Building thermal control: Hierarchical design from limited data using gray-box or black-box internal models for model predictive control

Yuqi Liu, Pauline Kergus, Fabien Claveau, Philippe Chevrel, Bruno Lacarrière

Journal of Building Engineering 110 (2025) 112936

Contents lists available at ScienceDirect

Journal of Building Engineering

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

Building thermal control: Hierarchical design from limited data using gray-box or black-box internal models for model predictive control

Yuqi Liu [a,b,*], Pauline Kergus [c], Fabien Claveau [a,b], Philippe Chevrel [a,b], Bruno Lacarrière [d,e]

[a] IMT Atlantique - Département Automatique, Productique et Informatique, 4, rue Alfred Kastler, Nantes, 44307, France
[b] Laboratoire des Sciences du Numérique de Nantes, CNRS, 2, rue de la Houssinière, Nantes, 44322, France
[c] LAPLACE, Université de Toulouse, CNRS, INPT, UPS, 2 rue Charles Camichel, Toulouse, 31000, France
[d] IMT Atlantique - Département Systèmes Energétiques et Environnement, 4, rue Alfred Kastler, Nantes, 44307, France
[e] Laboratoire de génie des procédés, environnement, agroalimentaire, CNRS, 3 rue du Maréchal Joffre, Nantes, 44041, France

ARTICLE INFO

Keywords:
Building thermal management
MPC
Gray-box modeling
VRFT
Hierarchical control

ABSTRACT

This paper presents a hierarchical model predictive control (MPC) framework designed to accommodate the diversity of heating equipment in building energy systems. The proposed architecture consists of a supervisory MPC for power planning and tracking controllers for device-level regulation. The framework systematically compares gray-box system identification, using an equivalent Resistance–Capacitance (RC) model, against black-box Subspace State Space Identification (4SID), evaluating their performance with limited datasets. At the tracking level, Virtual Reference Feedback Tuning (VRFT) offers a model-free approach to equipment control, eliminating the need for a detailed Heating, Ventilation, Air Conditioning system model. The proposed approach is validated using a multi-zone residential test case from the Building Optimization Testing Framework. Results show that although both models achieve similar open-loop prediction accuracy, their performance diverges under closed-loop control. Under different weighting configurations, the proposed scheme using the RC model achieves approximately a 3%–12% reduction in energy consumption while maintaining comparable or even lower levels of thermal discomfort compared to the black-box 4SID model. The study concludes with a discussion of practical considerations and the potential for broader deployment.

1. Introduction

Reports from the International Energy Agency show that buildings account for 30%–40% of energy consumption [1]. Within building operations, thermal management systems, encompassing Heating, Ventilation and Air Conditioning (HVAC) services, constitute approximately one-third of this energy utilization [2]. The Energy Performance of Buildings Directive thus introduces the concept of "Smart Buildings" and highlights thermal management efficiency as an important indicator of energy performance [3]. This regulatory evolution necessitates the development of sophisticated thermal management strategies that can effectively balance occupant comfort with energy efficiency objectives.

* Corresponding author at: IMT Atlantique - Département Automatique, Productique et Informatique, 4, rue Alfred Kastler, Nantes, 44307, France.
E-mail address: yuqi.liu@imt-atlantique.fr (Y. Liu).

https://doi.org/10.1016/j.jobe.2025.112936
Received 10 January 2025; Received in revised form 24 April 2025; Accepted 15 May 2025
Available online 2 June 2025

Nomenclature

Symbol	Description	Unit	Symbol	Description	Unit
System and state variables					
T_i	Zone i temperature	°C	y	Output vector	°C
P_i	Zone i power input	W	u	Input vector	W
G	Solar radiation	W/m^2	d	Disturbance vector	–
T_{ext}	External temperature	°C	$\bar{y}$	Measured output	°C
v_i	Valve opening i	–	$\bar{u}$	Measured input	W
$\bar{d}$	Measured disturbance	–	$\hat{y}$	Predicted output	°C
$\hat{d}$	Forecast disturbance	–	$\bar{e}$	Modeling error	–
e	Temperature error	°C	e_u	Power tracking error	W
ε	Controller output error	–	n	Number of zones	–
Model parameters					
$R_{i,j}$	Inter-zone resistance	K/W	$R_{i,e}$	External resistance	K/W
C_i	Zone capacitance	J/K	S_i	Solar surface	m^2
N_{all}	Set of all zones	–	$N_{ad,i}$	Adjacent zones to i	–
$Ad_{i,j}$	Adjacency indicator	–	θ	Parameter vector	–
θ^*	Optimal parameters	–	n_{id}	Data points count	–
Time parameters					
t_k	Time step k	s	$T_{dataset}$	Sampling time	s
Δt	Derivative step	s	T_{MPC}	MPC time step	s
T_{PID}	PID time step	s	H	Prediction horizon	h
System matrices and functions					
$A(\theta)$	RC state matrix	–	A_{4SID}	4SID state matrix	–
$B_1(\theta)$	RC input matrix	–	$B_{1,4SID}$	4SID input matrix	–
$B_2(\theta)$	RC disturbance matrix	–	$B_{2,4SID}$	4SID disturbance matrix	–
f_{RC}	RC state function	–	f_d	Discrete model	–
$f^*_{d,RC}$	Optimal RC model	–	$f_{d,4SID}$	4SID model	–
Control parameters					
$K_{P,i}$	Proportional gain	–	$K_{I,i}$	Integral gain	–
$K_{D,i}$	Derivative gain	–	$K_{aw,i}$	Anti-windup gain	–
n_H	Prediction steps	–	τ	Model time constant	s
τ_{act}	Actual time constant	s	u_{ref}	Power reference	W
$u_{ref,vr}$	Virtual reference	W	y_{ref}	Temperature reference	°C
M_{ref}	Reference transfer fn	–	$M_{ref,i}$	Zone i ref model	–
$M^{\dagger}_{ref}$	Model left inverse	–	C_{PID}	PID transfer fn	–
C_{aw}	Anti-windup fn	–	θ_{PID}	PID parameters	–
θ_{aw}	Anti-windup params	–	A_{in}	Internal loop matrix	–
A^*_{in}	Optimal loop matrix	–	D_{in}	Training set	–
MPC and Performance parameters					
Q	State error weight	–	R	Input weight	–
w	Weight scalar	–	w_0	Base weight	–
$\underline{u}$	Lower power limit	W	$\bar{u}$	Upper power limit	W
S_{all}	Building surface	m^2	n_{data}	Data point count	–
q	Shift operator	–	β_0, β_1	Model coefficients	–
a	Power increase	W	$T_{set,i}$	Temperature setpoint	°C
$I_{pred,i}$	Zone i prediction quality	°C^2	I_{dis}	Discomfort metric	K h/zone
I_{ener}	Energy consumption	kW h/m^2			

The implementation of Model Predictive Control (MPC) in building thermal management has demonstrated significant potential for optimizing system performance. MPC leverages predictions of a building's climate evolution, allowing systems to use the building's thermal inertia effectively. Studies have shown that MPC outperforms traditional Rule-Based Control (RBC) strategies in multi-objective planning [4,5]. Documented performance improvements include energy consumption reductions ranging from 15%–50% while maintaining thermal comfort [6–9].

The performance of MPC is fundamentally linked to the accuracy of their underlying prediction models. Notable contributions [10–12] developed precise simulation models for the entire building energy system in Modelica for MPC. In contrast, other investigations [13–17] adopted data-based modeling solutions to reduce modeling effort. In practice, two challenges emerge: (1) With the diversity of building thermal equipment, most control design methodologies is application-specific. Difficulty arises in the

development and application of the complete thermal management solutions. (2) Obtaining open-loop data on the building's thermal response can be expensive and time-consuming, which discourages the application of data-driven methods.

1.1. Equipment diversity

HVAC systems encompass diverse solutions to maintain indoor thermal conditions and air quality. In large-scale buildings, this often involves Air Handling Unit and Variable Air Volume systems for centralized climate control. For residential and smaller commercial applications, conventional solutions include hydraulic circulation with radiators, heat pumps, and thermal energy storage systems. The diversity of HVAC equipment and their control requirements presents a significant challenge in developing unified management strategies. This challenge stems from the complex and time-consuming nature of HVAC system modeling, which varies significantly with equipment mechanisms, thermal inertia, usage patterns, and circuit topology. Moreover, these systems typically exhibit nonlinear behaviors and are subject to operational constraints that further complicate their control.

Previous research has explored various thermal control methodologies. For instance, [18] built an environment with indoor climate and HVAC models for hierarchical reinforcement learning. [19] used data generated from physical model to train data-based models for adaptive MPC. Although these approaches provide valuable insights, they typically result in system-specific solutions. This specificity necessitates the development of unique device models for each HVAC configuration, which is a barrier to widespread implementation.

In contrast to the methods mentioned above, which seek to model the thermal inertia of buildings with HVAC actuators, a number of researches assume an ideal response of HVAC to prevent engaging with its modeling [20–24]. These works reformulate the control problem to operate directly with heating and cooling power as control variables, abstracting away the underlying actuation mechanisms such as valves, dampers in system like radiators, air conditioners. These approaches maintains independence from the particular type of HVAC system installed, offering a more generalized solution. While this strategy significantly reduces the complexity of model development for predictive control design, it introduces potential challenges in practical implementation, particularly in translating the computed power signals into actual equipment control commands that achieve the desired thermal output.

This gap and limitations of model-dependent approaches have motivated investigation into model-free control methodologies [25, 26]. One of the limit in these approaches lies in their analytical complexity and the difficulty in providing deterministic system behavior. This limitation highlights the need for methodologies that balance model independency with system interpretability.

The Virtual Reference Feedback Tuning (VRFT) [27] is a model-free control tuning method that accommodates the diversity of the HVAC system at a lower modeling cost. Using I/O data of the system, VRFT tunes controller parameters to match a desired closed-loop reference model. VRFT is appropriate for controller classes that are simple to analyze and verify, such as PID. The value of this approach has been corroborated by recent work on water reservoir control systems, where a similar methodology demonstrated success in managing complex, nonlinear systems [28]. This paper adapts the approach to building thermal management to tackle equipment-dependent aspects of control without requiring detailed system models.

1.2. Data condition

Data plays a critical role in building control design. Open-loop experiments with persistently exciting inputs, such as pseudo-random binary sequences or white/colored noise, are known to provide reliable data-driven models [29]. However, such experiments are often impractical in real building thermal systems. Instead, historical operational and maintenance data, though limited in amplitude and frequency by closed-loop operation, are more readily available. Under these constraints, the choice between gray-box and black-box modeling becomes essential, as each relies differently on data quality, quantity, and characteristics.

Gray-box modeling reduces reliance on large amounts of data by informing prior knowledge of the model structure [4,30,31]. It usually entails a vast simplification and linearization of system physics. Equivalent Resistance–Capacitance (RC) model is a typical gray-box model. It translates building thermodynamic to electric circuit topology. Its capability for modeling indoor climate is extensively studied [32–35].

Black-box models learn system dynamics directly from measured data without physical assumptions. Linear models, such as autoregressive-based models and Subspace State Space Identification (4SID) models, offer computational efficiency and interpretability but are limited in capturing complex nonlinear dynamics [14,15,36]. A comparison of these models is provided in [15], showing that 4SID offers better approximation performance and requires less modeling effort than autoregressive-based models. Nonlinear models such as Artificial Neural Networks and Gaussian Processes often outperform linear models, providing flexible input–output mapping and uncertainty quantification. However, they can lead to complex optimization problems and require more data [16,17]. The effective use of black-box models in system modeling and control involves careful consideration of data quality and quantity, as well as thoughtful application to ensure reliability [37].

