Titre: Title:	Decentralized Mixed-Integer Temperature Optimization for Multi-Zone Commercial Buildings
Auteur: Author:	Étienne Tremblay
Date:	2025
Type:	Mémoire ou thèse / Dissertation or Thesis
Référence: Citation:	Tremblay, É. (2025). Decentralized Mixed-Integer Temperature Optimization for Multi-Zone Commercial Buildings [Master's thesis, Polytechnique Montréal]. PolyPublie. https://publications.polymtl.ca/68390/

URL de PolyPublie: PolyPublie URL:	https://publications.polymtl.ca/68390/
Directeurs de recherche: Advisors:	Antoine Lesage-Landry
Programme: Program:	Génie électrique

POLYTECHNIQUE MONTRÉAL
affiliée à l'Université de Montréal

Decentralized Mixed-Integer Temperature Optimization for Multi-Zone Commercial Buildings

ÉTIENNE TREMBLAY
Département de génie électrique

Mémoire présenté en vue de l'obtention du diplôme de *Maîtrise ès sciences appliquées*
Génie électrique

Août 2025

POLYTECHNIQUE MONTRÉAL
affiliée à l'Université de Montréal

Ce mémoire intitulé :

Decentralized Mixed-Integer Temperature Optimization for Multi-Zone Commercial Buildings

présenté par **Étienne TREMBLAY**
en vue de l'obtention du diplôme de *Maîtrise ès sciences appliquées*
a été dûment accepté par le jury d'examen constitué de :

Hanane DAGDOUGUI, présidente
Antoine LESAGE-LANDRY, membre et directeur de recherche
Michaël KUMMERT, membre

DEDICATION

Life moves pretty fast.
If you don't stop and look
around once in a while,
you could miss it.

ACKNOWLEDGEMENTS

I would like to acknowledge the collaboration of BrainBox AI (BBAI) and the support of the MITACS program in making this work possible. A very special thank you goes to the Bourse d'Excellence Maurice Brisson – BBA for their generous and continued support throughout my postgraduate studies.

Thank you to Marie-Christine for breaking the ice on this project and helping me find my rhythm.

I'm deeply grateful to everyone who makes working enjoyable—the people in the lab, my former classmates now thriving in their respective fields, and my professors. I also want to thank all those who have nothing to do with engineering, school, or research: my friends and family, who remind me that everything I do outside of work gives meaning to the time I spend at work. Thank you to everyone who supported me just by being there. A special shoutout to M. Rémi Le Ny.

Thank you to sports—every single one of them: hockey, baseball, running, cycling, swimming, hiking, football, Australian football, soccer, climbing, and many more. They have kept me going and entertained, and I know they won't let me down—just like the people I share them with. Special thanks to Côte-des-Neiges and the Saint Joseph's Oratory for being a warm and grounding home.

I want to meaningfully and sincerely acknowledge the contribution of Prof. Antoine Lesage-Landry to this work—not only for the tangible effort he put in, but for his trust in me, his belief in the process, and his unwavering conviction that there is always a way forward.

Finally, I want to acknowledge the uncertainty of where this journey will lead. I entered research to see if it was something I would enjoy and as a way to give back to the incredible higher education system in Québec. I've worked hard, chasing an uncertain goal, but I've truly enjoyed the journey, and I'm glad I did it. At the risk of sounding vain, I want to recognize the effort I've put into this work and say, that I'm proud of it.

RÉSUMÉ

Les bâtiments commerciaux représentent une part importante de la consommation énergétique mondiale, les systèmes de chauffage, ventilation et climatisation (CVC) en constituant une portion substantielle. L'optimisation du fonctionnement des systèmes CVC permet non seulement d'améliorer l'efficacité énergétique, mais aussi de réduire les pointes de demande sur le réseau électrique, notamment lorsqu'ils sont intégrés à des programmes de gestion de la demande de puissance. Toutefois, ces gains doivent être réalisés sans compromettre le confort des occupants, car une baisse de confort pourrait nuire à la participation à ces programmes. La modélisation précise de la dynamique thermique d'un bâtiment est essentielle à une commande efficace, chaque bâtiment présentant un comportement unique qui nécessite une identification soignée du système. Par ailleurs, les bâtiments commerciaux utilisent fréquemment des unités de toit (*roof-top unit* RTU), qui fonctionnent selon des modes discrets, par exemple, plusieurs niveaux de chauffage/refroidissement et un ventilateur à fonctionnement tout ou rien, ce qui complique l'utilisation de stratégies de commande continues. Dans ce mémoire, nous proposons une approche décentralisée de commande prédictive à variables discrètes pour des bâtiments commerciaux à zones multiples. Cette stratégie de contrôle repose sur un modèle linéaire guidé par la physique et entraîné à partir de données simulées ou historiques, ainsi que de prévisions des conditions météorologiques. Grâce à sa formulation décentralisée à variables discrètes, l'approche est à la fois deployable à grande échelle et compatible avec les systèmes RTU existants, évitant ainsi des mises à niveau d'infrastructure coûteuses. Enfin, la méthode proposée est validée dans un environnement de simulation de bâtiment, en tenant compte des structures tarifaires à tarification selon l'heure d'utilisation (Time-of-Use), à tarification critique (Critical Peak Pricing) et à tarification en temps réel (Real-Time Pricing). Les résultats démontrent que le contrôleur basé sur l'optimisation, dans ses versions centralisée et décentralisée appliquées au bâtiment à deux zones, surpasse la stratégie de contrôle tout ou rien de référence. Bien que le contrôleur centralisé appliqué au bâtiment à cinq zones présente une capacité limitée à s'adapter à des configurations à grande échelle, il offre tout de même une performance supérieure à celle du contrôleur tout ou rien.

.

ABSTRACT

Commercial buildings represent a major share of global energy consumption, with heating, ventilation, and air conditioning (HVAC) systems accounting for a substantial portion of that use. Optimizing HVAC operations not only improves energy efficiency but also reduces peak loads on the electrical grid when integrated with demand response programs. However, these benefits must be achieved without compromising occupant comfort, as discomfort can hinder program participation. Accurately modelling a building's thermal dynamics is essential for effective control, yet each building exhibits unique behaviour that requires careful system identification. Morevover, commercial buildings often rely on rooftop unit (RTU) HVAC systems, which operate in discrete modes ,e.g., multiple heating/cooling stages and an ON/OFF fan control, posing challenges for continuous control strategies. In this Master's thesis, we propose a decentralized mixed-integer model predictive control (MPC) framework for multi-zone commercial buildings Heating, ventilation, and air conditioning (HVAC) optimization. The control strategy is built upon a data-driven, physics-informed linear model trained using simulated or historical data, as well as forecasted external conditions. By leveraging the mixed-integer decentralized MPC formulation, the approach is scalable and compatible with existing RTU systems, avoiding the need for costly infrastructure upgrades. Finally, the proposed method is validated in a building simulation environment, considering Time-of-Use, Critical Peak Pricing, and Real-Time Pricing rate structures. The results demonstrate that the optimization-based controller, in both its centralized and decentralized forms for the two-zone building, outperforms the benchmark bang-bang control strategy. While the centralized controller applied to the five-zone building shows limited scalability, it still yields improved performance compared to the bang-bang baseline.

TABLE OF CONTENTS

LIST OF TABLES

LIST OF FIGURES

LIST OF SYMBOLS AND ABBREVIATIONS

ADMM	alternating direction method of multipliers
ANN	artificial neural network
ARMAX	autoregressive moving average with external inputs
ARX	autoregressive with external inputs
AST	actual solar time
BBAI	BrainBox AI
BESS	battery energy storage systems
COHORT	Coordination Of HeterOgeneous Residential Thermostatically controlled loads
CPP	Critical peak pricing
DHI	diffuse horizontal irradiance
DNI	direct hormal irradiance
DR	demand response
DSM	demand-side management
GHG	greenhouse gas
GHI	global horizontal irradiance
HVAC	Heating, ventilation, and air conditioning
IETF	Internet Engineering Task Force
kW	kilowatt
kWh	kilowatt-hour
LORER	Laboratoire d'Optimisation des Réseaux Électriques Renouvelables
MILP	mixed-integer linear program
MINLP	mixed-integer nonlinear programming
MIP	mixed-integer programming
MIQP	mixed-integer quadratic programming
MLR	multiple linear regression

MPC	model predictive control
MST	mean solar time
NN	neural network
OSI	Open Systems Interconnection
PCM	phase change material
PID	Proportional-Integral-Derivative
PTR	Peak Time Rebates
PWM	pulse width modulation
RC	resistance–capacitance
RL	reinforcement learning
RTP	real-time pricing
RTU	rooftop unit
TCL	thermostatically controlled load
TES	thermal energy storage
TOU	Time-of-use
RMSE	root mean square error

LIST OF APPENDICES

CHAPTER 1 INTRODUCTION

Global energy and electricity demand is steadily increasing, driven by both population growth and sectoral shifts such as the electrification of transportation [2, 3]. This rising demand places additional stress on electrical grid infrastructure, increasing, e.g., the risk of equipment overloading and reliability issues. Electric grid operators are increasingly challenged to either expand distribution and/or transmission infrastructure or implement more intelligent and adaptive strategies for managing electricity consumption, especially during peak demand periods [4]. Constructing new infrastructure is a capital-intensive undertaking, with construction costs continuously rising [5]. Moreover, such developments can have significant environmental consequences, including habitat disruption, threats to local biodiversity, and greenhouse gas (GHG) emissions associated with expanded transmission networks [6].

Power grid flexibility strategies, such as demand response (DR), play a critical role in mitigating peak power demand within electrical networks—a process commonly referred to as peak-shaving—which in turn reduces or delays the need for additional infrastructure investment [7]. DR mechanisms, including incentive-based electricity pricing, aim to influence consumer behaviour, promoting shifts in electricity usage patterns and effectively enabling load-shifting. In addition, DR provides an opportunity for consumers to lower their electricity expenditures by adjusting their consumption in response to price signals or incentive mechanisms [8].

The building sector has a tremendous impact on global energy consumption [9] and a major part of this of energy, about 40%, is used for temperature control, i.e., by heating, ventilation, and air conditioning systems (HVAC) [10]. By leveraging the thermal mass of large commercial buildings [11] and by enrolling them to DR programs, we aim to manage power use more efficiently and thus offsetting and reduce electricity consumption compared to generic rule-based controller. This can lead to reduced peak loads [12], which in turn lowers the demand for additional electric infrastructure, and reduced cost for building operators [13].

In this Master's thesis, we propose a model predictive control approach—which is embedded with a data-driven, physics-informed thermal model—that participates to DR events for HVAC systems in multi-zone commercial buildings. In line with the low deployment cost argument for DR, the control method supports mixed-integer controls to ensure compatibility with existing systems [14, 15]. The method is also scalable to different building sizes thanks to its decentralized hence parallelizable computation infrastructure and is streamlined by dint of to its data-driven modelling pipeline.

1.1 Context

The primary scope of this work in terms of buildings is large-scale, open-space structures, such as warehouses and large retail stores. These types of building offer a great potential for taking advantage of DR if their HVAC controls are optimized adequately. Because of their large thermal mass, temperature fluctuation in the buildings can be anticipated and leveraged [11], and can act as an energy storage that is then employed as a tool to offset energy usage. Given their important HVAC energy usage and its corresponding operation cost [16], optimizing their control patterns to take advantage of incentive-based pricing offers both a building-level economic and a grid-level efficiency opportunity. Thus, optimizing large-scale, open-space building HVAC controls can have a significant impact on electrical grids' peak load and power consumption pattern while generating a substantial financial reward to building operators.

Large-scale, open-space buildings typically use rooftop units (RTUs) as their HVAC systems, which usually mainly admit discrete controls [14, 15]. To develop a method that can be applied immediately, without requiring retrofitting existing HVAC systems, we aim to design a mixed-integer HVAC optimization strategy specifically for large-scale open-space buildings.

Moreover, due to the wide diversity of building types and the considerable complexity and effort required to model each of them individually, a streamlined and efficient data-driven pipeline for building operation is suggested. This is possible, provided certain hardware and data availability conditions are met. Therefore, we focus on a scenario where ample historical data is available, and assume that control systems are already installed on all HVAC units in the building. In the case where no individual control systems are installed on each HVAC system, our approach can easily be used by the means of a multicore central unit in a decentralized manner to maintain the efficiency gains of parallel computing.

In our work, we employ a state-of-the-art simulator to test different control methods thanks to our collaboration with BrainBox AI (BBAI). BBAI is a company that specializes in energy consumption optimization via temperature management in buildings. Our collaboration involves using their simulator to develop a mixed-integer optimization algorithm that is computationally efficient enough for real-world use in large-scale commercial buildings. In practice, the control decision must be computed in under five minutes, which sets the upper bound for our computation time. The simulator is implemented in Python, so our control agents will be developed in the same language. While other languages, such as C++, could offer faster computation, they are not considered in this work. Finally, this project focuses on heating controls and will target only winter climates like in Montréal, Québec, Canada. However,

this approach can be extended to summer conditions or other climatic environments.

1.2 Background

We now introduce several key concepts necessary for this Master's thesis. The following sections overview demand response (DR), general optimization concepts, HVAC optimization and temperature modelling inside a building. All four of the subjects are central to this work.

1.2.1 Demand response

DR refers to demand-side management (DSM) methods that aim to modulate users consumption at a specific moment by sending incentive signals to support grid operations [7]. DSM helps grid operators to gain a new degree of freedom for power balancing operations.

DR as the potential to become an important aspect of electrical grid operations. Electrical grids have a limited amount of power generation that needs to be dispatched in the right amount, in different part of the grid, and at all times to ensure power adequacy. The growth of recent technologies such as renewable power generation and electromobility makes power balancing between generation and the load, harder than before [8]. DR is an important tool to provide flexibility to the grid and support operations [7]. DR can influence demand patterns in beneficial ways to assist grid operations, but it does not necessarily reduce global energy consumption [8].

As shown in Figure 1.1, typical power demand curves in both summer and winter exhibit high peaks that could potentially be offset through better demand-side management.

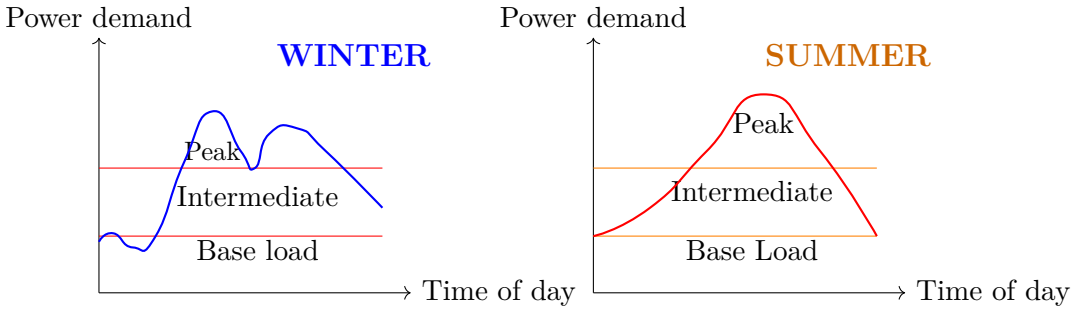

Figure 1.1 Power demand profiles for typical winter and summer days showing base, intermediate, and peak loads (source).

DSM methods like DR produce positive incomes for both users and grid operators. It provides an additional source of flexibility to the grid operators and can improve the grid's efficiency

and reliability [7]. Moreover, it gives a precious opportunity to costumers to reduce their power consumption costs [8].

DSM techniques can be distinguished by either having direct or indirect control over the load. In direct load control, grid operators have remote control to certain appliance, while indirect load control refers to programs where monetary incentives are distributed to clients willing to induce change in their consumption during specific moments or events, i.e. DR [17].

DR can be classified by incentive and time based DR [8]. In incentive-based DR, system operators usually initiate the consumption change. Examples of this include : critical peak pricing where grid operators send a signal and the user can react willingly to alter their load in exchange for some monetary incentives and capacity market programs where clients agree contractually to curtail or shed their consumption when the grid requires it. Time-based DR refers to programs that set a fixed electricity pricing schedule dependent on time instead of specific events. These type of programs are more consumer-based, i.e., that the impact on the grid comes mostly from the behaviour of the consumer and less from the grid operators. Therefore, the pricing schedule is set in a way to influence the consumers' behaviour, to offset electricity usage.

In this Master's thesis, we consider both settings. We leverage incentive and time-based electricity pricing in combination with a decentralized optimization technique and the thermal mass of large-scale, open-space commercial buildings to contribute to the power grid efficiency, e.g., by and reducing peak load demand. This approach also helps building operators reduce their operational costs.

Specifically, we consider the following DR programs [8]:

- Time-of-use (TOU) : As seen in Figure 1.2, TOU prices are predetermined and vary by time block (e.g., peak, mid-peak, off-peak), often based on historical patterns;
- Critical peak pricing (CPP) : Similar approach to TOU, but with adjustments that place greater emphasis on critical hours of the day, i.e., reducing the number of pricing tiers and price variation and introducing sharper price increases during critical events;
- Real-time pricing (RTP) : Prices fluctuate on hour, minute or even shorter timescale and are tied to the wholesale electricity market. Consumers receive signals or forecasts to adjust consumption accordingly.

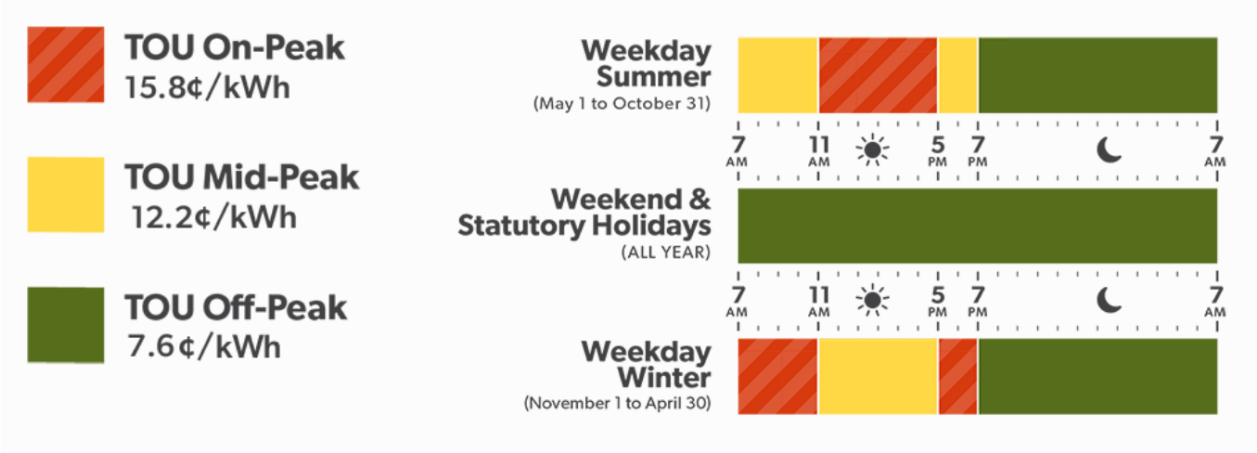

Figure 1.2 TOU pricing in Ontario as published in 2024 by Hydro One [1].

CPP is deployed in Québec, e.g., with Hydro-Québec's Rate Flex D [18] in the residential sector. During the winter period, Hydro-Québec applies an incentive-based pricing structure designed to promote peak-shaving. Electricity consumed outside of peak demand events is charged at a lower rate of 4.774¢ per kilowatt-hour (kWh), up to a daily threshold of 40 kWhs multiplied by the number of days in the billing period. Any additional consumption beyond this threshold is billed at a higher rate of 8.699¢ per kWh. Moreover, electricity used during designated peak demand events, which could be between 6 a.m. and 9 a.m. or 4 p.m. and 8 p.m., is subject to a significantly higher rate of 45.088¢ per kWh, providing a strong financial incentive for consumers to reduce or shift their energy usage during these critical periods.

1.2.2 General optimization concepts

Optimization is a tool at the core of this Master's thesis. In optimization, a decision-maker seeks to minimize or maximize an objective function subject to constraints on the decision they can make. Optimization problems provide us with the means to take the best possible decision considering an objective while ensuring that the decision is part of a set dependent on the constraints. It is a widespread approach in many fields such as portfolio management, manufacturing, machine learning [19], and power management [20].

A standard optimization problem takes the form of (1.1), where $\mathbf{x} \in \mathbb{R}^n$, $n \in \mathbb{N}$, is the optimization variable, $f_0 : \mathbb{R}^n \to \mathbb{R}$ the objective function and $g_i : \mathbb{R}^n \to \mathbb{R}, \quad i = 1, 2, ..., m, \quad m \in \mathbb{N}$, are the constraints.

$$\begin{aligned} \min_{x} \quad & f_0(x) \\ \text{subject to} \quad & g_i(x) \leq 0, \quad i = 1, \ldots, m. \end{aligned} \tag{1.1}$$

Optimization can require significant computing resources depending on how many variables are considered, their nature, the complexity of the objective function and constraints, and the sparsity of the problem [19]. In this specific work, two main factors impact the computational burden: the number of variables due to the rolling horizon we use and their mixed-integer nature.

Rolling horizon optimization, often called model predictive control (MPC), is an optimization technique that is ideal for HVAC system control [21]. It requires solving the problem for numerous steps ahead to take into account dynamic changes in the environment and forecasted information. In (1.2) Problem (1.1) is solved for a horizon of time ahead of length $W \in \mathbb{H}$ steps at each time steps t.

$$\begin{aligned} \min_{x} \quad & \sum_{t=\tau}^{\tau+W} f_0(x_t) \\ \text{s.t.} \quad & \sum_{t=\tau}^{\tau+W} g_i(x_t) \leq 0, \quad i = 1, \ldots, m. \end{aligned} \tag{1.2}$$

The decision-maker then applies the optimal decision corresponding to the first step within the horizon, and then restart the process at the next decision round with updated information about the problem states.

The solution considers future states of the problem's environment while making sure that the current state corresponds to the constraints. It, however, necessitates a new solution for the whole horizon of time at each time step in addition to detailed models of the environment. The former leads to a major increase of computing requirements because the original problem needs to be solved with $W = 5$, 10, 50 or even more times at each time step.

Mixed-integer optimization is central in this work because many HVAC systems are controlled by mixed-integer controls, viz, RTU-HVAC which equips many commercial buildings [14, 15]. It requires dedicated resolution approaches which are more computationally intensive because mixed-integer optimization problem are NP-hard and do not have an efficient polynomial-time resolution method [22]. In other words, the problem is *inherently* intractable [23] and that either an approximation, heuristic or relaxation is necessary to solve the problem. As emphasized by Fischetti, Lodi, and Salvagnin [24], techniques such as branching strategies, cut generation, and heuristic methods have been successfully applied to solve mixed-integer programming (MIP) problems. Nevertheless, the practical resolution of MIP problems remains constrained by the problem size and the complexity of the associated constraints.

Optimization may lead to a local optimum or even a saddle point instead of the global optimum. As shown in Figure 1.3, situations may occur where the combination of objective

function and feasible set has local and global maxima. Depending on the initial point of the solver used, we very well may converge to the local maximum instead of the global maximum. Because optimization problems typically aim to identify a global optimum, the presence of non-convexities renders the problem NP-hard.

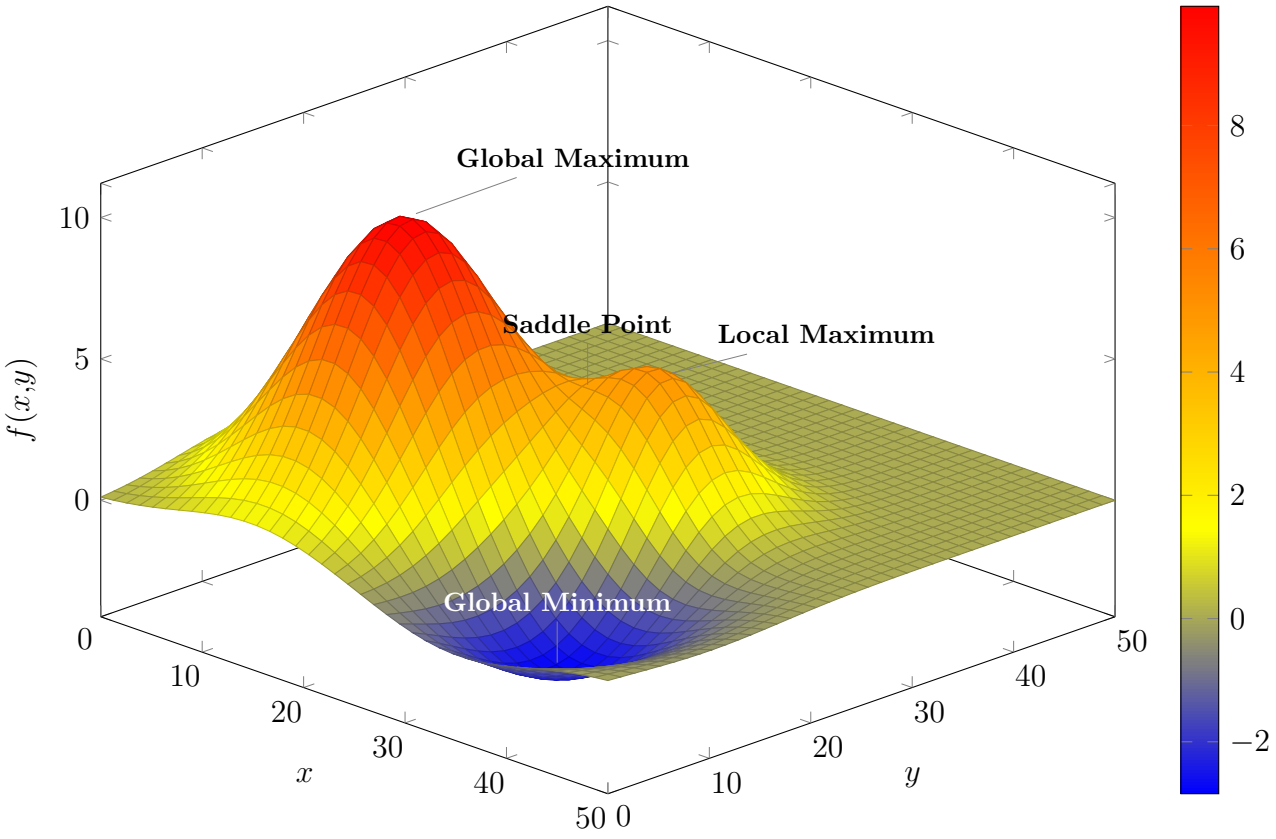

Figure 1.3 Local and global optima

The issue is avoided in the case of convex optimization problems. Convex optimization problem are defined by having convex objective function and constraints [19]. In (1.1), $f_0(\mathbf{x})$ and $g_i(\mathbf{x})$ must be convex to yield a convex optimization problem. The convexity of a problem ensures that any extrema is a global optimum. It is also interesting to note that if a problem has a linear objective function and affine constraints, then it is necessarily convex [19].

Finally, optimization can be performed in a centralized, decentralized, or distributed fashion. Decentralizing or distributing an optimization process results in parallelizing the resolution of a large problem between multiple agents. Each agent is assigned a reduced-size local optimization problem, which is solved independently, leading to local solutions shown to match the original centralized problem's one. This approach requires the deployment of additional computing agents; however, it significantly reduces the computational burden on each indi-

vidual agent compared to the processing power required to solve the centralized version of the same problem. Thus, decentralized and distributed approaches reduce computation time needed to solve a problem and help increase the scalability of optimization approaches [25].

In the context of HVAC control, decentralized or distributed optimization can be implemented using a star network topology, where each HVAC unit solves its own local sub-problem, or through a single high-performance computing agent employing parallelization. In both cases, these approaches enhance the scalability of the overall control and optimization process [26].

The difference between distributed and decentralized optimization resides in the need for a central collector. In decentralized optimization, the process needs to aggregate information at a central collector that will send back new information to all the agents at each iteration, potentially also processing the information. In distributed optimization, this process is not necessary, and all computations are performed in a completely individual manner by the agents [27], [28]. A communication of nonsensible information remains between agents. Nevertheless, communication requirements are reduced and privacy is kept between each agent.

1.2.3 HVAC optimization

The HVAC optimization problem can be viewed from the perspective of a thermal energy storage (TES) problem. TESs are used to store energy in the form of heat and then move it through time [29]. The thermal energy can be stored in various medium, e.g., the ground, a water tank, phase change material [30], or, in our case, the air mass in our buildings.

Problem (1.3) below provides a general optimization problem for TES [29] :

$$\min \sum_{t=1}^{T} c(t)e(t) + \beta \max_{1 \leq t \leq T} \left(P(t) \right), \tag{1.3}$$

where the objective focuses on two aspects : (i) to minimize the operation cost, given by the product of the energy cost $c(t)$ and the energy consumption $e(t)$, and (ii) to minimize the peak demand cost, which is the maximum power demand $P(t)$ multiplied by the peak power β. Energy cost can vary from jurisdiction to jurisdiction.

An example of energy pricing would be Hydro-Québec's CPP for commercial buildings that consume at least 50 kilowatts during one of the months of the billing period : Rate Flex M [18]. Hydro-Québec's Rate Flex M is a time-varying electricity pricing structure applied under an annual contract, consisting of a fixed monthly charge of $17.573 per kilowatt of billing demand, plus energy charges that vary by season. During the winter period, electricity

is billed at 3.820¢/kWh outside of peak demand events and 60.262¢/kWh during peak events. In the summer period, the rate is 6.061¢/kWh for the first 210,000 kWh and 4.495¢/kWh for any additional consumption.

In this work, we consider a modified version of (1.3) suited for building temperature control where we replace the energy consumption of the system by the control of a HVAC system of a building's specific zone, i, at time t: $u_i(t) \in \mathbb{R}$. We take into account multiple HVAC systems and denote the set of all HVACs as $\mathcal{N} \subset \mathbb{N}$. We also use $c_i(t) \in \mathbb{R}$ to model the energy price at t and thus to incorporate different DR programs in the problem. The notation i, added to the cost vector, $c_i(t)$, is only a formality because all zones in a building are subject to the same electricity price. In (1.4), we assume that the power consumption of the HVAC is constant for a given control and define the actual power usage in the numerical cases. We omit the peak demand term from the objective because we do not assume access to the building power consumption. Using a rolling horizon formulation, the problem is:

$$\min_{u_i} \sum_{t=\tau}^{\tau+W} \sum_{i \in \mathcal{N}} c_i(t) u_i(t), \tag{1.4}$$

where $W \in \mathbb{N}$ is the rolling window and $u_i(t)$ is the continuous control for HVAC i at time t. Problem (1.4) is subject to different constraints. To be defined thoroughly, they depend on the DR settings, the HVAC system, the operator's requirements, etc. From [29], we can translate the standard constraints on the controls $u_i(t)$ and the zones' temperature $T_i(t)$ that applies to this work into : HVAC control maximum and a minimum output, $\overline{u}$ and $\underline{u}$, the varying temperature bounds, $\overline{T}_i(t)$ and $\underline{T}_i(t)$, and the building dynamics function, f, describing the temperature evolution which we discuss in Section 1.2.4.

This leads to the constraints:

$$\underline{u_i} \leq u_i(t) \leq \overline{u_i}, \tag{1.5}$$

$$\underline{T}_i(t) \leq T_i(t) \leq \overline{T}_i(t), \tag{1.6}$$

$$T_i(t+1) = f(T_i(t), T_{-i}(t), u_i(t)), \quad \forall i \in \mathbb{N}, \tag{1.7}$$

where the notation $-i$ refers to all zones except the current zone i.

