Hierarchical Model Predictive Control for Building Thermal Management : A mixed model-based and data-based approach

Yuqi LIU[1,2], Pauline KERGUS[3], Fabien CLAVEAU[1,2], Philippe CHEVREL[1,2]

Abstract— The diversity of building thermal systems poses significant challenges for control-oriented modeling. This paper presents a hierarchical control framework for building thermal management that addresses these challenges at both supervisory and tracking levels. At the supervisory level, a model predictive controller employs a lumped physics-based model, independent of the installed equipment, and limitations of instantaneous thermal power response are accounted for as constraints. At the tracking level, PI controllers with anti-windup are designed using Virtual Reference Feedback Tuning, a model-free approach that circumvents the need for explicit thermal system modeling. This hierarchical control strategy effectively manages thermal system characteristics without requiring detailed thermal equipment models. The proposed framework is illustrated using the Building Optimization Testing Framework, a standardized virtual testbed, demonstrating its relevance and performance for building climate control.

I. Introduction

The International Energy Agency reports that buildings account for 30-40% of global energy consumption, with one-third attributed to building thermal management [1]. Advanced control methods, particularly Model Predictive Control (MPC), have shown promising results in this field, potentially reducing energy consumption while maintaining comfort [2].

The performance of MPC in this context largely depends on the predictive model's accuracy. Some approaches develop physic-based models for the entire building energy system [3], [4], while others adopt data-based modeling solutions to reduce modeling effort [5], [6]. Both strategies aim to construct comprehensive building system models, simulating indoor climate alongside Heating, Ventilation and Air Conditioning (HVAC) devices. In this paper, HVAC stand for a generic class of building thermal service equipment.

However, both approaches often result in control designs that are highly system-specific, necessitating the construction of a HVAC model for each building. This can be time-consuming, particularly given the diversity of HVAC technologies currently in use. Moreover, it poses a significant challenge in terms of scalable deployment as the HVAC modeling procedure would need to be done for every single building.

Analyzing the system complexity reveals a natural separation in modeling requirements:

1) The thermal characteristics of buildings can be modeled using relatively simple and performant physics-based prediction models, capturing essential elements like daylight conditions, insulation properties, and thermal capacity.
2) Modeling the specific nature of HVAC systems can be complex and time-consuming, depending on the equipment mechanisms, usage patterns, and circuit topology implemented. Furthermore, these systems often exhibit nonlinear behaviors and operational constraints that make them challenging to control.

To address these challenges, this paper proposes a hierarchical control framework that systematically decouples the indoor climate model from HVAC-specific dynamics while explicitly addressing both equipment diversity and data limitations. The key novelty lies in the integration of a grey-box MPC with model-free Virtual Reference Feedback Tuning (VRFT) at distinct hierarchical levels, enabling scalable deployment across diverse HVAC configurations without requiring detailed equipment models.

At the supervisory level, an MPC strategy based on a lumped physics-based model of the building's thermal dynamics is employed to handle 1. This model predicts multi-zone thermal evolution based on thermal power inputs, which will be provided by the HVAC devices at the tracking level, regardless of the specific equipments installed. The MPC incorporates a input rate constraint to account for the limitation of the instantaneous thermal power response assumption at the tracking level, thereby leveraging the thermal inertia in HVAC device.

At the tracking level, a set of PID controllers with anti-windup is implemented, tuned using the VRFT approach [7]. Inspired by the hierarchical combination of MPC and VRFT in [8] [9], this paper adapts the approach to building thermal management to tackle 2.

The rest of the paper is organized as follows: Section II introduces the case study and details the building thermal modeling approach. Section III presents the hierarchical control scheme, including the supervisory level MPC and the tracking level VRFT approach. Section IV discusses the results of the proposed methodology applied to the Building Optimization Testing Framework (BOPTEST) [10]. Finally, Section V concludes the paper and outlines future research directions.

