data, Electricity Mix, https://ourworldindata.org/electricity-mix.

[3] Bundesamt für Energie, Energiestrategie 2050, https://www.bfe.admin.ch/bfe/de/home/politik/energiestrategie-2050.html.

[4] European Comission, A European Green Deal, https://commission.europa.eu/strategy-and-policy/priorities-2019-2024/european-green-deal_en.

[5] BloombergNEF, Lithium-ion battery pack prices rise for first time to an average of $151/kWh, 2022, https://about.bnef.com/blog/lithium-ion-battery-pack-prices-rise-for-first-time-to-an-average-of-151-kwh/.

[6] Lazard, 2023 levelized cost of energy, 2023, https://www.lazard.com/research-insights/2023-levelized-cost-of-energyplus/.

[7] M. Fichtner, K. Edström, E. Ayerbe, M. Berecibar, A. Bhowmik, I.E. Castelli, S. Clark, R. Dominko, M. Erakca, A.A. Franco, A. Grimaud, B. Horstmann, A. Latz, H. Lorrmann, M. Meeus, R. Narayan, F. Pammer, J. Ruhland, H. Stein, T. Vegge, M. Weil, Rechargeable batteries of the future—the state of the art from a BATTery 2030+ perspective, Adv. Energy Mater. 12 (17) (2022) 2102904, http://dx.doi.org/10.1002/aenm.202102904.

[8] P. Grünewald, J. Torriti, Demand response from the non-domestic sector: early UK experiences and future opportunities, Energy Policy 61 (2013) 423–429, http://dx.doi.org/10.1016/j.enpol.2013.06.051.

[9] H.C. Gils, Economic potential for future demand response in Germany – Modeling approach and case study, Appl. Energy 162 (2016) 401–415, http://dx.doi.org/10.1016/j.apenergy.2015.10.083.