In the HVAC control problem, the output (i.e., the decision to be made) is the control action applied to the HVAC system, u_i, where i represents the specific temperature controlled zone i. The control scheme for an HVAC system can be either continuous or discrete. In this work, the control will be presented as discrete $u_i(t) \in \mathbb{N}$, thus, it can take the form of heating, cooling, fan operation, or no action (idling). The heating and cooling actions are typically

applied at specific output levels. Discrete controls are, e.g.,

$$u \in \{-2,\ -1,\ 0,\ 1,\ 2,\ 3\}, \tag{1.8}$$

where

$-2, 2$: level 2 cooling or heating;

$-1, 1$: level 1 cooling or heating;

0 : no action (null);

3 : idle or fan.

The inputs to this problem can vary widely, but primarily consist of exogenous data from environmental sensors (e.g., indoor temperature, humidity), weather data (e.g., outdoor temperature, sun position) and behavioural patterns of the users, i.e., building occupancy [16,21,29].

The HVAC optimization problem is constrained not only by the above thermal, control actions, and physical input limitations, but also by time, hardware, and software or communication-related factors. In particular, software and computational restrictions arise due to the limited time available for computing control actions, i.e, solving the optimization problem. This is especially critical in problems such as the HVAC's, where solutions must often be computed in real-time, thereby placing strict bounds on computing time and thus solver performance. As highlighted in [21], the application of MPC can be constrained by the computational resources available to the controller, due to the strict time requirements of real-time operations. In addition to the time step length, the overall prediction horizon also significantly influence the quality of the decision-making process. A shorter horizon typically results in lower prediction error but reduces the long-term foresight capabilities of the MPC controller. Moreover, the use of distributed architectures introduces communication delays and coordination steps, further complicating both the problem formulation and the solution process [31]. However, communication-related delays are not considered in this work.

In HVAC operation, hardware constraints may include cooldown times between control actions or ramping constraints. For the sake of the system's longevity and proper operation, it is not advisable to *skip* control levels [32]. For example, if the desired control is heating at power level 2, the system must first be set to power level 1 for a full time step. Thus, a ramping constraint can be added [33], as :

$$|u_i(t) - u_i(t-1)| \leq \Delta u_{\max}, \tag{1.9}$$

where, $\Delta u_{\max}$ is the maximum fluctuation in control between two time steps and $u_i(t)$ is

the control at time t. This constraint is not considered in this work. Ramping a toggling constraints are a topic for future work. Our primary interest lies in demonstrating that a mixed-integer control method can be implemented at large scale in a decentralized fashion in real time.

1.2.4 Temperature modelling inside a building

The temperature dynamic of a building can be thought as resulting from a complex function. It depends on several uncertain environmental and non-environmental factors such as the outdoor temperature, humidity, occupancy, and solar irradiation [34]. Furthermore, the non-linearity of the HVAC controls and the coupling between zones add another layer of complexity to the temperature's behaviour. Modelling of indoor temperature is thus a demanding but crucial part of this work because it plays a major role in an efficient MPC formulation [34]. Prior works have approximated the temperature evolution within a building using both linear and nonlinear models [35–37].

The linear model (1.10) from [35] represents the indoor temperature at time t, as a function of the indoor $T_{\text{in}}(t-1)$ and outdoor $T_{\text{out}}(t-1)$ temperatures at the previous step, the wind speed $v(t)$ and the solar irradiance $I(t)$. The other values such as the surface area and heat capacities are not time dependent or forecasted and are, therefore, parameters that characterize the environment and need to be known or learned. This particular model however does not account for the impact of a heating or cooling system. The linear model is:

$$T_{\text{in}}(t) = T_{\text{in}}(t-1) + \frac{\alpha_I A_I}{C} I(t-1)\Delta + \frac{b_1 v(t-1) + b_0 A_e}{C} \left(T_{\text{out}}(t-1) - T_{\text{in}}(t-1)\right)\Delta, \quad (1.10)$$

where, we define:

$T_{\text{in}}(t)$: indoor temperature at time t [K];
$T_{\text{in}}(t)$: indoor temperature at t [K];
$I(t)$: solar irradiance at t [W m^{-2}];
$T_{\text{out}}(t)$: outdoor temperature at t [K];
Δ : time step duration [s];
A_{I} : effective area exposed to solar radiation [m^2];
α_{I} : solar radiation factor [$-$];
A_{e} : effective area for heat exchange with the exterior [m^2];
$v(t)$: wind speed at t [m s^{-1}];
b_i : coefficients for wind influence on heat transfer [$\text{W m}^{-2}\,\text{K}^{-1}\,(\text{m s}^{-1})^{-1}$]*;
C : heat capacity (thermal capacitance) [J K^{-1}].

A linear model such as (1.10) can be extended to better capture the temperature dynamics by using a non-linear representation for one of the parameters and incorporate the impact of a HVAC system, e.g., the model from [37] yields,

$$\begin{aligned} T_{\text{in}}(t+1) - T_{\text{in}}(t) = &-\frac{\Delta}{\tau_3} T_{\text{in}}(t) \\ &+ \frac{\Delta}{\tau_3}\Big[T_{\text{ie}}(t) + (T_{\text{out}}(t) - T_{\text{in}}(t))(b_1 v(t) + b_2 v^2(t)) \\ &+ (a_1 I(t) + a_2 I^2(t) + a_3 I^3(t)) + \dot{Q}_{\text{occ}}\xi_{\text{occ}}(t) + \dot{Q}_{\text{HVAC}} u_{\text{bin}}(t)\Big], \end{aligned} \tag{1.11}$$

where, the new parameters are defined as :

$T_{\text{ie}}(t)$: interior wall surface temperature at time step t;
$\tau_3 = C_{\text{air}} R_{\text{air}}$: time constant of the space air node;
a_1, a_2, a_3 : coefficients for solar radiation effect;
b_1, b_2, b_3 : coefficients for wind impact;
$\xi_{\text{occ}}(t)$: occupancy indication (1 if occupied, 0 if not) at time k;
$u_{\text{bin}}(t)$: HVAC system control input (1 if ON, 0 if OFF) at time k;
$\dot{Q}_{\text{occ}}$: gain related to internal heat gains (e.g., occupancy);
$\dot{Q}_{\text{HVAC}}$: gain representing HVAC system heating/cooling capacity.;

The non-linear approach in (1.11) emphasizes on the wind speed (v) and global horizontal solar irradiation (I), allowing for greater flexibility and precision of the model.

A last example of thermal modelling comes from [36] and is provided in (1.12). This model takes the form of a differential equation and can go into as much if not more depth than its counterpart:

$$\dot{Q}_{\mathrm{c}} + \dot{Q}_{\mathrm{h}} = M \frac{\mathrm{d}T_{\mathrm{in}}}{\mathrm{d}t}, \tag{1.12}$$

where the rate of change of the indoor temperature ($\frac{\mathrm{d}T_{\mathrm{in}}}{\mathrm{d}t}$) depends on the heating power $\dot{Q}\mathrm{h}$ or cooling power $\dot{Q}_{\mathrm{c}}$ applied to the zone, as well as the thermal mass of the zone M. The thermal mass is determined by the mass of the walls, floor, and all other elements within the room, including the air, together with their respective heat capacities.

The heating and cooling powers can be defined in various ways. For example, in [36], the heating power is represented as the sum of the heating contributions from all the walls $\dot{Q}_{\mathrm{w}}$ in the room:

$$\dot{Q}_{\mathrm{h}} = \sum_{\mathrm{w} \in \mathrm{walls}} \dot{Q}_{\mathrm{w}}. \tag{1.13}$$

The heat transfer through each wall is then defined using a set of physical properties, such as thickness, surface area, and thermal conductivity, and a resistance–capacitance (RC) model. RC modelling is a common approach to indoor temperature modelling that helps simplify thermal load calculations [36]. As its name suggest, it is analogous to an RC electrical circuit. Instead of current, voltage, and resistance, it uses temperature, heat transfer, and thermal resistance, respectively.

In sum, there are multiples ways to express a thermal model depending on the desired level of accuracy, which in turn imposes different levels of complexity. No model is perfect yet, and the model selection varies depending on the usage of the model. In our case, we favour a linear model for its compatibility with optimization.

1.3 Research objective

This research focuses on mixed-integer HVAC optimization in large-scale commercial buildings. The goal is to develop a control method that is scalable, easy-to-implement, and capable of incorporating forecasted data and the thermal coupling between zones. To achieve this the sub-objective are :

1. **Real-time operations :** The control method must operate in real-time with respect to the building timescale. Thus, we impose a computation time limit of 5 minutes.

2. **Mixed-integer optimization :** The control method must be compatible with existing HVAC systems, and therefore must support integer control decisions.

3. **Scalability :** The control method must scale to multi-zone buildings. To achieve this, it must be designed in a decentralized manner so that adding a zone increases the computation time only minimally.

4. **Thermal coupling between zones :** In open-space, multi-zone buildings, thermal interactions between zones are significant and must be taken into account.

5. **Demand response :** The control method must be able to provide demand response services via incentive-based pricing to reduce the building's impact on the grid and benefit its users.

6. **Data-driven :** The controller needs to be able to predict the zones' future states (i.e., the temperature) and to maintain it in the defined deadband. To allow this method to be applied on multiple different buildings, a data-driven implementation pipeline that does not require extensive field work is preferable.

In this work, we propose a decentralized mixed-integer MPC method for HVAC control in multi-zone commercial buildings. Our approach employs a data-driven, physics-informed linear prediction temperature model. It ensures thermal comfort as defined by the building deadband while optimizing power consumption with respect to DR events.

1.4 Thesis outline

This Master's thesis is composed of 6 chapters. The first chapter, the current one, introduces the content of this work and provides some background. Chapter 2 covers the literature review as well as the current state of the art of HVAC control methods. Chapter 3 and 4 introduce our methodology. In the first part of the methodology, we detail the physics-informed, data-driven linear temperature model used to make predictive HVAC control. The second par of the methodology discusses the optimization based control approach as well as the decentralization of this process. Chapter 5 discusses numerical results as well as their detailed analysis. Finally, Chapter 6 concludes the thesis and proposes future lines of work.

CHAPTER 2 LITERATURE REVIEW

The objective of this literature review is to survey the state-of-the-art and to highlight gaps in the existing body of research. The scope of this Master's thesis is focused on mixed-integer programming, and distributed or decentralized optimization, in the context of, building temperature modelling and HVAC system control. We now review the literature on these topics, starting with the literature on established methods for HVAC control.

2.1 Current control methods

Control strategies for HVAC systems typically fall under categories such as rule-based, predictive, and data-driven methods, though hybrid and alternative approaches also exist. Each method is presenting specific advantages and limitations, as discussed below.

2.1.1 Rule-based control

Rule-based control methods rely on predefined set points or parameters [38]. The controller aims to follow these reference values as closely as possible. Binary ON/OFF control and Proportional-Integral-Derivative (PID) control are two common approaches to rule-based control. In the case of the ON/OFF control, the HVAC system is either activated or deactivated based on preset threshold bounding the thermal comfort or the power consumption [39]. PID control is a more refined approach that will continuously adjust the HVAC output to follow a target temperature defined by set parameters [39]. For example, in [40], the following PID is applied to building temperature control :

$$u(t+1) = u(t) + \alpha_P \Delta T + \frac{\alpha_I}{t} \int_0^t \Delta T(\tau)\, \mathrm{d}\tau + \alpha_D \left(\frac{\mathrm{d}(\Delta T)}{\mathrm{d}t} \right), \tag{2.1}$$

where,

T_{sp}: temperature set point (desired);

T_{in}: indoor temperature;

$\Delta T = T_{\mathrm{in}} - T_{\mathrm{sp}}$;

α_P: proportional gains;

α_I: integral gains;

α_D: derivative gains;

t: current time (for integral scaling).

These controllers are relatively simple to design, computationally efficient, and cost-effective. However, they do not utilize predictive information such as weather forecasts and are generally not suitable for integration with DR events. Moreover, they offer limited flexibility in handling multiple, potentially conflicting control objectives, and they lack adaptability to changing environments [39].

2.1.2 Model predictive control (MPC)

MPC improves on rule-based control method and offers a way to account for anticipated future states of the environment while considering its current state. It is currently a leading method to optimize for thermal comfort and power consumption [39].

The core components of a MPC framework include [38] :

- **Model**: Describes the evolution of future states based on the current state and forecasted data;
- **Prediction**: Utilize the model and incorporates forecasts such as weather conditions, occupancy levels, or electricity prices to predict future states;
- **Control**: Solves an optimization problem over a rolling horizon using the model and forecasts to determine optimal control actions.

MPC is particularly well-suited for applications involving time series data, as it explicitly incorporates forecasts of future system behaviour to optimize control actions over a prediction horizon. However, its effectiveness is highly dependent on the accuracy of the underlying model. A well-calibrated model enables reliable predictions and effective control. Conversely, an inaccurate model may lead to poor predictions and suboptimal decisions—potentially performing worse than basic controllers. Model development is thus critical, though it can be time-consuming, complex, and requires domain expertise [21].

Moreover, MPC performance can degrade due to the accumulation of model and forecast errors over the prediction horizon, as each predicted state depends on the accuracy of previous predictions, potentially compounding errors over time. To mitigate this, a *receding horizon* strategy is employed: after solving the optimization problem over the full prediction horizon, only the control inputs for the first time step are implemented. The horizon then advances by one step, and the process is repeated at each time step. This iterative optimization process increases the computational burden, as the problem must be solved repeatedly over the entire horizon [21].

A basic MPC formulation to control HVAC systems over a prediction horizon of length $W \in \mathbb{N}$ is presented in (2.2) for a three-zone building [41]. The objective function h is left unspecified, but typically represents the total energy consumption of the HVAC systems weighted by electricity cost. The variable $x_i(\tau)$ denotes the state (e.g., temperature) of zone i at time τ, and $u_i(\tau)$ is the control input for the corresponding HVAC unit. The vector $\mathbf{d}$ represents exogenous disturbances, such as the outdoor temperature. The first constraint captures a simplified temperature dynamics model for each room, accounting for interactions with neighbouring zones and HVAC operations, while the second and third constraints bound the temperature in each zone as well as the possible controls applicable by the HVAC systems. The parameters $\mathbf{A}$, $\mathbf{B}$, $\mathbf{C}$ and, $\mathbf{D}$ are specific parameters that describe respectively the heat transfer between zones, the impact of each HVAC system and the environmental impacts.

$$
\begin{aligned}
&\min_{u} \sum_{\tau=0}^{W} \sum_{i=1}^{3} h_i(u_i(t+\tau)) \\
&\text{s.t.} \quad \begin{pmatrix} x_1(t+\tau+1) \\ x_2(t+\tau+1) \\ x_3(t+\tau+1) \end{pmatrix} = \mathbf{A}^{-1} \left(\mathbf{C} \begin{pmatrix} x_1(t+\tau) \\ x_2(t+\tau) \\ x_3(t+\tau) \end{pmatrix} + \mathbf{B} \begin{pmatrix} u_1(t+\tau) \\ u_2(t+\tau) \\ u_3(t+\tau) \end{pmatrix} + \mathbf{D}\mathbf{d}(t+\tau) \right), \qquad (2.2) \\
&\quad u_{i,\min} \leq u_i(t+\tau) \leq u_{i,\max}, \\
&\quad x_{i,\min} \leq x_i(t+\tau) \leq x_{i,\max}, \quad i = 1,2,3, \quad \forall \tau.
\end{aligned}
$$

MPC is a well-established and an extensively researched control strategy. As demonstrated in [42], MPC has been applied across a wide range of domains with significant impact. More recently, [43] highlights the effectiveness of MPC in building control applications, particularly in its ability to coordinate with the electrical grid.

2.1.3 Data-driven Control

Data-driven control methods are control algorithms trained on large datasets with the aim of taking the best possible actions without prior knowledge of the building model [38]. These methods are essentially machine learning algorithms. In the case of temperature control, reinforcement learning (RL) is of particular interest [38, 39].

RL seeks to compute the optimal sequence of controls of an unknown environment modelled as a Markov decision process through successive interactions [44]. In other words, RL provides delayed feedback to the controller in the form of rewards based on the outcome of its control at a given state. The controller's goal is to maximize the total reward over time [38].

In RL, four main components define the system: the policy, reward signal, value function, and optionally, a model of the environment. The policy maps states to controls and dictates the agent's behaviour. The reward signal provides delayed feedback, guiding the agent toward controls that yield higher long-term returns. The value function estimates the expected future reward from each state, helping the agent make more farsighted decisions. Together, these elements shape how the agent learns and adapts within its environment [44, 45].

Data-driven methods such as RL are advantageous for HVAC control because they can implicitly model buildings' thermal dynamics environment while being efficient enough for real-time operation [45]. Furthermore, they do not require an expert-designed model. However, they come with several drawbacks: they typically require a large volume of training data, long training times, and substantial computational resources prior to deployment [46]. Moreover, these methods are generally not robust or do not provide out of the box safety guarantees. They are a *blackboxes*, meaning that the internal decision-making process is not easily interpretable by the system designer. This lack of transparency makes the control design process less intuitive and raises safety concerns, particularly in critical applications [47].

2.1.4 Summary of control methods

Table 2.1 compares the above-mentioned control methods.

Table 2.1 Comparison of HVAC Control Methods

Feature	**Rule-based Control**	**Predictive Control (MPC)**	**Data-driven Control (RL)**
Approach	Uses predefined rules or set points	Uses a model and forecasts to optimize control over a time horizon	Learns control policies from interaction with the environment
Model Requirement	None	Requires accurate model of the system	No model needed (model-free)
Forecast Usage	Not used	Required (e.g., weather, occupancy, electricity prices)	Optional (but can be included)
Computation Time After Deployment	Very low	High (due to optimization at each time step)	Very low
Data Requirement for Training	None	Requires model parameters	Requires large training datasets
Ease of Design	Easy	Complex (requires domain-specific expertise)	Requires ML knowledge and training infrastructure
Constraint Satisfaction	High under known conditions	Medium (depends on model accuracy)	Low (blackbox, non-interpretable)
Transparency	Transparent box	White/Gray box	Blackbox
Model Accuracy	No model	Good	Good
Main Drawback	Ignores future and system dynamics	Heavy computation and model dependency	Requires significant training data, low interpretability

2.2 Related work

Following the introduction on HVAC control paradigms, we now delve into the existing literature.

2.2.1 HVACs' impact on electrical grids

An essential consideration for this work is that optimizing HVAC control can meaningfully impact the electrical grid. HVACs must represent a significant portion of electricity demand so that the integration of DSM strategies with HVAC control methods result in a positive effect on the grid consumption patterns. More specifically, it is important to establish that HVACs in multi-zone, commercial buildings have a significant impact on power consumption at the grid level.

Proof of this impact can be found in [48] and [2]. These references are building energy data sets provided by the International Energy Agency, a comprehensive global and national statistics agency. They focus on energy consumption, end-uses, and emissions in residential and commercial buildings. They serve as a key reference to measure the impact of energy consumption in buildings. As seen in [48], in 2022, the operations of buildings was responsible for 30% of global final energy consumption and almost half of the energy consumed towards the building sector was used for space and water heating. Furthermore, [2] presents that heat pumps (HVACs) are expected to be a major driver of electricity demand growth through 2026 in the European Union, representing a substantial share of the projected increase in total consumption.

Other works, such as [10], reinforce the notion that HVAC systems have a significant impact on the electrical grid whilst noting the important role of commercial buildings. In this recent review, the authors emphasize that buildings account for a major share of global energy consumption and CO_2 emissions, largely driven by factors such as population growth, urbanization, and the increasing demand for comfort services like HVAC, which alone represents approximately 12% of global final energy use. Additionally, they highlight that in the United States, 44% of building energy consumption originates from the tertiary (service) sector, which includes commercial buildings.

References [13] and [49] provide concrete examples demonstrating the positive impact of DSM strategies on the electrical grid. From these results, one can reasonably infer that applying DSM techniques to HVAC control could yield similarly beneficial outcomes. They both proposes DSM strategies to reduce energy cost saving. In [13], Alimohammadisagvand et al. evaluate the life cycle cost optimization of thermal energy storage integrated with

ground source heat pumps in Finnish detached houses under various tank sizes and set point temperatures. Three DR control strategies were tested, with the MPC algorithm achieving up to 12% energy and 10% cost savings. The optimal configuration was a 0.3m^3 tank with a 60°C set point. Thermal mass had minimal impact on DR performance when adequate storage was available, and increasing heating set points was not economically beneficial. In [49], Li et al. propose an integrated DR framework for optimizing the design and control of battery energy storage systemss (BESSs) using accurate medium-term electricity demand forecasts from 72 buildings. By combining MPC and mixed-integer nonlinear programming (MINLP), the approach achieved significant cost savings (C$311k) and emissions reductions (471 t CO_2-e) over a winter period.

While energy cost saving does not automatically equate to a positive impact on the electrical grid, works such as [8], [12], and [7] show that DSM strategies related savings are a key tool for load management, have a positive impact on electrical grids' stability, and can delay the need to build new distribution infrastructure.

Other works such as [41], [45], [50], and [51] apply DSM directly to HVAC control in buildings and manage to obtain significant gains in energy cost saving over other techniques. These works will be discussed in detail in the following sections.

In sum, as discussed briefly in the introduction and stated in [49], commercial buildings have a major impact on power consumption patterns and GHG emissions and HVAC systems are responsible for a major part of that consumption. Buildings consume 30 % of the global power [48] and 38 % of a building consumption comes from these HVAC systems [10]. Furthermore, DSM strategies are the solution to leverage the energy usage in buildings to reduce peak load in power systems [49]. Therefore, developing an easily implementable control strategy that integrates DSM techniques for HVAC systems in commercial buildings is both relevant and promising, as it has the potential to positively influence electrical grid consumption patterns.

2.2.2 Basic HVAC operations

Basic HVAC operations can be considered a non-optimization-based control method. It typically refers to rule-based strategies, which serve as the foundational approaches that optimization techniques seek to improve or replace. While the background of such methods is already discussed in Section 2.1.1, it is nevertheless valuable to revisit the topic and highlight additional studies that have explored these approaches in practice.

Recent papers that introduce newer control techniques, such as [38] and [39], as well as

earlier works like [52], provide valuable insights into traditional control strategies. These studies frequently reference basic controllers, such as PID and ON/OFF, in their introductions or literature reviews, positioning them as foundational approaches in the evolution of HVAC control. As such, they serve as useful entry points for understanding the fundamental principles of HVAC operations.

Other studies, such as [53], investigate these methods in greater detail by comparing various types of ON/OFF controllers for HVAC systems with more advanced control techniques. In the context of classical control, earlier works like [54] focus on the design of robust PID controllers tailored for HVAC applications.

Rule-based control strategies are generally simple, easy to implement, and widely used in practice. However, they can fall short in performance when compared to well-tuned predictive or data-driven controllers [38].

2.2.3 Continuous optimization for HVAC control

A literature review focused on the optimization of continuous control strategies is now provided to understand the current state of the art in optimization for HVAC operations. We remark that continuous control simply defines control methods that communicates their control decisions by the mean of float values, and not necessarily integers.

2.2.3.1 Centralized continuous optimization

Prívara et al. [52] implemented a MPC strategy to manage the heating system of a real university building during the 2009–2010 heating season with seven control blocks. The authors developed a multiple-input, multiple-output linear state-space model to describe the thermal dynamics of the building. This model was identified using real operational data collected from the building management system. The MPC controller used weather forecasts and the identified thermal model to predict future indoor temperature trajectories and optimize heating control decisions over a defined prediction horizon. The optimization objective is to minimize heating energy consumption while maintaining indoor thermal comfort within predefined temperature bounds. The predictive controller operates in closed-loop, receiving updated measurements and re-optimizing control inputs at regular intervals, i.e., in a receding horizon fashion. To ensure adaptability and real-world feasibility, the authors incorporated constraints on control input (e.g., valve positions) and thermal comfort ranges directly into the optimization problem. Their method was benchmarked against a standard weather-compensated control strategy, which are control strategies where the output is ad-

justed based on the outdoor air temperature. Results demonstrated energy savings of up to 24%, confirming the practical benefits of MPC in real building environments and highlighting the importance of accurate model identification and reliable forecast data. A key limitation of this work lies in the challenge of obtaining an accurate building model. Standard system identification techniques often fail due to the poor statistical properties of typical building data. While first-principle, *white-box*, models, based on fundamental physical laws, can provide high accuracy, their development is time-consuming and impractical for large-scale or commercial applications. Nevertheless, the authors demonstrate that reliable models can be achieved through well-designed identification experiments combined with modified identification algorithms, provided that numerical stability is carefully addressed. Additionally, the study considers only outdoor temperature forecasts as exogenous inputs and operates at the building block level, without extending control to individual zones or rooms.

The author of [55] propose a hybrid control framework that combines nonlinear moving horizon estimation (MHE) with MPC to manage building HVAC systems with unknown dynamics. The building's thermal behaviour is modelled using a first-order RC model, while the HVAC equipment dynamics, often complex and difficult to model analytically, are captured through a data-driven model using a neural network. The building model analyzed in this study consists of a five-zone commercial structure served by a centralized cooling system. Thermal conditioning is delivered through a variable air volume distribution network equipped with reheat boxes in each zone, enabling individualized temperature control. To adapt the model in real time, the authors employ MHE, which estimates both unmeasured thermal states and unknown model parameters over a sliding time window using only data from standard building management systems. This updated model is then integrated into a MPC framework that optimizes continuous control actions such as supply air temperature and zone-level reheat commands. The controller operates in a receding horizon fashion to minimize energy consumption while maintaining occupant comfort. The methodology is validated through simulation, demonstrating robust performance without requiring detailed knowledge of the HVAC system or access to long-term historical data, thereby highlighting its potential for scalable real-world deployment.

In a pair of real-world field trials, West et al. [56] implemented a supervisory MPC system in two multi-zone commercial office buildings in Australia over 51 and 10 days, respectively. The objective is to minimize energy costs and GHG emissions while maintaining acceptable occupant thermal comfort. The system employs a data-driven *grey-box* building model, developed by combining knowledge of physical building dynamics with parameter estimation from up to three months of historical sensor data. The model structure is based on a multiple-input linear time-invariant system, capturing aggregate thermal behaviour. Although control

is ultimately executed at the individual zone level, the building model is constructed based on aggregate zone conditions and whole-building energy consumption. This modelling choice is motivated by practical limitations in submetering infrastructure and the influence of factors like air leakage and occupant movement across zones. As a result, fewer energy meters are needed, reducing installation and operational costs. Zone-level conditions are subsequently used to disaggregate control commands from the central model to each zone. The centralized MPC controller uses weather forecasts, occupant schedule profiles, and an online comfort feedback tool to dynamically adjust zone-level setpoints. The controller is designed to balance three competing objectives: energy cost, carbon emissions, and thermal comfort. Field trial results demonstrated energy savings of 19% and 32% in the two buildings, respectively, compared to baseline rule-based control systems. Importantly, these savings were achieved without negatively impacting occupant comfort, underscoring the practical benefits of MPC for energy-efficient HVAC operation in commercial buildings.

Centralized continuous optimization is largely based on MPC strategies. The existing literature on this topic is exhaustive, and the effectiveness of such control methods is proven and highly dependent on the accuracy of both the building model and the forecasted data [32].

2.2.3.2 Decentralized and Distributed Continuous Optimization

To enable scalable control strategies [41] and to alleviate the computational burden associated with predictive HVAC control [21], decentralized and distributed control methods have gained increasing attention.

In [41], the authors propose a scalable distributed control framework for coordinating thermostatically controlled loads (TCLs) using the alternating direction method of multipliers (ADMM) algorithm [57]. The method decomposes a global grid-level optimization problem into local subproblems while ensuring individual comfort and operational constraints. In real-world testing, [41] achieves an average peak load reduction of 12.5% without causing occupant discomfort and is shown to integrate easily with commercial smart thermostats, making it suitable for large-scale deployment. A discretization of the continuous control solution using pulse width modulation (PWM) is introduced to enable ON/OFF actuation. However, the method does not consider inter-zone dependencies, and the PWM-based discretization fails to enforce hardware cooldown constraints.

In [50], the authors propose a distributed control algorithm that integrates ADMM with deep learning-based models to regulate temperature setpoints and manage power consumption in multi-zone HVAC systems. The optimization problem under consideration in this work is inherently non-convex, whereas the ADMM algorithm traditionally requires convexity to

ensure convergence. Nonetheless, under certain Lipschitz continuity assumptions and with an appropriate selection of the penalty parameter ρ, convergence of the algorithm can still be guaranteed even in the non-convex setting. The method operates under a global coupled power constraint, distributing power allocations across local controllers without assuming a specific type of heating or cooling equipment. The building dynamics are learned entirely from simulated data, and the framework is validated on an 18-zone residential building simulated with EnergyPlus [58]. The approach achieves promising prediction accuracy—4.16% mean error using a state space model and 6.84% with a recurrent state space model—for one-hour ahead heating power forecasts. During DR events, both deterministic and stochastic variants of the algorithm successfully maintain power consumption below target thresholds, distributing the control effort among zones. While the model offers greater model accuracy than [41], it still does not account for inter-zone interactions and does not provide a method for translating continuous outputs into discrete control actions.

The study in [59] proposes a distributed MPC algorithm based on proximal the Jacobian ADMM [60, 61] for coordinating temperature control across multiple thermal zones in large open-space buildings. The approach relies on a linear model that captures inter-zone thermal coupling and demonstrates its applicability through a case study involving three zones. While the method is generalizable to a broader class of building control problems, it does not incorporate considerations such as time-of-use pricing or demand response events. Moreover, the optimization formulation only considers continuous variables and does not permit the use of integer ones, limiting its applicability to systems with discrete control requirements.

The work of Yu et al. in [62] proposes a distributed real-time HVAC control strategy, referred to as the cost-aware distributed real-time algorithm (CDRA), for multi-zone commercial buildings operating under a smart grid environment. The method aims to minimize the time-average expected total cost of HVAC operation, which includes both energy consumption and thermal discomfort, without relying on forecasts of electricity prices or external conditions. Using a Lyapunov optimization framework, the authors introduce virtual queues to regulate zone temperatures in real time, enabling system stability and bounded thermal deviation. The control is executed in a distributed manner: each zone determines its own control actions using only local measurements and limited communication with neighbouring zones, improving scalability and preserving privacy. Simulation results using real-world data show that the algorithm achieves significant cost savings while maintaining comfort. However, CDRA has limitations: it focuses only on long-term average performance and may not address short-term comfort violations; it requires manual tuning of control parameters; and it assumes perfect knowledge of local thermal dynamics and access to real-time data that may not be readily available in practice.

Zhuang et al. [25] present a noteworthy example of decentralized continuous control applied to standard HVAC configurations, including chillers, cooling towers, pumps, pipes, and air handling units. A key distinction of this work lies in its explicit modelling of system-wide coupling effects within the HVAC internal workings, alongside the simultaneous regulation of indoor temperature, humidity, and CO_2 concentration. This integrated approach makes it one of the few control strategies that jointly address both thermal comfort and indoor air quality. Although the method is not specifically designed for multi-zone commercial buildings, it remains relevant to this review due to its consideration of internal HVAC dynamics and the inclusion of exogenous environmental parameters. The primary objective of the control strategy is to minimize energy consumption while maintaining acceptable indoor environmental conditions. To address the computational complexity typically associated with such multi-objective control problems, the authors adopt an ordinal optimization approach [63]. This method seeks to identify sufficiently good—though not necessarily globally optimal—solutions, thereby improving scalability and enabling more practical implementation. Control actions are computed locally within each subsystem, which then exchanges information with neighbouring subsystems to evaluate relative performance changes. Based on this limited coordination, each unit determines whether its local solution can be integrated into a feasible global strategy. While the decentralized architecture allows for reduced computational overhead and avoids the complexities of centralized optimization, the use of ordinal optimization inherently sacrifices global optimality, and the coordination mechanism does not guarantee convergence to a consistent system-wide solution.