[1]IMT Atlantique, LS2N, UMR CNRS 6004, F-44307 Nantes, France
[2]LAPLACE, CNRS, INPT, Université de Toulouse.
Contact:yuqi.liu@imt-atlantique.fr

sys.png sys.bb

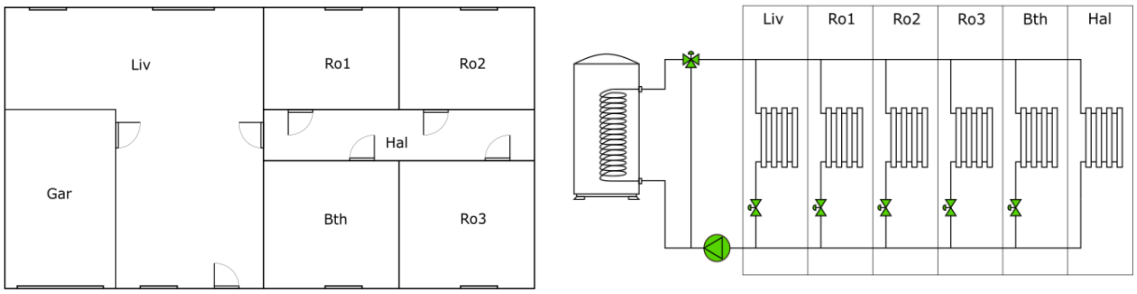

Fig. 1: Adjacency of the test case and hydronic circuit of heating system.

II. Problem formulation

To validate the proposed methodology, we employ the BOPTEST, a standardized virtual testbed for evaluating building control strategies.

A. System description

As illustrated in Figure 1, the 'Multi-zone residential hydronic' case comprises multiple heating zones, including an attic (not shown in the figure). Hydraulic radiators are installed in all zones except the attic and garage. Flow regulation is achieved through individual valve control for each radiator, with the exception of the hallway radiator, which serves to balance the distribution circuit pressure. The heating system consists of a gas boiler for hot water supply, a main pump for total flow control, and a three -way valve for supply temperature regulation through mixing of boiler hot water and return water.

The baseline controllers provided with the benchmark consist of multiple PID and Rule-Based Control to achieve the following objectives:

- Maintain boiler water temperature within safety boundaries.
- Track supply water temperature set-point according to the outside temperature.
- Minimizes the temperature error in the living room through main pump control.
- Tracking indoor temperature set-point.

The test case is modeled after a French residential dwelling compliant with the 2012 French thermal regulations. The system includes various nonlinear elements, such as valve flow characteristics, heat transfer processes, and solar radiation through glazing surfaces. The thermal dynamics involve multiple time scales, influenced by factors including building thermal mass and radiator heat storage. The system operates under changing occupancy patterns and internal heat loads, along with climate factors.

Simulations use climate data from Bordeaux, France, recorded in 2014.

B. Considered control architecture

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 methodology proposed in this paper, a hierarchical control scheme for valve systems is implemented. It segments the thermal management system based on power signals, aligning with the thermal response model. As shown in Figure 2, the architecture consists of two primary levels:

- A supervisory level employing MPC to handle long-term objectives and slow dynamics, planning energy

scheme.png scheme.bb

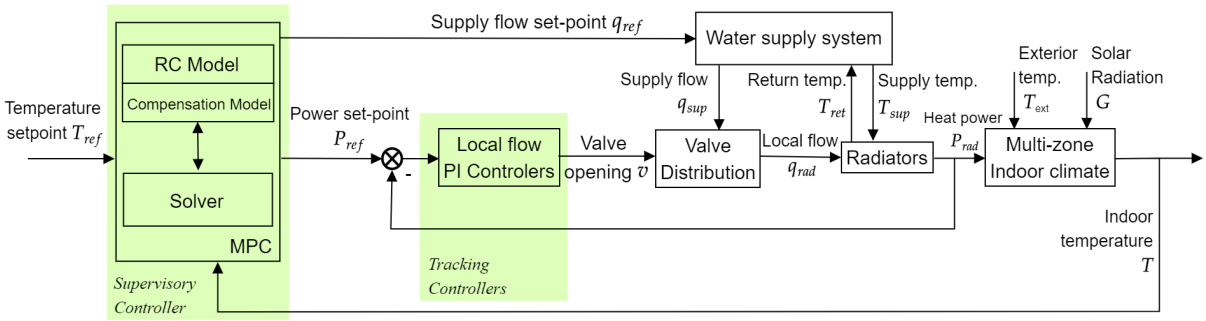

Fig. 2: Diagram of 'Multi-zone residential hydronic' test case with proposed solution. Boiler and mixture circuit are simplified as water supply system block. The hierarchical control is deployed on the flow distribution system. In this paper, system variables $y_{ref}, u_{ref}, v, u, y$ are composed of signals $T_{ref}, P_{ref}, v, P, T$ respectively.