[10] Swissgrid, Principles of ancillary services products, 2022, https://www.swissgrid.ch/dam/swissgrid/customers/topics/ancillary-services/as-documents/D220824-AS-Products-V19-en.pdf.
[11] Enerdata, Residential buildings: energy efficiency & consumption evolution in Europe, 2021, https://www.enerdata.net/publications/executive-briefing/households-energy-efficiency.html.
[12] Bundesamt für Energie, Energieverbrauch nach Verwendungszweck, https://www.bfe.admin.ch/bfe/de/home/versorgung/statistik-und-geodaten/energiestatistiken/energieverbrauch-nach-verwendungszweck.html.
[13] US Energy Information Administration, U.S. Energy Facts Explained - Consumption and Production, https://www.eia.gov/energyexplained/us-energy-facts/.
[14] US Energy Information Administration, Use of Energy in Commercial Buildings in Depth, https://www.eia.gov/energyexplained/use-of-energy/commercial-buildings-in-depth.php.
[15] US Energy Information Administration, Use of Energy in Homes, https://www.eia.gov/energyexplained/use-of-energy/homes.php.
[16] O. Abedrabboh, M. Koç, Y. Biçer, Sustainability performance of space-cooling technologies and approaches, Energy Sources Part A: Recover. Util. Environ. Eff. 44 (2022).
[17] E. Elnagar, S. Pezzutto, B. Duplessis, T. Fontenaille, V. Lemort, A comprehensive scouting of space cooling technologies in Europe: Key characteristics and development trends, Renew. Sustain. Energy Rev. 186 (2023) 113636, http://dx.doi.org/10.1016/j.rser.2023.113636.
[18] UK Department of Energy and Climate Change, The future of heating: meeting the challenge, 2013, https://www.gov.uk/government/publications/the-future-of-heating-meeting-the-challenge.
[19] Bundesamt für Statistik, Bau- und Wohnungswesen, Gebäude, Energiebereich, https://www.bfs.admin.ch/bfs/de/home/statistiken/bau-wohnungswesen/gebaeude/energiebereich.html.
[20] International Energy Agency, The future of heat pumps – analysis, 2022, https://www.iea.org/reports/the-future-of-heat-pumps.
[21] B. Parrish, P. Heptonstall, R. Gross, B.K. Sovacool, A systematic review of motivations, enablers and barriers for consumer engagement with residential demand response, Energy Policy 138 (2020) 111221, http://dx.doi.org/10.1016/j.enpol.2019.111221.
[22] Y. Chen, P. Xu, J. Gu, F. Schmidt, W. Li, Measures to improve energy demand flexibility in buildings for demand response (DR): A review, Energy Build. 177 (2018) 125–139, http://dx.doi.org/10.1016/j.enbuild.2018.08.003.
[23] F. Pallonetto, M. De Rosa, F. D'Ettorre, D.P. Finn, On the assessment and control optimisation of demand response programs in residential buildings, Renew. Sustain. Energy Rev. 127 (2020) 109861, http://dx.doi.org/10.1016/j.rser.2020.109861.
[24] H. Li, Z. Wang, T. Hong, M.A. Piette, Energy flexibility of residential buildings: A systematic review of characterization and quantification methods and applications, Adv. Appl. Energy 3 (2021) 100054, http://dx.doi.org/10.1016/j.adapen.2021.100054.
[25] H. Wang, S. Wang, R. Tang, Development of grid-responsive buildings: Opportunities, challenges, capabilities and applications of HVAC systems in non-residential buildings in providing ancillary services by fast demand responses to smart grids, Appl. Energy 250 (2019) 697–712, http://dx.doi.org/10.1016/j.apenergy.2019.04.159.
[26] Z.E. Lee, Q. Sun, Z. Ma, J. Wang, J.S. MacDonald, K. Max Zhang, Providing grid services with heat pumps: a review, ASME J. Eng. Sustain. Build. Cities 1 (1) (2020) http://dx.doi.org/10.1115/1.4045819.
[27] V. Chinde, Y. Lin, M.J. Ellis, Data-enabled predictive control for building HVAC systems, J. Dyn. Syst. Meas. Control. http://dx.doi.org/10.1115/1.4054314.
[28] H. Wolisz, T.M. Kull, D. Müller, J. Kurnitski, Self-learning model predictive control for dynamic activation of structural thermal mass in residential buildings, Energy Build. 207 (2020) 109542, http://dx.doi.org/10.1016/j.enbuild.2019.109542.
[29] F. D'Ettorre, M. Banaei, R. Ebrahimy, S.A. Pourmousavi, E.M.V. Blomgren, J. Kowalski, Z. Bohdanowicz, B. Łopaciuk-Gonczaryk, C. Biele, H. Madsen, Exploiting demand-side flexibility: state-of-the-art, open issues and social perspective, Renew. Sustain. Energy Rev. 165 (2022) 112605, http://dx.doi.org/10.1016/j.rser.2022.112605.