Lastly, Yang et al. in [64] introduce a two-level distributed control framework for multi-zone HVAC systems that simultaneously regulates thermal comfort and indoor air quality (IAQ)—specifically CO_2 concentration—while optimizing energy cost. The upper-level controller computes optimal zone-level air mass flow rates based on temperature targets and cost minimization, whereas the lower-level controller adjusts both ventilation and flow rates to ensure compliance with IAQ constraints, without significantly compromising energy efficiency. Both controllers operate in a distributed fashion, enhancing scalability, computational efficiency, and robustness relative to centralized strategies. Simulation results demonstrate that this method achieves near-centralized cost performance and surpasses existing distributed and demand-controlled ventilation strategies, delivering an 8–10% reduction in energy cost while maintaining acceptable IAQ along with thermal conditions. While effective in balancing energy cost and IAQ, the proposed two-level framework is limited by its reliance on simplified CO_2-based air quality modelling and assumption of knowledge of accurate system dynamics and fixed comfort preferences, which may hinder adaptability in practical settings.

Decentralized and distributed continuous HVAC control strategies have been widely explored

in the literature, often demonstrating successful outcomes. These approaches offer improved scalability by enabling local optimization of energy consumption while maintaining thermal comfort. However, many existing methods overlook the thermal coupling between zones and are not inherently compatible with mixed-integer control formulations, which limits their practical ability to multi-zone commercial buildings.

2.2.4 Mixed-integer optimization for HVAC control

Mixed-integer control represents a less-explored area in the context of HVAC systems optimization. While mixed-integer programming (MIP) is a well-established field, it is known to be NP-hard and thus computationally intensive when the number of variables increases [23, 24]. Moreover, the decentralization and parallelization of MIP formulations remain in the early stages of development, posing additional challenges for large-scale or real-time applications.

2.2.4.1 Centralized mixed-integer programming

In [45], the authors propose a data-driven HVAC control strategy using deep reinforcement learning (DRL), evaluated via co-simulation with EnergyPlus [58], real weather data, and dynamic electricity pricing. The study considers buildings with one, four, and five thermal zones and compares binary (ON/OFF) and multi-level air flow control. Results show that DRL can maintain thermal comfort while reducing energy costs compared to baseline and Q-learning methods. To address the complexity of multi-zone systems with large action spaces, a heuristic DRL variant is introduced, yielding improved scalability and reduced temperature violations. Nonetheless, challenges remain: the method scales poorly with an increasing number of zone and is data-intensive—relying heavily on simulated environments, whereas real-world deployment would require substantial historical data.

In [51], the authors present a data-driven convex mixed-integer MPC framework for energy management in rooftop-unit-based HVAC systems within small commercial buildings. The method uses a discrete, data-driven thermal model, enabling straightforward implementation and supporting market-based strategies such as TOU pricing and critical peak rebates. Simulation results show a favourable balance between thermal comfort, energy savings, and cost reduction. Despite its effectiveness, the approach is not parallelizable, limiting its scalability to larger or more complex buildings due to the high computational demands of mixed-integer programming. The authors acknowledge this limitation and propose future extensions using distributed optimization techniques such as ADMM, along with methods for obtaining provably optimal solutions.

2.2.4.2 Decentralized mixed-integer linear programming

Falsone et al. [27] propose a decentralized iterative scheme for solving large-scale mixed-integer linear programs (MILPs) with a focus on preserving privacy in multi-agent systems. The approach leverages Lagrangian duality and builds on the framework introduced in [65], constituting a variant of the dual subgradient algorithm. By decomposing the global optimization problem into smaller subproblems, the method ensures finite-time convergence to a feasible solution and offers a less conservative performance guarantee compared to previous techniques. Its effectiveness is demonstrated in the context of optimal charging of plug-in electric vehicles, where it achieved improved performance. A fully distributed extension, which eliminates the need for a central coordinator through max-consensus and decentralized dual updates while preserving convergence properties, was later introduced in [28]. However, both methods require careful step-size tuning and may face challenges in achieving feasibility or optimality under stringent computational constraints. Recent research aims to establish upper bounds on the number of required iterations and expand the framework to nonconvex settings. To the best of our knowledge, these methods have not been applied to HVAC control. In this thesis, we adopt the decentralized mixed-integer optimization method and apply it for the first time to HVAC control.

In [66], the authors propose a distributed algorithm for MILPs over networks. The problem constraints are assigned to a network of agents, each with limited computational and memory resources, capable of communicating only with local neighbours. The approach is based on a cutting-plane and constraint exchange method that iteratively refines the convex hull of the feasible region by generating and exchanging linear constraints, enabling agents to converge to an optimal solution in a decentralized fashion. Notably, the algorithm supports asynchronous execution and is robust to communication delays and packet loss, requiring only that the communication graph is jointly strongly connected. Under mild assumptions, agents can also autonomously detect convergence and terminate the algorithm using a distributed stopping criterion. Although convergence is guaranteed in a finite number of communication rounds, generating effective cutting planes remains computationally intensive. Nevertheless, the method is scalable and provides a promising foundation for distributed MILP solving in constrained and unreliable environments. To the best of our knowledge, this method as not been used in the context of HVAC control as well.

2.2.4.3 Decentralized mixed-integer nonlinear programming

The literature on the decentralization of mixed-integer nonlinear programming (MINLP) is limited. To the best of our knowledge, there is no decentralized or distributed method

for MINLP that does not rely on heuristics, i.e., with a provable convergence guarantee to optimality. An example of such an approach can be found in [67]. Some work has been done on mixed-integer quadratic programming (MIQP), such as [68], which falls under the broader category of MINLP, but this appears to mark the boundary of current efforts in decentralized nonlinear programming. Other recent studies, such as [69], continue to push decentralized MIQP by demonstrating that the dual problem converges to a zero duality gap asymptotically. However, our findings suggest that, within MIP, achieving provable optimal performance is more accessible when the formulation is restricted to linear programming.

2.2.5 Data-driven control

Data-driven control method have been first introduced in Section 2.1.3, but we make a point of revisiting the subject to discuss their contributions to the current field.

As previously mentioned, [39] serves as a valuable entry point into the field of data-driven HVAC control, providing a focused review of relevant methodologies. Additionally, reference [38] offers a comprehensive study, highlighting RL's potential for real-world HVAC applications while considering its limitation caused by its lack of scalability.

Numerous other recent studies have explored the use of RL for HVAC control; however, most of these approaches face significant scalability challenges, particularly when applied to large or multi-zone building environments. One such study is [70], where Nguyen et al. present a deep reinforcement learning (DRL) approach to model and optimize building HVAC control strategies. The authors formulate the HVAC control problem as a Markov decision process and use a deep Q-network to learn optimal control policies that minimize energy consumption while maintaining thermal comfort. The DRL agent interacts with a building simulation environment, receiving feedback in the form of energy cost and comfort violations, to update its policy over time. The model is trained using simulation data from a typical building scenario and is compared against conventional rule-based and MPC strategies. Results show that the DRL-based controller can adapt to dynamic conditions and outperform traditional methods in terms of both energy efficiency and comfort maintenance. A key advantage of this method is its ability to handle complex, nonlinear environments without requiring an explicit system model. However, the approach requires extensive training data, simulation time, and careful reward function design, which may limit its immediate deployment in real-world settings.

Another example is [71]. The authors investigate the integration of predictive information into DRL for HVAC control in office buildings. The study introduces a predictive DRL framework that incorporates weather forecasts and occupancy schedules into the agent's state

representation to enhance control performance. A twin delayed deep Deterministic policy gradient (TD3) algorithm is employed to learn a continuous control policy that balances energy efficiency and occupant comfort. The proposed method is tested in a high-fidelity building simulation environment and is shown to significantly outperform both standard DRL and rule-based controllers, especially under variable and uncertain external conditions. The inclusion of predictive information allows the agent to anticipate future changes and proactively adjust control actions, reducing energy consumption and temperature deviations. This work highlights the potential of combining DRL with forecasting data for building control. However, the reliance on accurate forecasts and the computational burden of training remain challenges for real-world deployment.

2.2.6 Building modelling

As previously discussed, the performance of HVAC control is highly dependent on the accuracy of the underlying prediction model used to capture the building's dynamic thermal response. This dependency becomes especially critical in MPC strategies with long prediction horizons, where even moderate errors in forecasting can significantly degrade control performance and lead to suboptimal or unreliable decisions. Therefore, it is essential to review existing literature on building thermal modelling to better understand the modelling techniques and their implications for control accuracy.

2.2.6.1 Thermal modelling of buildings

The following works describe and compare different approaches for modelling indoor temperature in buildings. The literature offers a wide range of methods; the following are selected examples that were useful to the understanding of the key concepts. For an in depth review, the reader is referred to [72].

As previously mentioned, Patonai et al. [35] provide a comprehensive review of current building thermal modelling techniques, including physics-informed approaches, autoregressive moving average with external inputs (ARMAX), and artificial neural network (ANN). They also propose a multiple linear regression model to predict indoor temperature in mobile containers using inputs like outdoor temperature and solar radiation. The model is simple, interpretable, and suitable for control, though limited under nonlinear or rapidly changing conditions. In [36], a RC network model is developed based on a heat balance equation, designed for real-time applications. While it is applied to intensive cooling scenarios, such as walk-in freezers, the methodology is generalizable and serves as a solid foundation for RC-based modelling. They proposed a RC thermal model for real-time cooling load estimation

in HVAC systems. The model balances accuracy and computational efficiency, making it suitable for online control applications, though simplifications may limit precision in complex dynamics. Building on this, Wang et al. [37] introduce a second-order thermal network model for residential buildings, capturing key heat transfer dynamics with improved accuracy over first-order models. Their work emphasizes the interaction between thermal control and grid dynamics, using peak shaving and mitigation of the *duck curve* as motivating factors [73]. The approach remains computationally efficient, though it may require detailed parameter identification for different building types. Finally, Bünning et al. [74] present a linear, physics-informed MPC strategy tailored to residential buildings. A key contribution of their study is a detailed comparison between physics-based and machine learning models, showing that their linear regression model achieves comparable performance while remaining interpretable and computationally efficient. The modelling approach developed in [74] forms the basis for the control strategy adopted in this thesis.

2.3 Literature gap

HVAC systems optimization can contribute to grid efficiency and defer investments in electrical infrastructure through DR programs. In this chapter, we surveyed the most conventional control methods for HVAC, discussing their advantages and limitations. Additionally, we examined data-driven control approaches, including machine learning techniques, highlighting both their strengths and challenges. Optimization-based methods—both continuous and mixed-integer—were also reviewed, along with their computational implications. Given the intensive computational requirements of centralized approaches, we further explored decentralized and distributed control strategies. Finally, we reviewed relevant literature on building thermal modelling and the integration of weather forecast data.

From this review, a clear research gap emerges in the decentralized control of indoor temperature using mixed-integer optimization. While mixed-integer formulations have demonstrated success in small-scale HVAC systems, existing implementations are generally limited to fewer than five zones. Moreover, many of the proposed control techniques face practical limitations when applied in real-world settings, such as complex modelling requirements, extensive data dependency, and challenges in meeting real-time computational constraints. Therefore, the development of a scalable and easily implementable, with performance guarantee, decentralized control strategy capable of handling a larger number of thermal zones would constitute a significant contribution to the field.

To address this gap in the literature, this work proposes the development of a decentralized, data-driven controller for HVAC systems in multi-zone commercial buildings. The control

strategy is based on MPC with mixed-integer decision variables, enabling compatibility with real-time constraints and participation in DR programs. The predictive model within the MPC framework is a data-driven, physics-informed linear representation of building thermal dynamics, inspired by the modelling approach used in [27]. The decentralization technique draws on the architecture proposed in [74], enabling scalable control across multiple zones while maintaining computational efficiency. The overall approach also integrates insights from decentralized optimization strategies such as those presented in [59].

CHAPTER 3 PHYSICS-INFORMED LINEAR REGRESSION FOR THERMAL MODELLING

The first component of the methodology is the thermal prediction model. To this end, a linear, data-driven, physics-informed model inspired by [74] is developed next. This model offers several advantages: it is easily adaptable to different multi-zone buildings and weather conditions due to its data-driven nature; its linearity enables it to be embedded in a decentralized mixed-integer program; and it captures thermal coupling between zones. Moreover, when integrated into an MPC controller, it has been shown to achieve comparable performance to neural network-based models [74]. This modelling approach aligns with the requirements of the proposed control framework.

In the sequel, we formulate a physics-informed, linear thermal model for HVAC controlled multi-zone commercial building with parameter identification done via a linear regression. Let $T_i(t) \in \mathbb{R}$ be the temperature of a building's zone at time t. Let $\mathbf{u}_i(t) \in \{0,1\}^4$ be the discrete control in the form of a one-hot vector taken by the HVAC$_i$ at time t. Let $\dot{Q}_\mathrm{s} \in \mathbb{R}$ be the heat exchanges from sources and sinks impacting the zones' temperature. During the development of the model, building characteristics and unknown parameters are replaced by variables denoted as θ_s, which are intended to be learned through a linear regression process. Other parameters not defined by θ_s are either a known variable are a variable meant to be optimized. The details of the training procedure, including the data used for parameter identification, are presented in Section 5.1.2.

3.1 Physics-informed linear regression model

From [74] we get the physical model of the temperature dynamics in a HVAC controlled zone expressed as a differential equation:

$$mc_\mathrm{p}\frac{\mathrm{d}T}{\mathrm{d}t} = \dot{Q}_\mathrm{out} + \dot{Q}_\mathrm{n} + \dot{Q}_\mathrm{sol} + \dot{Q}_\mathrm{occ} + \dot{Q}_\mathrm{act}, \tag{3.1}$$

where each $\dot{Q}_\mathrm{s}$ represent the thermodynamical impact of an aspect of the environments, respectively, the impact of the outdoor temperature $\dot{Q}_\mathrm{out}$, neighbouring zones temperature $\dot{Q}_\mathrm{n}$, solar irradiance $\dot{Q}_\mathrm{sol}$, the occupancy $\dot{Q}_\mathrm{occ}$, and the actuators $\dot{Q}_\mathrm{act}$, m represents the mass, and c_p is the specific heat capacity of the mass. In our case, the mass considered is the air in a zone. Finally, $\frac{\mathrm{d}T}{\mathrm{d}t}$ describes the time derivative of the temperature, i.e., the temperature rate of change through time. Model (3.1) takes into account important and

impactful environmental factors, viz., outdoor temperature and solar irradiance while taking into account the coupling between neighbouring zones and the impact of the HVAC system. The occupancy term, $\dot{Q}_{\text{occ}}$, is neglected in this work because determining its value represents ongoing research [74].

3.1.1 Forward difference

We replace the continuous-time derivative in (3.1) with a form that is more suitable for numerical implementation. As detailed in [75], numerical differentiation techniques can be employed to approximate derivatives effectively. Using the definition of the derivative,

$$\frac{\mathrm{d}t_0}{\mathrm{d}t} = \lim_{h \to 0} \frac{f(t_0 + h) - f(t_0)}{h}, \tag{3.2}$$

we derive a discrete-time formulation that fits the needs of our discrete-time MPC. As seen in (3.2) the derivative can be approximated by the change in the function's value over a small time interval. This approximation becomes increasingly accurate as h decreases, or when the function varies slowly over time. In our application, temperature and forecast data are available at discrete intervals of five minutes, corresponding to one time step. While five minutes is not infinitesimally small, we assume that in a large open-space commercial building, temperature changes over such intervals are gradual. This justifies the use of the forward difference method, as presented in (3.3), to approximate the derivative of temperature in our model.

$$\frac{\mathrm{d}T}{\mathrm{d}t} \approx \frac{T(t_0 + 1 \text{ time step}) - T(t_0)}{1 \text{ time step}} = T(t_0 + 1) - T(t_0) \tag{3.3}$$

Using the forward difference approximation, we obtain:

$$\begin{aligned} mc_{\text{p}} \frac{(T(t_0 + 1) - T(t_0))}{h} &= \dot{Q}_{\text{out}} + \dot{Q}_{\text{n}} + \dot{Q}_{\text{sol}} + \dot{Q}_{\text{act}} \\ T(t_0 + 1) &= \frac{h}{mc_{\text{p}}} (\dot{Q}_{\text{out}} + \dot{Q}_{\text{n}} + \dot{Q}_{\text{sol}} + \dot{Q}_{\text{act}}) + T(t_0), \end{aligned} \tag{3.4}$$

which describes the evolution of the temperature within a given zone. To effectively utilize this equation, it is necessary to analyze the influence of each heat flow term $\dot{Q}$ and understand its contribution to the thermal dynamics of the system.

3.1.2 Outdoor and neighbouring temperature impact

The terms $\dot{Q}_{\text{out}}$ and $\dot{Q}_{\text{n}}$ are modelled following the definitions provided in [74]. As shown in (3.5) and (3.6), these heat transfer components are given by a linear relation between the

temperature difference, either with the outdoor temperature T_{out} or neighbouring zones T_{n}, and the temperature within the controlled zone T. The coefficients θ_{out} and θ_{n} serve as a weighting factor that captures the thermal properties of the specific environment in :

$$\dot{Q}_{\text{out}} = \theta_{\text{out}}(T_{\text{out}} - T), \tag{3.5}$$

$$\dot{Q}_{\text{n}} = \theta_{\text{n}}(T_{\text{n}} - T). \tag{3.6}$$

In the case of neighbouring zones, multiple adjacent areas may influence the thermal dynamics of the controlled zone. To account for their combined effect, we substitute (3.5) by (3.7), which aggregates the impact of all neighbouring temperatures:

$$\begin{aligned} &\dot{Q}_{\text{n}} = \sum_{i=1}^{N} \theta_{\text{n}_i}(T_{\text{n}_i} - T), \\ &T_{\text{n}_i} : \text{neighbours' temperature;} \\ &T : \text{temperature in your zone;} \\ &\theta_{\text{n}_i} : \text{thermal coupling between zones, weight to learn.} \end{aligned} \tag{3.7}$$

Furthermore, the relationship between the outdoor temperature and the zone temperature is slightly modified to reflect the characteristics of the simulation environment. In this work, we employ a building temperature simulator developed and provided by our collaborators at BBAI. The simulator is built upon open-source Modelica libraries. The simulation environment and configurations, i.e., the building model, weather data, and operational profiles are all sourced directly from BBAI. The simulation used in this work includes a variable, u_{eco}, which represents the proportion of air exchange between the indoor and outdoor environments. According to the simulator's documentation, u_{eco} is a continuous value between 0 and 1, where 0 corresponds to fully recirculated air with no outdoor intake, and 1 corresponds to complete reliance on outdoor air with no recirculation. In our model given by

$$\dot{Q}_{\text{out}} = u_{\text{eco}}\theta_{\text{out}}(T_{\text{out}} - T), \tag{3.8}$$

u_{eco} is used as a scaling factor for the parameter θ_{out}, thereby adjusting the influence of the outdoor temperature according to the level of ventilation. The modification shown in (3.8), captures the intuitive notion that higher u_{eco} values indicate increased exposure to outdoor conditions, analogous to the presence of open windows.

3.1.3 Solar impact

Solar radiation represents one of the most significant sources of heat gain for the building. Following the approach of [74], its impact is modelled by:

$$\dot{Q}_{\text{sol}} = A_{\text{win}} \sin(\alpha - \alpha_0) \frac{\cos(\beta)}{\sin(\beta)} I_{\text{hor}}, \tag{3.9}$$

and depends on the solar azimuth angle α, and elevation angle β, as well as the window surface area A_{win}, and the global horizontal irradiance (GHI), denoted as I_{hor}. In the following section, we examine these parameters in details to better understand (3.9) and identify appropriate sources for their values.

Before analyzing the individual parameters of (3.9), it is important to emphasize that the expression is inherently non-linear due to the presence of trigonometric functions. However, as will be shown in subsequent sections, both the solar elevation angle β, and the azimuth angle α, are known quantities determined by the sun's position, which is itself a function of time and geographic location. The azimuthal offset α_0, is treated as a learned parameter that accounts for the orientation of the building. These characteristics enable us to reasonably approximate (3.9) as linear in function in α and β for the purposes of modelling and control. We reformulate (3.9) using the trigonometric identity,

$$\sin(\alpha - \alpha_0) = \sin\alpha \cos\alpha_0 - \cos\alpha \sin\alpha_0, \tag{3.10}$$

where the $\cos\alpha_0$ and $\sin\alpha_0$ terms can now be treated as separate parameters to be learned as weight in a linear regression leading to

$$\dot{Q}_{\text{sol}} = A_{\text{win}} (\sin\alpha \cos\alpha_0 - \cos\alpha \sin\alpha_0) \frac{\cos(\beta)}{\sin(\beta)} I_{\text{hor}}. \tag{3.11}$$

Finally, (3.11) can be further simplified by grouping all constant terms into two distinct parameters, resulting in

$$\dot{Q}_{\text{sol}} = (\theta_{\text{sol}} \sin\alpha - \theta'_{\text{sol}} \cos\alpha) \frac{\cos(\beta)}{\sin(\beta)} I_{\text{hor}}. \tag{3.12}$$

3.1.3.1 Solar Irradiation

The model incorporates the GHI, I_{hor}, which represents the total solar radiation received per unit area on a horizontal surface at the Earth's surface [76]. Formally, solar irradiance received at the Earth's surface can be decomposed into three fundamental components: direct

normal irradiance (DNI), diffuse horizontal irradiance (DHI), and global horizontal irradiance (GHI) [76]:

- DNI measures the solar radiation received per unit area by a surface that is always perpendicular to the sun's rays. It represents unscattered, direct sunlight and is especially relevant for systems that track the sun or use solar concentration;
- DHI represents sunlight that has been scattered by the atmosphere and reaches a surface from all directions. It is measured on a horizontal plane and excludes direct sunlight;
- GHI is the total solar radiation incident on a horizontal surface, combining both the direct component and the diffuse component. It is commonly used in evaluating fixed solar panel installations.

The relation between the solar irradiation measurements and the zenith angle $\vartheta_\mathrm{z} \in [0, 90]^\circ$, defined as the angle between the sun and the vertical direction, can be seen in:

$$\mathrm{GHI} = \mathrm{DHI} + \mathrm{DNI}\cos(\vartheta_\mathrm{z}), \tag{3.13}$$

as presented in [77]. These four parameters are typically included in standard meteorological datasets. In our case, the simulation environment provides values for GHI, DHI, and DNI.

3.1.3.2 Elevation angle

The solar elevation angle $\beta \in [0, 90]^\circ$ indicates the height of the sun above the horizon as seen by an observer. It is the complementary to the zenith angle ϑ_z:

$$\beta = 90^\circ - \vartheta_\mathrm{z}. \tag{3.14}$$

If the zenith angle is not directly available, it can be estimated using (3.13), provided that all necessary irradiance components are known. To avoid potentially dividing by zero, we assume that when DNI is null, the zenith angle is set to zero degree. This corresponds to an elevation angle of 90°, resulting in a solar heat gain $\dot{Q}_\mathrm{sol}$ equal to zero. Alternatively, the

zenith angle can also be calculated using [77]:

$$\cos\vartheta_z = \cos\varphi\cos\delta\cos\omega + \sin\varphi\sin\delta,$$

where,

$$\begin{aligned} \varphi &: \text{Latitude;} \\ \delta &: \text{Declination angle;} \\ \omega &: \text{Hour angle.} \end{aligned} \tag{3.15}$$

In (3.15), the latitude, φ, represents the geographical coordinate of the location (e.g., 45.5088°N for Montréal, Canada). The hour angle ω denotes the angular displacement of the sun relative to the local meridian (73^{rd} meridian for Montréal), with negative values in the morning, positive values in the afternoon, and zero at the solar noon. The declination angle δ describes the solar position at solar noon and varies between $-23.45°$ and $23.45°$ over the course of the year. Its value can be computed using:

$$\delta = 23.45° \sin\left(\frac{360°(284+n)}{365}\right), \tag{3.16}$$

where n is the day number of the year starting from January 1^{st} (see [77, Table 1.6.1] for example values of n). Once the zenith angle is obtained, the elevation angle can be derived from (3.14).

3.1.3.3 Azimuth angle

The azimuth angle $\alpha \in [-180, 180]°$ represents the horizontal angle between the sun's position and a reference direction, typically south (or north, depending on the convention used). Together with the elevation and zenith angles, it provides a complete description of the sun's position in the sky. The azimuth angle is defined in:

$$\alpha = \text{sign}(\omega)\left|\arccos\left(\frac{\cos\vartheta_z\sin\varphi - \sin\delta}{\sin\vartheta_z\cos\varphi}\right)\right|, \tag{3.17}$$

using the same parameters as in the computation of the elevation angle. Because the inverse cosine function (arccos) returns values in the first or second quadrant, the sign of the hour angle is used to resolve the ambiguity and determine the correct quadrant. It is important to note that (3.17) assumes south as the reference direction; if north is preferred, an offset of 180° should be applied.

3.1.3.4 Hour angle

The hour angle $\omega \in [-180°, +180°]$ is a key parameter used in the calculation of both the azimuth angle α and the elevation angle β and represents the angular position of the sun relative to the local meridian. This angle is directly dependent on the solar time. Solar time can be expressed either as actual solar time (AST) or mean solar time (MST). AST reflects the actual interval between two solar noons, which varies slightly throughout the year, while MST standardizes this interval to exactly 24 hours. In both systems, solar noon corresponds to a solar time of 12:00, when the sun reaches its highest point in the sky. For practical purposes, MST is commonly used as the solar hour due to its consistency and ease of calculation. This work uses MST as solar time, as represented in:

$$\omega = 15°(\text{MST} - 12). \tag{3.18}$$

MST is defined as a function of the local clock time, CT, and geographic longitude (e.g., 73.561668°W for Montréal, Canada) [78]:

$$\text{MST} = \text{CT} + \frac{\Delta L}{15°}, \tag{3.19}$$

in decimal hours, where ΔL represents the longitudinal difference (in degrees) between the location of interest and the standard meridian for the local time zone, with positive values indicating an easterly direction. For Montréal, the standard meridian for the UTC-5 time zone is located at 75°W, while the city's actual longitude is 73.561668°W. Therefore, the longitudinal offset is given by $\Delta L = 75° - 73.5617° = 1.4383°$.

3.1.4 Actuator's impact

There are several approaches to model the thermal impact of the HVAC system. We employ the formulation, adapted from [74]:

$$\dot{Q}_{\text{act}} = \dot{m}_{\text{f}} c_{\text{p,f}} (T_{\text{sup}} - T_{\text{ret}}), \tag{3.20}$$

where $\dot{m}_{\text{f}}$ represents the mass flow rate of the fluid (air, in this case), and $c_{\text{p,f}}$ is its specific heat capacity. The two temperatures involved are the supply temperature T_{sup} delivered by the HVAC system, and the return temperature T_{ret} which corresponds to the air returning from the zone. In air-based heating or cooling systems, the return temperature is often approximated by the current zone temperature. To simplify the model, the product $\dot{m}_{\text{f}} c_{\text{p,f}}$ can be approximated by a constant parameter θ_{act}, yielding a simpler form based on a learned

parameter as seen in [74]:

$$\dot{Q}_{\text{act}} \approx \theta_{\text{act}}(T_{\text{sup}} - T). \tag{3.21}$$

The only remaining component to be defined is the supply temperature delivered to the zone. The supply temperature delivered to the zone is not directly observable and may vary depending on factors such as the system's capacity, configuration, and control logic. Nevertheless, it is reasonable to assume that the supply temperature is determined by the control signal applied to the HVAC system. In this work, the control strategy is modelled as a mixed-integer formulation. Consequently, we represent the control input using a one-hot encoded vector,

$$\mathbf{u} = \begin{bmatrix} u_{\text{none}} \\ u_{\text{fan}} \\ u_{\text{heat}_1} \\ u_{\text{heat}_2} \end{bmatrix} \in \{0,1\}^4, \tag{3.22}$$

where $u_{\text{none}} + u_{\text{fan}} + u_{\text{heat}_1} + u_{\text{heat}_2} = 1$, i.e., a single control can be deployed at a time, and each control action corresponds to a parameter representing the associated supply temperature. In the simplified case considered here, which focuses exclusively on heating, the supply temperature can be defined as:

$$T_{\text{sup}} = \boldsymbol{\theta}_{\text{sup}}^{\top}\mathbf{u} = [\theta_{\text{sup}} \quad \theta'_{\text{sup}} \quad \theta''_{\text{sup}} \quad \theta'''_{\text{sup}}] \begin{bmatrix} u_{\text{none}} \\ u_{\text{fan}} \\ u_{\text{heat}_1} \\ u_{\text{heat}_2} \end{bmatrix}, \tag{3.23}$$

where the possible control actions include: no action (none), activation of the fan without temperature change (fan), and two levels of heating (heat_1 and heat_2), which lead to the actuator heat flow model given by:

$$\dot{Q}_{\text{act}} = \theta_{\text{act}}([\theta_{\text{sup}} \quad \theta'_{\text{sup}} \quad \theta''_{\text{sup}} \quad \theta'''_{\text{sup}}] \begin{bmatrix} u_{\text{none}} \\ u_{\text{fan}} \\ u_{\text{heat}_1} \\ u_{\text{heat}_2} \end{bmatrix} - T). \tag{3.24}$$

However, a significant issue arises when using (3.24) as the actuator model. Specifically, the

formulation results in a non-zero heat flow even when the control signal corresponds to the *none* action, which is physically inconsistent. To address this, we considered two alternatives: setting the corresponding parameter, $\boldsymbol{\theta}_{\text{sup}}$, equal to the zone temperature T, or incorporating the *none* control as a multiplicative factor for the entire function, e.g.,

$$\dot{Q}_{\text{act}} = (1 - u_{\text{none}})([\theta'_{\text{sup}} \quad \theta''_{\text{sup}} \quad \theta'''_{\text{sup}}] \begin{bmatrix} u_{\text{fan}} \\ u_{\text{heat}_1} \\ u_{\text{heat}_2} \end{bmatrix} - T). \tag{3.25}$$

However, both approaches introduce bilinear terms and induce non-separable constraints that complicate the optimization process, viz., that is not compatible with our decentralized approach.

A suitable solution is to adopt the simplified formulation of (3.24) proposed in [74]:

$$\dot{Q}_{\text{act}} = \boldsymbol{\theta}_{\text{act}}^{\top} \mathbf{u} = [\theta_{\text{act}} \quad \theta'_{\text{act}} \quad \theta''_{\text{act}}] \begin{bmatrix} u_{\text{fan}} \\ u_{\text{heat}_1} \\ u_{\text{heat}_2} \end{bmatrix}, \tag{3.26}$$

where $u_{\text{fan}} + u_{\text{heat}_1} + u_{\text{heat}_2} \in \{0, 1\}$ and $u_{\text{fan}} = u_{\text{heat}_1} = u_{\text{heat}_2} = 0$ corresponds to $u_{\text{none}} = 1$ in (3.25). In this model, the actuator's heat contribution is represented in a more physically consistent manner. While this approach neglects the explicit influence of the temperature difference between the zone and the supply air, it ensures that no heat is added when the control signal is inactive. Despite its simplifications, this method is sufficiently accurate and well-suited for the requirements of our control-oriented application.