When applying data-driven model to MPC, data is even more crucial. The training of efficient predictive models depends on data collection conditions — including sampling time, prediction horizon, and operating points — which should closely match the conditions under which the models will be used [38]. Historical building data typically reflect occupant comfort constraints, whereas MPC optimization may explore a broader range of scenarios to identify the optimum. As a result, predictive models trained solely on operational data may be insufficient. It is therefore essential to validate the proposed control strategies using only historical operational data, in order to demonstrate the robustness of the methodology under the practical conditions considered.

1.3. Contribution

This article proposes a hierarchical MPC framework that address the two challenges mentioned above. In the hierarchy, VRFT is used to design direct controllers for heating equipment without the need for equipment model modeling. The MPC as a top-level controller is embedded with a special rate constraint to cope with the asymmetric power rise and fall rates of large thermal inertia devices. This work also presents a detailed analysis of the RC and 4SID models as prediction models within the proposed framework, using historical datasets. The study evaluates their performance in both open-loop and closed-loop.

This paper's novelty lies in its systematic approach to addressing the device diversity challenge within a hierarchical control design. Integrating VRFT with MPC in control design opens new possibilities. Moreover, by leveraging the International Building Performance Simulation Association benchmark, the proposed framework and model comparisons provide valuable insights to this expanding open-source community.

The following sections cover all aspects of the proposed hierarchical control strategy: model identification, controller formulation and tuning. The rest of the paper is organized as follows: Section 2 introduces the control hierarchy and provides the relevant preliminaries and practical insights of system identification, rate-constrained MPC, and VRFT approach. Section 3 apply the proposed techniques to testbed for validation. Conclusion and outlooks are given in Section 4.

2. Methodology

This section details the proposed scheme structure, along with relevant preliminaries and practical insights. It begins with an overview of the hierarchical control architecture, then focuses on the supervisory level by describing the system identification process for predictive modeling, followed by the formulation of the MPC. The section concludes with the VRFT-based design approach applied at the tracking level.

2.1. Control hierarchy

The proposed hierarchical control architecture handles the challenges above in two levels:

- Supervisory level MPC plans energy distribution across zones to balance thermal comfort and consumption. It predicts temperature evolution over the prediction horizon and produces power references for the tracking level. This long-term optimization addresses slow dynamics and system-wide objectives.
- The tracking level controllers focus on short-term control. It uses a set of decentralized PID controllers with anti-windup to manage power demands from MPC and regulate local actuator, such as valve openings, pump flows, and turbines, ensuring they track the power reference signals from the supervisory level.

The proposed hierarchical control scheme separates the thermal management system through power signals. This approach has been successfully employed in previous studies [10,32,34]. These two level correspond to the two nested loops in control block diagram, as shown in Fig. 1.

VRFT tunes the tracking level controllers with power and actuator data. While VRFT optimizes controller parameters, the inherent limitations of HVAC systems may prevent achieving ideal tracking performance. This is particularly evident in hydronic heating systems, where the actuation principle exhibits inherent asymmetry: power increase occurs through active heat transfer from the hot water circulation, while power decrease relies on passive convective cooling to the indoor environment. Heating power increasing capabilities significantly exceed decreasing rates, creating a nonlinear system response that standard control approaches may not adequately address.

To account for this asymmetric behavior, we propose a Rate-constrained MPC formulation, which augments the thermal response model with an input rate constraint. This formulation explicitly identify and account for the physical limitations in power reduction through a lower bound on the rate of decrease. The rate constraint parameters are identified directly from operational data, preserving the approach's generality while capturing system-specific thermal characteristics. By combining the predictive capabilities of MPC with the adaptability of VRFT-tuned controllers, our method aims to strike a balance between model simplicity and control effectiveness, potentially offering a more widely applicable solution for building thermal management.

Implementation of this control scheme can take two forms: a supervisory control and data acquisition system [39] with programmable logic controllers at the local level, or an integrated single-chip solution.

The design workflows for the two control levels are independent of each other, as shown in Fig. 2. The supervisory level design requires thermal power and indoor temperature measurements to identify the prediction model used as a constraint in the MPC optimization. The tracking level relies on the actuator actions and thermal power outputs from the HVAC system as datasets to optimize the local controller parameters.

The hierarchical approach enables effective management of disturbances at different time scales. At the supervisory level, MPC can account for slower, predictable disturbances like weather changes or occupancy patterns. Meanwhile, tracking level controllers quickly respond to fast, local disturbances such as sudden changes in flow rates or supply temperatures. This strategy maintains stable operation without compromising the overall optimization goals.

While thermal power is not always directly measurable for local control, it can be estimated using readily available measurements such as water flow rates and supply/return temperatures [40]. This approach allows for effective power-based control at the local level without requiring additional sensors.

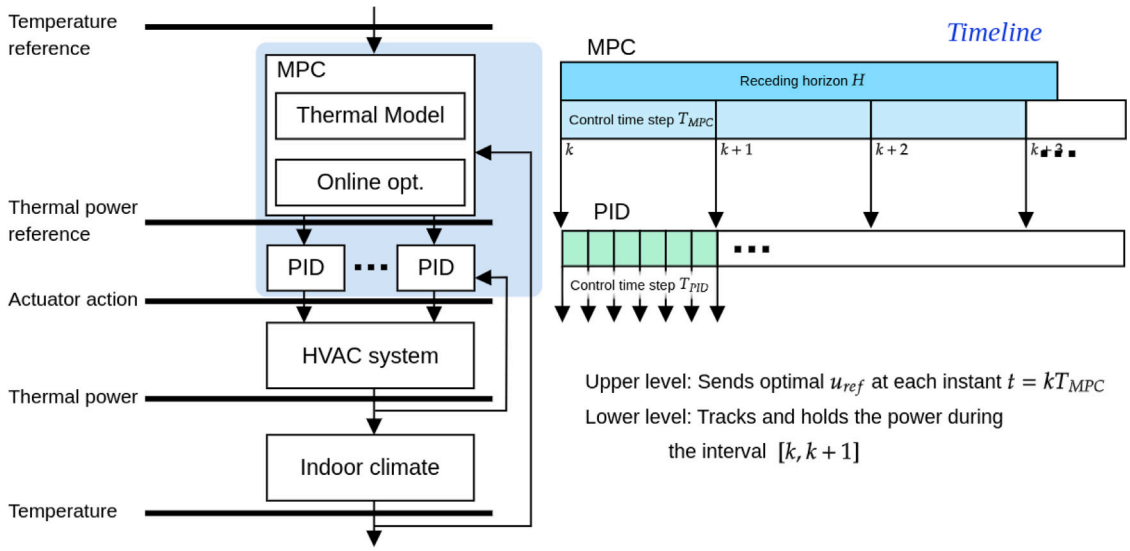

Fig. 1. Hierarchical control scheme (control solution is marked in blue) with data coordination. The supervisory-level MPC perform a prediction over the horizon H with time step T_{MPC}. The thermal power reference is updated by MPC at each time step. During MPC's time step, discrete PID regulate the actuators to follow power reference with time step T_{PID}.

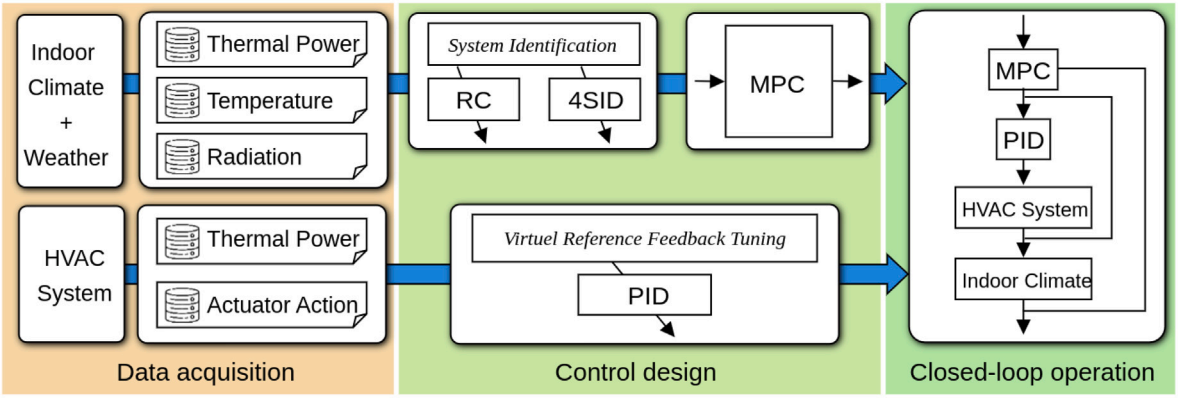

Fig. 2. Implementation work flow of the proposed methodology.

This paper is illustrated with the Building Optimization Test Framework (BOPTEST), which is dedicated to performance evaluation and benchmarking of building control strategies [41]. This framework, developed through an international collaboration led by IBPSA Project 1 researchers, provides standardized virtual testbeds constructed using well-established Modelica-based libraries for building energy modeling.

The comparison of the RC and 4SID models contains both open-loop and closed-loop scenarios. In the open-loop part, the sensitivity of the two models to the prediction window and the amount of data is compared. The similar open-loop performance of the two is a validation of the fairness of the closed-loop comparison. In the closed-loop part, the optimal prediction trajectories and thermal management indicators are compared. The analysis further shows the influence of historical data on both models when applied with MPC.

2.2. System identification

For a multi-zone thermal system, this paper adopts the centralized modeling approach used in [33]. The 1st-Order Resistance–Capacitance Multi-Zone (1RCMZ) model represents the indoor climate of a multi-zone building. Each zone is modeled as a first-order RC circuit.

A typical multi-zone thermal model can be described as:

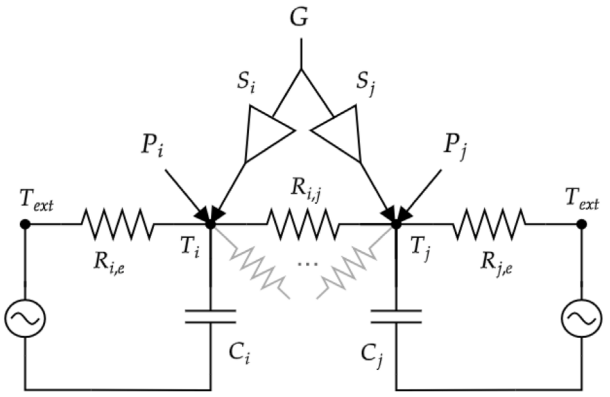

Fig. 3. Equivalent RC circuit of multi-zone gray-box modeling, showing two sample rooms as an example. Triangle shapes in the diagram represent equivalent surfaces, indicating the gain from solar radiation to absorbed radiative power.

- **Multiple Zones**: The total number of zones is n. The system consists of a set of zones represented as $N_{all} := \{1, 2, \dots, n\}$. Each zone $i \in N_{all}$ has an equivalent thermal capacitance C_i, a temperature T_i, a thermal power source P_i.
- **Thermal Exchange**: Each zone i engages in thermal exchange with its adjacent zones $j \in N_{ad,i}$. There exists a thermal resistance $R_{i,j}$ between each pair of adjacency.
- **External Walls and Floor**: Assume that all zones i either have external walls or in contact with the external ground. In equivalent circuit, the connection to exterior are simplified into a constant voltage source T_{ext} and a resistance $R_{i,e}$ in series.
- **Solar Radiation Input**: All zones are subject to solar radiation. The radiation intensity is consistent across the system, denoted as G. Each zone has an equivalent incident area for solar radiation labeled S_i.