[30] M. Killian, M. Kozek, Ten questions concerning model predictive control for energy efficient buildings, Build. Environ. 105 (2016) 403–412, http://dx.doi.org/10.1016/j.buildenv.2016.05.034.
[31] E.T. Maddalena, Y. Lian, C.N. Jones, Data-driven methods for building control — A review and promising future directions, Control Eng. Pract. 95 (2020) 104211, http://dx.doi.org/10.1016/j.conengprac.2019.104211.
[32] G. Serale, M. Fiorentini, A. Capozzoli, D. Bernardini, A. Bemporad, Model predictive control (MPC) for enhancing building and HVAC system energy efficiency: problem formulation, applications and opportunities, Energies 11 (3) (2018) 631, http://dx.doi.org/10.3390/en11030631.
[33] J. Arroyo, F. Spiessens, L. Helsen, Comparison of optimal control techniques for building energy management, Front. Built Environ. 8 (2022) http://dx.doi.org/10.3389/fbuil.2022.849754.
[34] L. Di Natale, Y. Lian, E.T. Maddalena, J. Shi, C.N. Jones, Lessons learned from data-driven building control experiments: contrasting Gaussian process-based MPC, bilevel DeePC, and deep reinforcement learning, 2022, http://dx.doi.org/10.48550/arXiv.2205.15703, arXiv.
[35] F. Bünning, B. Huber, A. Schalbetter, A. Aboudonia, M. Hudoba de Badyn, P. Heer, R.S. Smith, J. Lygeros, Physics-informed linear regression is competitive with two machine learning methods in residential building MPC, Appl. Energy 310 (2022) 118491, http://dx.doi.org/10.1016/j.apenergy.2021.118491.
[36] B. Ajib, S. Lefteriu, A. Caucheteux, S. Lecoeuche, Predicting the air temperature of a building zone by detecting different configurations using a switched system identification technique, J. Build. Eng. 31 (2020) 100995, http://dx.doi.org/10.1016/j.jobe.2019.100995.
[37] M.H. Benzaama, L.H. Rajaoarisoa, B. Ajib, S. Lecoeuche, A data-driven methodology to predict thermal behavior of residential buildings using piecewise linear models, J. Build. Eng. 32 (2020) 101523, http://dx.doi.org/10.1016/j.jobe.2020.101523.
[38] G.C. Chasparis, T. Natschläger, Supervisory system identification for bilinear systems with application to thermal dynamics in buildings, in: 2014 IEEE International Symposium on Intelligent Control, ISIC, 2014, pp. 832–837, http://dx.doi.org/10.1109/ISIC.2014.6967608.
[39] G.C. Chasparis, T. Natschläger, Supervisory output prediction for bilinear systems by reinforcement learning, IET Control Theory Appl. 11 (10) (2017) 1514–1521, http://dx.doi.org/10.1049/iet-cta.2016.1400.
[40] M. Killian, B. Mayer, M. Kozek, Effective fuzzy black-box modeling for building heating dynamics, Energy Build. 96 (2015) 175–186, http://dx.doi.org/10.1016/j.enbuild.2015.02.057.
[41] M. Killian, B. Mayer, M. Kozek, Cooperative fuzzy model predictive control for heating and cooling of buildings, Energy Build. 112 (2016) 130–140, http://dx.doi.org/10.1016/j.enbuild.2015.12.017.
[42] B. Mayer, M. Killian, M. Kozek, Hierarchical model predictive control for sustainable building automation, Sustainability 9 (2) (2017) 264, http://dx.doi.org/10.3390/su9020264.
[43] P. Fazenda, P. Lima, P. Carreira, Context-based thermodynamic modeling of buildings spaces, Energy Build. 124 (2016) 164–177, http://dx.doi.org/10.1016/j.enbuild.2016.04.068.
[44] E.M. Burger, H.E. Perez, S.J. Moura, Piecewise linear thermal model and recursive parameter estimation of a residential heating system, Escholarship (2016).
[45] T.T. Gorecki, L. Fabietti, F.A. Qureshi, C.N. Jones, Experimental demonstration of buildings providing frequency regulation services in the Swiss market, Energy Build. 144 (2017) 229–240, http://dx.doi.org/10.1016/j.enbuild.2017.02.050.
[46] B. Ning, S. Schiavon, F.S. Bauman, A classification scheme for radiant systems based on thermal timeconstant, 2015, URL https://escholarship.org/uc/item/1sx88662.
[47] C. Vallianos, J. Candanedo, A. Athienitis, Thermal modeling for control applications of 60,000 homes in North America using smart thermostat data, Energy Build. 303 (2024) 113811, http://dx.doi.org/10.1016/j.enbuild.2023.113811.