distribution across zones to balance thermal comfort and consumption.
- A tracking level using PID controllers with anti-windup to manage immediate power demands and regulate local actuators, such as valve openings, pump flows, and turbines.

This hierarchical structure enables the system to effectively handle both long-term optimization and rapid local adjustments, adapting to slow system changes and quick disturbances while maintaining overall performance goals.

For this study, we assume perfect knowledge of external temperature and solar irradiation for the duration of the prediction horizon. This idealized approach allows us to focus on the core performance of the control strategy without the added complexity of forecast uncertainties. It's worth noting that in real-world applications, weather forecasting and solar irradiation modeling have been extensively studied and sophisticated prediction methods are available [11], [12].

III. Hierarchical control scheme

This section details the proposed hierarchical control scheme. Section III-A presents the identification of prediction model, a physics-based lumped RC-model. Then Section III-B and Section III-C show the model-based supervisory level and data-driven tracking level controllers, respectively.

A. Building thermal modeling

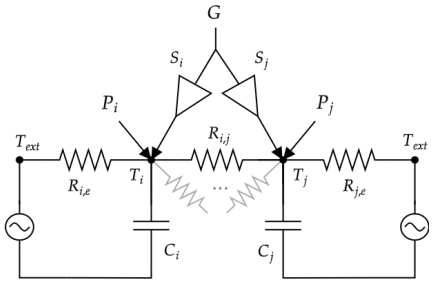

Fig. 3: Equivalent RC circuit of multi-zone modeling, showing two sample rooms as an example. Triangle shapes in the diagram represent equivalent surfaces, indicating the gain from solar radiation to absorbed radiative power.

This study adopts a centralized modeling approach for multi-zone thermal systems, as proposed by [13]. The 1st Order Resistance-Capacitance (RC) Multi-Zone model is employed to represent the indoor climate of a multi-zone building, with each zone modeled as a first-order RC circuit, as shown in Figure 3

The multi-zone RC model can be expressed in a continuous state-space form as:

$$\frac{\mathrm{d}y}{\mathrm{d}t} = f_{RC}(\theta, y, u, d) = A(\theta)y + B_1(\theta)u + B_2(\theta)d \quad (2)$$

where:

- $y \in \mathbb{R}^n$ is the output (as well as state) vector representing zone temperatures: $y = [T_1, T_2, \ldots, T_n]^\top$
- $u \in \mathbb{R}^n$ is the input vector of thermal powers: $u = [P_1, P_2, \ldots, P_n]^\top$
- $d \in \mathbb{R}^2$ is the disturbance vector: $d = [G, T_{\text{ext}}]^\top$, where G is the solar radiation and T_{ext} is the external temperature
- θ is the parameter vector to be estimated:

$$\theta = [R_{1,2}, \ldots, R_{n,n-1}, R_{1,e}, \ldots, R_{n,e}, C_1, \ldots, C_n, S_1, \ldots, S_n]^\top \quad (3)$$

where $R_{i,j}$ are thermal resistances between zones, $R_{i,e}$ are thermal resistances to the exterior, C_i are thermal capacitances, and S_i are equivalent solar radiation surfaces.

The matrices $A(\theta)$, $B_1(\theta)$, and $B_2(\theta)$ are defined as Equation (1) according to lumped 1st-order thermal dynamic.

In this formulation, $Ad_{i,j}$ indicates the adjacency between zones i and j, defined as:

$$Ad_{i,j} = \begin{cases} 0 & \text{if } j \notin N_{ad,i} \\ 1 & \text{if } j \in N_{ad,i} \end{cases} \quad (4)$$

where $N_{ad,i}$ is the set of zones adjacent to zone i. Historical data of the testcase baseline solution is used for system identification. According to Equation (1), the estimation of parameter vector Equation (3) can be formulated to least square form.