For each zone node, an internal energy balance differential equation holds as shown in Eq. (1), Fig. 3.

$$C_i\dot{T}_i = \sum_{\forall j \in N_i} \frac{1}{R_{i,j}} Ad_{i,j}(T_j - T_i) + \frac{1}{R_{i,e}}(T_{ext} - T_i) + P_i + S_i G, \text{with } Ad_{i,j} = \begin{cases} 0 & , j \notin N_{ad,i} \\ 1 & , j \in N_{ad,i} \end{cases} \tag{1}$$

where $Ad_{i,j}$ indicates the adjacency from zone i to j.

From the perspective of system identification, the parameter vector to be estimated is

$$\theta = [R_{1,2}, \dots, R_{n,n-1}, R_{1,e}, \dots, R_{n,e}, C_1, \dots, C_n, S_1, \dots, S_n]^\top. \tag{2}$$

Input vector is

$$u = \left[P_1, P_2 \dots, P_n\right]^\top, \tag{3}$$

disturbance (weather factors) vector is

$$d = \left[G, T_{\text{ext}}\right]^\top, \tag{4}$$

output vector (state is measured as output) is

$$y = \left[T_1, T_2, \dots, T_n\right]^\top. \tag{5}$$

The 1RCMZ model can be rewritten in a continuous state space form as shown in Eq. (6).

$$f_{RC}(\theta, y, u, d) = A(\theta)y + B_1(\theta)u + B_2(\theta)d \tag{6a}$$

$$A(\theta)_{(n,n)} = \begin{bmatrix} -(\frac{1}{R_{1,e}C_1} + \Sigma_{i\in N_{ad,1}} \frac{1}{R_{1,i}C_1}) & \frac{Ad_{1,2}}{R_{1,2}C_1} & \cdots & \frac{Ad_{1,n}}{R_{1,n}C_1} \\ \frac{Ad_{2,1}}{R_{2,1}C_2} & -(\frac{1}{R_{2,e}C_2} + \Sigma_{i\in N_{ad,2}} \frac{1}{R_{2,i}C_2}) & \cdots & \frac{Ad_{2,n}}{R_{2,n}C_2} \\ \vdots & \vdots & \ddots & \vdots \\ \frac{Ad_{n,1}}{R_{n,1}C_n} & \frac{Ad_{n,2}}{R_{n,2}C_n} & \cdots & -(\frac{1}{R_{n,e}C_n} + \Sigma_{i\in N_{ad,n}} \frac{1}{R_{n,i}C_n}) \end{bmatrix} \tag{6b}$$

$$B_1(\theta)_{(n,n)} = \begin{bmatrix} \frac{1}{C_1} & 0 & \cdots & 0 \\ 0 & \frac{1}{C_2} & \cdots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \cdots & \frac{1}{C_n} \end{bmatrix}, \quad B_2(\theta)_{(n,2)} = \begin{bmatrix} \frac{S_1}{C_1} & \frac{1}{R_{1,e}C_1} \\ \frac{S_2}{C_2} & \frac{1}{R_{2,e}C_2} \\ \vdots & \vdots \\ \frac{S_n}{C_n} & \frac{1}{R_{n,e}C_n} \end{bmatrix} \tag{6c}$$

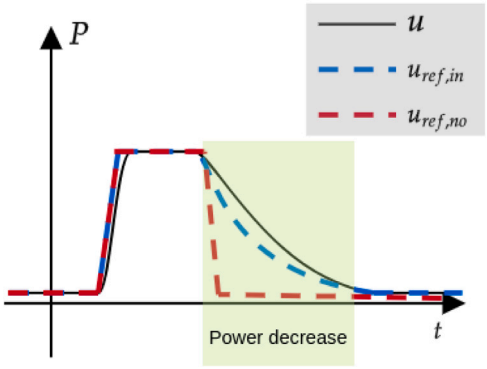

Fig. 4. Power response of heating equipment. A nominal MPC suggests power reference $u_{ref,no}$ that waits for power decrease process. An internal loop model contributes to a closer prediction to local prediction, like $u_{ref,in}$.

The training data set is denoted $\bar{y}(t_k)$, $\bar{u}(t_k)$, $\bar{d}(t_k)$, $k \in \{1, \dots, n_{id}\}$ and is measured in discrete time. Training data satisfy the following equation.

$$\frac{\mathrm{d}\bar{y}(t_k)}{\mathrm{dt}} = A(\theta)\bar{y}(t_k) + B_1(\theta)\bar{u}(t_k) + B_2(\theta)\bar{d}(t_k) + \bar{e}(t_k, \theta) \tag{7}$$

where $\bar{e}$ is the modeling error term. $A(\theta)$ and $B(\theta)$ is defined in Eqs. (6b) and (6c).

Finally, the identification problem comes down to finding the optimal set of parameters θ^* in the sense of the following optimal problem

$$\min_{\theta} \sum_{k=1}^{n_{id}} \|\bar{e}(t_k, \theta)\|^2 \tag{8a}$$

$$s.t. \quad \theta \geq 0 \tag{8b}$$

As in [33], once rewritten in matrix form, it can be solved by the least-squares method under linear constraints. The Matlab Optimization Control Toolbox provides the `lsqlin()` function for solving linear quadratic problems [42].

This optimization problem contains convex objective function and convex constraints. This avoids the risk of converging to a local optimum. The simplicity of this formulation (low dimension and convexity) is an important asset for scalable deployment in real world applications.

Once the optimal solution θ^* has been found, the predictive model is defined as

$$\frac{\mathrm{d}\hat{y}}{\mathrm{dt}} = A(\theta^*)\hat{y} + B_1(\theta^*)u + B_2(\theta^*)\hat{d} \tag{9}$$

where $\hat{d}$ is forecast disturbance.

The MPC optimization problem must integrate the evolution equation of the state representation Eq. (9). To this end, the latter is discretized using the Runge–Kutta 4 integration method leading to Eq. (10).

$$\hat{y}(t_{k+1}) = f^*_{d,RC}(\hat{y}(t_k), u(t_k), \hat{d}(t_k)) \tag{10}$$

The system identified by the RC structure has two intrinsic properties :

- Every row in state space model corresponds to a node in thermal graph. Similar to electric circuit, nodal analysis here implies that the sum of heat flow is zero. The multi-zone model respects thus energy conservation when taking interaction into account.
- The construction adjacency is informed to A matrix. It instructs energy distribution in multi-zone.

Remark. Using a first-order model for each zone avoids the problems of observability and identifiability. It ensures that the identification process is a convex optimization. There is no longer need to make initial guesses on the optimization variables. On the other hand, higher order models lead to better approximation. See more model structures in [34]. Higher order RC models require temperature measurements of additional thermal capacitances, such as walls. The alternative is to estimate these temperatures as system states along with the RC model parameters by methods such as extended Kalman filtering. However the estimated values of these states lose their interpretability like the parameters. Therefore using higher order RC without additional measurements leads to a partial black-boxing of the model.

To compare with RC model, a linear black-box model is implemented in this work. With same system order, linear black-box and gray-box have the same level of mathematical complexity. Also, their convexity leads to a same computational effort of online optimization. The 4SID prediction model is

$$\frac{\mathrm{d}\hat{y}}{\mathrm{dt}} = A_{4SID}\hat{y} + B_{1,4SID}u + B_{2,4SID}\hat{d} \tag{11}$$

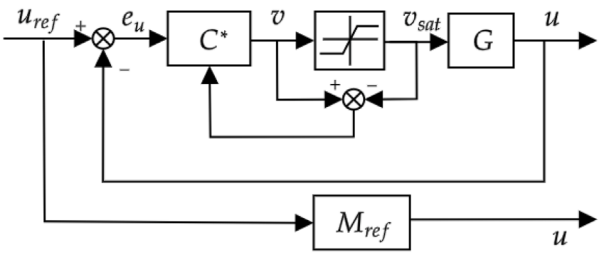

Fig. 5. Block diagram of VRFT process. C^* represents an ideal controller that drives the closed-loop to achieve the desired performance M_{ref}..

where A_{4SID}, $B_{1,4SID}$, $B_{2,4SID}$ are full matrix with no sparsity. With the same treatment to Eq. (9), Eq. (11) is discretized to

$$\hat{y}(t_{k+1}) = f_{d,4SID}(\hat{y}(t_k), u(t_k), \hat{d}(t_k)) \tag{12}$$

Mathwork developed the system identification toolbox for fast modeling [43]. The `N4SID()` function can quickly deploy singular value analysis and identification process.

2.3. Supervisory level design: MPC approach

Assuming an ideal tracking level control, the typical optimization solved by MPC at this level is

Nominal MPC

$$\min_{u_{ref}} \sum_{t=t_k}^{t_{k+n_H}} (e(t_{k+1})^T Q e(t_{k+1}) + u_{ref}(t_k)^T R u_{ref}(t_k)) \tag{13a}$$

$$e(t_{k+1}) = y_{ref}(t_{k+1}) - \hat{y}(t_{k+1}) \tag{13b}$$

$$\hat{y}(t_{k+1}) = f_d(\hat{y}(t_k), u_{ref}(t_k), \hat{d}(t_k)) \tag{13c}$$

$$\underline{u} \le u_{ref}(t_k) \le \bar{u} \tag{13d}$$

At each time t_k, the MPC solves the optimization defined in Eq. (13). It searches for the optimal power reference trajectory u_{ref} within the lower limits $\underline{u}$ and upper limits $\bar{u}$. The first term of the quadratic cost function Eq. (13a) quantifies the thermal comfort through the norm of the temperature error signal e over the prediction horizon H. The second term quantifies the consumption through the thermal power requested. Q and R are positive definite diagonal weight matrices. As the thermal requirements are common to each zone, Q is chosen as unitary (identity matrix I) and R is defined as wI, where w is a weighting parameter. $\hat{y}$ here represents the predicted trajectories. f_d is the indoor climate model defined in Eq. (10) or Eq. (12). $\hat{d}$ is the forecast disturbance, namely radiation and exterior temperature. In this paper, the formulation in Eq. (13) is called 'Nominal MPC'.

Nominal MPC provides a general mathematical framework for control problems, but assumes perfect tracking of power references at the local level — an assumption that may not hold in practice. Fig. 4 illustrates the thermal power response characteristics of the internal control loop. The graph demonstrates that both the blue and red thermal power reference trajectories can result in the same actual power response, represented by the black curve. In the nominal MPC approach, the controller would request an abrupt change in power demand (red dashed curve) and then wait for the actual power to decrease. This behavior stems from the operating mechanism of the equipment. For instance, water-based HVAC exhibit thermal inertia that can result in a gradual power fading process due to the indoor convective heat transfer mechanisms, potentially spanning several MPC time steps. In contrast, an MPC with a input rate constraint (blue dashed curve) can fully account for the internal loop dynamics.