3.2 Physics-informed linear thermal model

The physics-informed linear thermal model is constructed by starting from the general heat balance in (3.1), applying the forward difference approximation as shown in (3.4), and incorporating the definitions of heat flow given in (3.7), (3.8), (3.12), and (3.26). This process

leads to the linear model defined by:

$$\begin{aligned} T(t+1) - T(t) = \quad & u_{\text{eco}}(t)\theta_{\text{out}}(T_{\text{out}}(t) - T(t)) + \sum_{i=1}^{N} \theta_{\text{n}_i}(T_{\text{n}_i}(t) - T(t)) \\ & + (\theta_{\text{sol}} \sin(\alpha(t)) - \theta'_{\text{sol}} \cos(\alpha(t))) \frac{\cos(\beta(t))}{\sin(\beta(t))} I_{\text{hor}}(t) \\ & + [\theta_{\text{act}} \quad \theta'_{\text{act}} \quad \theta''_{\text{act}}] \begin{bmatrix} u_{\text{fan}}(t) \\ u_{\text{heat}_1}(t) \\ u_{\text{heat}_2(t)} \end{bmatrix} + \theta_{\text{bias}}, \end{aligned} \tag{3.27}$$

where the term $\frac{h}{mc_{\text{p}}}$, which appears in (3.4) has been absorbed into the corresponding θ parameters for each heat flow component and θ_{bias} is an added bias term used to render the learning process more flexible.

3.3 Linear System Model Formulation

We aim to create a model that generalizes across varying numbers of zones and time steps. To this end, we begin by analyzing the dynamics of a two-zone building in order to identify structural patterns or rules that can support the formulation of a scalable and recurrent model express in matrix form.

Applying (3.27) to a two-zone building and denoting as $T_1(t)$ and $T_2(t)$ zone 1's and 2's temperature at time t, the $t = 1$ temperature is expressed as function of the neighbouring zone temperature at $t = 0$. Specifically, we have:

$$\begin{aligned} T_1(1) = \quad & u_{\text{eco}_1}(0)\theta_{\text{out}_1}(T_{\text{out}}(0) - T_1(0)) + \theta_{\text{n}_1}(T_2(0) - T_1(0)) \\ & + (\theta_{\text{sol}_1} \sin(\alpha(0)) - \theta'_{\text{sol}_1} \cos(\alpha(0))) \frac{\cos(\beta(0))}{\sin(\beta(0))} I_{\text{hor}}(0) \\ & + [\theta_{\text{act}_1} \quad \theta'_{\text{act}_1} \quad \theta''_{\text{act}_1}] \begin{bmatrix} u^1_{\text{fan}}(0) \\ u^1_{\text{heat}_1}(0) \\ u^1_{\text{heat}_2}(0) \end{bmatrix} + \theta_{\text{bias}_1} + T_1(0) \end{aligned} \tag{3.28}$$

$$
\begin{aligned}
T_2(1) = \quad & u_{\mathrm{eco}_2}(0)\theta_{\mathrm{out}_2}(T_{\mathrm{out}}(0) - T_2(0)) + \theta_{\mathrm{n}_2}(T_1(0) - T_2(0)) \\
& + (\theta_{\mathrm{sol}_2}\sin(\alpha(0)) - \theta'_{\mathrm{sol}_2}\cos(\alpha(0)))\frac{\cos(\beta(0))}{\sin(\beta(0))}I_{\mathrm{hor}}(0) \\
& + \begin{bmatrix} \theta_{\mathrm{act}_2} & \theta'_{\mathrm{act}_2} & \theta''_{\mathrm{act}_2} \end{bmatrix} \begin{bmatrix} u^2_{\mathrm{fan}}(0) \\ u^2_{\mathrm{heat}_1}(0) \\ u^2_{\mathrm{heat}_2}(0) \end{bmatrix} + \theta_{\mathrm{bias}_2} + T_2(0)
\end{aligned} \tag{3.29}
$$

Because the initial temperature $T_1(0)$ and $T_2(0)$ are known, it does not add complexity to the model. However, when predicting the temperature at $t = 2$, the previously known temperature becomes a variable—specifically, the temperature at $t = 1$—which significantly increases the complexity of the equations and introduces strong interdependencies between zones, thereby complicating model separation and decentralization.

It becomes evident that if we would substitute $T_2(1)$ with the full expression from (3.29) in $T_2(2)$ in

$$
\begin{aligned}
T_1(2) = \quad & u_{\mathrm{eco}_1}(1)\theta_{\mathrm{out}_1}(T_{\mathrm{out}}(1) - T_1(1)) + \theta_{\mathrm{n}_2}(T_2(1) - T_1(1)) \\
& + (\theta_{\mathrm{sol}_1}\sin(\alpha(1)) - \theta'_{\mathrm{sol}_1}\cos(\alpha(1)))\frac{\cos(\beta(1))}{\sin(\beta(1))}I_{\mathrm{hor}}(1) \\
& + \begin{bmatrix} \theta_{\mathrm{act}_1} & \theta'_{\mathrm{act}_1} & \theta''_{\mathrm{act}_1} \end{bmatrix} \begin{bmatrix} u^1_{\mathrm{fan}}(1) \\ u^1_{\mathrm{heat}_1}(1) \\ u^1_{\mathrm{heat}_2}(1) \end{bmatrix} + \theta_{\mathrm{bias}_1} + T_1(1)
\end{aligned} \quad , \tag{3.30}
$$

that the latter would become a more entangled and hard to follow formulation and that this pattern of dependency and substitution will only compound with each additional prediction step, resulting in rapidly growing model complexity over time.

To address the increasing complexity of the problem and to develop a scalable recurrent formulation that grows with both the prediction horizon and the number of zones, the thermal system model is reformulated as a linear matrix equation of the form:

$$
\sum_{i=1}^{N} \mathbf{A}_i \mathbf{x}_i = \mathbf{b}.
$$

A systematic procedure for constructing this recurrent system is presented, enabling its generation in a modular and automated manner. It is important to emphasize that this linear recurrent formulation does not inherently reduce the entanglement or computational complexity of the time-coupled system. However, it significantly improves its manipulability, facilitates implementation, and aligns with the structure required by our resolution frame-

work.

In this representation, each matrix $\mathbf{A}_i \in \mathbb{R}^{NW \times 6W+2}$, where N is the number of zones and W the control horizon, encodes the system dynamics and thermal coefficients, $\mathbf{x}_i \in \mathbb{R}^{6W+2}$ denotes the time-indexed input variables (e.g., outdoor temperature, solar irradiance, and control actions), and $\mathbf{b} \in \mathbb{R}^{NW}$ corresponds to the predicted zone temperatures over the forecast horizon. This compact formulation improves clarity, supports automation, and is well-suited to the proposed optimization-based control strategy.

An example of the structure of the input vector $\mathbf{x}_i$ is provided in (3.31), which contains the variables for a two-step prediction horizon for a given zone i. This vector can be generalized to accommodate any prediction horizon length and is constructed independently of the total number of zones in the building. Specifically, the input vector $\mathbf{x}_i$ requires only the current temperature at t and the control inputs for the zone, with the control sequence specified for the entire prediction horizon. In addition, it includes the forecasted solar irradiance and outdoor temperature values, which are considered exogenous parameters over the horizon. These values would be supplied by the forecast model used in the MPC. The final entry in the vector accounts for a bias term of (3.27). A systematic approach can thus be used to create easily the input vectors for different configuration.

$$\mathbf{x}_i(t) = \begin{bmatrix} T_i(t) \\ T_{\text{out}}(t) \\ T_{\text{out}}(t+1) \\ \dots \\ \sin(\alpha(t)))\frac{\cos(\beta(t))}{\sin(\beta(t))} I_{\text{hor}}(t) \\ -\cos(\alpha(t)))\frac{\cos(\beta(t))}{\sin(\beta(t))} I_{\text{hor}}(t) \\ \sin(\alpha(t+\tau)))\frac{\cos(\beta(t+\tau))}{\sin(\beta(t+\tau))} I_{\text{hor}}(t+\tau) \\ -\cos(\alpha(t+\tau)))\frac{\cos(\beta(t+\tau))}{\sin(\beta(t+\tau))} I_{\text{hor}}(t+\tau) \\ \dots \\ \mathbf{u}_i(t) \\ \mathbf{u}_i(t+1) \\ \dots \\ 1 \end{bmatrix}. \tag{3.31}$$

The structure of the matrices $\mathbf{A}_1$ and $\mathbf{A}_2$ for a two-zone configuration with a two-step prediction horizon starting at t are:

$$\mathbf{A}_1 = \left[\begin{array}{cccccccccccccc} \kappa_1 & \theta_{\text{out}_1} u_{\text{eco}_1} & 0 & \theta_{\text{sol}_1} & \theta'_{\text{sol}_1} & 0 & 0 & \theta_{\text{act}_1} & \theta'_{\text{act}_1} & \theta''_{\text{act}_1} & 0 & 0 & 0 & b_1 \\ \kappa_1^2 + \theta_{\text{n}_1}\theta_{\text{n}_2} & \kappa_1\theta_{\text{out}_1} u_{\text{eco}_1} & \theta_{\text{out}_1} u_{\text{eco}_1} & \kappa_1\theta_{\text{sol}_1} & \kappa_1\theta'_{\text{sol}_1} & \theta_{\text{sol}_1} & \theta'_{\text{sol}_1} & \kappa_1\theta_{\text{act}_1} & \kappa_1\theta'_{\text{act}_1} & \kappa_1\theta''_{\text{act}_1} & \theta_{\text{act}_1} & \theta'_{\text{act}_1} & \theta''_{\text{act}_1} & \kappa_1 b_1 + b_1 \\ \hdashline \theta_{\text{n}_2} & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ \kappa_1\theta_{\text{n}_2} + \kappa_2\theta_{\text{n}_2} & \theta_{\text{n}_2}\theta_{\text{out}_1} u_{\text{eco}_1} & 0 & \theta_{\text{n}_2}\theta_{\text{sol}_1} & \theta'_{\text{sol}_1}\theta_{\text{n}_2} & 0 & 0 & \theta_{\text{act}_1}\theta_{\text{n}_2} & \theta'_{\text{act}_1}\theta_{\text{n}_2} & \theta''_{\text{act}_1}\theta_{\text{n}_2} & 0 & 0 & 0 & \theta_{\text{n}_2} b_1 \end{array}\right], \tag{3.32}$$

where $\kappa_1 = 1 - \theta_{\text{out}_1} u_{\text{eco}_1} - \theta_{\text{n}_1}$ and,

$$\mathbf{A}_2 = \left[\begin{array}{cccccccccccccc} \theta_{\text{n}_1} & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ \kappa_1\theta_{\text{n}_1} + \kappa_2\theta_{\text{n}_1} & \theta_{\text{n}_1}\theta_{\text{out}_2} u_{\text{eco}_2} & 0 & \theta_{\text{n}_1}\theta_{\text{sol}_2} & \theta'_{\text{sol}_2}\theta_{\text{n}_1} & 0 & 0 & \theta_{\text{act}_2}\theta_{\text{n}_1} & \theta'_{\text{act}_2}\theta_{\text{n}_1} & \theta''_{\text{act}_2}\theta_{\text{n}_1} & 0 & 0 & 0 & \theta_{\text{n}_1} b_2 \\ \hdashline \kappa_2 & \theta_{\text{out}_2} u_{\text{eco}_2} & 0 & \theta_{\text{sol}_2} & \theta'_{\text{sol}_2} & 0 & 0 & \theta_{\text{act}_2} & \theta'_{\text{act}_2} & \theta''_{\text{act}_2} & 0 & 0 & 0 & b_2 \\ \kappa_2^2 + \theta_{\text{n}_1}\theta_{\text{n}_2} & \kappa_2\theta_{\text{out}_2} u_{\text{eco}_2} & \theta_{\text{out}_2} u_{\text{eco}_2} & \kappa_2\theta_{\text{sol}_2} & \kappa_2\theta'_{\text{sol}_2} & \theta_{\text{sol}_2} & \theta'_{\text{sol}_2} & \kappa_2\theta_{\text{act}_2} & \kappa_2\theta'_{\text{act}_2} & \kappa_2\theta''_{\text{act}_2} & \theta_{\text{act}_2} & \theta'_{\text{act}_2} & \theta''_{\text{act}_2} & \kappa_2 b_2 + b_2 \end{array}\right], \tag{3.33}$$

where $\kappa_2 = 1 - \theta_{\text{out}_2} u_{\text{eco}_2} - \theta_{\text{n}_2}$.

The matrices are separated by horizontal lines to denote the contribution to each zones. In this case, the upper part of (3.32) and (3.33) are associated with zone 1 while the lower part is associated with zone 2. To highlight how the system complexity increases even after a single additional step, the fully expanded matrices for $\mathbf{A}_1$ and $\mathbf{A}_2$ under a three-step horizon starting at time t are provided in Appendix A.

The construction of the $\mathbf{A}_i$ matrix can be systematically derived by following the patterns observed in the two and three time steps examples. This recurrence is embedded in the matrix structure and can be implemented by applying the following steps for each new row corresponding to a future time step:

1. Incorporate the coefficients associated with the forecasted outdoor temperature, using the product $\theta_{\text{out}_i} u_{\text{eco}_i}$, and the solar heat gains, represented by $\boldsymbol{\theta}_{\text{sol}_i}$;
2. Add the coefficients corresponding to the control actions, denoted by $\boldsymbol{\theta}_{\text{act}_i}$;
3. Apply a factor of κ_i to zone i previous temperature;
4. Apply a factor of $\boldsymbol{\theta}_{\text{n}_{-i}}$ to the previous temperatures of all neighbouring zones;
5. Repeat steps 1 to 4.

3.4 Other factors in the model

The core components required to construct the physics-informed linear model have been presented. However, some additional refinements can be introduced to enhance both the model's accuracy and its performance within a time-series optimization framework.

3.4.1 ARMAX

Although the use of an autoregressive moving average with exogenous inputs (ARMAX) [79] approach is commonly recommended in the literature for modelling indoor temperature dynamics [74, 79, 80], our analysis showed that the inclusion of autoregressive terms had a negligible impact on the overall model performance. Hence, to simplify the model, they have been neglected.

3.4.2 Artificial biases

For the centralized implementation of the physics-informed linear model, artificial biases were introduced to improve the optimization performance. Although the model generally captured temperature dynamics well, it occasionally exhibited systematic deviations—being either overly optimistic or pessimistic in certain zones. To address this, zone-specific bias terms is added to correct for these discrepancies. The decentralized implementation did not require the addition of such biases, as its structure inherently allowed for more localized adaptation and accuracy.

3.4.3 Slack variables and back-up controller

Given that our control framework is optimization-based and operates within a dynamic and uncertain environment, it is possible for the optimization problem to become infeasible—particularly when model prediction errors push zone temperatures beyond the bounds that the available control actions can correct. Such infeasibility prevents the controller from generating a valid solution. A common approach to address this issue is the introduction of slack variables into the temperature constraints, allowing controlled violations at the cost of penalization in the objective function. However, the inclusion of slack variables significantly increases the complexity of the problem, as it requires extensive tuning to balance the trade-off between constraint violation, e.g., temperature bounds, and the corresponding optimization objectives. Moreover, the added variables increase the computational burden. To avoid these drawbacks, we opt to implement a backup controller that activates when the optimization problem becomes infeasible. The backup controller operates by selecting between

two discrete actions: either no heating or maximum heating. When the zone temperature falls below a predefined threshold—specifically, the minimum allowable temperature plus a safety margin—the controller activates maximum heating. Otherwise, it remains inactive. This control logic is based on the assumption that, outside of the threshold region, the system will be capable of recovering the desired temperature in subsequent control steps without requiring immediate intervention. This approach provides a simpler and more robust alternative, avoiding the tuning sensitivity and processing overhead associated with slack-based formulations. The design of the backup controller is detailed in a subsequent section.

3.5 Conclusion

We have formulated a physics-informed linear model that captures the temperature dynamics in multi-zone commercial buildings. This model includes a set of unknown parameters θ_{s}, which will be identified through linear regression using simulation data. The model is structured as :

$$\sum_{i=1}^{N} \mathbf{A}_i \mathbf{x}_i = \mathbf{b},$$

in such a way that the input vector $\mathbf{x}_i$ and the coefficient matrix $\mathbf{A}_i$ associated with each zone i are independent of the others. This structure enables the decentralization of the optimization problem, which will be addressed in the following section.

CHAPTER 4 DECENTRALIZED MIXED-INTEGER OPTIMIZATION FOR MULTI-ZONE COMMERCIAL BUIDING HVACs

In this chapter, we address the decentralization of the optimal control process. We begin by presenting the centralized optimization-based controller, followed by the development of a decentralized algorithm that enables its decentralized implementation.

4.1 Centralized optimization problem

We continue the work done in Section 1.2.3 and further explain the optimization problem tackled in this work. We formulate the control problem as a MPC problem. As previously introduced in Section 2.1.2, MPC determines optimal control actions based on a prediction model of the system dynamics together with forecast information of exogenous parameters over a predefined prediction horizon. At each time step, only the first control action of the optimized sequence is applied to the system. The environment then provides updated state information (e.g., zone temperature), which serves as the new initial condition for the subsequent optimization. This receding horizon strategy enables the controller to account for both current and forecasted states while mitigating the impact of modelling inaccuracies and forecast uncertainty. In this work, the prediction model is given by the linear regression function presented in Chapter 3, where each θ parameter is learned via a linear regression based on historical building data. Our problem can take the same form has the one described in [27] given by

$$\begin{aligned} \min_{\mathbf{x}_i} \quad & \sum_{i=1}^{N} \mathbf{c}_i^\top \mathbf{x}_i \\ \text{s.t.} \quad & \sum_{i=1}^{N} \mathbf{A}_i \mathbf{x}_i \leq \mathbf{b}, \\ & \mathbf{x}_i \in \mathcal{X}_i, \quad i = 1, \ldots, N, \end{aligned} \tag{4.1}$$

where $\mathcal{X}_i$ represents the feasible set for the input vector $\mathbf{x}_i$. Most components of (4.1) have already been introduced in Section 3.3, where a recurrent approach is used to define both the system matrix $\mathbf{A}_i$ and the input vector $\mathbf{x}_i$. We note that the majority of the elements in $\mathbf{x}_i$ are either forecasts or fixed parameters. Therefore, the only elements that must be explicitly constrained within $\mathcal{X}_i$ are the control inputs, denoted by $\mathbf{u}_i$. By definition, in the context of this work, $\mathbf{u}_i$ is modelled as a one-hot vector, meaning it is composed of boolean values with only one active element at a time. The two remaining elements that have not yet been described are the thermal limit constraint vector $\mathbf{b}$ and the cost vector $\mathbf{c}_i$, which

will be introduced in the following section.

4.1.1 Thermal limit vector

As previously discussed, each matrix $\mathbf{A}_i$ is defined per zone. In a two-zone configuration, for instance, the upper half of both $\mathbf{A}_1$ and $\mathbf{A}_2$ corresponds to zone 1, while the lower half corresponds to zone 2. Consequently, the combined expression $\mathbf{A}_1\mathbf{x}_1 + \mathbf{A}_2\mathbf{x}_2$ yields a vector in which the first half contains the predicted values over the entire horizon for zone 1, and the second half contains those for zone 2. Accordingly, the constraint vector $\mathbf{b}$ must match this structure and is therefore defined in $\mathbb{R}^{NW}$, where N is the number of zones and W the length of the time horizon. However, this formulation only allows a single temperature threshold to be enforced at a time. To accommodate both upper and lower temperature bounds, the problem must be reformulated, as shown in (4.2).

$$\begin{aligned} \min_{\mathbf{x}_i} \quad & \sum_{i=1}^{N} \mathbf{c}_i^\top \mathbf{x}_i \\ \text{s.t.} \quad & \sum_{i=1}^{N} \begin{bmatrix} -\mathbf{A}_{i,\min} \\ \mathbf{A}_{i,\max} \end{bmatrix} \mathbf{x}_i \leq \begin{bmatrix} -\mathbf{b}_{\min} \\ \mathbf{b}_{\max} \end{bmatrix}, \\ & \mathbf{u}_i \in \{0,1\}^{3W}, \quad i = 1, \ldots, N, \end{aligned} \tag{4.2}$$

The specific values assigned to the constraint vector, now $\mathbf{b} \in \mathbb{R}^{2NW}$, are determined by the desired upper and lower temperature bounds enforced within the control problem. The first half of $\mathbf{b}$ encodes the negative values of the lower temperature limits, while the second half encodes the corresponding upper bounds. These values are defined for all zones across the entire prediction horizon. In the specific case of a two-zone building with predefined upper and lower temperature bounds set to $\overline{T}(t)$ and $\underline{T}(t)$ °C respectively, and considering a prediction horizon of three time steps, the resulting thermal constraint vector takes the form

$$\begin{aligned} \mathbf{b} = [\; & -\underline{T}_1(t{+}1) \; -\underline{T}_1(t{+}2) \; -\underline{T}_1(t{+}3) \; -\underline{T}_2(t{+}1) \; -\underline{T}_2(t{+}2) \; -\underline{T}_2(t{+}3) \\ & \overline{T}_1(t{+}1) \; \overline{T}_1(t{+}2) \; \overline{T}_1(t{+}3) \; \overline{T}_2(t{+}1) \; \overline{T}_2(t{+}2) \; \overline{T}_2(t{+}3)]^\top. \end{aligned} \tag{4.3}$$

4.1.2 Cost vector

The cost vector $\mathbf{c}_i$ spans the entire prediction horizon and is primarily composed of zeros, except at positions corresponding to active control inputs $\mathbf{u}_i$. The control vector at time t is

defined as

$$\mathbf{u}_i(t) = \begin{bmatrix} u^i_{\text{fan}_1}(t) \\ u^i_{\text{heat}_1}(t) \\ u^i_{\text{heat}_2}(t) \end{bmatrix},$$

where each entry represents the activation status of a specific action in zone i. The cost associated with these actions is weighted according to their relative energy consumption. The instantaneous operating cost of HVAC unit i at time t is computed as

$$c_i(t) = \Delta c_P(t) \begin{bmatrix} P_{\text{fan}} \\ P_{\text{heat}_1} \\ P_{\text{heat}_2} \end{bmatrix}^\top \mathbf{u}_i(t), \tag{4.4}$$

where P_{s} are the respective power consumption values in kW for each actuator, Δ is the control time step duration in hours, and $c_P(t)$ is the energy price in ¢/kWh at time t. In vector form, this corresponds to an element-wise product between power usage, duration, dynamic pricing, and the one-hot structure of the control signal.

The power consumption values P_{s} are obtained directly from the simulator, and the resulting cost values reflect temporal variations in electricity pricing. In Section 5.1.5 of the the numerical study, we consider three representative pricing schemes offered by Hydro-Québec: TOU, CPP, and RTP.

4.1.3 Centralized problem resolution

The optimization problem defined in (4.2) can be efficiently solved using off-the-shelf commercial solvers, as the formulation is mixed-integer convex. However, the computational complexity increases significantly as the control horizon is extended, leading to larger coefficient matrices $\mathbf{A}_i$ and a greater number of binary variables within the input vector $\mathbf{x}_i$. Similarly, the addition of more zones further expands the size and number of these matrices.

4.2 Decentralized optimization problem

Finally, we adopt the decentralized algorithm for MILPs proposed in [27] to deploy our approach in an efficient and scalable way. This method is particularly well-suited for our problem structure, as it allows each zone to solve its own subproblem while coordinating with neighbouring zones through dual variables, thus maintaining inter-zone consistency.

This method guarantees finite-time feasibility and provides a performance bound on the deviation from the optimal solution under certain assumptions, given by,

$$\sum_{i=1}^{m} \mathbf{c}_i^\top \mathbf{x}_i(k) - J_{\mathcal{P}}^\star \leq \bar{\gamma} + \frac{\|\tilde{\rho}\|_\infty}{N\zeta}\,\tilde{\gamma}, \tag{4.5}$$

where the error–defined as the difference between the optimal objective $J_{\mathcal{P}}^\star$ and the actual objective $\sum_{i=1}^{m} \mathbf{c}_i^\top \mathbf{x}_i(k)$ is bounded by three elements: (i) the number of agents N, (ii) the finite-time asymptotic error $\bar{\gamma}$ at iteration k, representing the maximum gap between worst-case and best-case outcomes extrapolated to the finite-time horizon, and (iii) the absolute theoretical maximum difference between worst-case and best-case objective values $\tilde{\gamma}$, with $\bar{\gamma} \leq \tilde{\gamma}$. The latter is defined as

$$\tilde{\gamma} = N \max_{i \in \{1,\ldots,m\}} \left\{ \max_{\mathbf{x}_i \in X_i} \mathbf{c}_i^\top \mathbf{x}_i \;-\; \min_{\mathbf{x}_i \in X_i} \mathbf{c}_i^\top \mathbf{x}_i \right\}. \tag{4.6}$$

Lastly, $\|\tilde{\rho}\|_\infty$ represents a worst-case bound on constraint variability. Specifically, it quantifies the maximum fluctuation of the j-th aggregated constraint across agents due to differences in local feasible decisions, and is defined component-wise as

$$[\tilde{\rho}]_j = N \max_{i \in \{1,\ldots,m\}} \left\{ \max_{\mathbf{x}_i \in X_i} [\mathbf{A}_i]_j \mathbf{x}_i \;-\; \min_{\mathbf{x}_i \in X_i} [\mathbf{A}_i]_j \mathbf{x}_i \right\}, \tag{4.7}$$

for all jth row of the matrix $\mathbf{A}_i$. The adopted approach is detailed in Algorithm 1, which describes the iterative coordination mechanism used to reach a feasible and near-optimal solution in a distributed fashion, as shown in [27].

Algorithm 1 is a modified dual subgradient method, adapted for decentralized optimization of mixed-integer linear problems. This work is primarily based on the algorithmic framework presented in [65], which also served as the foundation for our understanding of dual subgradient algorithms.

Lines 1−4 of the algorithm serve as initialization steps. Specifically, we initialize the dual variable $\boldsymbol{\lambda} \in \mathbb{R}^{2NW}$, the iteration counter k, and two auxiliary variables, $\underline{\mathbf{s}} \in \mathbb{R}^{2NW}$ and $\overline{\mathbf{s}} \in \mathbb{R}^{2NW}$, which are used to track the local lower and upper bounds for constraint violation across iterations.

Line 7 corresponds to the local update phase, which is solved independently by each agent.

Algorithm 1 Decentralized MILP [27]

$\boldsymbol{\lambda}(0) = 0$
$\overline{\mathbf{s}}_i(0) = -\infty, \quad i = 1, \ldots, N$
$\underline{\mathbf{s}}_i(0) = +\infty, \quad i = 1, \ldots, N$
$k = 0$
repeat
  for $i = 1$ to N **do**
    $\mathbf{x}_i(k+1) \leftarrow \arg \min_{\mathbf{x}_i \in \text{vert}(\mathcal{X}_i)} (\mathbf{c}_i^\top + \boldsymbol{\lambda}(k)^\top \mathbf{A}_i)\mathbf{x}_i$
  end for
  $\overline{\mathbf{s}}_i(k+1) = \max\{\overline{\mathbf{s}}_i(k), \mathbf{A}_i\mathbf{x}_i(k+1)\}, \quad i = 1, \ldots, N$
  $\underline{\mathbf{s}}_i(k+1) = \min\{\underline{\mathbf{s}}_i(k), \mathbf{A}_i\mathbf{x}_i(k+1)\}, \quad i = 1, \ldots, N$
  $\boldsymbol{\rho}_i(k+1) = \overline{\mathbf{s}}_i(k+1) - \underline{\mathbf{s}}_i(k+1), \quad i = 1, \ldots, N$
  $\boldsymbol{\rho}(k+1) = N \max\{\boldsymbol{\rho}_1(k+1), \ldots, \boldsymbol{\rho}_N(k+1)\}$
  $\boldsymbol{\lambda}(k+1) = \left[\boldsymbol{\lambda}(k) + \alpha(k)\left(\sum_{i=1}^{N} \mathbf{A}_i\mathbf{x}_i(k+1) - \mathbf{b} + \boldsymbol{\rho}(k+1)\right)\right]_+$
  $k \leftarrow k+1$
until some stopping criterion is met.

This step represents the inner minimization problem of the dual problem of (4.1) given by:

$$\begin{aligned} \sup_{\boldsymbol{\lambda}} \quad & -\boldsymbol{\lambda}^\top \mathbf{b} + \sum_{i \in N} \min_{\mathbf{x}_i \in \mathcal{X}_i} (\mathbf{c}_i^\top \mathbf{x}_i + \boldsymbol{\lambda}^\top \mathbf{A}_i \mathbf{x}_i) \\ \text{s.t} \quad & \boldsymbol{\lambda} \geq 0. \end{aligned} \tag{4.8}$$

If the duality gap can be reduced to zero, the optimal solutions of the primal and dual problems coincide [19, 61]. In the algorithm, solving only the inner problem,

$$\min_{\mathbf{x}_i \in \mathcal{X}_i} (\mathbf{c}_i^\top \mathbf{x}_i + \boldsymbol{\lambda}^\top \mathbf{A}_i \mathbf{x}_i), \tag{4.9}$$

allows us to decompose the overall objective into subproblems that can be handled independently by each agent. This structure is key to enabling decentralized optimization. Moreover, the authors of [65] have shown that the optimal solutions of (4.8) differ from those of the primal problem (4.1) in only a small, finite number of components—significantly fewer than the total number of variables in the problem. The work of [27] further strengthens this result by demonstrating that the bound on the difference between the primal and dual optimal solutions is even tighter than the one provided in [65]. It is important to note that in the algorithm, the inner minimization is performed over the set of *vertices* of the feasible polyhedron, i.e., $\text{vert}(\mathcal{X}_i)$. This is justified by the fact that the objective function is linear, and for linear programs over convex polytopes, the optimum is always attained at an extreme point.

Thus, restricting the feasible set to the vertices of $\mathcal{X}_i$ does not exclude any potential optimal solutions.

Lines 9−12 of the algorithm are executed by the central coordinator after collecting the local solutions from each agent. This implies that all agents must first communicate their respective optimal solutions $\mathbf{x}_i(k+1)$ before the central update can proceed. The goal of these steps is to monitor the impact of each agent's solution on the global constraint and to adaptively tighten that constraint, if necessary. Specifically, Line 9 and 10 update the historical maximum and minimum contributions to the coupling constraint, denoted as $\bar{\mathbf{s}}_i$ and $\underline{\mathbf{s}}_i$, respectively. These correspond to the maximum and minimum values observed for $\mathbf{A}_i\mathbf{x}_i(k+1)$. Line 11 then computes the individual spread $\boldsymbol{\rho}_i$, which quantifies the variability of each agent's contribution to the global constraint. In Line 12, the global tightening margin $\boldsymbol{\rho}$ is determined as the product of the number of agents N and the largest of the $\boldsymbol{\rho}_i$ values. This global tightening margin is crucial in a decentralized setting, because the aggregation of locally optimal solutions may not respect the original global constraint. The tightening margin introduces a controlled level of conservativeness to compensate for this potential infeasibility. This adjustment is applied indirectly through the update of the dual variable $\boldsymbol{\lambda}(k+1)$ in Line 13, which incorporates the tightening term $\boldsymbol{\rho}(k+1)$. As a result, the algorithm gradually enforces a more conservative version of the original constraint, promoting convergence toward a globally feasible solution despite the decentralized nature of the optimization.

Line 13 corresponds to an update of the dual variable as a function of both the global tightening margin $\boldsymbol{\rho}$ and the step size $\alpha(k)$. The choice of step size is essential to the convergence properties of the algorithm. It must satisfy the standard diminishing step size conditions [27]:

$$\lim_{k\to+\infty} \alpha(k) = 0 \quad \text{and} \quad \sum_{k=0}^{+\infty} \alpha(k) = +\infty. \tag{4.10}$$

These conditions ensure that the algorithm takes progressively smaller updates, allowing for finer adjustments as it converges, while still making sufficient progress to reach a solution. This line also applies a projection onto the non-negative orthant $\mathbb{R}^m_+$ to make sure that the dual variables remains positive.