Once the optimal solution θ^* has been determined, the discret predictive model is defined as:

$$\begin{aligned} \hat{y}(t_{k+1}) &= f_d(\hat{y}(t_k), u(t_k), \hat{d}(t_k)) \\ &= \textstyle\int_{t_k}^{t_{k+1}} (A(\theta^*)\hat{y}(t) + B_1(\theta^*)u(t) + B_2(\theta^*)\hat{d}(t))dt \end{aligned} \quad (5)$$

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

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

where $\hat{y}$ represents the predicted zone temperatures and $\hat{d}$ is the forecast disturbance.

B. Supervisory level design: MPC approach

Assuming an ideal tracking level control, the typical optimization problem solved by MPC at this level is Equation (6)

Nominal MPC

$$\min_{u_{ref}} \sum_{t=t_k}^{t_{k+H}} (e(t_{k+1})^T Q e(t_{k+1}) + u_{ref}(t_k)^T R u_{ref}(t_k)) \tag{6a}$$

$$e(t_{k+1}) = y_{ref}(t_{k+1}) - \hat{y}(t_{k+1}) \tag{6b}$$

$$\hat{y}(t_{k+1}) = f_d(\hat{y}(t_k), u_{ref}(t_k), \hat{d}(t_k)) \tag{6c}$$

$$\underline{u} \leq u_{ref}(t_k) \leq \bar{u} \tag{6d}$$

At each time step t_k, the MPC solves Equation (6), optimizing the power reference trajectory u_{ref} within limits $\underline{u}$ and $\bar{u}$. The cost function in Equation (6a) balances thermal comfort (first term) and energy consumption (second term) over the prediction horizon H. Here, e is the temperature error, Q and R are weighting matrices (with $Q = I$ and $R = wI$, where w is a scalar parameter), $\hat{y}$ represents predicted temperatures, and f_d is the thermal model from Equation (5). The disturbance $\hat{d}$ includes forecasted radiation and exterior temperature. This formulation is referred to as 'Nominal MPC'.

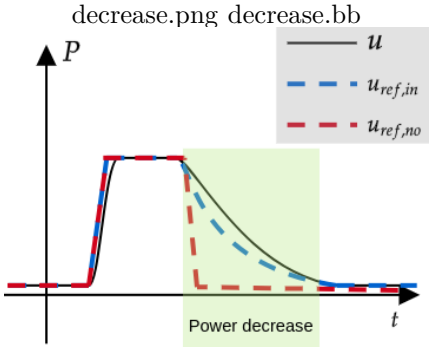

Fig. 4: Conceptual illustration of hydronic heating equipment thermal power response. A nominal MPC suggests power reference $u_{ref,no}$ that waits for power decrease process. An rate constraint contributes to a closer prediction to local response, like $u_{ref,in}$

Nominal MPC approaches often assume rapid power response capabilities in heating systems, expecting actuators to achieve abrupt changes in thermal power output. However, in water-based radiator systems like the considered system, this assumption proves problematic due to significant thermal inertia.

Figure 4 illustrates this challenge by showing how the system's actual power response (black curve) evolves when given different reference trajectories. A nominal MPC generating steep power decrease commands (red dashed curve) must wait for the radiator's gradual cooling process, which can span several MPC time steps. In contrast, an MPC with enhanced prediction capabilities (blue dashed curve) fully accounts for this dynamic. Even if both MPC formulations result in the same actual power response, the latter allows earlier reactions to setpoint changes and optimizes residual energy utilization, potentially improving overall energy efficiency.

The thermal inertia of the HVAC devices is then accounted for as a rate constraint [14] in the proposed MPC supervisory control. The input rate is lower bounded, as shown in Equation (7)

$$u_{ref}(t_{k+1}) - u_{ref}(t_k) \geq \Delta u_{min} \tag{7a}$$

$$\Delta u_{min} = A_{in} u_{ref}(t_k) - u_{ref}(t_k) \tag{7b}$$

$$\Rightarrow u_{ref}(t_{k+1}) - A_{in} u_{ref}(t_k) \geq 0 \tag{7c}$$

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. Equation (7c) 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 increment of thermal power remain non-constrained.