By modeling and constraining the gradual thermal power evolution, the MPC can react to set-point change earlier, thereby optimizing the utilization of residual energy in the system and potentially improving overall energy efficiency. An internal loop model is then proposed as a rate constraint:

$$u_{ref}(t_{k+1}) - u_{ref}(t_k) \ge A_{in} u_{ref}(t_k) - u_{ref}(t_k) \tag{14a}$$

$$\Rightarrow u_{ref}(t_{k+1}) - A_{in} u_{ref}(t_k) = a(t_k) \ge 0 \tag{14b}$$

where A_{in} is a stable diagonal matrix to be estimated (elements smaller than 1), representing the thermal inertia of the heating device. This model approximates the internal loops as discrete first-order systems. Eq. (14b) is incorporated into the optimization as a dynamic constraint for the MPC. It establishes a lower bound on the rate of decrease for the optimization variable, the reference power vector u_{ref}. The variable a allows for power increases, acting as a non-negative slack variable. When $a = 0$, $u_{ref}(t_{k+1})$ reaches its minimum value of $A_{in} u_{ref}(t_k)$. a is not upper-bounded, reflecting the rapid power increase capability of the system. Consequently, the optimization always permits $u_{ref}(t_{k+1})$ to exceed $u_{ref}(t_k)$, subject to the maximum power limit of the device.

Only the data relating to the reduction in power is required to identify the model in Eq. (14b). The time index of the training set is

$$D_{in} = \{t_k | (\bar{u}(t_k) - \bar{u}(t_{k+1})) \geq 0, \forall t_k \in \{1, \dots, n_{id}\}\} \tag{15}$$

A_{in}^* is obtained by solving the following optimization problem:

$$\min_{A_{in}} \sum_{t_k \in D_{in}} \|u(t_{k+1}) - A_{in} u(t_k)\|^2 \tag{16}$$

In the end, optimization can be formalized as follows:

Rate-constrained MPC

$$\min_{u_{ref}} \sum_{t=I_k}^{t_{k+n_H}} (e(t_{k+1})^T Q e(t_{k+1}) + u_{ref}(t_k)^T R u_{ref}(t_k)) \tag{17a}$$

$$e(t_k) = y_{ref}(t_k) - \hat{y}(t_k) \tag{17b}$$

$$\hat{y}(t_{k+1}) = f_d(\hat{y}(t_k), u_{ref}(t_k), \hat{d}(t_k)) \tag{17c}$$

$$u_{ref}(t_{k+1}) = A_{in}^* u_{ref}(t_k) + a(t_k) \tag{17d}$$

$$a(t_k) \geq 0 \tag{17e}$$

$$\underline{u} \leq u_{ref}(t_k) \leq \bar{u} \tag{17f}$$

where Eqs. (17d) and (17e) are the internal loop model corresponding to Eq. (14b). This optimization problem is solved using CasADi [44], which enables efficient implementation through its automatic differentiation and mathematical-like syntax for problem definition. The solution employs multiple shooting to enhance both feasibility and computational speed of the optimization process.

Remark. An alternative is to add input incremental cost terms to the cost function to penalize variations of inputs. However, to handle the asymmetry, a conditional constraint that determines this cost based on the positive or negative increment is necessary. This would introduce nonlinearity and non-convexity in the optimization, thus making the solution more difficult.

2.4. Tracking level design: VRFT approach

In this level, a set of PID controllers is tuned by the VRFT approach. This approach is particularly interesting for handling the wide variety of equipment encountered in building systems. Considering that actuators always have physical saturation, [45] developed VRFT tuning for anti-windup PID controllers. This data-driven approach supports very diverse control schemes by designing controller arrays and reference models. Another advantage of VRFT is the usage of a reference model, which quantifies the desired closed-loop performance.

The typical scenario of power tracking control is as follows:

- The HVAC system G is a multiple-input multiple-output system with an input data vector $v = [v_1, \dots, v_n]^\top$ and an output data vector $u = [P_1, \dots, P_n]^\top$.
- A set of discrete controllers C is expected to achieve the power tracking goal. Its transfer function is given by Eq. (18).

$$\begin{aligned} v(t_k) &= C_{PID}(q^{-1})\theta_{PID} e_u(t_k) + C_{aw}(q^{-1})\theta_{aw}\Delta v_{sat}(t_k) \\ &= \left[1, \frac{T_{PID}}{2}\frac{1+q^{-1}}{1-q^{-1}}, \frac{2}{T_s}\frac{1-q^{-1}}{3-q^{-1}}\right]\theta_{PID} e_u(t_k) + \left(\frac{T_{PID}}{2}\frac{1+q^{-1}}{1-q^{-1}}\right)\theta_{aw}\Delta v_{sat}(t_k) \end{aligned} \tag{18}$$

  where q denotes the shift operator. $e_u = [e_{u,1}, \dots, e_{u,n}]^\top$ is the input vector to the controllers, a power error vector. C_{PID} is the discrete transfer function vector of the PID controller. T_{PID} is the discrete time step. The parameter vector of the PID controller array is $\theta_{PID} = \left[\theta_{PID,1}, \dots, \theta_{PID,n}\right]^\top$. $\theta_{PID,i} = [K_{P,i}, K_{I,i}, K_{D,i}]^\top$ are the gains of the proportional, integral, and derivative terms respectively. C_{aw} is the transfer function of the anti-windup branch, which includes a set of gains $\theta_{aw} = [K_{aw,1}, \dots, K_{aw,n}]$ to be tuned. Δv_{sat} is the difference of controller output before and after the saturation.
- Reference models M_{ref} represent the desired response model of the internal closed-loop. $M_{ref}^\dagger$ is the left inverse of M_{ref}, i.e., $M_{ref}^\dagger M_{ref} = I$.

Fig. 5 illustrates the VRFT method. This approach optimizes a given controller class (e.g., PID with anti-windup) to approximate an ideal controller C^* using only input–output data from the system G. The ideal controller C^* is defined as the controller that, if inserted in the closed-loop during the experiment in which the data was collected, would have provided exactly the desired behavior. The ideal controller VRFT generates a virtual reference $u_{ref,vr}$ by reverse-simulating the desired closed-loop performance M_{ref}. This virtual reference allows VRFT to construct the necessary controller input data. By using this virtual reference and the actual system output, VRFT tunes the controller parameters to match the ideal controller's performance without requiring an explicit system model.

The optimal parameter θ_{PID} and θ_{aw} can be obtained by the optimization below:

$$\min_{\theta_{PID},\theta_{aw}} \sum_{k=1}^{n_{id}} \|\varepsilon(t_k)\|^2 \tag{19a}$$

$$\text{s.t.} \quad \varepsilon(t_k) = v(t_k) - C_{PID}(q^{-1})\theta_{PID}e_u(t_k) - C_{aw}(q^{-1})\theta_{aw}\Delta v_{sat}(t_k), \tag{19b}$$

$$u(t_k) = M_{ref}(q^{-1})u_{ref,vr}(t_k) \tag{19c}$$

$$e_u(t_k) = u_{ref,vr}(t_k) - u(t_k). \tag{19d}$$

Then the error in Eq. (19d) can be expressed as

$$e_u(t_k) = (M_{ref}^{\dagger}(q^{-1}) - 1)u(t_k) \tag{20}$$

When generating the dataset in open-loop, inputs should exceed the saturation to make Δv_{sat} data informative. What is tuned by parameter estimation is actually the linear part of the controller, and the saturation is compensated into the system. Quadratic programming solvers, such as the '`quadprog()`' function in Matlab Optimization Toolbox, are capable of solving Eq. (19).

Remark. The time constant is a critical parameter that balances the response speed of the MPC with the inherent dynamics of the radiator. It must be fast enough to support the quasi-static power response assumption of the MPC, yet slow enough to be achievable by the physical system. As an initial guideline, setting the time constant to 1/10 of the MPC time step is a good starting point for further refinement. This ratio represents a balance between meeting the supervisory level's power tracking requirements and considering actuator longevity, as overly aggressive response times could accelerate equipment wear and tear.

In (20), $M_{ref}^{\dagger}$ is used to backward-simulate the virtual reference to compute e. If the reference model has more poles than zeros, $M_{ref}^{\dagger}$ will become a non-causal system, making simulation infeasible. Therefore, an approximate approach is to add a zero to the transfer function of the reference model with a time constant much faster than the poles. Experiments show that a zero 10 times faster is sufficient. For example, M_{ref} is a diagonal transfer matrix consisting of several transfer functions $M_{ref,i}$ in Eq. (21)

$$M_{ref,i}(q^{-1}) = ZOH\left(\frac{\frac{1}{10}\tau s + 1}{\tau s + 1}\right) = \frac{\beta_0 + \beta_1 q^{-1}}{1 - q^{-1}e^{-T_{PID}/\tau}} \tag{21}$$

where τ refers to the desired gain and time constant. T_{PID} is the time step of the discrete PID controller. β_0 and β_1 are obtained through ZOH discretization.

In the proposed structure, local-level variables such as zone temperature, supply water temperature, and supply water flow rate can be considered as disturbances to the tracking level control. However, these variables are typically outputs of other closed-loop systems within the building's thermal management setup, which naturally constrains their variability. When generating the dataset for controller tuning, incorporating these realistic operating conditions helps ensure the resulting controller's effectiveness. The controller parameters should be re-tuned when the operation point of the tracking control are changed. Additionally, the PI controller structure inherently possesses disturbance rejection capabilities suitable for building services. This combination of constrained disturbances and the robust structure results in a reliable power regulation.

3. A case study

3.1. BOPTEST testcase: Multi-zone residential hydraulic

BOPTEST is a framework for building control performance benchmarking [41]. It proposes a series of test benchmarks covering commercial/residential, single-zone/multi-zone, and multiple fluid scenarios. These building systems are constructed using Modelica-based libraries such as IDEA [46] and Buildings [47]. The run-time environment of these test cases is deployed using Docker and a RESTful HTTP API.

The 'Multi-zone residential hydronic' test case was selected to validate the proposed methodology. As shown in Fig. 6, the test case includes multiple heating zones, with an attic above all zones that is not shown in the diagram. All zones, except the attic and garage, have hydraulic radiators. Each radiator, except the one in the hallway, is controlled by a valve to regulate flow. The hallway branch balances the distribution circuit pressure. On the supply side, a gas boiler provides hot water, a main pump controls the total flow, and a three-way valve mixes boiler hot water with return water to achieve supply temperature control. The baseline controllers provided with the benchmark consist of multiple PID and RBC. Their objectives are as follows:

- The boiler water temperature remains within security boundaries.
- The supply water temperature tracks a set-point related to the outside temperature.
- The pumping system minimizes the temperature error in the living room.
- Indoor temperature set-point tracking.

The test case represents a residential building designed in accordance with the French thermal regulation RT2012, which established stringent requirements for building energy performance. The system's complexity manifests through multiple nonlinear phenomena: valve characteristics exhibit nonlinear flow-pressure relationships, heat transfer processes incorporate both radiative

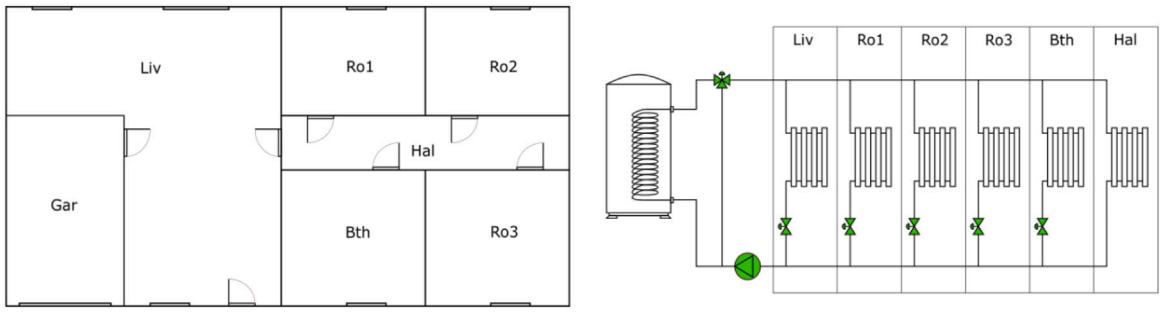

Fig. 6. Adjacency of the test case and hydronic circuit of heating system.