The dual variables can be interpreted as penalties proportional to the constraint violations [81], as illustrated in Line 7. The magnitude of a dual variable reflects the extent to which the corresponding constraint is violated. In applications where the optimization space is smooth, the updates of the dual variables follow the gradient of the dual function, thereby guiding the algorithm toward the optimal dual solution [82]. In non-smooth applications, such as Algorithm 1, the update direction does not necessarily align with the true optimal

direction. Nevertheless, over successive iterations, the updates still provide a mechanism for approaching optimality [83]. In both gradient and subgradient methods, dual variables can thus be understood as dynamic penalty coefficients that balance the minimization of the objective function with the enforcement of constraints. In decentralized optimization frameworks such as the one considered in this work, dual variables serve as coordination signals: they enable implicit communication between agents by conveying information about constraint violations, without requiring direct exchange of primal decision variables [61].

Finally, the algorithm terminates when a predefined stopping criterion is satisfied. In the context of feasibility-focused applications, [27] suggests stopping the iterations once two or three consecutive feasible solutions have been identified. This strategy is effective when the primary objective is to guarantee constraint satisfaction. However, in applications where optimal performance is prioritized—such as minimizing energy costs or maximizing system efficiency—more refined stopping conditions can be employed.

In summary, the structure of the algorithm can be interpreted as follows:

- **Lines 1–4:** Initialization;
- **Line 7:** Solution of the local sub-problem for each zone i;
- **Lines 9–13:** Update of the dual variables in accordance with the constraint violations, scaled by the step size;
- **Line 15:** Evaluation of the stopping criterion to determine convergence.

4.2.1 Performance-oriented algorithm

A performance-oriented version of Algorithm is the proposed in [27]. This approach involves a trade-off, as it requires sharing additional information with a central coordinator, thereby requiring additional computational resources. The main changes consists of retaining the best-performing solution in memory throughout the iterative process. Specifically, the algorithm stores the best solution obtained so far and continues to perform as many iterations as possible within a predefined time constraint. Although this strategy increases the computational effort, it is well-suited for scenarios like ours, where strict minimization of computation time below five minutes is critical.

As illustrated in Algorithm 2, Lines 5−8 are dedicated to saving the current solution, conditional on the evaluation in Line 11, which verifies whether it outperforms the previously stored solution. The variable $J(k)$ is employed to record the best objective value obtained

throughout the iterations. The stopping criterion is defined based on elapsed computation time rather than on iteration count or convergence thresholds. In this work, we use Algorithm 2 to solve MPC problem (4.2) in an efficient decentralized fashion.

Algorithm 2 Performance-oriented version

$\boldsymbol{\lambda}(0) = 0$
$\underline{\mathbf{s}}_i(0) = -\infty, \quad \overline{\mathbf{s}}_i(0) = +\infty, \quad i = 1, \ldots, N$
$J(0) = +\infty, \quad \delta(0) = 0, \quad k = 0$
repeat
  for $i = 1$ to N **do**
    if $\delta(k) = 1$ **then**
      $\hat{\mathbf{x}}_i(k) \leftarrow \mathbf{x}_i(k)$
    end if
    $\mathbf{x}_i(k+1) \leftarrow \arg \min_{\mathbf{x}_i \in \mathrm{vert}(X_i)} (\mathbf{c}_i^\top + \boldsymbol{\lambda}(k)^\top \mathbf{A}_i)\mathbf{x}_i$
  end for
  if $\sum_{i=1}^{N} \mathbf{A}_i \mathbf{x}_i(k+1) \leq \mathbf{b}$ **and** $\sum_{i=1}^{N} \mathbf{c}_i^\top \mathbf{x}_i(k+1) < J(k)$ **then**
    $J(k+1) \leftarrow \sum_{i=1}^{N} \mathbf{c}_i^\top \mathbf{x}_i(k+1)$
    $\delta(k+1) \leftarrow 1$
    $\hat{\mathbf{x}}_i(k+1) \leftarrow \mathbf{x}_i(k+1), \quad i = 1, \ldots, N$
  else
    $\delta(k+1) \leftarrow 0$
  end if
  $\underline{\mathbf{s}}_i(k+1) = \max\{\underline{\mathbf{s}}_i(k), \mathbf{A}_i \mathbf{x}_i(k+1)\}, \quad i = 1, \ldots, N$
  $\overline{\mathbf{s}}_i(k+1) = \min\{\overline{\mathbf{s}}_i(k), \mathbf{A}_i \mathbf{x}_i(k+1)\}, \quad i = 1, \ldots, N$
  $\boldsymbol{\rho}_i(k+1) = \overline{\mathbf{s}}_i(k+1) - \underline{\mathbf{s}}_i(k+1), \quad i = 1, \ldots, N$
  $\boldsymbol{\rho}(k+1) = p \max\{\boldsymbol{\rho}_1(k+1), \ldots, \boldsymbol{\rho}_N(k+1)\}$
  $\boldsymbol{\lambda}(k+1) = \left[\boldsymbol{\lambda}(k) + \alpha(k) \left(\sum_{i=1}^{N} \mathbf{A}_i \mathbf{x}_i(k+1) - \mathbf{b} + \boldsymbol{\rho}(k+1) \right)\right]_+$
  $k \leftarrow k+1$
until time is over

4.3 Back-up controller design

As the control advances through the prediction horizon, accumulated errors in state estimation and forecast data may degrade the model's predictive accuracy. This can eventually render the optimization problem infeasible, preventing the computation of a valid control input that satisfies all constraints. To ensure robustness in such situations, a fallback or back-up controller must be implemented to maintain operational continuity, as the optimization-based

controller may be unable to generate a feasible solution.

In the event that the optimization-based controller fails to return a feasible solution, a simple rule-based back-up controller is employed to ensure continued operations. This controller follows a conservative logic: if the temperature in a given zone falls below the minimum acceptable comfort threshold, the heating control is immediately set to level 2:

$$\mathbf{u}_i(t) = \begin{cases} \begin{bmatrix} 0 & 0 & 1 \end{bmatrix}^\top & \text{if } T_i(t) < \underline{T}_i(t) + 0.5, \\ \begin{bmatrix} 0 & 0 & 0 \end{bmatrix}^\top & \text{otherwise.} \end{cases} \tag{4.11}$$

This choice reflects the assumption that infeasibility is likely to occur under conditions of high thermal stress, such as a sudden drop in temperature, and that a strong corrective action is required. An analogous rule is applied in the case of excessive temperatures, where the cooling control would be set to its maximum level. This fallback strategy guarantees that thermal comfort constraints are respected, even when the primary controller is unable to produce a valid control sequence.

CHAPTER 5 NUMERICAL CASE STUDY

We now illustrate the performance of our approach in a numerical case study employing a two-zone and a five-zone commercial building.

5.1 Setting

We next detail the global simulation environment, the datasets used for training and evaluation, the algorithm's parameters, and the optimization problem settings—including temperature constraints and the cost function definition. These elements collectively define the experimental framework and pave the way for the analyses of the results, which are presented in the subsequent subsection.

5.1.1 Simulation environment

As previously mentioned, this work leverages a building simulator developed by BBAI, which is based on open-source Modelica libraries [84]. The simulator operates on a five-minute time step and implements discrete control actions at each interval. Specifically, it receives zone-specific control inputs $\mathbf{u}_i(t)$ and, in return, simulates and outputs the temperature of each zone at the subsequent time step. The internal dynamics are influenced by the building model and environmental conditions, such as outdoor temperature, humidity, and solar irradiation.

The environmental conditions are specified using EnergyPlus Weather (EPW) files [58], also provided by BBAI. In this study, we use weather files corresponding to two locations in Québec: Montréal and Mirabel.

The physical and thermal characteristics of the building are described using functional mock-up units (FMUs), also provided by BBAI. FMUs are standardized, tool-independent packages used to encapsulate dynamic models for use across different simulation platforms [85]. Defined by the functional mock-up interface (FMI) standard, FMUs enable both model exchange and co-simulation. A typical FMU contains a structured `XML` description of the model, compiled binaries for various platforms, and optional resource files such as parameters and input data. This format facilitates seamless integration of complex physical models into external control or optimization workflows.

In this work, we employ two FMU files corresponding to buildings with two and five thermal zones, respectively. These models serve as the simulation environment for evaluating the performance of our proposed control strategies and are based on the Modelica Building

Library [86]. Control agents developed in this study interact with the simulator by sending control actions computed via decentralized MILP in response to the current state of the environment, thereby closing the loop between simulation and control.

5.1.2 Data sets and function learning

To develop the physics-informed linear model used in both the two-zone and five-zone configurations, we employed standard linear regression techniques using the `scikit-learn` library [87], which minimizes the squared L2-norm of the residuals, i.e., the residual sum of squares. A bias term is included in the regression to enhance model flexibility and account for unmodelled dynamics. The training dataset is composed of one year of simulation data generated using a bang-bang controller, as well as an additional year of data incorporating randomly generated control signals applied twice per hour. The inclusion of random control inputs served to diversify the state space by inducing variability in zone temperatures, thereby simulating more realistic environmental conditions and enriching the dataset for model learning. For the purpose of this study, only data corresponding to the winter period—defined as January to March and October to November—was utilized. This seasonal selection is consistent with our focus on heating performance under cold weather conditions. All simulations were based on historical environmental data from the year 2021. Numerical results in relation with model training can be seen in Section 5.2.1. In a real-world application, the building operator would typically have access to historical data encompassing environmental conditions, indoor temperature measurements, and the control actions executed by each HVAC system. This data could be used to identify and model the physical dynamics of the actual building, removing the need for simulated environments and historical weather datasets.

The data used for model training and parameter tuning was extracted from the weather file `CAN_PQ_Montreal.Intl.AP.716270_CWEC.epw` [58], which represents typical meteorological conditions for the Montréal–Trudeau International Airport area. To ensure generalization and avoid overfitting, final testing was performed using a distinct weather file, `CAN_PQ_Montreal.Mirabel.716278_CWEC.epw` [58], corresponding to the Mirabel region (approximately 40 km north-west of Montreal) and serving as out-of-sample testing data. Although both files cover similar climate profiles, they are not identical and introduce variability in the temperature and weather conditions. In both training and testing phases, the year 2021 is used to maintain temporal consistency while validating model performance on unseen data. To simulate forecast uncertainty during testing, the input data is perturbed with a synthetic noise. Specifically, a Gaussian noise with a standard deviation ratio of 0.05 is added to the solar irradiation forecasts, while the temperature forecasts are perturbed using normally

distributed noise with a standard deviation of 0.5°C. This procedure is applied consistently to both the two-zone and five-zone building scenarios.

5.1.3 Controllers

Both the centralized and decentralized approaches are employed in this study to enable a comparative analysis of their performance. The centralized formulation serves as a benchmark, as it has access to global information and can jointly optimize all zones, which typically results in improved constraint satisfaction and lower objective values. However, this centralized method is inherently less scalable due to its computational complexity and the need for global coordination. In contrast, the decentralized approach trades some performance for increased scalability and practical deployment feasibility in larger systems. A bang-bang controller developed by BBAI is also used as a benchmark to represent traditional rule-based control strategies. This controller operates as a simple ON/OFF mechanism, with control actions triggered based on a predefined temperature dead band.

5.1.3.1 Centralized approach

To solve the centralized optimization problem described in (4.2), a controller was implemented using the `CVXPY` parser [88] with the `GUROBI` solver [89]. This controller receives, at each time step, the current temperatures of all zones from the simulator, along with forecasted values of exogenous parameter like the solar irradiation and outdoor temperature. To better reflect real-world forecasting uncertainty and avoid relying on perfect knowledge, Gaussian noise is added to the simulator-provided inputs. This introduces stochastic variability consistent with the nature of short-term weather forecasting in the lookahead component of the MPC. Figure 5.1 illustrates the closed-loop interaction between the controller and the building simulator. In this framework, $\mathbf{T}(t)$ denotes the vector of indoor temperatures across all zones at time step t, while $\boldsymbol{\chi}(t)$ represents the exogenous input data, including forecasts of outdoor temperature and solar irradiance over the prediction horizon. The vector $\mathbf{u}(t)$ corresponds to the control actions computed by the controller for all zones at the current time step. Although the controller optimizes the control sequence over the entire prediction horizon, only the control input for the first time step is applied to the simulator, following a receding horizon strategy. This control loop is executed iteratively at each time step throughout the simulation horizon. We remark that the decentralized controller also operates through a multi-step procedure, as detailed in Algorithm 2.

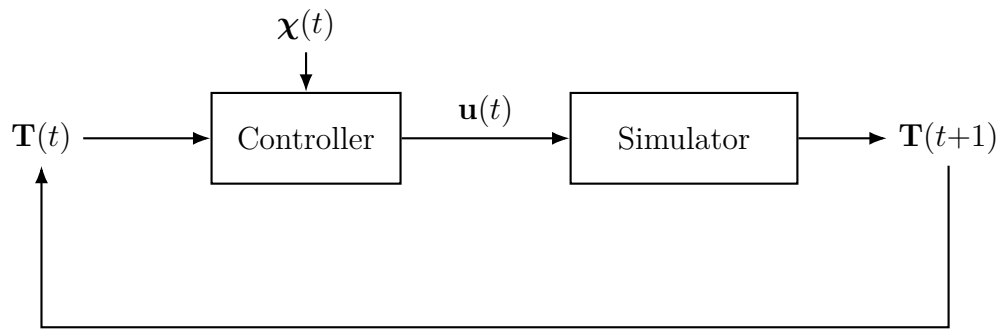

Figure 5.1 Closed-loop control architecture.

The controller computes an optimal sequence of control actions over a prediction horizon of 1.5 hours, corresponding to 18 future time steps (at 5-minute intervals), i.e, $W = 18$. Then, only the first control action of the sequence is applied to the simulation environment, in accordance with the receding horizon principle of MPC. The simulation advances in 5-minute increments, resulting in 288 control decisions per day. While longer horizons could theoretically improve performance, a 1.5-hour window is chosen to maintain tractable computation times given the complexity of the optimization problem. In cases where the optimization problem becomes infeasible, the back-up controller discussed in Section 4.3 is triggered to ensure uninterrupted simulation execution.

Finally, in the case of the centralized controller, we introduce an additional bias to the existing thermal model which is beneficial, as discussed in Section 3.4.2. To determine an appropriate value for this additional bias, a grid search is performed using the first three days of January under Montréal weather data. The bias value that yields the best overall objective performance is then selected for subsequent deployments.

5.1.3.2 Decentralized approach

To address the problem in a decentralized manner, we adopt the approach outlined in Algorithm 2, using a receding horizon of $W = 18$ steps, consistent with Section 5.1.3.1. As with centralized approach, if no feasible solution is found at the end of the algorithm, the back-up controller insures uninterrupted simulation execution.

In order to satisfy the constraints imposed on $\alpha(k)$, as defined in (4.10), we employ the following decaying formulation:

$$\alpha(k) = \theta_\alpha \frac{1}{1+k}, \tag{5.1}$$

where $\theta_\alpha > 0$ is a tunable parameter selected through a grid search. The grid search was con-

ducted using the first three days of January, under Montréal-based weather data. The value of θ_α that yields the best objective performance is retained for all subsequent simulations.

Another key parameter in Algorithm 2 is the number of optimization steps—or equivalently, the maximum computation time—allowed per decision round t. Several values are tested, and it is found that allowing 100 optimization steps provided a favourable trade-off between solution quality and computational feasibility. Specifically, this configuration enabled the algorithm to remain within the five-minute control interval required for real-time deployment in both the two-zone and five-zone scenarios. While increasing the number of steps, e.g., to 500, can enhance performance, empirical results showed only marginal improvements beyond the 100-step threshold.

5.1.4 Temperature constraints

The temperature bounds used in this work are inspired by the study presented in [90], which models thermal control in a primary school located in Montreal using time-dependent setpoints. In that study, the heating setpoint is maintained at 21°C during operational hours (6:00 AM to 9:00 PM) and lowered to 16°C during off-hours. Similarly, the cooling setpoint is set to 24°C during the day and increased to 27°C at night. Comparable setpoints are also commonly used throughout the literature, as shown in [91–93]. However, drawing from these references, and to ensure sufficient flexibility in our control scheme, we adopt a conservative lower setpoint bound of 20°C for winter conditions in our simulation framework. This choice allows us to better showcase the responsiveness and effectiveness of the proposed control strategy, as illustrated in Figure 5.2.

5.1.5 Cost function definition

To evaluate the proposed control strategies under various realistic operating conditions, the electricity price $c_{\mathrm{P}}(t)$ is defined according to several representative pricing schemes. In particular, the performance of the controllers is assessed under three commonly used DR programs: TOU, CPP, and RTP. Each pricing structure introduces distinct incentives and constraints on energy usage, thereby enabling a thorough evaluation of the controllers' adaptability and effectiveness across different economic scenarios.

To compute the entries of the cost vector $\mathbf{c}_i$ associated with each control action $\mathbf{u}_i(t) \in \{0,1\}^3$, both the instantaneous electricity price $c_P(t)$ and the corresponding power consumption P_s of each actuator are required. The specific power ratings used in this study are as follows:

- Fan: 637.76 W;

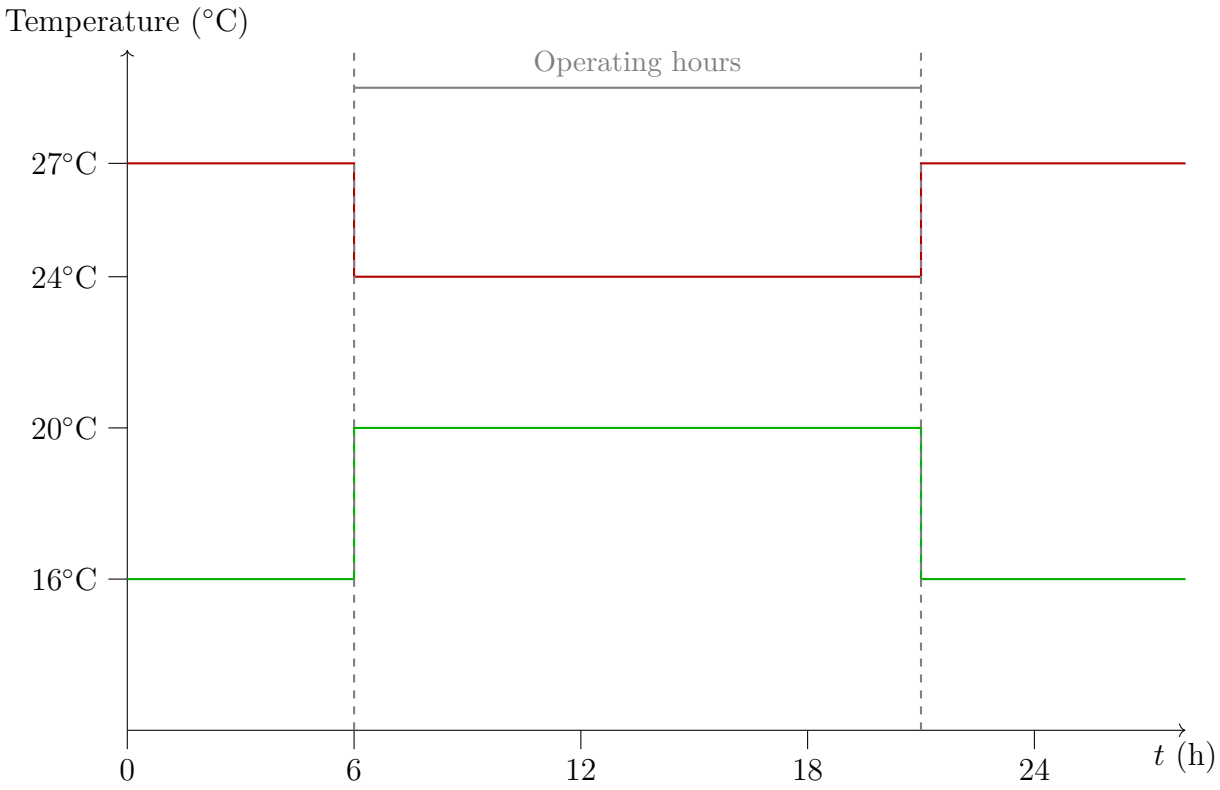

Figure 5.2 Daily heating and cooling temperature limits

- Heating level 1 ($heat_1$): 2910.48 W;
- Heating level 2 ($heat_2$): 5183.21 W;

for all RTU-HVAC and zone. Moreover, the simulation time step Δ is set to 5 minutes, corresponding to 0.083 hours. This value is used to convert instantaneous power consumption into energy cost, ensuring consistency across all cost calculations.

5.1.5.1 Hydro Québec's TOU

This time-based incentive program considers three distinct pricing periods: off-peak (night), mid-peak, and peak, as defined by Hydro-Québec [18] and illustrated in Figure 5.3. Off-peak periods offer the lowest electricity prices, while mid-peak pricing is set below that of peak periods. This pricing structure is designed to encourage customers to shift their electricity usage away from periods of high demand—typically from 6 a.m. to 10 a.m. and from 12 p.m. to 8 p.m.—toward times of lower overall system usage. The primary objective is to alleviate stress on the electric grid by reducing the overall peak power demand.

5.1.5.2 Hydro Québec's CPP

The CPP scheme, previously introduced in Section 1.2.3 as Rate Flex M [18], defines specific periods during which electricity prices are significantly increased to reflect peak demand conditions. According to Hydro-Québec, peak periods can occur from 6 a.m. to 9 a.m. and from 4 p.m. to 8 p.m., with a maximum of two events per day. Each event lasts three hours in the morning and four hours in the evening and must be separated by at least seven hours. The total number of peak hours is limited to 100 hours during the winter season.

In this study, we model a total of 10 CPP events for the month of January at random moments. Five of theses events are set to in the morning and five in the afternoon. Furthermore, at least one day must contain two events.

Rate Flex M consists of two components: a fixed demand charge of \$17.573 per kilowatt (\$/kW) based on the maximum monthly power demand, and a variable energy charge that depends on whether consumption occurs during a peak event. During the winter period, energy consumed outside peak events is charged at 3.820¢/kWh, whereas energy consumed during a peak event is charged at 60.262¢/kWh. In this work, we focus exclusively on the variable energy component to assess the controller's response to price signals. An example of a day featuring two CPP events is shown in Figure 5.4.

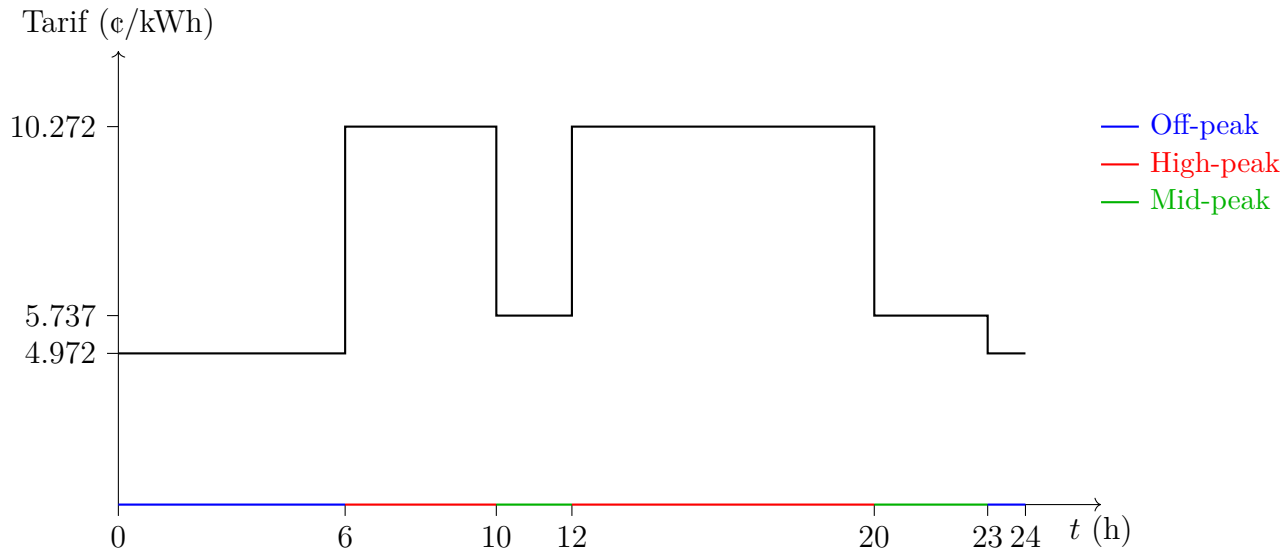

Figure 5.3 Hourly electricity rate as a function of TOU periods (off-peak, high-peak, mid-peak) in winter for Hydro Québec

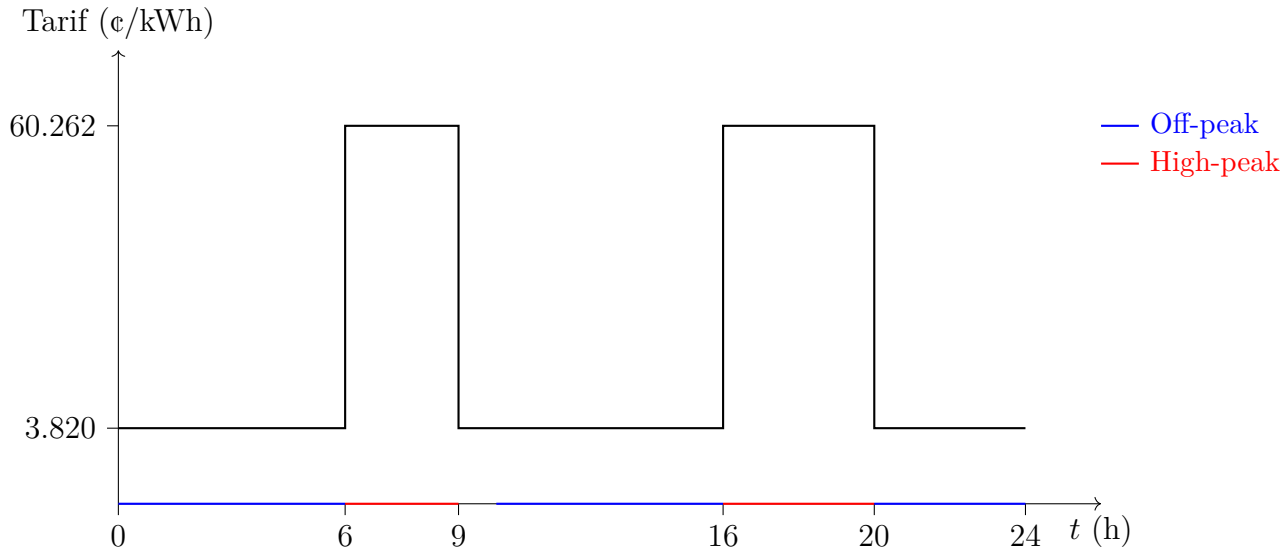

Figure 5.4 Hourly electricity rate on a two-event day considering CPP in winter for Hydro Québec

5.1.5.3 Real time pricing

RTP is a dynamic electricity pricing mechanism in which the cost of electricity varies on an hourly basis according to the fluctuations in the wholesale electricity market [94]. This scheme can be likened to a stock market, where prices respond in real time to supply and demand conditions. In general, higher overall consumption leads to higher electricity prices, while lower consumption results in lower prices. To allow users to respond to these signals, electricity providers typically communicate the price schedule to consumers one day in advance, enabling them to plan and adjust their usage accordingly. In this work, the RTP signal is modelled as inversely proportional to the consumption profile of a typical winter day. This approach reflects the principle that prices tend to rise when demand is high and fall when demand is low, and is used here to simulate economically motivated behaviour under real-time pricing.

The load profile of a typical winter day is previously illustrated in Figure 1.1. To construct an approximated RTP price signal, we rely on historical load profiles. We chose to source our data locally. Specifically, we used the publicly available historical electricity demand data published by Hydro-Québec [95], which provides hourly electricity consumption records from 2019 to 2024. Given that the RTP price signal is defined on an hourly basis, the temporal resolution of this dataset is well suited to our modelling needs.

We selected the data from January 2023 to reflect the characteristics of winter energy demand in Québec, which is representative of the climate and seasonal conditions relevant to this study. The resulting demand curve is shown in Figure **??** considering only the first day of the month. The data, expressed in megawatts (MW), corresponds to the total hourly demand across the Hydro-Québec grid.

To approximate the principle behind RTP—namely, that electricity prices should increase with higher demand and decrease with lower demand—we transformed the electricity demand curve by applying a power of three to each data point and dividing the result by a factor of 3×10^{12}. This transformation serves two purposes: first, to bring the resulting values into a range more comparable with typical ¢/kWh pricing; and second, to amplify the variability in demand, thereby emphasizing changes in consumption behaviour. As illustrated in Figure **??**, the RTP price curve retains the general shape of the original demand profile, while exhibiting sharper ramps. This reflects a stronger pricing signal during periods of high demand, reinforcing the incentive to reduce consumption when system load is elevated. We remark that this only serve as a stylized example of RTP for our numerical case study.

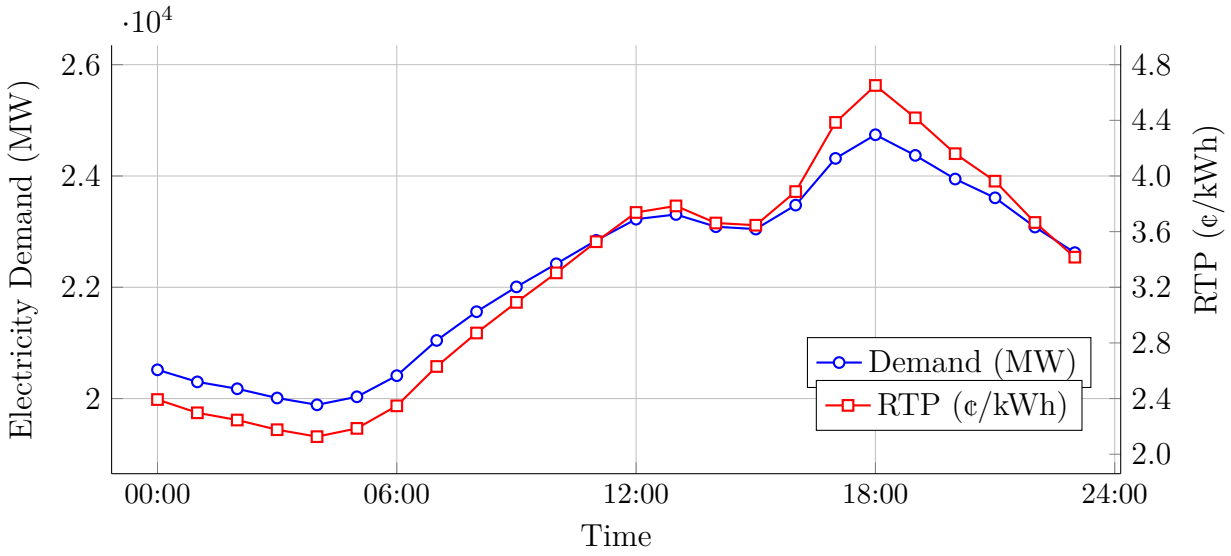

5.2 Results and analysis

This section presents the results of the linear regression used to identify temperature prediction models for both the two-zone and five-zone buildings. It also includes the outcomes of a one-month simulation performed under each of the previously described DR strategies. We compare the performance of three control approaches—namely, a rule-based controller, a centralized controller, and a decentralized controller—based on two key metrics: energy cost and thermal comfort. Because of its limited scalability, the centralized approach is only utilized for the two-zone building.