Only the data relating to the reduction in power is required to identify the model in Equation (7c). The time index of the training set is

$$D_{in} = \{t_k | (\bar{u}(t_k) - \bar{u}(t_{k+1}) \geq 0, \forall t_k \in \{1, ..., n_{id}\}\} \tag{8}$$

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

$$\min_{A_{in}} \sum_{t_k \in D_{in}} ||u(t_{k+1}) - A_{in} u(t_k)||^2 \tag{9}$$

In the end, optimization can be formalized as follows:

Rate-constrained MPC

$$\min_{u_{ref}} \sum_{t=t_k}^{t_{k+H}} (e(t_{k+1})^T Q e(t_{k+1}) + u_{ref}(t_k)^T R u_{ref}(t_k)) \tag{10a}$$

$$e(t_k) = y_{ref}(t_k) - \hat{y}(t_k) \tag{10b}$$

$$\hat{y}(t_{k+1}) = f_d(\hat{y}(t_k), u_{ref}(t_k), \hat{d}(t_k)) \tag{10c}$$

$$u_{ref}(t_{k+1}) - A^*_{in} u_{ref}(t_k) \geq 0 \tag{10d}$$

$$\underline{u} \leq u_{ref}(t_k) \leq \bar{u} \tag{10e}$$

where Equation (10d) is the internal loop constraint corresponding to Equation (7c).

The CasADi package [15] is used to implement and solve the eq. (10) efficiently. Multiple shooting is adopted in this study, due to its advantages in feasibility and computational speed [16].

This Rate-constrained MPC formulation is effective for hydronic systems like the the considered system, providing a balance between model simplicity and effective control. It has potential applicability to more HVAC type. The specific constraints and dynamics in Equation (7c) can be modified for thermal power increase rapidity constraint, according to HVAC mechanism. For instance, gas boiler's power rising stage is possible to be slower than MPC time step as well.

C. Tracking Level Design: VRFT Approach

At the tracking level, a set of PID controllers with anti-windup is implemented and tuned using the VRFT approach [7]. VRFT is a data-driven controller tuning method that avoids the need for explicit system modeling, making it particularly suitable for the diverse equipment encountered in building systems.

The VRFT method optimizes controller parameters to approximate an ideal controller C^* using only input-output data from the system G (Figure 5). It generates a virtual reference u_{ref} by reverse-simulating the desired closed-loop performance M_{ref}. This virtual reference allows VRFT to construct the necessary controller input data and tune the parameters to match the ideal controller's performance without requiring an explicit system model.

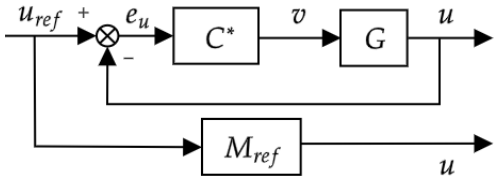

Fig. 5: Block diagram of VRFT process. C^* represents an ideal controller that drives the closed-loop to achieve the desired performance M_{ref}.

The optimal PID parameters θ_{PID} are obtained by solving the optimization problem:

$$\begin{aligned}\min_{\theta_{PID}} \sum_{k=1}^{n_{id}} \|\varepsilon(t_k)\|^2 \\ \text{s.t. } \varepsilon(t_k) = v(t_k) - C_{PID}(q^{-1})\theta_{PID} e_u(t_k), \\ u(t_k) = M_{ref}(q^{-1}) u_{ref,vr}(t_k) \\ e_u(t_k) = u_{ref,vr}(t_k) - u(t_k).\end{aligned} \tag{11}$$

where $\varepsilon(t_k)$ is the error between the actual and desired controller outputs, C_{PID} is the PID controller transfer function, and e_u is the tracking error.