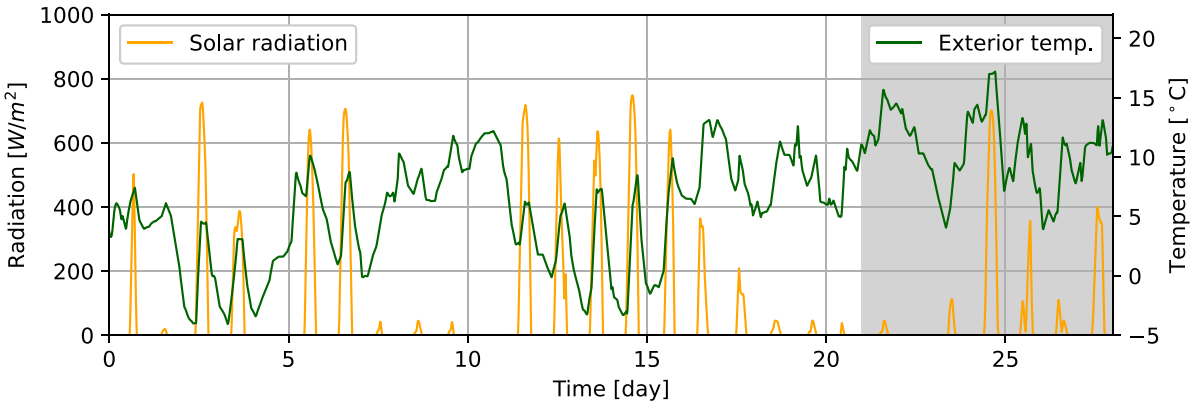

Fig. 7. Weather data used in benchmark test case. Validation dataset is marked in gray.

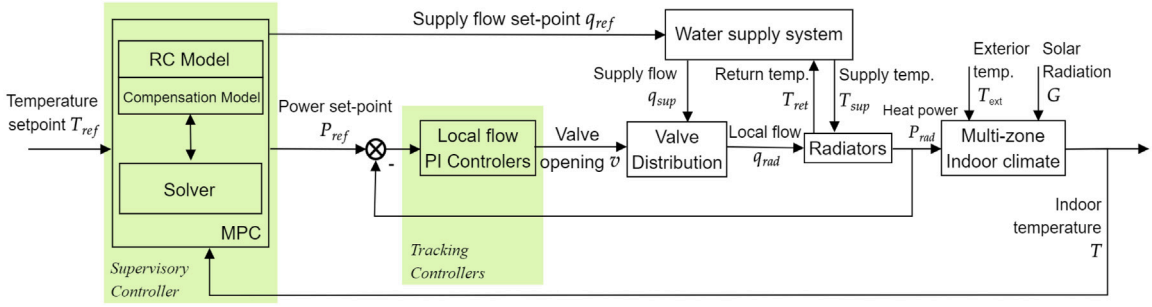

Fig. 8. Diagram of 'Multi-zone residential hydronic' test case with proposed solution. Boiler and mixture circuit are simplified as water supply system block. The hierarchical control is deployed on the flow distribution system. In this paper, system variables $y_{ref}, u_{ref}, v, u, y$ are composed of signals $T_{ref}, P_{ref}, v, P, T$ respectively.

and convective mechanisms, and solar gains through glazed surfaces follow complex geometric and material-dependent patterns. The simulation uses a climate file of weather data from 2014 in Bordeaux, France. The outdoor temperature and solar radiation are shown in Fig. 7.

In this work, the validation retains the baseline controllers for the boiler and mixture. The demand of supply flow is computed through the sum of thermal power command. According to the control methodology proposed in this paper, a hierarchical control scheme for valve systems is implemented, as shown in Fig. 8.

This study assumes perfect forecasting of external temperature and solar irradiation over the prediction horizon, allowing evaluation of the control methodology's fundamental capabilities without the confounding effects of prediction uncertainties. While this represents an idealization of operational conditions, established methods for weather prediction and solar irradiation forecasting are available for practical implementations [48–50].

The Rate-constrained MPC is responsible for centrally optimizing the power demand of the 5 zones and generates the total reference flow based on these demands. The predictive model embedded in the MPC law is the 1RCMZ model. Note that according

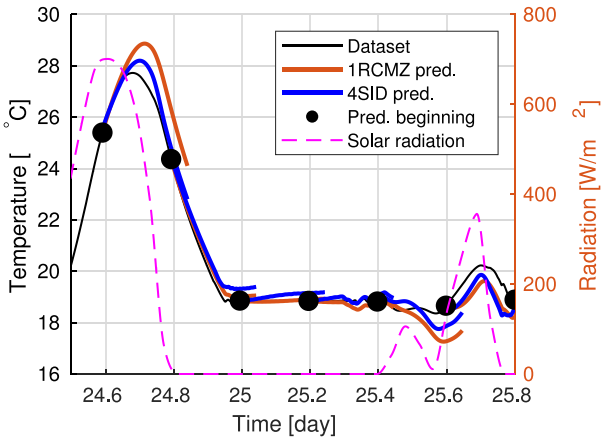

Fig. 9. Data validation of indoor climate models predictions. Temperature of bedroom 1 (Ro1.) is shown as an example. Prediction horizon is 6 h.

to Eq. (17c), the internal loop dynamic is taken into consideration. The five valve controllers are PI controllers with anti-windup design, tuned by the VRFT approach. The simulation uses OpenModelica's default DASSL solver, which is packaged in a Functional Mock-up Unit (FMU). The FMU co-simulates with a CasADi-based controller class, which is coded in Python.

3.2. Result

The Results section evaluates the proposed hierarchical MPC framework for building thermal management. It is structured in three parts:

1. Validation of indoor climate models (1RCMZ and 4SID).
2. Performance of tracking level controllers tuned by VRFT.
3. Overall performance of the hierarchical control scheme and model comparison in closed-loop.

3.2.1. Indoor climate model data validation

The indoor climate model is trained and validated using baseline control historical data in the test case. The dataset is divided into a training set, covering the period from January 1 to January 21, and a validation set, spanning January 22 to January 28. Both the gray-box (1RCMZ) and black-box (4SID) models are trained using data regulated by baseline controllers. The dataset sampling time for building multi-zone model identification is $T_{dataset} = 60\,\text{s}$. Temperature derivatives $\frac{\mathrm{d}\tilde{y}(t_k)}{\mathrm{dt}}$ are approximated by differentiation with $\Delta t = 20\,\text{s}$. To ensure fairness in comparison, the order of the 4SID model is set to 8, the same as the RC model according to the zone number. Fig. 9 presents the prediction error over a 6-h horizon for both the 1RCMZ and 4SID models in data validation. The 1RCMZ model demonstrates good approximation during periods of low solar radiation ($t = 24.8$ and $t = 25$). However, its error noticeably increases during strong solar radiation hours ($t = 24.6$ and $t = 25.6$). This highlights a weakness in the 1RCMZ model's handling of solar radiation. The actual radiation process is far more complex than the simple static gain and direct irradiance multiplication used in 1RCMZ, involving time-variant and non-linear factors like sunspot. In contrast to 1RCMZ, the 4SID model shows lower error during high solar radiation periods and higher error when solar radiation is low.

To better evaluate prediction quality, an indicator I_{pred} is introduced. This indicator, defined in Eq. (22), measures the average Mean Squared Error (MSE) of temperature predictions across all MPC control time steps in the simulation. For each zone i, $I_{\text{pred},i}$ is calculated by first computing the MSE between predicted temperatures $\hat{T}_i$ and actual measured temperatures T_i over the prediction horizon at each MPC time step. These MSE are then averaged over all MPC time steps in the simulation.

$$I_{\text{pred},i} = \frac{1}{n_{\text{MPC}}} \sum_{k=1}^{n_{\text{MPC}}} \left(\frac{1}{n_H} \sum_{p=1}^{n_H} (T_i(t_{k+p}) - \hat{T}_i(t_{k+p}))^2 \right) \tag{22}$$

Fig. 10 illustrates the approximation capabilities of the two models, focusing on bedroom 1 (Ro1.), 2 (Ro2.), and 3 (Ro3.) as representative examples. All these three zones are the same in size, but Ro3. is on the opposite side of the house. Therefore, Ro1. and Ro2. share the same radiation condition, Ro3. is different. The figure displays the I_{pred} of three zones predicted by 1RCMZ and 4SID in different forecast horizons.

I_{pred} increases with the length of the forecast horizon for both models, reflecting the accumulation of errors over time. For Ro3. which has different solar radiation conditions, its predicted quality differs from that of Ro1. and Ro2. This confirms that the 1RCMZ modeling error is sensitive to solar radiation. Within the first 4–6 h, both models exhibit comparable prediction accuracies, ensuring that their similar open-loop performance validates the fairness of the closed-loop comparison.

The sensitivity to training set length for both models is further studied in Fig. 11. The results reveal significant differences in the training stage. The 1RCMZ model converges to its best performance within one day's data. In contrast, the 4SID model requires

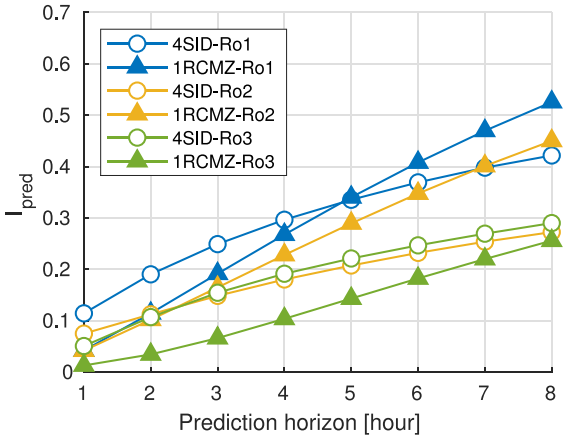

Fig. 10. Error propagation with growth of prediction horizon.

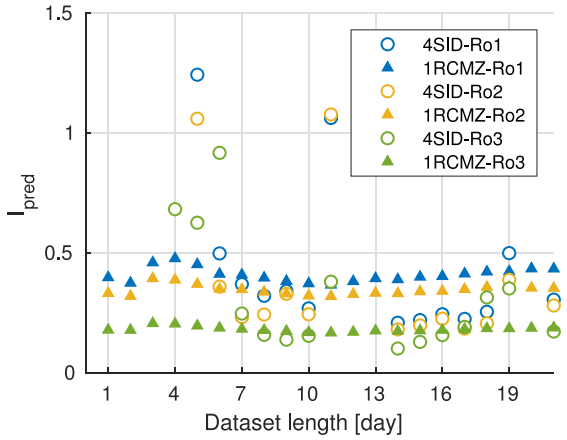

Fig. 11. Identification convergence and training dataset length. In this figure, 6 h prediction quality is validated in zone Ro1.,Ro2. and Ro3.

nearly 10 days to converge to a good prediction model. Notably, the MSE for days 1, 2, and 3 are so high they fall outside the graph's range. 1RCMZ demonstrates strong consistency, with its MSE consistently floating between 0.15 and 0.4 across different dataset lengths. The 4SID model also shows more variability. While the overall trend of 4SID's MSE is decreasing with more data, it produces outlier MSE (day 11, 19) with certain training sets. The poor performance on certain datasets can be mitigated by adjusting the order of the 4SID model or modifying the data sampling time. However, these adjustments lead to heavy workload in scalable deployment. These results confirm the 4SID method's higher sensitivity to the sampling configuration compared to 1RCMZ.