5.2.1 Linear regression training results

As described previously, the dataset used for model training comprises one year of bang-bang control data and one year of randomized bang-bang control data, with both datasets limited to the winter months. To assess model performance, we perform a 5-fold cross-validation (k-split $= 5$), allowing us to evaluate the model's generalizability. Model accuracy is primarily assessed using the root mean square error (RMSE), as it most directly reflects the impact of prediction errors on temperature control.

5.2.1.1 Two-zone linear regression

Table 5.1 reports the RMSE and associated standard deviation for each zone in the two-zone environment. In both zones, the average prediction error remains below 0.5,°C per time step. According to ASHRAE (American Society of Heating, Refrigerating and Air-Conditioning Engineers) standards [96], temperature sensors used in HVAC systems typically have an

accuracy of around ±0.5,°C, which sets a practical lower bound on achievable model precision. Thus, maintaining prediction errors within this range is generally considered sufficiently precise for control applications. However, it should be noted that consecutive occurrences of near-maximum errors could still result in significant deviations from the actual temperature trajectory.

We evaluate the predictive performance of the identified model using a separate control dataset, consisting of one year of bang-bang generated data at the Mirabel location. To ensure consistency between simulation and ground truth, the model is initialized with the same temperature values as the dataset and subjected to the same control inputs observed in the data. The predicted temperature trajectories are then compared to the actual values, as illustrated in Figure 5.5. The results show that the prediction error varies significantly over time. In some instances, the model exhibits a high degree of accuracy, with errors below 1°C. However, at other times, the error can reach up to 4°C, indicating that while the model captures general dynamics effectively, it may struggle under certain operating conditions or state transitions.

Figure 5.6 illustrates the model's performance over an extended prediction horizon spanning three consecutive days. The results suggest that the prediction error does not necessarily accumulate over time, indicating a level of stability in the model's dynamics. Moreover, a cyclical pattern is observed in the error behaviour, with improved accuracy during periods of elevated indoor temperature. This trend is likely attributable to increased solar gains during daylight hours, which appear to enhance the model's predictive reliability. We may also suggest that samples with high solar irradiation are overrepresented in the training dataset. Incorporating more data from low-irradiation periods could help produce a more balanced and generalizable model. It is also important to acknowledge that the prediction model is constrained to a linear form, whereas temperature dynamics are inherently non-linear [34, 74, 97]. This limitation implies that some degree of approximation error is inevitable, and improved accuracy in one operating region may come at the cost of reduced precision in another.

Table 5.1 RMSE for linear regression learning of two-zone model

Zone	**RMSE (°C)**	**Std RMSE (°C)**	**Error range (°C)**
1	0.486	0.004	[0.482, 0.490]
2	0.431	0.003	[0.428, 0.434]

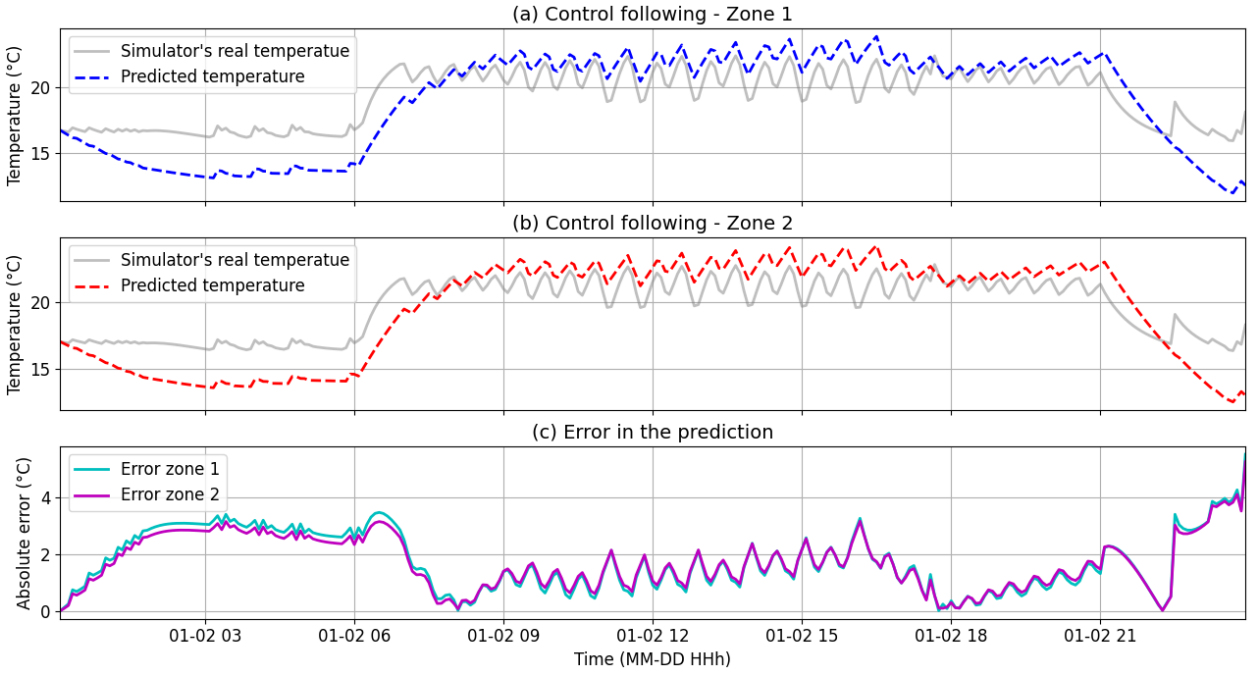

Figure 5.5 Comparison between predicted and actual indoor temperature on January 2nd 2021 in the two-zone building.

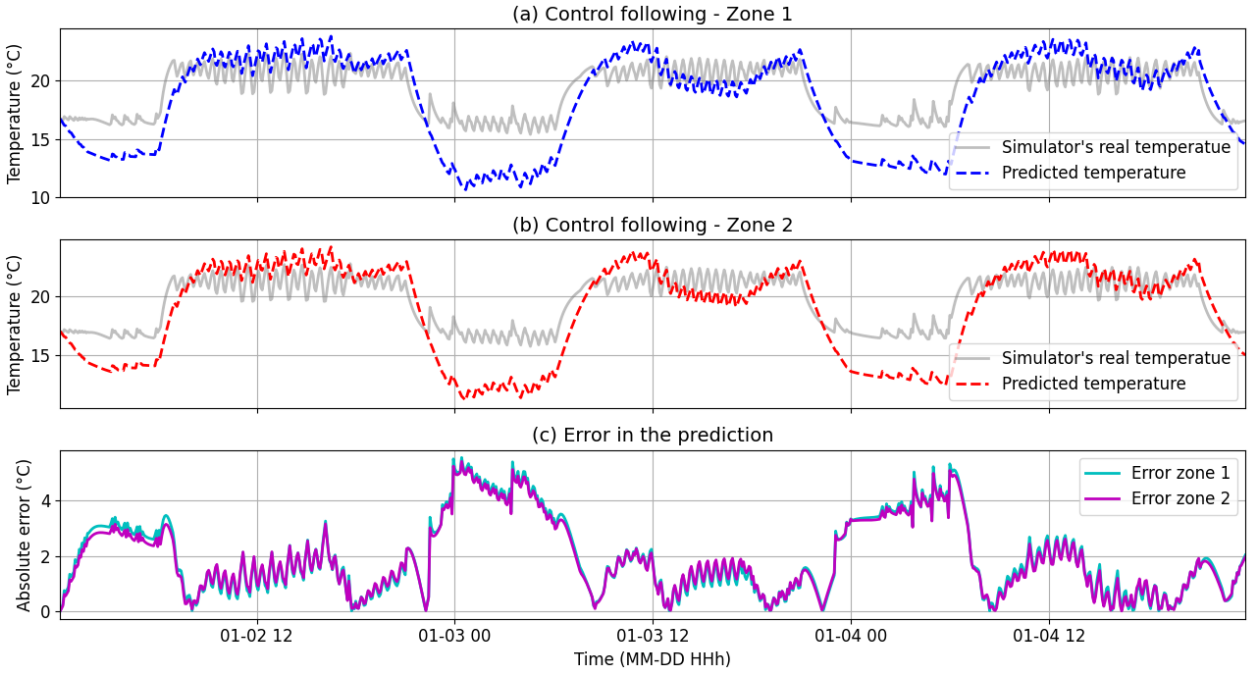

Figure 5.6 Comparison between predicted and actual indoor temperature from January 2nd to 4th 2021 in the two-zone building.

5.2.1.2 Five-zone linear regression

We now evaluate the performance of the five-zone model using the same methodology applied to the two-zone case. Table 5.2 presents the model's performance metrics, disaggregated by zone, in terms of the RMSE. Figure 5.7 provides a visual comparison of the predicted and actual temperatures across all five zones. However, due to the density of information in this figure, it may not be optimal for detailed analysis. For clearer insights, the reader is referred to Appendix B, which contains individual, full scale plots of temperature predictions and corresponding errors for each zone.

Similar results are observed for the five-zone model as in the two-zone case. The RMSE remains below 0.5 in all zones, indicating satisfactory prediction accuracy. A correlation appears to exist between prediction quality and the magnitude of solar impact—instance with higher solar exposure tend to exhibit lesser prediction errors. Nevertheless, the errors do not appear to accumulate or grow over time, suggesting stable model behaviour throughout the prediction horizon.

5.2.2 Simulation results

In this section, we present a performance comparison between the proposed decentralized controller, and compare it to the benchmark approaches when possible. The analysis is conducted through both tabular and graphical representations. The tables report the total energy cost and the comfort rating over the evaluation period. Energy cost is expressed in dollars and directly corresponds to the objective function of the optimization problem. The comfort rating, measured in °C·h, quantifies the cumulative constraint violation time across all building zones and serves as an indicator of thermal comfort performance and constraint violation.

5.2.2.1 TOU results

We present here the results of the simulation conducted under the TOU pricing scheme. Table 5.3 summarizes the total operating cost over a one-month period for each of the three controllers, applied to both two-zone and five-zone building configurations. As shown, the rule-based controller consistently yields the highest operational cost, followed by the decentralized controller. The centralized controller achieves the lowest cost, indicating superior cost-efficiency under the TOU pricing structure.

Table 5.3 also reports the comfort rating associated with each control strategy. While this metric may not be immediately intuitive, it quantifies the total cumulative temperature de-

Table 5.2 RMSE for linear regression learning of five-zone model

Zone	RMSE (°C)	Std RMSE (°C)	Error range (°C)
1	0.486	0.004	[0.482, 0.490]
2	0.431	0.003	[0.428, 0.434]
3	0.473	0.002	[0.471, 0.475]
4	0.447	0.002	[0.445, 0.449]
5	0.413	0.002	[0.411, 0.415]

Best performing zone:	Zone 5 (RMSE: 0.413 °C)
Worst performing zone:	Zone 1 (RMSE: 0.486 °C)
Performance difference:	0.073 °C

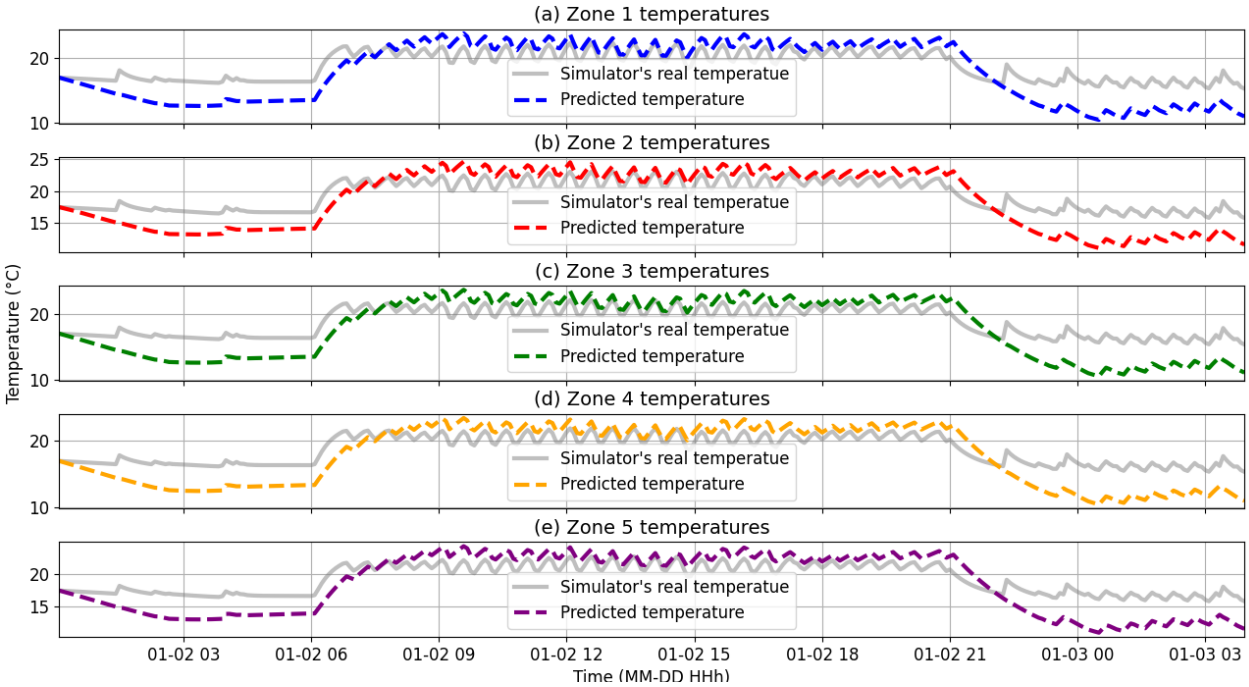

Figure 5.7 Comparison between predicted and actual temperature evolution on January 3rd 2021 in the five-zone environment.

viation from the comfort bounds across all zones and time steps. To facilitate interpretation, we added the mean step comfort rating. Consider an example: a comfort rating of 10,000°C in a two-zone building over a 30-day simulation period (with 12 steps per hour in each zone) corresponds to a deviation of 10,000°C over 17,280 five-minute intervals. This translates to an average temperature violation of approximately 0.58°C per time step.

From the results in Table 5.3, we observe that the rule-based controller performs the worst in terms of comfort, followed by the decentralized controller. The centralized controller yields the best comfort rating. This is seen consistently in subsequent results: the centralized approach tends to outperform the decentralized in term of comfort but not in price. In general, reducing energy costs typically comes at the expense of thermal comfort, and vice versa. We could say that the centralized approach enforces constraints more strictly from the outset compared to the decentralized approach. In contrast, the decentralized formulation gradually tightens its constraints through iterative updates, which can be interpreted as a relaxed version of the centralized problem. Increasing the number of iterations may allow the decentralized method to converge toward the centralized solution. This behaviour could explain why the decentralized approach is more prone to constraint violations, while still achieving lower operational costs.

The following results also reveal a consistent trend: both the centralized and decentralized controllers outperform the rule-based controller not only in terms of cost minimization but also with respect to occupant comfort. This outcome is somewhat unexpected. Typically, a rule-based controller is explicitly designed to enforce thermal comfort constraints and serve as a reference standard for maintaining acceptable indoor conditions. However, while its parameters are set to satisfy these constraints, the rule-based approach lacks adaptability and foresight in dynamic and heterogeneous environments such as the one considered in this study. Although the optimization-based controller yields satisfactory results, it does not necessarily generate a truly optimal control sequence. If we were to compare the controller's decisions with those obtained under perfect foresight, a clear performance gap would emerge. This discrepancy arises from two fundamental challenges: (i) the building environment is not easily tractable with linear approximations, and (ii) significant uncertainties exist in the forecasts of exogenous variables such as weather conditions and occupancy.

In contrast, optimization-based controllers are also highly effective at enforcing comfort constraints. These controllers dynamically adjust to changing conditions thanks to the receding horizon, enabling them to maintain thermal comfort more reliably in complex scenarios. While they may, in some cases, risk comfort violations in pursuit of lower energy costs, their ability to adapt and anticipate system behaviour often allows them to outperform rigid

rule-based strategies, particularly when the latter are poorly tuned for extreme or variable conditions. This likely explains the superior performance observed. In addition to the tabulated results, we present time-series plots to further illustrate the behaviour of the building under different control strategies. Figures 5.8 and 5.9 show the temperature trajectories for each zone in the five-zone configuration over a representative three-day period. These plots correspond to the, respectively, rule-based and decentralized controllers operating under the TOU pricing scheme.

Alongside the zone temperatures, we also include the outdoor temperature profile and the TOU pricing signal. These visualizations help to contextualize how each controller responds to external disturbances and price fluctuations, and they provide additional insight into the observed differences in cost and comfort performance.

The plots illustrate that the controllers generally attempt to maintain the indoor temperature within the specified thermal limits. Among them, the most effective strategy appears to operate near the lower bound of the acceptable temperature range, thereby minimizing power consumption. Minimizing energy cost requires maintaining the temperature as close as possible to its lower bound, which increases the risk of violating temperature constraints due to uncertainties in the controller and prediction model. Figures 5.11 and 5.12 illustrate the control signals produced by the two proposed optimization-based controllers for the two-zone building. These plots are specifically intended to highlight the influence of the back-up controller on system behaviour. The results correspond to the first day of January 2021. For clarity and to enable a more focused analysis, only the temperature and control trajectories of Zone 1 are presented. The control signals produced by the rule-based controller are shown in Figure 5.10 as a reference. In Figures 5.11 and 5.12, it can be observed that when the zone temperature drops below the defined lower limit, a control signal with a value of 2 is issued by the back-up controller. The figures also reveal the threshold behaviour that governs the activation of this controller—specifically, the conditions under which it chooses to engage heating at level 2. Notably, reaching the temperature threshold does not necessarily trigger the back-up controller. In some instances, the optimization-based controller is still able to maintain the temperature within acceptable bounds without intervention. Conversely, the back-up controller may occasionally activate even when the temperature is above the lower limit. This may be attributed to thermal inaccuracies or limited computation time. In such cases, the back-up controller operates strictly based on the predefined temperature thresholds, regardless of the optimization strategy.

Table 5.3 Controller comparison over one month under TOU pricing for both 2-zone and 5-zone configurations.

2-Zone Configuration				
Controller	**Price ($)**	**Price (%)**	**Comfort Rating (°C·h)**	**Mean Comfort Rating (°C/step)**
Rule-based	443.91	100.00	765.37	0.531
Centralized	317.74	71.6	619.87	0.430
Decentralized	305.45	68.8	758.30	0.527
5-Zone Configuration				
Controller	**Price ($)**	**Price (%)**	**Comfort Rating (°C·h)**	**Mean Comfort Rating (°C/step)**
Rule-based	1,127.02	100.00	2,351.63	0.653
Centralized	856.99	76.0	1987.04	0.552
Decentralized	826.60	73.3	2,264.48	0.629

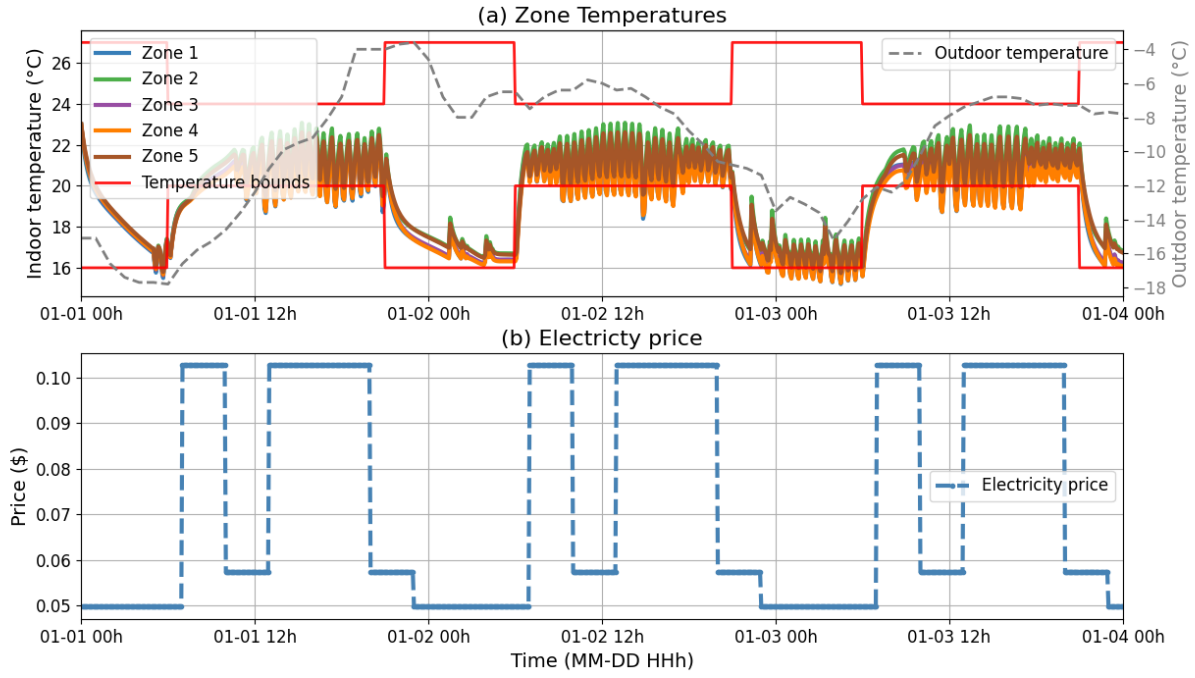

Figure 5.8 Temperature trajectories under rule-based control for a three-day period with TOU pricing.

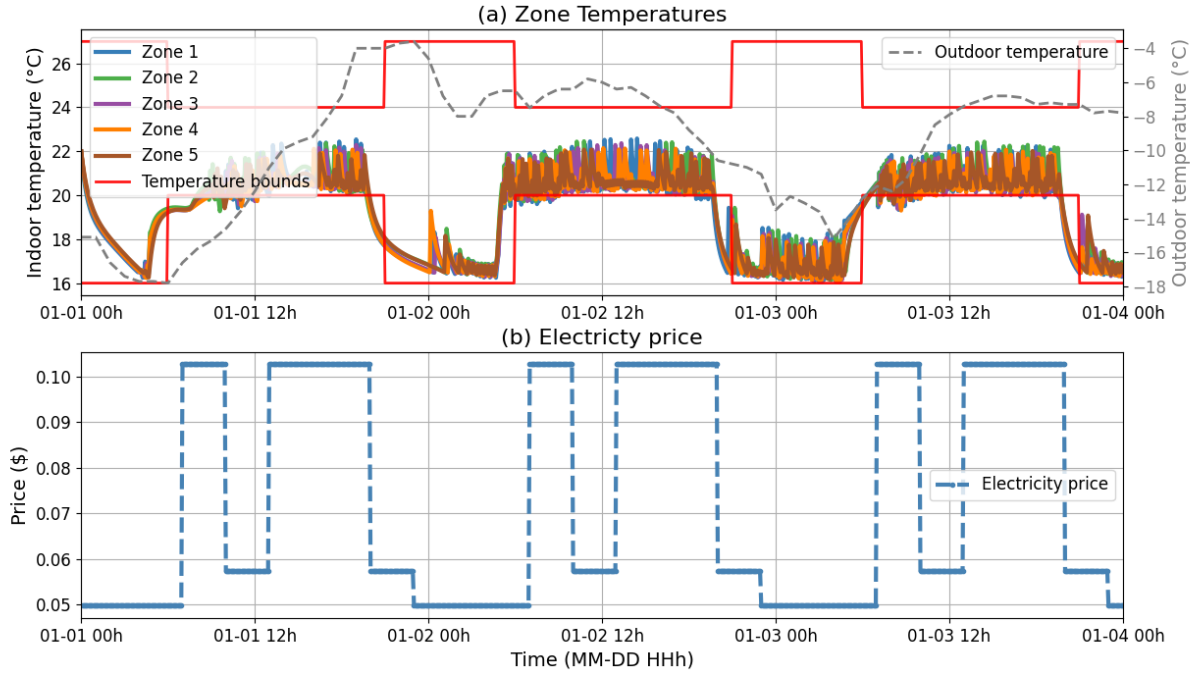

Figure 5.9 Temperature trajectories under decentralized control for a three-day period with TOU pricing.

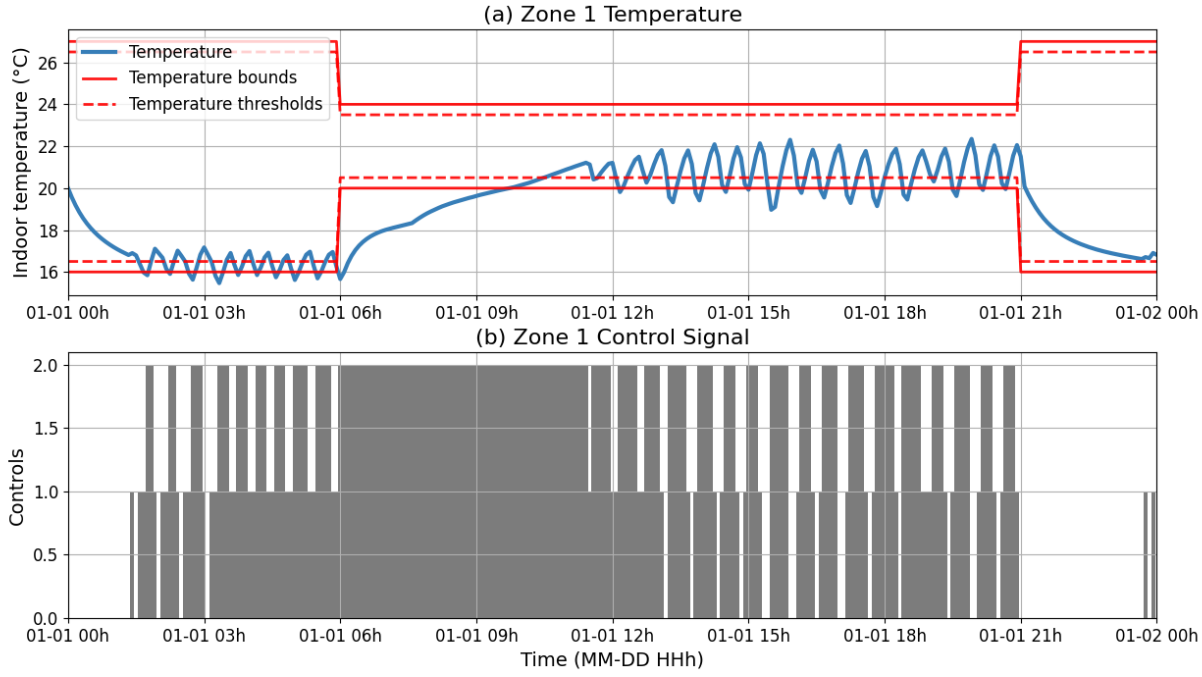

Figure 5.10 Temperature trajectories and controls of zone 1 under rule-based control for a one-day period with TOU pricing.

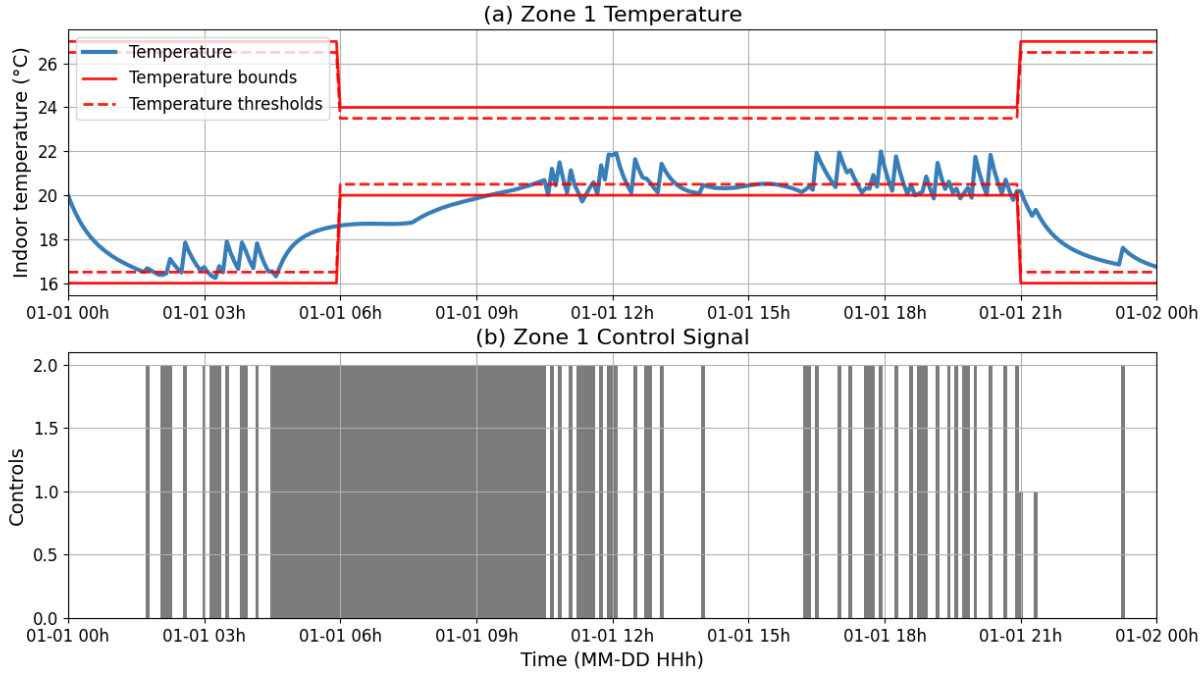

Figure 5.11 Temperature trajectories and controls of zone 1 under centralized control for a one-day period with TOU pricing.

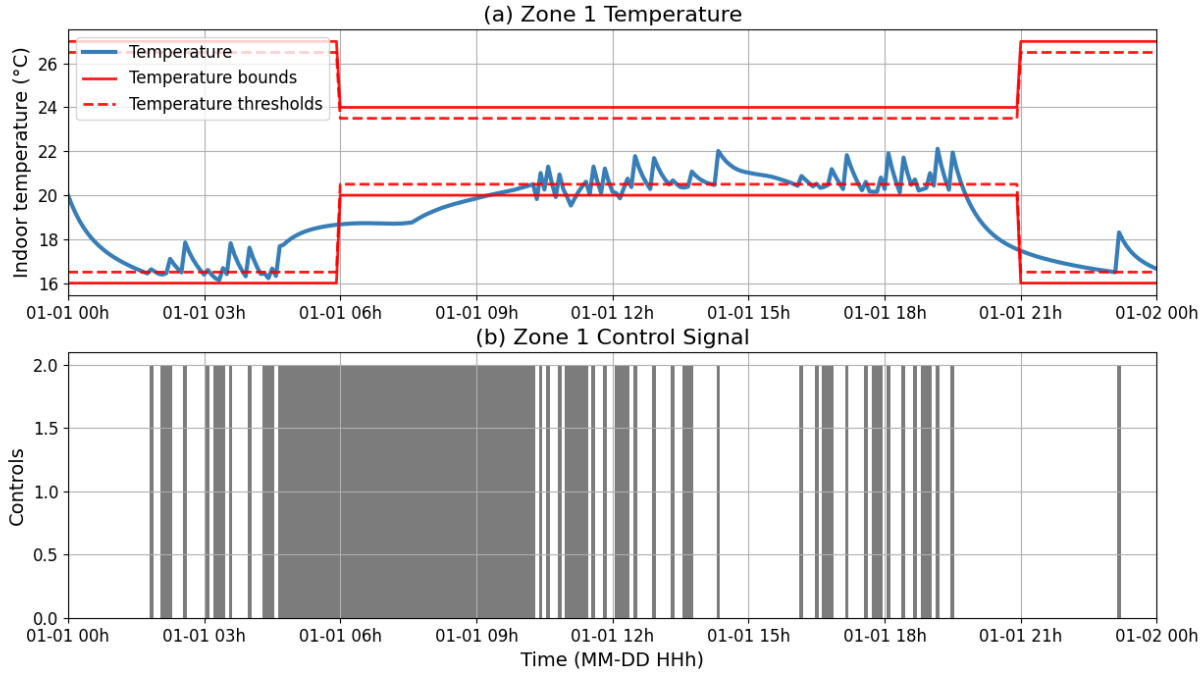

Figure 5.12 Temperature trajectories and controls of zone 1 under decentralized control for a one-day period with TOU pricing.