To account for actuator saturation, an anti-windup design is incorporated [17]. The optimization problem is expanded to include the anti-windup parameters:

$$\begin{aligned}\min_{\theta_{PID},\theta_{aw}} \sum_{k=2}^{n_{id}} \|\bar{v}(t_k) - C_{PID}(q^{-1})\theta_{PID}(M_{ref}^{\dagger}(q^{-1}) - 1)\bar{u}(t_{k-1}) \\ - C_{aw}(q^{-1})\theta_{aw}\Delta\bar{v}_{sat}(t_{k-1})\|^2\end{aligned} \tag{12}$$

where $\Delta\bar{v}_{sat}$ is the difference between the controller output before and after saturation, C_{aw} is the anti-windup compensator transfer function, and $\theta_{aw} = [K_{aw,1}, ..., K_{aw,n}]$ are the anti-windup gains to be tuned. When generating the dataset in open-loop, inputs should exceed the saturation limits to ensure that $\Delta\bar{v}_{sat}$ data is informative.

This formulation allows for simultaneous tuning of both the PID controller parameters and the anti-windup gains. The optimization problem can be solved using quadratic programming solvers, such as the quadprog() function in MATLAB's Optimization Toolbox.

The reference model M_{ref} is chosen as a first-order system with a zero added to ensure causality in the optimization process:

$$M_{ref,i}(q^{-1}) = ZOH(\frac{\frac{1}{10}\tau s + 1}{\tau s + 1}) = \frac{\beta_0 + \beta_1 q^{-1}}{1 - q^{-1}e^{-T_{PID}/\tau}} \tag{13}$$

where τ is the desired time constant and T_{PID} is the discrete time step of the PID controller.

This VRFT approach provides a flexible and efficient method for tuning the tracking level controllers, adapting to various heating and cooling equipment without requiring detailed system models. The resulting PID controllers with anti-windup ensure robust performance across the diverse range of actuators typically found in building thermal management systems.

IV. RESULTS

This section presents the performance evaluation of our proposed hierarchical control scheme. The analysis is based on data generated from the test case in the BOPTEST framework, using the baseline control plan from January 1 to January 21. Our evaluation encompasses the test of VRFT-tuned PI controllers, isolated validation of these controllers, and full system validation in BOPTEST. We examine the parameters obtained for different zones, assess the tracking performance of the controllers, and evaluate the overall performance of our hierarchical control scheme, including the effectiveness of rate constraint in the MPC and temperature control under different prediction horizons. All simulations, except for the isolated VRFT validation, were conducted within the BOPTEST benchmark environment. The validation period spans from January 22 to January 29. Table I

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	$4.0e^{-15}$	$1.1e^{-16}$

TABLE I: Each zone PI parameters tuned by VRFT

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.

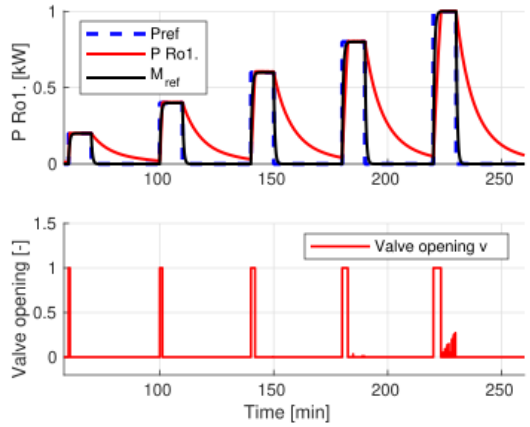

Fig. 6: Internal loop simulation that validate VRFT tuned controllers. Zone Ro1. is shown as a sample. Hot water temperature, zone temperature and water flow is set to constant. Power reference is square wave with 10 minutes heating and 40 minutes stand-by.

Figure 6 validates the VRFT-tuned controllers in an isolated closed-loop simulation. The power response shows fast and accurate tracking during power increase. Power decrease is slower than the reference model M_{ref} desired. During decrease, the valve opening saturates at 0. The antiwindup mechanism effectively prevents continuing integration, avoiding delays in subsequent responses. For this study, we chose a reference model time constant τ of 1 minute, which aligns with the guideline of being approximately 1/10 of the MPC time step, see Table II. 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.