3.2.2. Tracking level control performance

Table 1 shows the PI controller parameters for different zones, tuned by VRFT. The nominal power indicates radiator size, with the Living room (Liv.) having the largest at 1.8 kW. A clear trend emerges: for larger radiators, K_P is smaller while K_I and K_{aw} are larger. This relationship aligns with PI control principles. A smaller K_P for larger radiators reduces the risk of overshoot, as these radiators have more thermal inertia. The larger K_I compensates for the reduced proportional action, ensuring steady-state errors are eliminated over time. The increased K_{aw} helps prevent integral windup, which is more critical for larger systems with potentially longer saturation periods.

Fig. 12 validates the VRFT-tuned controllers in closed-loop operation. The power response shows fast and accurate tracking during power increase. The valve primarily operates binary-like states (almost fully open or closed) due to the high supply water temperature. Radiator can reach the power reference with a fast water intake and maintain the power with a small water flow. When it maintain a higher power reference (like the fifth square wave), it requires a more obvious valve opening to continuously replenish hot water. Although valve operations in real systems exhibit smoother behavior rather than impulse signals shown in the figure, this stand-alone test justify the tracking quality of VRFT-tuned controllers. More complex power dynamics will be present in next subsection.

Power decrease is slower than the reference model M_{ref} desired. During decrease, the valve opening saturates at 0. The anti-windup mechanism effectively prevents continuing integration, avoiding delays in subsequent responses. For this study, we chose

Table 1
Each zone PI parameters tuned by VRFT.

Zone	Nominal power	K_P	K_I	K_{aw}
Liv.	1.8 kW	0.0651	$1.9e^{-13}$	$3.4e^{-13}$
Bth.	0.8 kW	2.9579	$3.9e^{-15}$	$8.5e^{-17}$
Ro1.	1 kW	1.9813	$4.1e^{-15}$	$1.1e^{-16}$
Ro2.	0.9 kW	3.0351	$3.2e^{-15}$	$7.2e^{-17}$
Ro3.	1.1 kW	1.9825	$4.0e^{-15}$	$1.1e^{-16}$

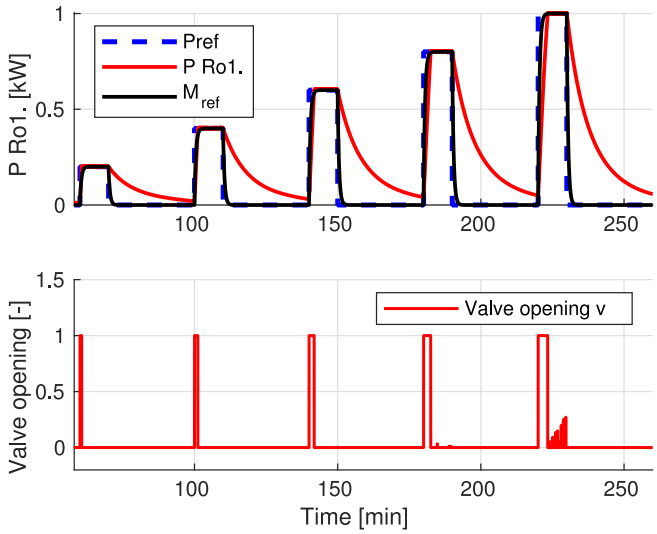

Fig. 12. Internal loop simulation that validate VRFT-tuned controllers. Zone Ro1. is shown as a sample. Hot water temperature, zone temperature and water flow is set to constant. Power reference is square wave with 10 min heating and 40 min stand-by.

Table 2
Rate-constrained MPC parameters used in test case.

Parameters	Values	Description
Model	1RCMZ/4SID	Prediction model
T_{MPC}	10 min	MPC control time step
H	4/6 h	MPC prediction horizon
Q	$I_{(8,8)}$	Weighting on state error vector
R	$wI_{(6,6)}$	Weighting on power input vector
Solver	'qpoases' [51]	Solver in MPC
w	$w_0 = 1.6 \times 10^{-4}$	Weighting on objectives

a reference model time constant τ of 1 min, which aligns with the guideline of being approximately 1/10 of the MPC time step, see Table 2. The actual power rising time τ_{act} is longer than the reference model's τ, with τ_{act}/τ ratios ranging from 1.4 to 2.7. This ratio increases with higher demanded power. Despite response is slower than intended, it remains sufficiently fast to support supervisory level control.

3.2.3. Hierarchical control performance

In the test case, performance of hierarchical control schemes are further evaluate. In this stage, 1RCMZ and 4SID are integrated into the Rate-constrained MPC formulation. Parameters tuned by VRFT in Table 1 is adopted for local controllers. Parameter setting of Rate-constrained MPC is shown in Table 2.

Fig. 13 illustrates the temperature evolution across three zones from January 22 to January 29, using the same climate file, solar radiation, and exterior temperature as the validation dataset. Despite the 4SID model shows a comparable approximation capabilities to 1RCMZ in Fig. 10, it has poor performance in the closed-loop validation. Notably, deviations from the set-point are evident: the Living Room is under-heated, while Room 2 is overheated. While the lower-level controller rejects disturbances in power tracking, it operates without thermal state awareness — it precisely executes MPC-generated power references regardless of their thermodynamic validity. These discrepancies highlight the 4SID model's sensitivity to training data limitation. When environmental perturbations exceed its extrapolation boundaries, the resulting biased power references become thermodynamically inconsistent yet get accurately implemented, ultimately manifesting as temperature deviations. In contrast, the RC model demonstrates consistent reliability when integrated into our framework, as its informed knowledge maintain prediction fidelity across wider operating conditions.

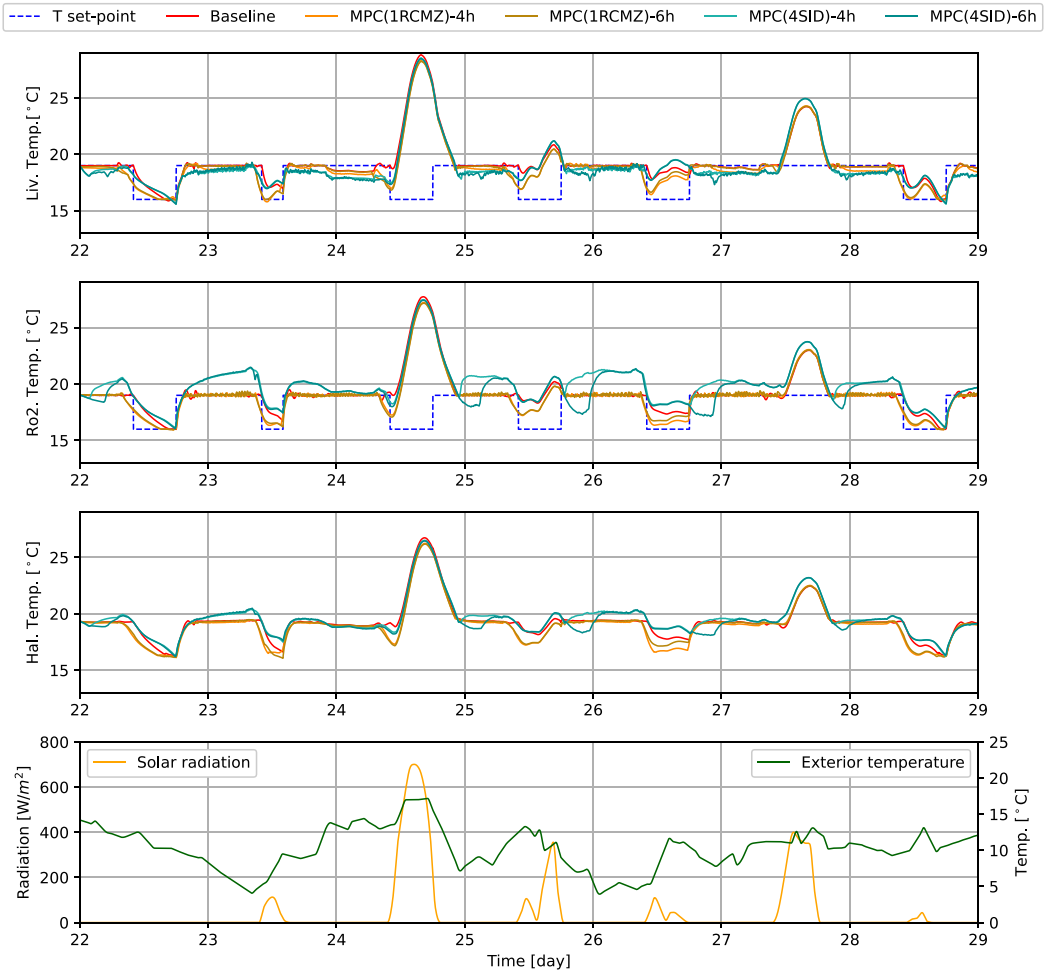

Fig. 13. Comparison of 1RCMZ and 4SID performance in proposed control scheme. The top three figures show the temperatures of the living room (liv.), bedroom (Ro2.), and hallway (hal.). The hallway is not a controlled zone but is still relevant to consumption. Both models were tested with 4-h and 6-h prediction horizons. The bottom figure shows disturbance variation.

Fig. 14 provides a one-day zoom-in to compare the prediction quality of 1RCMZ and 4SID models. Fig. 14(a) shows that 1RCMZ predicts well during occupied periods, with most predictions aligning with actual temperatures, except at midday. The MPC consistently underestimates indoor temperatures during solar radiation hours, but predictions realign with actual temperatures as solar radiation fades.

In Fig. 14(b), 4SID predictions deviate significantly from reality. At $t = 24.9$, the model overestimates temperature decrease when there is no solar radiation or radiator power. At $t = 25.2$, despite Ro2. being overheated, the MPC still demands maximum radiator power $\bar{u} = 800$ W, as indicated by u_{ref} Ro2. This inaccurate prediction in Ro2 also affects Liv. Actual radiator power u Liv. is not enough to achieve the reference u_{ref} Liv. due to excessive demand in Ro2, as Ro2 consumes a disproportionate share of power in the hot water circulation.

Baseline scheme historical data, i.e., training set is very concentrated around the set-point. It is similar to red curve in Fig. 13. The modeled gradient is calibrated over a narrow range. However, the MPC searches for optimal control inputs and state trajectories over a broader range. The optimization process may identify solutions within the 4SID state space that lie outside the reliable region but yield a lower cost in the objective function. Consequently, while the predicted trajectories converge well to the set-point, the power reference generated by the MPC may be inconsistent.

Fig. 15 compares overall prediction quality of 1RCMZ and 4SID model. The 4SID model produces higher errors than the 1RCMZ model throughout the simulation. The figure also reveals that the mismatch in Liv. is higher than in other rooms. For 1RCMZ, as explained, the insufficiency lies in the radiation model. The living room has a larger window surface, so it is more sensitive than other zones. For 4SID, predicted trajectories always converge to the set-point, error is thus larger when temperatures rise in solar radiation hours. Since living room temperature in non-occupied time is always higher than other zones, it has bigger deviations.