5.2.2.2 CPP results

We now analyze the results obtained under the CPP scheme. The simulation setup remains identical, covering 30 days of building operations. Overall, energy costs are lower under CPP compared to TOU pricing, primarily due to the reduced frequency of high price periods. As observed previously, the rule-based controller consistently underperforms relative to the optimized strategies. In the two-zone building, the decentralized controller achieves the lowest energy cost but delivers poorer thermal comfort compared to the centralized controller that explains the lower cost. As in the case of the TOU pricing. We now present time-series results for the same three-day period considered in the TOU analysis, this time under the CPP scheme and for the two-zone configuration. Because the rule-based controller does not adapt its behaviour to the price signal, its temperature trajectories remain unchanged across different pricing scenarios. The results for the rule-based, centralized, and decentralized controllers are shown in Figures 5.13, 5.14, and 5.15.

An interesting and desirable behaviour can be observed in both plots with optimization based controllers. Prior to the first high-price period, both controllers proactively raise the temperature in each zone. This preemptive action—more pronounced in the centralized case but still noticeable in the decentralized one—demonstrates a form of thermal pre-charging. By intentionally heating the zones before the onset of expensive electricity periods, the controllers reduce the need for active HVAC operation during peak pricing. This behaviour is aligned with demand-side management objectives and highlights the controllers' ability to anticipate and react to pricing dynamics in an intelligent manner.

It is also important to note that even an optimal solution may not always satisfy all system constraints under this experimental setup. Due to forecasting errors and the interplay between the strength of the HVAC modules and the thermal dynamics of the building model, the optimization-based controller may generate decisions that will eventually cause the violation of certain constraints. However, such decisions can still be considered optimal when accounting for prediction uncertainties and the physical limitations of the system. Thus, our optimal controls are limited by prediction errors and the modelled environment.

In our view, the most effective strategy for the controller is to pre-heat the building before a scheduled price increase, raising the temperature either to the upper threshold or to the level necessary to maintain comfort throughout the price spike without reactivating the HVAC modules. During high-price periods, the controller should then minimize energy consumption while ensuring that indoor temperatures remain above the lower comfort bound.

Table 5.4 Controller comparison over one month under CPP for both 2-zone and 5-zone configurations.

2-Zone Configuration				
Controller	**Price ($)**	**Price (%)**	**Comfort Rating (°C)**	**Mean Comfort Rating (°C/step)**
Rule-based	409.22	100.00	765.37	0.531
Centralized	311.81	76.2	607.89	0.422
Decentralized	301.33	73.6	759.31	0.527
5-Zone Configuration				
Controller	**Price ($)**	**Price (%)**	**Comfort Rating (°C)**	**Mean Comfort Rating (°C/step)**
Rule-based	1,035.43	100.00	2,351.63	0.653
Centralized	405.01	80.4	1982.44	0.551
Decentralized	801.15	77.4	2,265.59	0.629

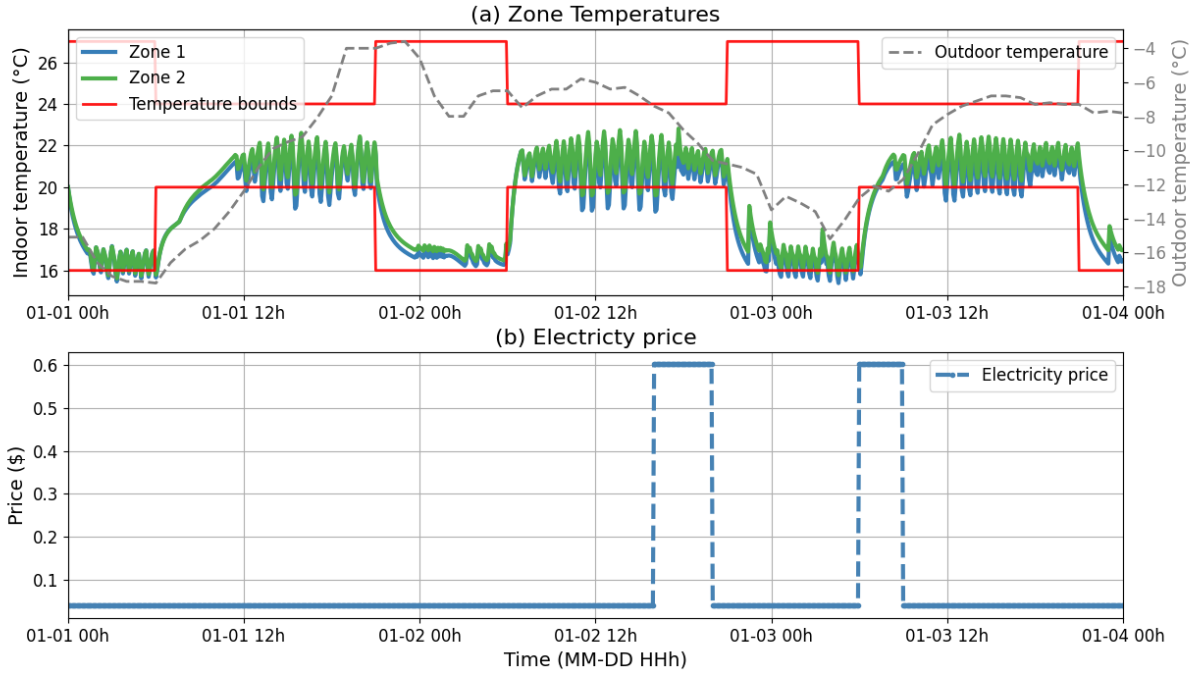

Figure 5.13 Temperature trajectories under rule-based control for a three-day period with CPP.

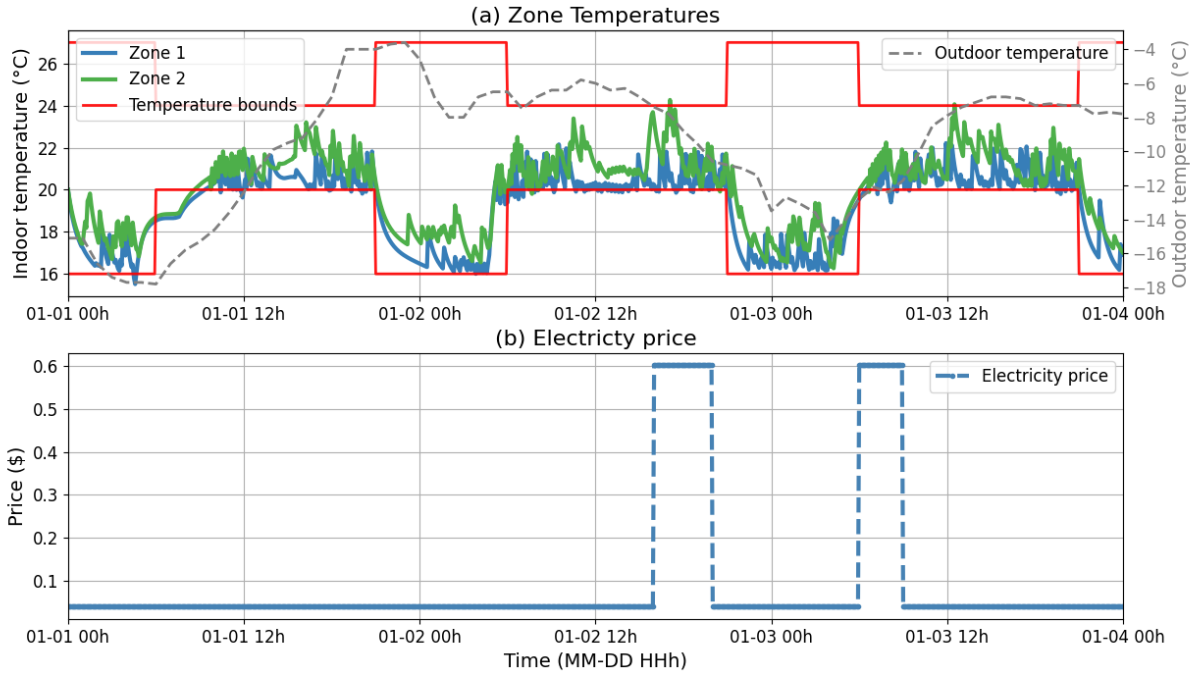

Figure 5.14 Temperature trajectories under centralized control for a three-day period with CPP.

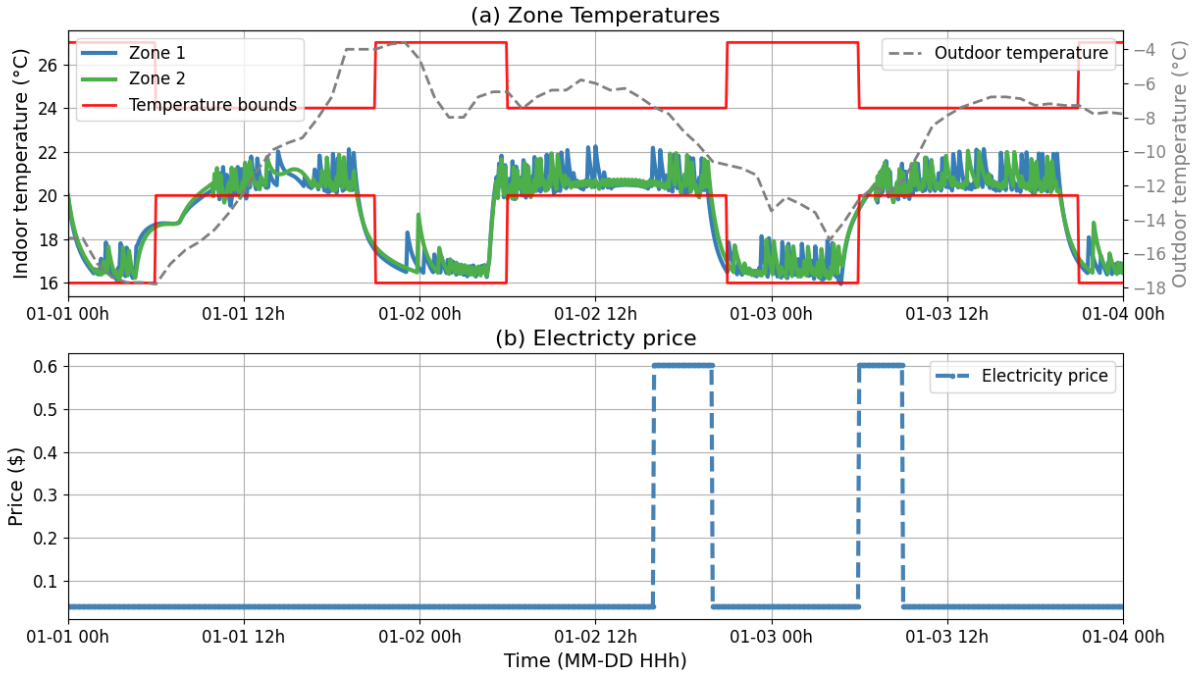

Figure 5.15 Temperature trajectories under decentralized control for a three-day period with CPP.

5.2.2.3 RTP results

We now evaluate the performance of the proposed controllers under the RTP scheme. As presented in Table 5.5, the results exhibit similar patterns to those observed under the TOU and CPP schemes (Tables 5.3 and 5.4, respectively). Notably, for the two-zone building the centralized controller achieves better thermal comfort but incurs slightly higher costs compared to the decentralized controller.

We also observe that in all three pricing schemes, the five-zone controllers perform worst in price percentage and mean comfort rating compared to the two-zone building. The reduction in mean comfort performance in multi-zone scenarios is expected. As the number of zones increases, the underlying thermodynamic interactions become more complex, leading to a higher likelihood of prediction errors. This, in turn, results in more frequent violations of the temperature bounds. Consequently, both centralized and decentralized controllers are more prone to deviations in larger-scale configurations, which explains the observed drop in comfort performance.

The final set of graphical results presents the temperature trajectories over the entire 30-day simulation period under the RTP scheme, using the two-zone configuration. Figures 5.16, 5.17, and 5.18 illustrate the performance of the rule-based, centralized, and decentralized controllers, respectively.

Across all three controllers, we observe persistent constraint violations during extremely cold periods, particularly when outdoor temperatures fall below $-20\,^{\circ}\mathrm{C}$. Even the rule-based controller—which is designed to prioritize comfort and constraint satisfaction over cost—is unable to maintain indoor temperatures within the defined bounds under such conditions. This behaviour highlights limitations inherent in the thermal model and simulation environment underscoring that optimization-based controllers must be formulated around an appropriate problem definition in order to yield control actions that are truly relevant and optimal in real-world conditions. Specifically, the inability to satisfy comfort constraints suggests that the simulated building lacks sufficient thermal insulation or that the HVAC system lacks the necessary capacity to compensate for extreme cold. In a real-world context, such results would point to either poor envelope performance or undersized heating equipment.

Table 5.5 Controller comparison over one month under RTP for both 2-zone and 5-zone configurations.

2-Zone Configuration				
Controller	**Price ($)**	**Price (%)**	**Comfort Rating (°C)**	**Mean Comfort Rating (°C/step)**
Rule-based	445.41	100.00	765.37	0.531
Centralized	325.02	73.0	605.52	0.420
Decentralized	312.03	70.1	758.03	0.526
5-Zone Configuration				
Controller	**Price ($)**	**Price (%)**	**Comfort Rating (°C)**	**Mean Comfort Rating (°C/step)**
Rule-based	1,133.38	100.00	2,351.63	0.653
Centralized	878.81	77.5	1983.86	0.551
Decentralized	843.17	74.4	2,264.18	0.629

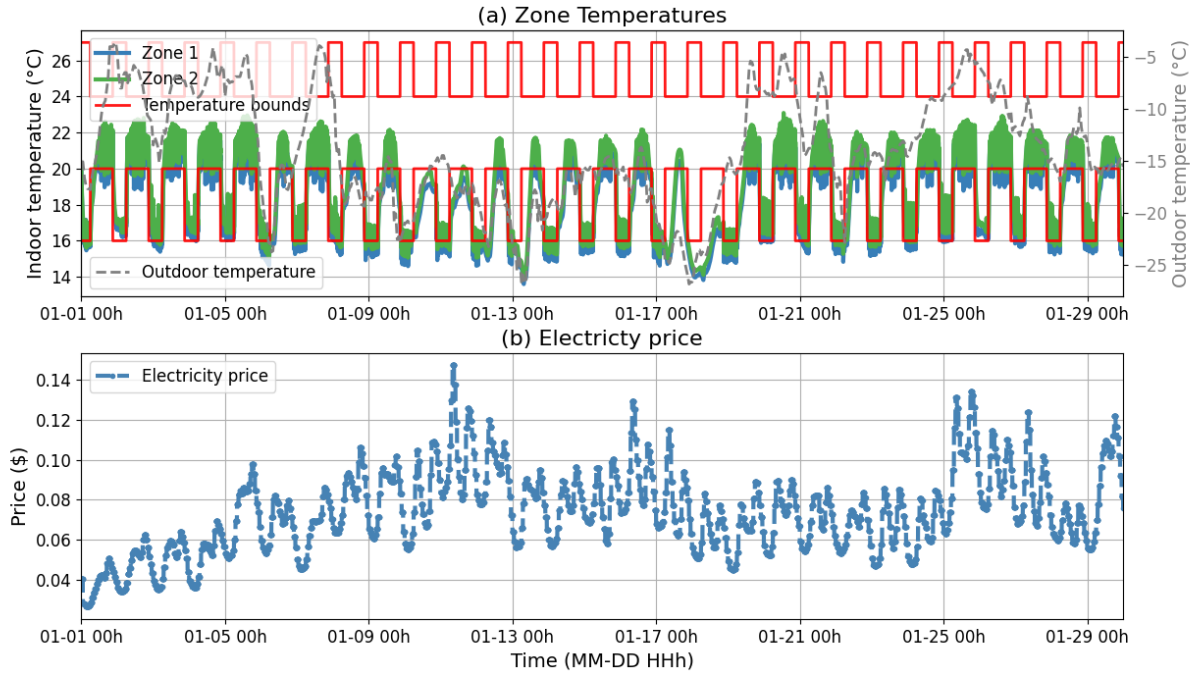

Figure 5.16 Temperature trajectories under rule-based control for a 30-day period with RTP.

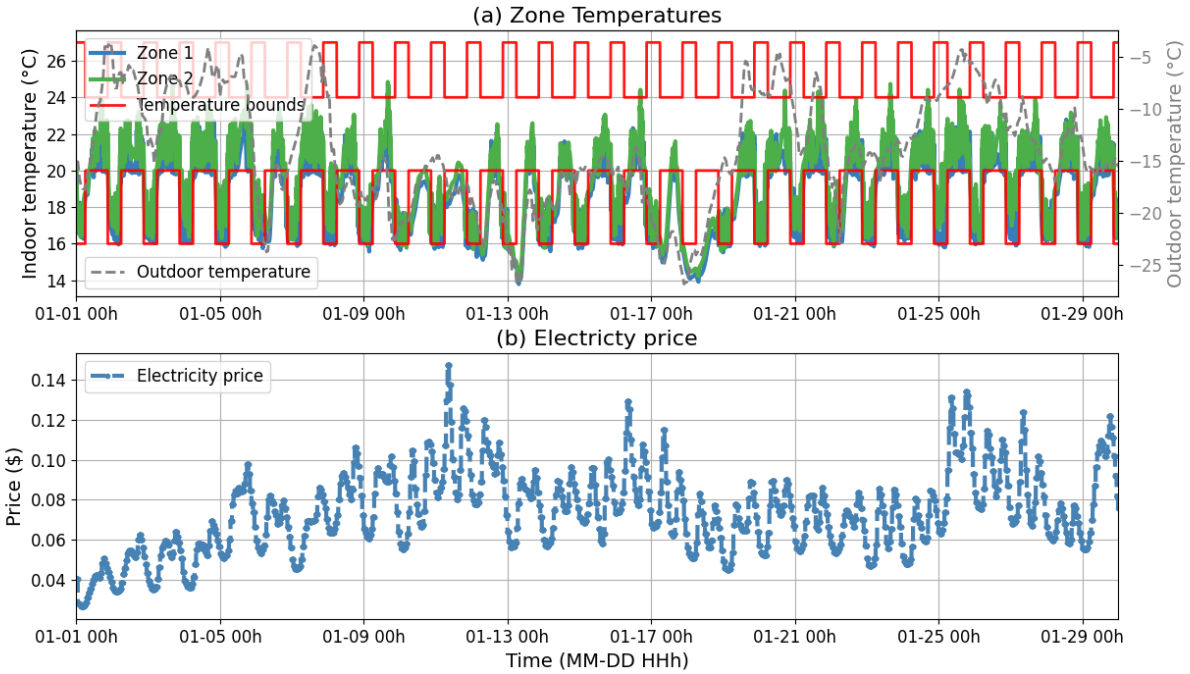

Figure 5.17 Temperature trajectories under centralized control for a 30-day period with RTP.

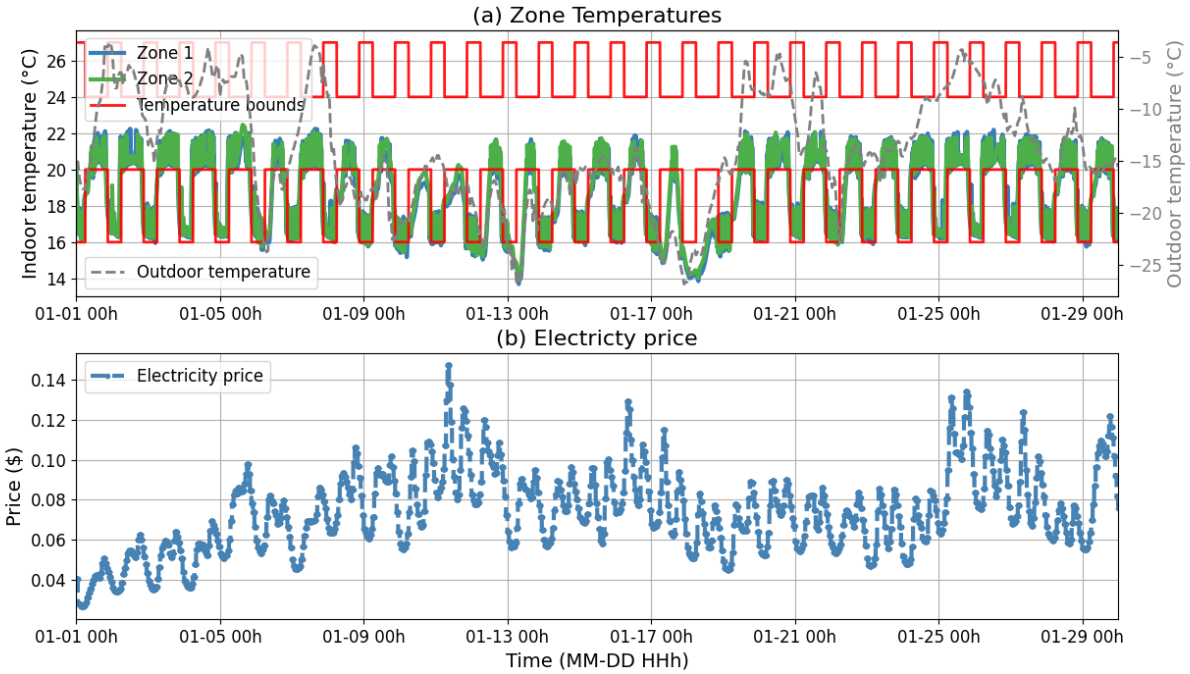

Figure 5.18 Temperature trajectories under decentralized control for a 30-day period with RTP.

CHAPTER 6 CONCLUSION

Finally, this chapter presents the conclusion of this Master's thesis. The work conducted throughout the preceding chapters is summarized, its limitations are discussed, and potential directions for future research are proposed.

6.1 Summary of Works

This work proposes decentralized temperature optimization with respect to DR strategies in multi-zone commercial buildings. A comprehensive literature review is conducted, highlighting current temperature control methodologies, the challenges of decentralization, and the importance of HVAC electricity consumption in the context of power system operations. From this review, we identify a gap in the literature concerning mixed-integer decentralized optimization for temperature regulation in multi-zone buildings.

To address this gap, we propose a decentralized optimization-based, predictive control method built upon a data-driven, physics-informed linear model. This method is tailored for multi-zone commercial buildings and is designed to operate under various DR incentive schemes. We evaluate the proposed control strategy using a simulation framework that replicates real-world conditions across three different types of DR programs. Simulation results illustrates that the decentralized control method achieves performance comparable to that of a centralized approach while offering improved scalability through parallel or local computations.

Ultimately, decentralized control strategies that incorporate DR programs show promise in supporting electrical grid flexibility and reliability. By reducing peak demand and enhancing responsiveness, such methods can contribute to deferring costly infrastructure investments and improving overall grid efficiency and reduce costs for building operators.

6.2 Limitations

This work is subject to several limitations that warrant further investigation.

- **Simulation horizon.** The evaluation is limited to a 30-day simulation period. Extending this horizon—or ideally validating the approach in a real building—would provide more robust insights into its long-term effectiveness and practical applicability. Furthermore, in certain situations, the strategy relies on a backup controller, which may lead to suboptimal decisions.

- **Building's simulator model sizing.** The simulation environment employs a relatively simplistic building model that does not fully reflect the thermal behaviour of real buildings. For example, under extreme outdoor temperatures, even maximum heating is sometimes insufficient to maintain comfort—a situation that is less likely in real commercial buildings, which typically benefit from greater thermal mass and more advanced envelope designs.

- **Baseline controller weaknesses.** The rule-based controller used for comparison serves as a basic benchmark and is relatively weak. Although it is designed to enforce comfort constraints, it often fails to maintain acceptable temperature levels. In practice, rule-based systems are typically tuned to strictly prioritize constraint satisfaction, but in this study, the baseline consistently underperformed across all test cases. It should be noted with more extensive tuning, a well-designed rule-based controller could potentially achieve performance that is comparable to, or even better than, that of the proposed optimization-based approach but would not be applicable to a vast ranges of environment or scenarios, like our approach.

- **Computational and hardware assumptions.** The computational performance was evaluated using a standard desktop computer. In real-world applications, controllers would likely operate on embedded processors within each zone's HVAC. Although preliminary estimates suggest computation times would remain under five minutes, this has not been validated on actual embedded hardware. We also assumes perfect communication and synchronization among local agents and any central coordinator. Factors such as communication delays, message loss, or asynchronous updates are not considered. These issues, however, could pose serious challenges to feasibility and real-time performance in practical deployments. Moreover, the controller assumes that control inputs are applied continuously, neglecting practical actuator limitations such as discrete steps, ON/OFF switching constraints, or actuation delays.

- **Robustness to uncertainty and disturbances.** The controller's robust approach to model mismatch, external disturbances, sensor noise, and actuator faults is not explicitly analyzed. In practice, these uncertainties are unavoidable and can significantly impair performance. While the use of MPC can mitigate some effects of imperfect data, a formal robustness analysis remains absent from this study.

6.3 Future Research

There are numerous promising avenues for future research and development. Improvements can be made across various dimensions, including the optimization algorithm, the modelling framework, the data utilized, and the testing environment.

First, the optimization process could be extended from a decentralized to a fully distributed formulation. Such an approach would eliminate the need for a central coordinator, thereby reducing communication overhead and further enhancing scalability. For readers interested in this direction, we refer to the distributed optimization framework presented in [28]. Additionally, alternative decentralized mixed-integer optimization techniques could be explored, such as cutting-plane methods like the one proposed in [66].

Second, the predictive model itself offers opportunities for enhancement. Incorporating occupancy data into the temperature dynamics could improve model accuracy. Moreover, exploring more complex and potentially nonlinear models—and devising methods to decentralize their control—could provide further performance gains. Another direction involves extending the model and control scheme to include active cooling systems, enabling a more comprehensive thermal management strategy throughout the year.

Finally, the scalability of the control method could be further evaluated by applying it to environments with more than five zones. This would require a simulation platform capable of handling larger and more complex building configurations which were not available to us at the present time. Ultimately, a critical next step would be to validate the control strategy in a real-world building environment, assessing its practical effectiveness under operational constraints.

REFERENCES

[1] Hydro One, "Electricity Pricing and Costs," https://www.hydroone.com/rates-and-billing/rates-and-charges/electricity-pricing-and-costs, 2024, accessed: 2025-06-09.

[2] International Energy Agency, "Electricity 2024: Executive Summary," 2024, accessed: 2024-10-02. [Online]. Available: https://www.iea.org/reports/electricity-2024/executive-summary

[3] F. Li, I. Kocar, and A. Lesage-Landry, "Mitigating Equipment Overloads due to Electric Vehicle Charging Using Customer Incentives," in *2023 IEEE Power & Energy Society General Meeting (PESGM)*, 2023, pp. 1–5.

[4] G. A. Orfanos, P. S. Georgilakis, and N. D. Hatziargyriou, "Transmission Expansion Planning of Systems With Increasing Wind Power Integration," *IEEE Transactions on Power Systems*, vol. 28, no. 2, pp. 1355–1362, 2013.

[5] T. Liu, Z. Wu, C. Chen, H. Chen, and H. Zhou, "Carbon Emission Accounting during the Construction of Typical 500 kV Power Transmissions and Substations Using the Carbon Emission Factor Approach," *Buildings*, vol. 14, no. 1, 2024.

[6] L. Donida Biasotto and A. Kindel, "Power lines and impacts on biodiversity: A systematic review," *Environmental Impact Assessment Review*, vol. 71, 07 2018.

[7] H. Yang, X. Zhang, Y. Ma, R. Yang, and L. Zhang, "Critical peak rebate strategy and application to demand response," *Protection and Control of Modern Power Systems*, vol. 6, no. 1, p. 28, 2021.

[8] P. Palensky and D. Dietrich, "Demand Side Management: Demand Response, Intelligent Energy Systems, and Smart Loads," *IEEE Transactions on Industrial Informatics*, vol. 7, no. 3, pp. 381–388, 2011.

[9] J. K. Calautit and H. N. Chaudhry, "Sustainable Buildings: Heating, Ventilation, and Air-Conditioning," *Energies*, vol. 15, no. 21, p. 8208, 2022.

[10] M. González-Torres, L. Pérez-Lombard, J. F. Coronel, I. R. Maestre, and D. Yan, "A review on buildings energy information: Trends, end-uses, fuels and drivers," *Energy Reports*, vol. 8, pp. 626–637, 2022.

[11] E. Roque, R. Vicente, R. M. S. F. Almeida, and V. M. Ferreira, "The Impact of Thermal Inertia on the Indoor Thermal Environment of Light Steel Framing Constructions," *Energies*, vol. 15, no. 9, p. 3061, 2022.

[12] Z. Wang, B. Lu, B. Wang, Y. Liu, N. Qian, F. Song, Y. Zeng, and D. Guan, "Incentive based emergency demand response effectively reduces peak load during heatwave without harm to vulnerable groups," *Nature Communications*, vol. 14, p. 6202, 2023.

[13] B. Alimohammadisagvand, J. Jokisalo, S. Kilpeläinen, M. Ali, and K. Sirén, "Cost-optimal thermal energy storage system for a residential building with heat pump heating and demand response control," *Applied Energy*, vol. 174, pp. 275–287, 2016.

[14] R. D. McAllister and J. B. Rawlings, "Advances in mixed-integer model predictive control," in *2022 American Control Conference (ACC)*, 2022, pp. 364–369.

[15] M. J. Risbeck, C. T. Maravelias, J. B. Rawlings, and R. D. Turney, "A mixed-integer linear programming model for real-time cost optimization of building heating, ventilation, and air conditioning equipment," *Energy and Buildings*, vol. 142, pp. 220–235, 2017.

[16] A. Khalilnejad, R. H. French, and A. R. Abramson, "Data-driven evaluation of HVAC operation and savings in commercial buildings," *Applied Energy*, vol. 278, p. 115505, 2020.

[17] A. N. M. M. Haque, M. Nijhuis, G. Ye, P. H. Nguyen, F. W. Bliek, and J. G. Slootweg, "Integrating Direct and Indirect Load Control for Congestion Management in LV Networks," *IEEE Transactions on Smart Grid*, vol. 10, no. 1, pp. 741–751, 2019.

[18] Hydro-Québec, "Hydro-Québec's electricity rates for its electricity distribution activities, effective April 1, 2025," 2025, established in accordance with Schedule 1 of the Hydro-Québec Act. Legal deposit, Bibliothèque et Archives nationales du Québec, 2nd quarter 2025.

[19] S. Boyd and L. Vandenberghe, *Convex Optimization*. Cambridge: Cambridge University Press, 2004, available online.

[20] J. A. Taylor, *Convex Optimization of Power Systems*. Cambridge University Press, 2015.

[21] M. Killian and M. Kozek, "Ten questions concerning model predictive control for energy efficient buildings," *Building and Environment*, vol. 105, pp. 403–412, 2016.

[22] N. Mazyavkina, S. Sviridov, S. Ivanov, and E. Burnaev, "Reinforcement learning for combinatorial optimization: A survey," *Computers Operations Research*, vol. 134, p. 105400, 2021.