Figure 7 demonstrates the effect of the rate constraint in the BOPTEST environment. Rate-constrained MPC in Equation (10) has a lower bounded by a identified thermal inertia on equipment thermal power decrease, while nominal MPC in Equation (6) has no limit. Configuration of two formulation is shown in Table II

Rate-constrained MPC aligns better with the actual power response. At 10:00, when the residential building changes to non-occupied mode, the nominal MPC provides a power reference that quickly drops to 0. However, the heating equipment cannot follow immediately, forcing the MPC 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 and starting the power decrease process earlier than the nominal MPC. This approach fully manages the equipment's thermal inertia.

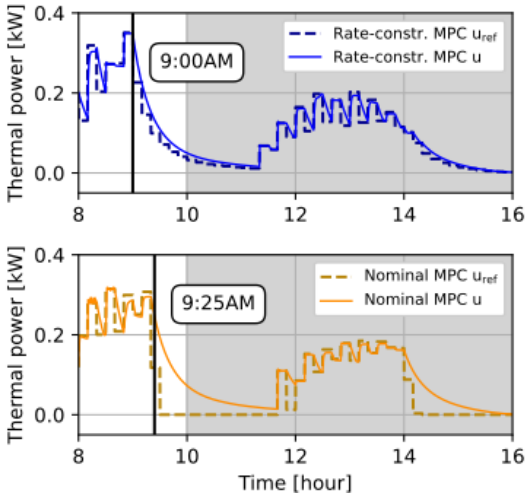

Fig. 7: Radiator power tracking performance. Two subplots compare Rate-constrained MPC in Equation (10) and nominal MPC in Equation (6). The grey zone represents non-occupied time requiring lower temperature demand.

Parameters	Values	Description
T_{MPC}	10 minutes	MPC control time step
H	6 hours	MPC prediction horizon
Q	$I_{(8,8)}$	Weight on state error
R	$wI_{(6,6)}$	Weight on power input
Solver	'qpoases'[18]	Solver in MPC
w	$w_0 = 1.6e^{-4}$	Weighting on objectives

TABLE II: MPC parameters used in test case

Figure 8 illustrates the temperature evolution in a bedroom (Ro2) under different prediction horizons, using BOPTEST climate data from January 22 to January 29. The results demonstrate the impact of prediction horizon on control performance. Notably, the 6-hour prediction horizon shows improved performance during daytime hours, particularly evident on days 26-27. The longer horizon allows the controller to anticipate temperature changes more effectively, resulting in less extreme temperature dips compared to the 4-hour horizon. This behavior highlights the controller's ability to balance comfort and energy efficiency more effectively with a longer prediction window.

V. CONCLUSION

This paper presents a novel hierarchical control framework for building thermal management, addressing challenges of equipment diversity and model complexity in

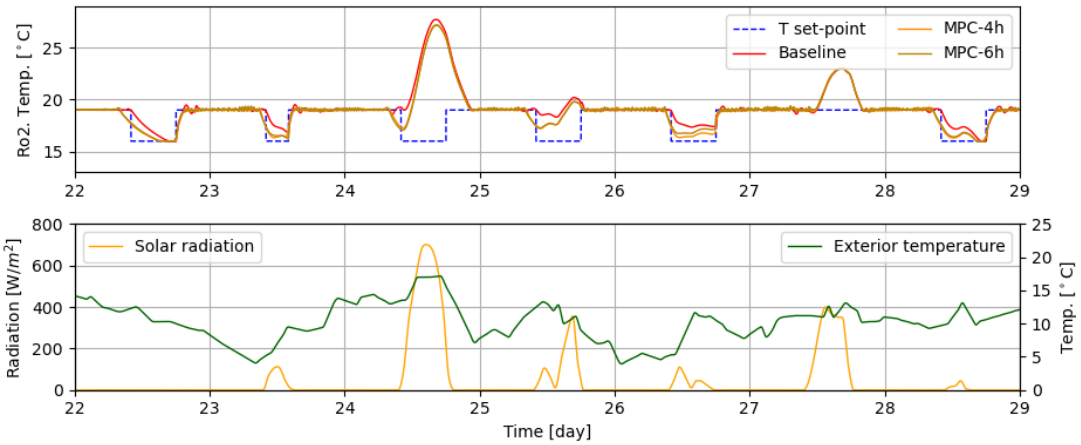

Fig. 8: Performance comparison of the proposed control scheme with 4-hour and 6-hour prediction horizons in the BOPTEST environment. The figure shows the temperature evolution in bedroom Ro2, along with the set-point, baseline performance, and disturbance variations (solar radiation and exterior temperature).