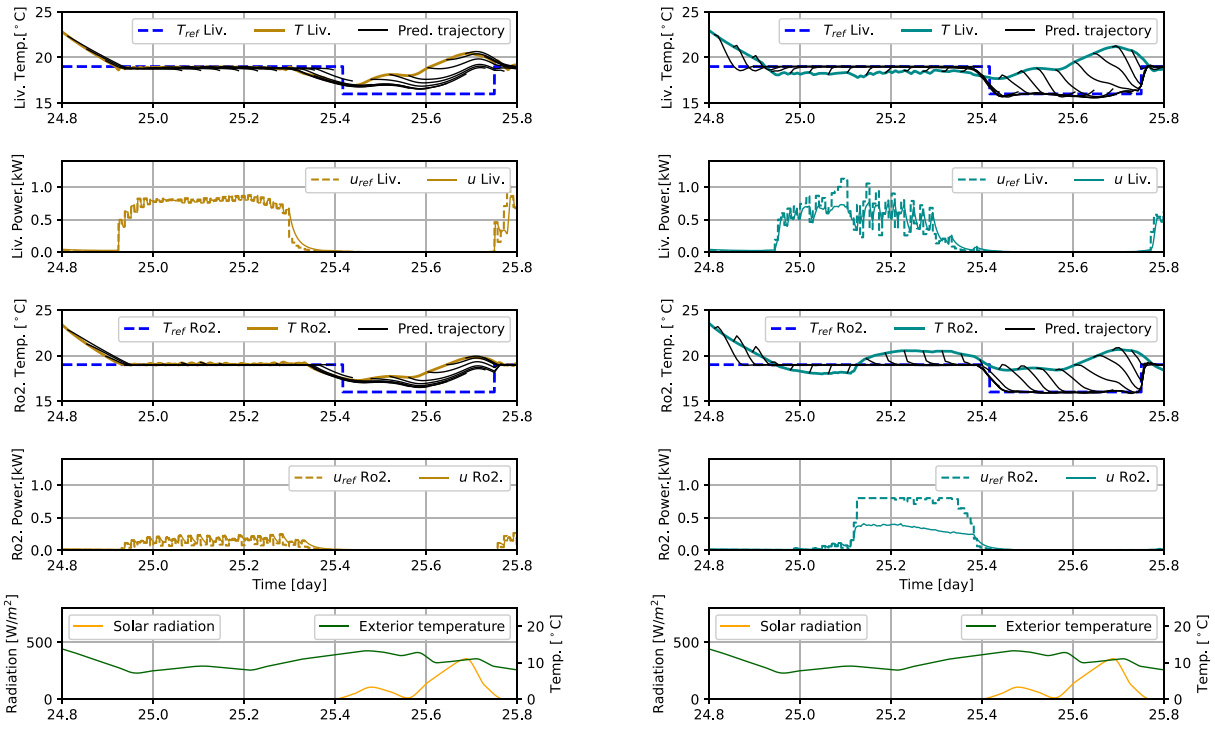

(a) Rate-constrained MPC-1RCMZ scheme with 6 hour prediction

(b) Rate-constrained MPC-4SID scheme with 6 hour prediction

Fig. 14. Zoom-in on the comparison of Rate-constrained MPC formulation with 1RCMZ and 4SID models in test case. Each scheme shows Liv. and Ro2. as samples. Figure exhibits temperature and power responses of each zone and also disturbances of corresponding period. Every hour ($6 * T_{MPC}$) a temperature prediction trajectory is shown as a black line labeled 'Pred. trajectory' in temperature subplots.

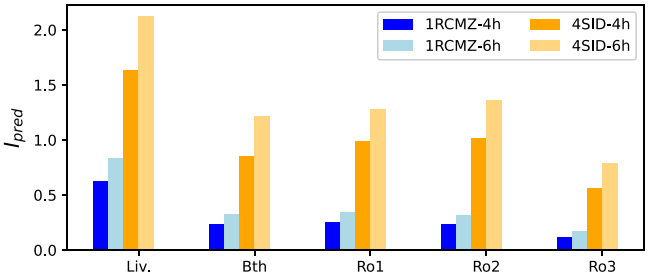

Fig. 15. Prediction quality comparison between different scheme. Expression of I_{pred} is Eq. (22).

Fig. 16 shows the effect of the power decrease constraint. Rate-constrained MPC in Eq. (17) is constrained by the internal loop model. Nominal MPC in Eq. (13) has no limit on equipment power decrease. Rate-constrained MPC aligns better with the actual power response. At 10:00, the residential building changes to non-occupied mode. The nominal MPC provides a power reference that quickly drops to 0. The heating equipment cannot follow immediately, MPC has to wait for the radiator to exhaust its power. Conversely, Rate-constrained MPC suggests a first-order-shaped curve with a slow descent, adequately planning the usage of residual heat, thereby starting the power decrease process earlier than the nominal MPC. Equipment thermal inertia is then fully managed.

In run-time environment, two Key Performance Indicators (KPI) are defined to evaluate the control scheme. These KPI rate the performance of energy consumption and discomfort according to [41]. The thermal discomfort indicator is

$$I_{dis} = \sum_{i=1}^{n} \sum_{k=1}^{n_{data}} max(T_{set,i}\,(t_k) - T_i\,(t_k), 0) \tag{23}$$

where $max()$ corresponds to a saturation on temperature error, only under-heating error is counted as discomfort. $T_{set,i}$ is the heating set-point temperature of each zone. T_i is the measured temperature of each zone. The unit of this indicator is converted to [Kelvin hour] according to simulation sampling time.

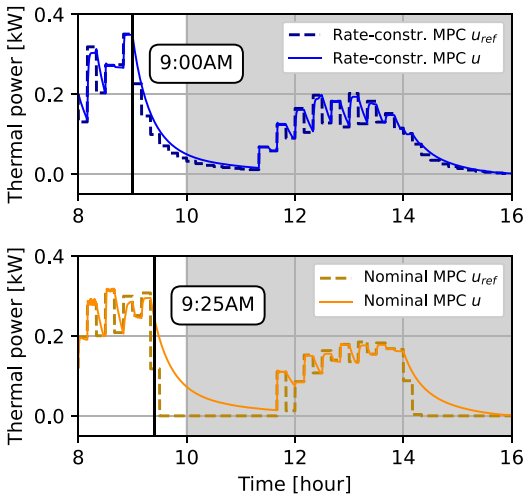

Fig. 16. Radiator power tracking performance. Two subplots compare Rate-constrained MPC-1RCMZ in Eq. (17) and nominal MPC-1RCMZ in Eq. (13). The gray zone represents non-occupied time requiring lower temperature demand.

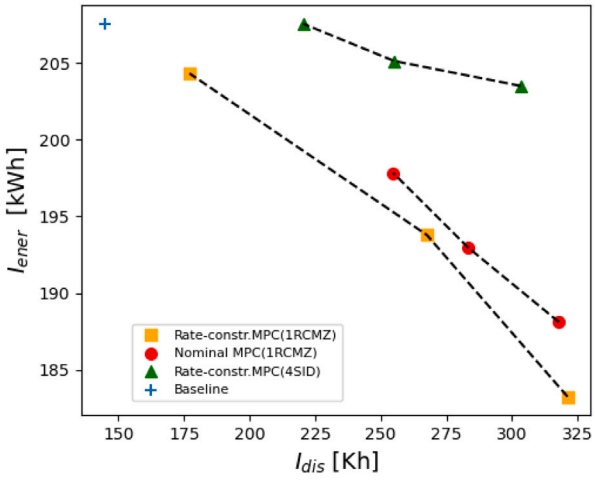

Fig. 17. KPI map for different schemes. I_{dis} is the thermal discomfort indicator, see Eq. (23). I_{ener} is the energy consumption indicator, see Eq. (24). For each scheme, three Weighting coefficient w in the MPC is tested. Corresponding value from left to right is $w \in \{100w_0, w_0, 0.1w_0\}$.

The consumption indicator is

$$I_{ener} = \sum_{i=1}^{n} \sum_{k=1}^{n_{data}} P_i\,(t_k) \tag{24}$$

where P_i is heating power of radiator in each zone. The unit of this indicator is converted to kWatt hour according to simulation sampling time.

Fig. 17 synthesizes and compares the KPI of several schemes. Including original baseline control scheme in the testcase, a set of hierarchical control scheme with Eqs. (13) and (17). The better performing scheme will be closer to the bottom left corner. However, compromises between the two objectives will shift the result from the upper left to the lower right, depending on the objective weights. It corresponds to the Pareto front of the optimization. The comparison between the nominal MPC and the Rate-constrained MPC is interesting. With a heavy weight on consumption ($w = 100w_0$), Rate-constrained MPC consumes less energy without significantly decreasing thermal comfort. It fully used the residue heat in equipment. When the consumption weight is relatively low ($w = w_0$), the results of the two models are comparable. Adding thermal power rate constraint does not absolutely imply a reduction in energy consumption. Model errors also fluctuate results. The residue heat released by equipment sometimes does not lead to a overheat, but covers the under-heating at beginning of non-occupied time. This happens in low radiation days, such as January 23 midday in Liv., see Fig. 13. Thus, compensated model will not result in noticeable improvements when the energy savings target is not important enough. With 4SID model, Rate-constrained MPC has more discomfort, as the 4SID model's

divergence leads to considerable overheating. The 1RCMZ model achieves approximately 3%–12% reduction in energy consumption, while having comparable or even lower levels of thermal discomfort compared to the 4SID model.

3.2.4. Summary

Both the 1RCMZ and 4SID models effectively capture the thermal behavior of individual zones within a building. The 1RCMZ model exhibited clear deterministic behavior in the dataset length sensitivity test. Although the data validation performance is comparable for both models, the RC and 4SID models demonstrate different closed-loop performances when integrated into the MPC framework. The RC model maintains good closed-loop performance when trained on historical datasets. The proposed methodology, combining MPC and VRFT, enables efficient, and even optimized, management of the produced power for thermal regulation purposes. Furthermore, the integration of power variation constraints improves energy management by enhancing the trade-off between energy consumption and occupant comfort.

4. Conclusion

The hierarchical regulation design methodology proposed in this article offers a global solution by providing answers to various well-known difficulties, linked to the diversity of equipment encountered, to the limitation of available data, to the fine management of the compromise between thermal comfort and energy cost. In this article, the methodology articulates different stages and solutions depending on the hierarchical level considered. This methodology responds to the two key challenges:

1. To handle HVAC system diversity and complexity, the methodology implements VRFT, a model-free approach, at the tracking level and a identified rate constraint at the supervisory level. No detailed HVAC system model is required in the proposed design procedure.
2. Regarding the challenges with the experimental data in practice, this work compares the performance of RC and 4SID models trained using building operation history data within the proposed framework. The results demonstrate that the RC model achieves a reduction in energy consumption of approximately 3%–12%, while maintaining comparable or even lower levels of thermal discomfort compared to the 4SID model.

The combination of a gray-box MPC for long-term optimization and VRFT-tuned PID controllers for local regulation strikes an effective balance between model simplicity and control effectiveness. This approach offers a more widely applicable solution for building thermal management, addressing the challenges of equipment diversity and limited historical data availability. The Rate-constrained MPC formulation further enhances the system's performance by accounting for equipment thermal inertia, demonstrating improved energy efficiency without compromising comfort. This showcases the framework's ability to adapt to specific system characteristics while maintaining its general applicability.

This study is based on assumptions that help isolate the core contributions but also define its scope. First, the use of ideal weather forecasts removes external uncertainty, which facilitates analysis but may differ from real operational conditions. Second, the benchmark testcase does not include measurement or actuation noise, allowing for focused evaluation of the control framework without additional disturbances. Third, the test case centers on heating services in a residential setting. Future studies may consider extending the approach to other services, such as cooling, or to larger-scale deployments, to further explore its applicability across a wider range of building types and operating environments.