[48] M. Schmelas, T. Feldmann, E. Bollin, Adaptive predictive control of thermo-active building systems (TABS) based on a multiple regression algorithm, Energy Build. 103 (2015) 14–28, http://dx.doi.org/10.1016/j.enbuild.2015.06.012.
[49] SwissGrid: Balancing, https://www.swissgrid.ch/en/home/operation/grid-data/balance.html.
[50] M.B. Saltık, L. Özkan, J.H.A. Ludlage, S. Weiland, P.M.J. Van den Hof, An outlook on robust model predictive control algorithms: Reflections on performance and computational aspects, J. Process Control 61 (2018) 77–102, http://dx.doi.org/10.1016/j.jprocont.2017.10.006.
[51] G.C. Calafiore, L. Fagiano, Robust model predictive control via scenario optimization, IEEE Trans. Autom. Control 58 (1) (2013) 219–224, http://dx.doi.org/10.1109/TAC.2012.2203054.
[52] A. Parisio, M. Molinari, D. Varagnolo, K.H. Johansson, A scenario-based predictive control approach to building HVAC management systems, in: 2013 IEEE International Conference on Automation Science and Engineering, CASE, 2013, pp. 428–435, http://dx.doi.org/10.1109/CoASE.2013.6654024.
[53] X. Zhang, K. Margellos, P. Goulart, J. Lygeros, Stochastic model predictive control using a combination of randomized and robust optimization, in: 52nd IEEE Conference on Decision and Control, 2013, pp. 7740–7745, http://dx.doi.org/10.1109/CDC.2013.6761118.
[54] X. Zhang, G. Schildbach, D. Sturzenegger, M. Morari, Scenario-based MPC for energy-efficient building climate control under weather and occupancy uncertainty, in: 2013 European Control Conference, ECC, 2013, pp. 1029–1034, http://dx.doi.org/10.23919/ECC.2013.6669664.
[55] L. Di Natale, B. Svetozarevic, P. Heer, C.N. Jones, Physically consistent neural networks for building thermal modeling: Theory and analysis, Appl. Energy 325 (2022) 119806, http://dx.doi.org/10.1016/j.apenergy.2022.119806.
[56] M. Koch, C.N. Jones, Comparison of behavioral systems theory and conventional linear models for predicting building zone temperature in long-term in situ measurements, 2023, http://dx.doi.org/10.48550/arXiv.2302.04063, arXiv:2302.04063.
[57] A. Wang, R. Li, S. You, Development of a data driven approach to explore the energy flexibility potential of building clusters, Appl. Energy 232 (2018) 89–100, http://dx.doi.org/10.1016/j.apenergy.2018.09.187.
[58] M. Hu, F. Xiao, Quantifying uncertainty in the aggregate energy flexibility of high-rise residential building clusters considering stochastic occupancy and occupant behavior, Energy 194 (2020) 116838, http://dx.doi.org/10.1016/j.energy.2019.116838.
[59] F. Plaum, R. Ahmadiahangar, A. Rosin, J. Kilter, Aggregated demand-side energy flexibility: A comprehensive review on characterization, forecasting and market prospects, Energy Rep. 8 (2022) 9344–9362, http://dx.doi.org/10.1016/j.egyr.2022.07.038.
[60] Y. Zhao, C. Zhang, Y. Zhang, Z. Wang, J. Li, A review of data mining technologies in building energy systems: Load prediction, pattern identification, fault detection and diagnosis, Energy Built Environ. 1 (2) (2020) 149–164, http://dx.doi.org/10.1016/j.enbenv.2019.11.003.
[61] L. Di Natale, M. Zakwan, P. Heer, G.F. Trecate, C.N. Jones, SIMBa: system identification methods leveraging backpropagation, 2023, http://dx.doi.org/10.48550/arXiv.2311.13889, arXiv:2311.13889.
[62] R. Schwan, Y. Jiang, D. Kuhn, C.N. Jones, PIQP: a proximal interior-point quadratic programming solver, 2023, http://dx.doi.org/10.48550/arXiv.2304.00290, arXiv:2304.00290.