MPC implementations. Our two-level approach is based on separating building thermal characteristics, which can be accurately described by simple yet accurate physics-based model, from the HVAC devices, for which the modelling task is highly system specific and is therefore a challenge in large-scale deployment of advanced control techniques. The proposed methodology is validated on the BOPTEST testbed. Results indicate that this approach achieves tight coordination between thermal power demand planning and local power tracking. Future research will aim to demonstrate the potential benefits and possibilities of the proposed hierarchical control methodology, including its scalability and adaptation to various building types.

References

[1] IEA, "World energy outlook 2015," IEA, 2015.

[2] F. Oldewurtel, C. N. Jones, A. Parisio, and M. Morari, "Stochastic model predictive control for building climate control," IEEE Transactions on Control Systems Technology, vol. 22, no. 3, pp. 1198–1205, 2014.

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

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

[5] F. Gauthier-Clerc, H. Le Capitaine, F. Claveau, and 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.

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

[7] M. C. Campi, A. Lecchini, and S. M. Savaresi, "Virtual reference feedback tuning: a direct method for the design of feedback controllers," Automatica, vol. 38, no. 8, pp. 1337–1346, 2002.

[8] P. Kergus, S. Formentin, M. Giuliani, and A. Castelletti, "Learning-based hierarchical control of water reservoir systems," IFAC Journal of Systems and Control, vol. 19, p. 100185, 2022.

[9] D. Piga, S. Formentin, and A. Bemporad, "Direct data-driven control of constrained systems," IEEE Transactions on Control Systems Technology, vol. 26, no. 4, pp. 1422–1429, 2017.

[10] 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," Journal of Building Performance Simulation, vol. 14, no. 5, pp. 586–610, 2021.

[11] E. Lorenz, J. Hurka, D. Heinemann, and H. G. Beyer, "Irradiance forecasting for the power prediction of grid-connected photovoltaic systems," IEEE Journal of selected topics in applied earth observations and remote sensing, vol. 2, no. 1, pp. 2–10, 2009.

[12] R. De Coninck and L. Helsen, "Practical implementation and evaluation of model predictive control for an office building in brussels," Energy and Buildings, vol. 111, pp. 290–298, 2016.

[13] C. Vallianos, A. Athienitis, and B. Delcroix, "Automatic generation of multi-zone rc models using smart thermostat data from homes," Energy and Buildings, vol. 277, 10 2022.

[14] D. Mayne, "Model predictive control theory and design," Nob Hill Pub, Llc, 1999.

[15] J. A. E. Andersson, J. Gillis, G. Horn, J. B. Rawlings, and M. Diehl, "CasADi – A software framework for nonlinear optimization and optimal control," Mathematical Programming Computation, vol. 11, no. 1, pp. 1–36, 2019.

[16] J. Drgoňa, J. Arroyo, I. Cupeiro Figueroa, D. Blum, K. Arendt, D. Kim, E. P. Ollé, J. Oravec, M. Wetter, D. L. Vrabie, and L. Helsen, "All you need to know about model predictive control for buildings," Annual Reviews in Control, vol. 50, pp. 190–232, 2020.

[17] V. Breschi and S. Formentin, "Direct data-driven control with embedded anti-windup compensation," in Learning for Dynamics and Control, pp. 46–54, PMLR, 2020.

[18] H. J. Ferreau, C. Kirches, A. Potschka, H. G. Bock, and M. Diehl, "qpoases: A parametric active-set algorithm for quadratic programming," Mathematical Programming Computation, vol. 6, pp. 327–363, 2014.