CRediT authorship contribution statement

Yuqi Liu: Writing – review & editing, Writing – original draft, Visualization, Validation, Software, Formal analysis, Data curation, Conceptualization. **Pauline Kergus:** Writing – review & editing, Validation, Supervision, Methodology, Formal analysis, Conceptualization. **Fabien Claveau:** Writing – review & editing, Validation, Supervision, Methodology, Formal analysis, Conceptualization. **Philippe Chevrel:** Writing – review & editing, Supervision, Resources, Project administration, Methodology, Investigation, Funding acquisition, Formal analysis, Conceptualization. **Bruno Lacarrière:** Writing – review & editing, Validation, Supervision.

Declaration of competing interest

The authors declare the following financial interests/personal relationships which may be considered as potential competing interests: Yuqi LIU reports financial support was provided by Regional Council of Pays de la Loire. 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

Data will be made available on request.

References

[1] World Energy Outlook 2024, IEA, Paris, 2024, https://www.iea.org/reports/world-energy-outlook-2024, licence: Cc by 4.0 (report); cc by nc sa 4.0 (annex a), IEA.

[2] U.S.S.O. Engineering, Phase Change Composite Materials for Energy Efficient Building Envelopes, Tech. Rep., UCLA Samueli School Of Engineering, 2018, URL https://www.seas.ucla.edu/pilon/PCMIntro.html.

[3] E.P. of Buildings Directive, Directive (eu) 2024/1275 of the european parliament and of the council of 24 april 2024 on the energy performance of buildings, 2024.

[4] J. Drgoňa, J. Arroyo, I. Cupeiro Figueroa, D. Blum, K. Arendt, D. Kim, E.P. Ollé, J. Oravec, M. Wetter, D.L. Vrabie, L. Helsen, All you need to know about model predictive control for buildings, Annu. Rev. Control. 50 (2020) 190–232, http://dx.doi.org/10.1016/j.arcontrol.2020.09.001.

[5] K. Bamdad, N. Mohammadzadeh, M. Cholette, S. Perera, Model predictive control for energy optimization of hvac systems using energyplus and aco algorithm, Buildings 13 (12) (2023) http://dx.doi.org/10.3390/buildings13123084.

[6] F. Oldewurtel, A. Parisio, C.N. Jones, D. Gyalistras, M. Gwerder, V. Stauch, B. Lehmann, M. Morari, Use of model predictive control and weather forecasts for energy efficient building climate control, Energy Build. 45 (2012) 15–27, http://dx.doi.org/10.1016/j.enbuild.2011.09.022.

[7] F. Oldewurtel, C.N. Jones, A. Parisio, M. Morari, Stochastic model predictive control for building climate control, IEEE Trans. Control Syst. Technol. 22 (3) (2014) 1198–1205, http://dx.doi.org/10.1109/TCST.2013.2272178.

[8] D. Sturzenegger, D. Gyalistras, M. Gwerder, C. Sagerschnig, M. Morari, R. Smith, Model predictive control of a swiss office building, Climate 2013 (2013) 3227–3236.

[9] Y. Ma, F. Borrelli, B. Hencey, B. Coffey, S. Bengea, P. Haves, Model predictive control for the operation of building cooling systems, IEEE Trans. Control Syst. Technol. 20 (3) (2011) 796–803.

[10] J. Arroyo, F. Spiessens, L. Helsen, Comparison of model complexities in optimal control tested in a real thermally activated building system, Buildings 12 (5) (2022) 539.

[11] F. Jorissen, D. Picard, K. Six, L. Helsen, Detailed white-box non-linear model predictive control for scalable building hvac control, in: Modelica Conferences, 2021, pp. 315–323.

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

[13] F. Gauthier-Clerc, H. Le Capitaine, F. Claveau, P. Chevrel, Comparing neural network and linear models in economic mpc: Insights from boptest for building temperature control, in: ECC 2024: European Control Conference, 2024.

[14] P. Van Overschee, B. De Moor, Subspace Identification for Linear Systems: Theory—Implementation—Applications, Springer Science & Business Media, 2012.

[15] S. Prívara, Z. Váňa, D. Gyalistras, J. Cigler, C. Sagerschnig, M. Morari, L. Ferkl, Modeling and identification of a large multi-zone office building, in: 2011 IEEE International Conference on Control Applications, CCA, IEEE, 2011, pp. 55–60.

[16] H. Huang, L. Chen, E. Hu, A neural network-based multi-zone modelling approach for predictive control system design in commercial buildings, Energy Build. 97 (2015) 86–97.

[17] H.Y. Noh, R. Rajagopal, Data-driven forecasting algorithms for building energy consumption, in: Sensors and Smart Structures Technologies for Civil, Mechanical, and Aerospace Systems 2013, Vol. 8692, SPIE, 2013, pp. 221–228.

[18] R. Yang, L. Wang, Multi-zone building energy management using intelligent control and optimization, Sustain. Cities Soc. 6 (2013) 16–21.

[19] D. Zhao, D. Watari, Y. Ozawa, I. Taniguchi, T. Suzuki, Y. Shimoda, T. Onoye, Data-driven online energy management framework for hvac systems: An experimental study, Appl. Energy 352 (2023) 121921.

[20] S. Zhan, A. Chong, Data requirements and performance evaluation of model predictive control in buildings: A modeling perspective, Renew. Sustain. Energy Rev. 142 (2021) 110835.

[21] S.R. West, J.K. Ward, J. Wall, Trial results from a model predictive control and optimisation system for commercial building hvac, Energy Build. 72 (2014) 271–279.

[22] Y. Zong, G.M. Böning, R.M. Santos, S. You, J. Hu, X. Han, Challenges of implementing economic model predictive control strategy for buildings interacting with smart energy systems, Appl. Therm. Eng. 114 (2017) 1476–1486.

[23] A. Mirakhorli, B. Dong, Market and behavior driven predictive energy management for residential buildings, Sustain. Cities Soc. 38 (2018) 723–735.

[24] J. Joe, P. Karava, A model predictive control strategy to optimize the performance of radiant floor heating and cooling systems in office buildings, Appl. Energy 245 (2019) 65–77.

[25] X. Liu, M. Ren, Z. Yang, G. Yan, Y. Guo, L. Cheng, C. Wu, A multi-step predictive deep reinforcement learning algorithm for hvac control systems in smart buildings, Energy 259 (2022) 124857.

[26] S. Zhou, A. Shah, P. Leung, X. Zhu, Q. Liao, A comprehensive review of the applications of machine learning for hvac, DeCarbon (2023) 100023.

[27] M.C. Campi, A. Lecchini, S.M. Savaresi, Virtual reference feedback tuning: a direct method for the design of feedback controllers, Automatica 38 (8) (2002) 1337–1346.

[28] P. Kergus, S. Formentin, M. Giuliani, A. Castelletti, Learning-based hierarchical control of water reservoir systems, IFAC J. Syst. Control. 19 (2022) 100185.

[29] L. Ljung, System identification, in: Signal Analysis and Prediction, Springer, 1998, pp. 163–173.

[30] P. Stoffel, L. Maier, A. Kümpel, T. Schreiber, D. Müller, Evaluation of advanced control strategies for building energy systems, Energy Build. 280 (2023) 112709, http://dx.doi.org/10.1016/j.enbuild.2022.112709.

[31] Grey-box modeling and application for building energy simulations - a critical review, Renew. Sustain. Energy Rev. 146 (2021) 111174, http://dx.doi.org/10.1016/j.rser.2021.111174.

[32] A. Lefort, R. Bourdais, G. Ansanay-Alex, H. Guéguen, Hierarchical control method applied to energy management of a residential house, Energy Build. 64 (2013) 53–61.

[33] C. Vallianos, A. Athienitis, B. Delcroix, Automatic generation of multi-zone rc models using smart thermostat data from homes, Energy Build. 277 (2022) http://dx.doi.org/10.1016/j.enbuild.2022.112571.

[34] J. Arroyo, F. Spiessens, L. Helsen, Identification of multi-zone grey-box building models for use in model predictive control, J. Build. Perform. Simul. 13 (4) (2020) 472–486.

[35] E. Biyik, A. Kahraman, A predictive control strategy for optimal management of peak load, thermal comfort, energy storage and renewables in multi-zone buildings, J. Build. Eng. 25 (2019) 100826.

[36] F. Langner, M. Frahm, W. Wang, J. Matthes, V. Hagenmeyer, Hierarchical-stochastic model predictive control for a grid-interactive multi-zone residential building with distributed energy resources, J. Build. Eng. 89 (2024) 109401.

[37] Y. Zhang, Z. O'Neill, B. Dong, G. Augenbroe, Comparisons of inverse modeling approaches for predicting building energy performance, Build. Environ. 86 (2015) 177–190.

[38] Z. Afroz, G. Shafiullah, T. Urmee, G. Higgins, Modeling techniques used in building hvac control systems: A review, Renew. Sustain. Energy Rev. 83 (2018) 64–84.

[39] A. Daneels, W. Salter, What is scada? 1999.

[40] D. Simon, Optimal State Estimation: Kalman, H Infinity, and Nonlinear Approaches, John Wiley & Sons, 2006.
[41] D. Blum, J. Arroyo, S. Huang, J. Drgoňa, F. Jorissen, H.T. Walnum, Y. Chen, K. Benne, D. Vrabie, M. Wetter, et al., Building optimization testing framework (boptest) for simulation-based benchmarking of control strategies in buildings, J. Build. Perform. Simul. 14 (5) (2021) 586–610.
[42] Matlab optimization toolbox. URL https://fr.mathworks.com/help/optim/index.html?s_tid=CRUX_lftnav.
[43] Matlab system identification toolbox. URL https://fr.mathworks.com/products/sysid.html.
[44] J.A.E. Andersson, J. Gillis, G. Horn, J.B. Rawlings, M. Diehl, CasADi – a software framework for nonlinear optimization and optimal control, Math. Program. Comput. 11 (1) (2019) 1–36, http://dx.doi.org/10.1007/s12532-018-0139-4.
[45] V. Breschi, S. Formentin, Direct data-driven control with embedded anti-windup compensation, in: Learning for Dynamics and Control, PMLR, 2020, pp. 46–54.
[46] R. Baetens, R. De Coninck, F. Jorissen, D. Picard, L. Helsen, D. Saelens, Openideas-an open framework for integrated district energy simulations, in: Building Simulation, 2015, pp. 345–354.
[47] M. Wetter, W. Zuo, T.S. Nouidui, X. Pang, Modelica buildings library, J. Build. Perform. Simul. 7 (4) (2014) 253–270.
[48] E. Lorenz, J. Hurka, D. Heinemann, H.G. Beyer, Irradiance forecasting for the power prediction of grid-connected photovoltaic systems, IEEE J. Sel. Top. Appl. Earth Obs. Remote. Sens. 2 (1) (2009) 2–10.
[49] E. Lorenz, J. Kühnert, D. Heinemann, Short term forecasting of solar irradiance by combining satellite data and numerical weather predictions, in: Proc. of 27th EUPVSEC, 2012, pp. 4401–4405.
[50] R. De Coninck, L. Helsen, Practical implementation and evaluation of model predictive control for an office building in brussels, Energy Build. 111 (2016) 290–298.
[51] H.J. Ferreau, C. Kirches, A. Potschka, H.G. Bock, M. Diehl, Qpoases: A parametric active-set algorithm for quadratic programming, Math. Program. Comput. 6 (2014) 327–363.