[23] M. R. Garey and D. S. Johnson, *Computers and Intractability: A Guide to the Theory of NP-Completeness*. San Francisco: W. H. Freeman and Company, 1979.

[24] M. Fischetti, A. Lodi, and D. Salvagnin, "Just MIP it!" Università di Padova and Università di Bologna, Tech. Rep., 2010, available upon request or from the authors.

[25] L. Zhuang, X. Chen, and X. Guan, "A Decentralized Method for Energy Conservation of an HVAC System," *Building Simulation*, vol. 13, pp. 155–170, 2020.

[26] Y. Yang, G. Hu, and C. J. Spanos, "HVAC Energy Cost Optimization for a Multizone Building via a Decentralized Approach," *IEEE Transactions on Automation Science and Engineering*, vol. 17, no. 4, pp. 1950–1960, 2020.

[27] A. Falsone, K. Margellos, and M. Prandini, "A Decentralized Approach to Multi-Agent MILPs: Finite-Time Feasibility and Performance Guarantees," *Automatica*, vol. 103, pp. 141–150, 2019.

[28] ——, "A Distributed Iterative Algorithm for Multi-Agent MILPs: Finite-Time Feasibility and Performance Characterization," *IEEE Control Systems Letters*, vol. 2, no. 4, pp. 563–568, 2018.

[29] R. Ooka and S. Ikeda, "A review on optimization techniques for active thermal energy storage control," *Energy and Buildings*, vol. 106, pp. 225–233, 2015, sI: IEA-ECES Annex 31 Special Issue on Thermal Energy Storage.

[30] S. Koohi-Fayegh and M. Rosen, "A review of energy storage types, applications and recent developments," *Journal of Energy Storage*, vol. 27, p. 101047, 2020.

[31] L. Ballotta, M. R. Jovanović, and L. Schenato, "Can Decentralized Control Outperform Centralized? The Role of Communication Latency," *IEEE Transactions on Control of Network Systems*, vol. 10, no. 3, pp. 1629–1640, 2023.

[32] A. Afram and F. Janabi-Sharifi, "Theory and applications of HVAC control systems– A review of model predictive control (MPC)," *Building and Environment*, vol. 72, pp. 343–355, 2014.

[33] C. Anuntasethakul and D. Banjerdpongchai, "Design of Supervisory Model Predictive Control for Building HVAC System With Consideration of Peak-Load Shaving and Thermal Comfort," *IEEE Access*, vol. 9, pp. 41 066–41 081, 2021.

[34] Y. Balali, A. Chong, A. Busch, and S. O'Keefe, "Energy modelling and control of building heating and cooling systems with data-driven and hybrid models—A review," *Renewable and Sustainable Energy Reviews*, vol. 183, p. 113496, 2023.

[35] Z. Patonai, R. Kicsiny, and G. Géczi, "Multiple Linear Regression Based Model for the Indoor Temperature of Mobile Containers," *Heliyon*, vol. 8, no. 12, p. e12098, 2022.

[36] M. A. Fayazbakhsh, F. Bagheri, and M. Bahrami, "A Resistance–Capacitance Model for Real-Time Calculation of Cooling Load in HVAC-R Systems," *Journal of Thermal Science and Engineering Applications*, vol. 7, no. 4, p. 041008, Dec. 2015.

[37] J. Wang, Y. Jiang, C. Y. Tang, and L. Song, "Development and Validation of a Second-Order Thermal Network Model for Residential Buildings," *Applied Energy*, vol. 306, 11 2021. [Online]. Available: https://www.osti.gov/biblio/1842207

[38] Z. Wang and T. Hong, "Reinforcement Learning for Building Controls: The Opportunities and Challenges," *Applied Energy*, vol. 269, p. 115036, jul 2020.

[39] K. Al Sayed, A. Boodi, R. Sadeghian Broujeny, and K. Beddiar, "Reinforcement learning for HVAC control in intelligent buildings: A technical and conceptual review," *Journal of Building Engineering*, vol. 95, p. 110085, 2024.

[40] A. Wemhoff, "Calibration of HVAC equipment PID coefficients for energy conservation," *Energy and Buildings*, vol. 45, pp. 60–66, 2012.

[41] B. Chen, J. Francis, M. Pritoni, S. Kar, and M. Bergés, "COHORT: Coordination of Heterogeneous Thermostatically Controlled Loads for Demand Flexibility," in *Proceedings of the 7th ACM International Conference on Systems for Energy-Efficient Buildings, Cities, and Transportation*, ser. BuildSys '20. New York, NY, USA: Association for Computing Machinery, 2020, pp. 31–40.

[42] S. Qin and T. A. Badgwell, "A survey of industrial model predictive control technology," *Control Engineering Practice*, vol. 11, no. 7, pp. 733–764, 2003.

[43] D. Lu, G. Augenbroe, and Z. Zeng, "An occupants' diversity-aware discussion on the economic benefits of model predictive control in buildings, journal = Energy and Buildings," vol. 337, p. 115668, 2025.

[44] R. S. Sutton and A. G. Barto, *Reinforcement Learning: An Introduction*, 2nd ed. Cambridge, MA: MIT Press, 2018, second Edition.

[45] T. Wei, Y. Wang, and Q. Zhu, "Deep Reinforcement Learning for Building HVAC Control," in *Proceedings of the 54th Annual Design Automation Conference 2017*, ser. DAC '17. New York, NY, USA: Association for Computing Machinery, 2017.

[46] D. Xu, F. Zhu, Q. Liu, and P. Zhao, "Improving exploration efficiency of deep reinforcement learning through samples produced by generative model," *Expert Systems with Applications*, vol. 185, p. 115680, 2021.

[47] Y. Qing, S. Liu, J. Song, H. Wang, and M. Song, "A Survey on Explainable Reinforcement Learning: Concepts, Algorithms, Challenges," 2023.

[48] International Energy Agency (IEA), "World Energy Statistics and Balances (Database)," 2022, accessed on 24 September 2024; cited on 17 June 2025.

[49] R. Li, Q. Tu, H. Feng, and Z. Zou, "Demand response based battery energy storage systems design and operation optimization," *Energy and Buildings*, vol. 338, p. 115738, 2025.

[50] V. Taboga and H. Dagdougui, "A Distributed ADMM-Based Deep Learning Approach for Thermal Control in Multi-Zone Buildings Under Demand Response Events," *IEEE Transactions on Automation Science and Engineering*, vol. 22, pp. 5994–6008, 2025.

[51] M.-C. Paré, V. Dermardiros, and A. Lesage-Landry, "Efficient Data-Driven Model Predictive Control for Demand Response of Commercial Buildings," in *IECON 2025-51st Annual Conference of the IEEE Industrial Electronics Society*, 2025, pp. 1–8, accepted.

[52] S. Prívara, J. Široký, L. Ferkl, and J. Cigler, "Model predictive control of a building heating system: The first experience," *Energy and Buildings*, vol. 43, no. 2, pp. 564–572, 2011.

[53] K. Chinnakani, A. Krishnamurthy, J. Moyne, A. Arbor, and F. Gu, "Comparison of energy consumption in HVAC systems using simple ON-OFF, intelligent ON-OFF and optimal controllers," in *2011 IEEE Power and Energy Society General Meeting*, 2011, pp. 1–6.

[54] M. Kasahara, T. Matsuba, Y. Kuzuu, and T. Yamazaki, "Design and Tuning of Robust PID Controller for HVAC Systems," *ASHRAE Transactions*, vol. 105, p. 154, 1999.

[55] S. Mostafavi, H. Doddi, K. Kalyanam, and D. Schwartz, "Nonlinear Moving Horizon Estimation and Model Predictive Control for Buildings with Unknown HVAC Dynamics," *IFAC-PapersOnLine*, vol. 55, no. 41, pp. 71–76, 2022, 4th IFAC Workshop on Cyber-Physical and Human Systems CPHS 2022.

[56] S. R. West, J. K. Ward, and J. Wall, "Trial results from a model predictive control and optimisation system for commercial building HVAC," *Energy and Buildings*, vol. 72, pp. 271–279, 2014.

[57] S. Boyd, N. Parikh, E. Chu, B. Peleato, and J. Eckstein, "Distributed Optimization and Statistical Learning via the Alternating Direction Method of Multipliers," *Foundations and Trends® in Machine Learning*, vol. 3, no. 1, pp. 1–122, 2011.

[58] U.S. Department of Energy, *EnergyPlus Engineering Reference: The Reference to EnergyPlus Calculations*, U.S. Department of Energy, 2023, version 23.1.0.

[59] X. Hou, Y. Xiao, J. Cai, J. Hu, and J. E. Braun, "A Distributed Model Predictive Control Approach for Optimal Coordination of Multiple Thermal Zones in a Large Open Space," in *Proceedings of the International High Performance Buildings Conference*. West Lafayette, IN, USA: Purdue University, School of Mechanical Engineering, 2016, paper 232.

[60] ——, "Distributed model predictive control via Proximal Jacobian ADMM for building control applications," in *2017 American Control Conference (ACC)*, 2017, pp. 37–43.

[61] S. Boyd, N. Parikh, E. Chu, B. Peleato, and J. Eckstein, "Distributed Optimization and Statistical Learning via the Alternating Direction Method of Multipliers," *Foundations and Trends® in Machine Learning*, vol. 3, no. 1, pp. 1–122, 2011.

[62] L. Yu, D. Xie, T. Jiang, Y. Zou, and K. Wang, "Distributed Real-Time HVAC Control for Cost-Efficient Commercial Buildings Under Smart Grid Environment," *IEEE Internet of Things Journal*, vol. 5, no. 1, pp. 44–55, 2018.

[63] Y.-C. Ho, Q.-C. Zhao, and Q.-S. Jia, *Ordinal Optimization: Soft Optimization for Hard Problems*. New York: Springer, 2007.

[64] Y. Yang, S. Srinivasan, G. Hu, and C. J. Spanos, "Distributed Control of Multizone HVAC Systems Considering Indoor Air Quality," *IEEE Transactions on Control Systems Technology*, vol. 29, no. 6, pp. 2586–2597, 2021.

[65] R. Vujanic, P. Mohajerin Esfahani, P. J. Goulart, S. Mariéthoz, and M. Morari, "A decomposition method for large scale MILPs, with performance guarantees and a power system application," *Automatica*, vol. 67, pp. 144–156, 2016.

[66] A. Testa, A. Rucco, and G. Notarstefano, "Distributed Mixed-Integer Linear Programming via Cut Generation and Constraint Exchange," *IEEE Transactions on Automatic Control*, vol. 65, no. 4, pp. 1456–1467, 2020.

[67] R. Ko and S.-K. Joo, "Stochastic Mixed-Integer Programming (SMIP)-Based Distributed Energy Resource Allocation Method for Virtual Power Plants," *Energies*, vol. 13, no. 1, 2020.

[68] V.-A. Le and A. A. Malikopoulos, "Distributed Mixed-Integer Quadratic Programming for Mixed-Traffic Intersection Control," 2025.

[69] X. Gu, S. Ahmed, and S. S. Dey, "Exact Augmented Lagrangian Duality for Mixed Integer Quadratic Programming," *SIAM Journal on Optimization*, vol. 30, no. 1, pp. 781–797, 2020.

[70] A. T. Nguyen, D. H. Pham, B. L. Oo, M. Santamouris, Y. Ahn, and B. T. Lim, "Modelling building HVAC control strategies using a deep reinforcement learning approach," *Energy and Buildings*, vol. 310, p. 114065, 2024.

[71] Y. Gao, S. Shi, S. Miyata, and Y. Akashi, "Successful application of predictive information in deep reinforcement learning control: A case study based on an office building HVAC system," *Energy*, vol. 291, p. 130344, 2024.

[72] X. Li and J. Wen, "Review of building energy modeling for control and operation," *Renewable and Sustainable Energy Reviews*, vol. 37, pp. 517–537, 2014.

[73] M. A. Rahman, I. Rahman, and N. Mohammad, "Implementing Demand Response in The Residential Energy System to Solve Duck Curve Problem," in *2020 2nd International Conference on Advanced Information and Communication Technology (ICAICT)*, 2020, pp. 466–470.

[74] F. Bünning, B. Huber, A. Schalbetter, A. Aboudonia, M. Hudoba de Badyn, P. Heer, R. S. Smith, and J. Lygeros, "Physics-Informed Linear Regression is Competitive with Two Machine Learning Methods in Residential Building MPC," *Applied Energy*, vol. 310, p. 118491, 2022.

[75] R. L. Burden and J. D. Faires, *Numerical Analysis*, 9th ed. Brooks/Cole, Cengage Learning, 2010.

[76] S. N. Laboratories, "Global Horizontal Irradiance (GHI) - PV Performance Modeling Collaborative," https://pvpmc.sandia.gov/modeling-guide/1-weather-design-inputs/irradiance-insolation/global-horizontal-irradiance/, 2025, accessed: 2025-02-18.

[77] *Solar Radiation*. John Wiley Sons, Ltd, 2013, ch. 1, pp. 3–42.

[78] R. U. Anderson, "Derivation of Solar Position Formulae," 2020. [Online]. Available: https://arxiv.org/abs/2009.07094

[79] Z. Jiang, J. Francis, A. K. Sahu, S. Munir, C. Shelton, A. Rowe, and M. Bergés, "Data-driven Thermal Model Inference with ARMAX, in Smart Environments, based on Normalized Mutual Information," in *2018 Annual American Control Conference (ACC)*, 2018, pp. 4634–4639.

[80] J. C.-M. Yiu and S. Wang, "Multiple ARMAX modeling scheme for forecasting air conditioning system performance," *Energy Conversion and Management*, vol. 48, no. 8, pp. 2276–2285, 2007. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S0196890407000854

[81] Q. Shi and M. Hong, "Penalty dual decomposition method for nonsmooth nonconvex optimization—part i: Algorithms and convergence analysis," *IEEE Transactions on Signal Processing*, vol. 68, pp. 4108–4122, 2020.

[82] D. Lee, "Dual gradient ascent (lecture notes)," KAIST / Convex Optimization Course, Tech. Rep., 2022, see also general dual-ascent lecture notes and course materials.

[83] Y. Nesterov and V. Shikhman, "Dual subgradient method with averaging for optimal resource allocation," *European Journal of Operational Research*, vol. 270, no. 3, pp. 907–916, 2018.

[84] M. Wetter, W. Zuo, T. S. Nouidui, and X. Pang, "Modelica Buildings library," *Journal of Building Performance Simulation*, vol. 7, no. 4, pp. 253–270, 2014.

[85] Modelica Association Project, "Functional Mock-up Interface (FMI) Standard Version 3.0," https://fmi-standard.org, 2023, accessed: 2025-07-24.

[86] Lawrence Berkeley National Laboratory Modelica Buildings Library, "Buildings.examples.vavreheat.ashrae2006: Variable air volume flow system with terminal reheat and five thermal zones," https://simulationresearch.lbl.gov/modelica/

releases/v12.1.0/help/Buildings_Examples_VAVReheat.html#Buildings.Examples.VAVReheat.ASHRAE2006, 2025, accessed: 2025-08-27.

[87] F. Pedregosa, G. Varoquaux, A. Gramfort, V. Michel, B. Thirion, O. Grisel, M. Blondel, P. Prettenhofer, R. Weiss, V. Dubourg, J. Vanderplas, A. Passos, D. Cournapeau, M. Brucher, M. Perrot, and É. Duchesnay, "Scikit-learn: Machine Learning in Python," *Journal of Machine Learning Research*, vol. 12, pp. 2825–2830, 2011.

[88] S. Diamond and S. Boyd, "Cvxpy: A python-embedded modeling language for convex optimization," *Journal of Machine Learning Research*, vol. 17, no. 83, pp. 1–5, 2016.

[89] L. Gurobi Optimization, *Gurobi Optimizer Reference Manual*, 2024, https://www.gurobi.com.

[90] S. Yang, "Assessment of Urban Microclimate and Its Impact on Outdoor Thermal Comfort and Building Energy Performance," PhD thesis, Concordia University, 2024, available at Spectrum Research Repository.

[91] R. Jevons, C. Carmichael, A. Crossley, and A. Bone, "Minimum indoor temperature threshold recommendations for English homes in winter – A systematic review," *Public Health*, vol. 136, pp. 4–12, 2016.

[92] S. Chandel and R. Aggarwal, "Thermal Comfort Temperature Standards for Cold Regions," *Innovative Energy Policies*, vol. 2, 01 2012.

[93] J. Hu, Y. He, X. Hao, N. Li, Y. Su, and H. Qu, "Optimal temperature ranges considering gender differences in thermal comfort, work performance, and sick building syndrome: A winter field study in university classrooms," *Energy and Buildings*, vol. 254, p. 111554, 2022.

[94] H. Zhang, D. Zhao, C. Gu, B. Liu, and Y. Xue, "Economic optimization of smart distribution networks considering real-time pricing," *Journal of Modern Power Systems and Clean Energy*, vol. 2, no. 4, pp. 350–356, 2014.

[95] Hydro-Québec, "Historique de la demande d'électricité au Québec," https://donnees.hydroquebec.com/explore/dataset/historique-demande-electricite-quebec/information/, 2024, données horaires de la demande d'électricité mesurée en mégawatts, mises à jour annuellement. Licence CC BY-NC 4.0. Données fournies à titre informatif seulement. Dernière modification: 27 novembre 2024.

[96] *ANSI/ASHRAE Standard 55-2020: Thermal Environmental Conditions for Human Occupancy*, American Society of Heating, Refrigerating and Air-Conditioning Engineers (ASHRAE) Std., 2020.

[97] Z. Wang and Y. Chen, "Data-driven modeling of building thermal dynamics: Methodology and state of the art," *Energy and Buildings*, vol. 203, p. 109405, 2019.

APPENDIX A FULL EXPRESSION OF THE LINEAR MATRIX SYSTEM

$$A_1 = \begin{bmatrix}
\kappa_1 & \theta_{\text{out}_1} u_{\text{eco}_1} & 0 & 0 & \theta_{\text{sol}_1} & \theta'_{\text{sol}_1} & 0 & 0 & 0 & 0 & \theta_{\text{act}_1} & \theta'_{\text{act}_1} & \theta''_{\text{act}_1} & 0 & 0 & 0 & 0 & 0 & 0 & b_1 \\
\kappa_1^2 + \theta_{\text{n}_1}\theta_{\text{n}_2} & \kappa_1\theta_{\text{out}_1} u_{\text{eco}_1} & \theta_{\text{out}_1} u_{\text{eco}_1} & 0 & \kappa_1\theta_{\text{sol}_1} & \kappa_1\theta'_{\text{sol}_1} & \theta_{\text{sol}_1} & \theta'_{\text{sol}_1} & 0 & 0 & \kappa_1\theta_{\text{act}_1} & \kappa_1\theta'_{\text{act}_1} & \kappa_1\theta''_{\text{act}_1} & \theta_{\text{act}_1} & \theta'_{\text{act}_1} & \theta''_{\text{act}_1} & 0 & 0 & 0 & \kappa_1 b_1 + b_1 \\
\kappa_1\left(\kappa_1^2 + \theta_{\text{n}_1}\theta_{\text{n}_2}\right) + \theta_{\text{n}_1}\left(\kappa_1\theta_{\text{n}_2} + \kappa_2\theta_{\text{n}_2}\right) & \kappa_1^2\theta_{\text{out}_1} u_{\text{eco}_1} + \theta_{\text{n}_1}\theta_{\text{n}_2}\theta_{\text{out}_1} u_{\text{eco}_1} & \kappa_1\theta_{\text{out}_1} u_{\text{eco}_1} & \theta_{\text{out}_1} u_{\text{eco}_1} & \kappa_1^2\theta_{\text{sol}_1} + \theta_{\text{n}_1}\theta_{\text{n}_2}\theta_{\text{sol}_1} & \kappa_1^2\theta'_{\text{sol}_1} + \theta'_{\text{sol}_1}\theta_{\text{n}_1}\theta_{\text{n}_2} & \kappa_1\theta_{\text{sol}_1} & \kappa_1\theta'_{\text{sol}_1} & \theta_{\text{sol}_1} & \theta'_{\text{sol}_1} & \kappa_1^2\theta_{\text{act}_1} + \theta_{\text{act}_1}\theta_{\text{n}_1}\theta_{\text{n}_2} & \kappa_1^2\theta'_{\text{act}_1} + \theta'_{\text{act}_1}\theta_{\text{n}_1}\theta_{\text{n}_2} & \kappa_1^2\theta''_{\text{act}_1} + \theta''_{\text{act}_1}\theta_{\text{n}_1}\theta_{\text{n}_2} & \kappa_1\theta_{\text{act}_1} & \kappa_1\theta'_{\text{act}_1} & \kappa_1\theta''_{\text{act}_1} & \theta_{\text{act}_1} & \theta'_{\text{act}_1} & \theta''_{\text{act}_1} & \kappa_1\left(\kappa_1 b_1 + b_1\right) + \theta_{\text{n}_1}\theta_{\text{n}_2} b_1 + b_1 \\
\theta_{\text{n}_2} & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\
\kappa_1\theta_{\text{n}_2} + \kappa_2\theta_{\text{n}_2} & \theta_{\text{n}_2}\theta_{\text{out}_1} u_{\text{eco}_1} & 0 & 0 & \theta_{\text{n}_2}\theta_{\text{sol}_1} & \theta'_{\text{sol}_1}\theta_{\text{n}_2} & 0 & 0 & 0 & 0 & \theta_{\text{act}_1}\theta_{\text{n}_2} & \theta'_{\text{act}_1}\theta_{\text{n}_2} & \theta''_{\text{act}_1}\theta_{\text{n}_2} & 0 & 0 & 0 & 0 & 0 & 0 & \theta_{\text{n}_2} b_1 \\
\kappa_2\left(\kappa_1\theta_{\text{n}_2} + \kappa_2\theta_{\text{n}_2}\right) + \theta_{\text{n}_2}\left(\kappa_1^2 + \theta_{\text{n}_1}\theta_{\text{n}_2}\right) & \kappa_1\theta_{\text{n}_2}\theta_{\text{out}_1} u_{\text{eco}_1} + \kappa_2\theta_{\text{n}_2}\theta_{\text{out}_1} u_{\text{eco}_1} & \theta_{\text{n}_2}\theta_{\text{out}_1} u_{\text{eco}_1} & 0 & \kappa_1\theta_{\text{n}_2}\theta_{\text{sol}_1} + \kappa_2\theta_{\text{n}_2}\theta_{\text{sol}_1} & \kappa_1\theta'_{\text{sol}_1}\theta_{\text{n}_2} + \kappa_2\theta'_{\text{sol}_1}\theta_{\text{n}_2} & \theta_{\text{n}_2}\theta_{\text{sol}_1} & \theta'_{\text{sol}_1}\theta_{\text{n}_2} & 0 & 0 & \kappa_1\theta_{\text{act}_1}\theta_{\text{n}_2} + \kappa_2\theta_{\text{act}_1}\theta_{\text{n}_2} & \kappa_1\theta'_{\text{act}_1}\theta_{\text{n}_2} + \kappa_2\theta'_{\text{act}_1}\theta_{\text{n}_2} & \kappa_1\theta''_{\text{act}_1}\theta_{\text{n}_2} + \kappa_2\theta''_{\text{act}_1}\theta_{\text{n}_2} & \theta_{\text{act}_1}\theta_{\text{n}_2} & \theta'_{\text{act}_1}\theta_{\text{n}_2} & \theta''_{\text{act}_1}\theta_{\text{n}_2} & 0 & 0 & 0 & \kappa_2\theta_{\text{n}_2} b_1 + \theta_{\text{n}_2}\left(\kappa_1 b_1 + b_1\right)
\end{bmatrix}, \tag{A.1}$$

where, $\kappa_1 = 1 - \theta_{\text{out}_1} u_{\text{eco}_1} - \theta_{\text{n}_1}$

$$A_2 = \begin{bmatrix}
\theta_{\text{n}_1} & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\
\kappa_1\theta_{\text{n}_1} + \kappa_2\theta_{\text{n}_1} & \theta_{\text{n}_1}\theta_{\text{out}_2} u_{\text{eco}_2} & 0 & 0 & \theta_{\text{n}_1}\theta_{\text{sol}_2} & \theta'_{\text{sol}_2}\theta_{\text{n}_1} & 0 & 0 & 0 & 0 & \theta_{\text{act}_2}\theta_{\text{n}_1} & \theta'_{\text{act}_2}\theta_{\text{n}_1} & \theta''_{\text{act}_2}\theta_{\text{n}_1} & 0 & 0 & 0 & 0 & 0 & 0 & \theta_{\text{n}_1} b_2 \\
\kappa_1\left(\kappa_1\theta_{\text{n}_1} + \kappa_2\theta_{\text{n}_1}\right) + \theta_{\text{n}_1}\left(\kappa_2^2 + \theta_{\text{n}_1}\theta_{\text{n}_2}\right) & \kappa_1\theta_{\text{n}_1}\theta_{\text{out}_2} u_{\text{eco}_2} + \kappa_2\theta_{\text{n}_1}\theta_{\text{out}_2} u_{\text{eco}_2} & \theta_{\text{n}_1}\theta_{\text{out}_2} u_{\text{eco}_2} & 0 & \kappa_1\theta_{\text{n}_1}\theta_{\text{sol}_2} + \kappa_2\theta_{\text{n}_1}\theta_{\text{sol}_2} & \kappa_1\theta'_{\text{sol}_2}\theta_{\text{n}_1} + \kappa_2\theta'_{\text{sol}_2}\theta_{\text{n}_1} & \theta_{\text{n}_1}\theta_{\text{sol}_2} & \theta'_{\text{sol}_2}\theta_{\text{n}_1} & 0 & 0 & \kappa_1\theta_{\text{act}_2}\theta_{\text{n}_1} + \kappa_2\theta_{\text{act}_2}\theta_{\text{n}_1} & \kappa_1\theta'_{\text{act}_2}\theta_{\text{n}_1} + \kappa_2\theta'_{\text{act}_2}\theta_{\text{n}_1} & \kappa_1\theta''_{\text{act}_2}\theta_{\text{n}_1} + \kappa_2\theta''_{\text{act}_2}\theta_{\text{n}_1} & \theta_{\text{act}_2}\theta_{\text{n}_1} & \theta'_{\text{act}_2}\theta_{\text{n}_1} & \theta''_{\text{act}_2}\theta_{\text{n}_1} & 0 & 0 & 0 & \kappa_1\theta_{\text{n}_1} b_2 + \theta_{\text{n}_1}\left(\kappa_2 b_2 + b_2\right) \\
\kappa_2 & \theta_{\text{out}_2} u_{\text{eco}_2} & 0 & 0 & \theta_{\text{sol}_2} & \theta'_{\text{sol}_2} & 0 & 0 & 0 & 0 & \theta_{\text{act}_2} & \theta'_{\text{act}_2} & \theta''_{\text{act}_2} & 0 & 0 & 0 & 0 & 0 & 0 & b_2 \\
\kappa_2^2 + \theta_{\text{n}_1}\theta_{\text{n}_2} & \kappa_2\theta_{\text{out}_2} u_{\text{eco}_2} & \theta_{\text{out}_2} u_{\text{eco}_2} & 0 & \kappa_2\theta_{\text{sol}_2} & \kappa_2\theta'_{\text{sol}_2} & \theta_{\text{sol}_2} & \theta'_{\text{sol}_2} & 0 & 0 & \kappa_2\theta_{\text{act}_2} & \kappa_2\theta'_{\text{act}_2} & \kappa_2\theta''_{\text{act}_2} & \theta_{\text{act}_2} & \theta'_{\text{act}_2} & \theta''_{\text{act}_2} & 0 & 0 & 0 & \kappa_2 b_2 + b_2 \\
\kappa_2\left(\kappa_2^2 + \theta_{\text{n}_1}\theta_{\text{n}_2}\right) + \theta_{\text{n}_2}\left(\kappa_1\theta_{\text{n}_1} + \kappa_2\theta_{\text{n}_1}\right) & \kappa_2^2\theta_{\text{out}_2} u_{\text{eco}_2} + \theta_{\text{n}_1}\theta_{\text{n}_2}\theta_{\text{out}_2} u_{\text{eco}_2} & \kappa_2\theta_{\text{out}_2} u_{\text{eco}_2} & \theta_{\text{out}_2} u_{\text{eco}_2} & \kappa_2^2\theta_{\text{sol}_2} + \theta_{\text{n}_1}\theta_{\text{n}_2}\theta_{\text{sol}_2} & \kappa_2^2\theta'_{\text{sol}_2} + \theta'_{\text{sol}_2}\theta_{\text{n}_1}\theta_{\text{n}_2} & \kappa_2\theta_{\text{sol}_2} & \kappa_2\theta'_{\text{sol}_2} & \theta_{\text{sol}_2} & \theta'_{\text{sol}_2} & \kappa_2^2\theta_{\text{act}_2} + \theta_{\text{act}_2}\theta_{\text{n}_1}\theta_{\text{n}_2} & \kappa_2^2\theta'_{\text{act}_2} + \theta'_{\text{act}_2}\theta_{\text{n}_1}\theta_{\text{n}_2} & \kappa_2^2\theta''_{\text{act}_2} + \theta''_{\text{act}_2}\theta_{\text{n}_1}\theta_{\text{n}_2} & \kappa_2\theta_{\text{act}_2} & \kappa_2\theta'_{\text{act}_2} & \kappa_2\theta''_{\text{act}_2} & \theta_{\text{act}_2} & \theta'_{\text{act}_2} & \theta''_{\text{act}_2} & \kappa_2\left(\kappa_2 b_2 + b_2\right) + \theta_{\text{n}_1}\theta_{\text{n}_2} b_2 + b_2
\end{bmatrix}, \tag{A.2}$$

where, $\kappa_2 = 1 - \theta_{\text{out}_2} u_{\text{eco}_2} - \theta_{\text{n}_2}$

APPENDIX B EXTENDED TEMPERATURE PREDICTION

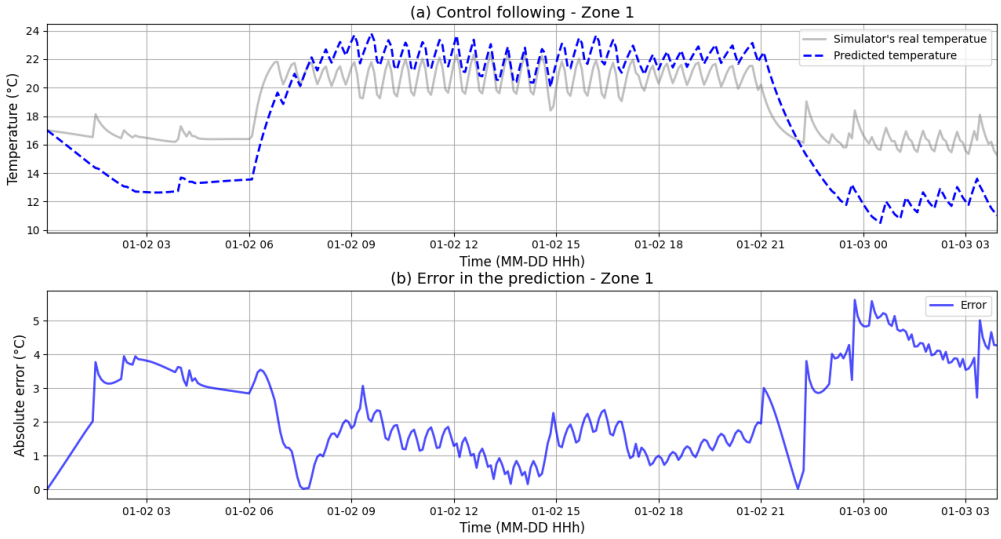

Figure B.1 Comparison between predicted and actual temperature evolution on January 3rd 2021 in zone 1 of the five-zone environment.

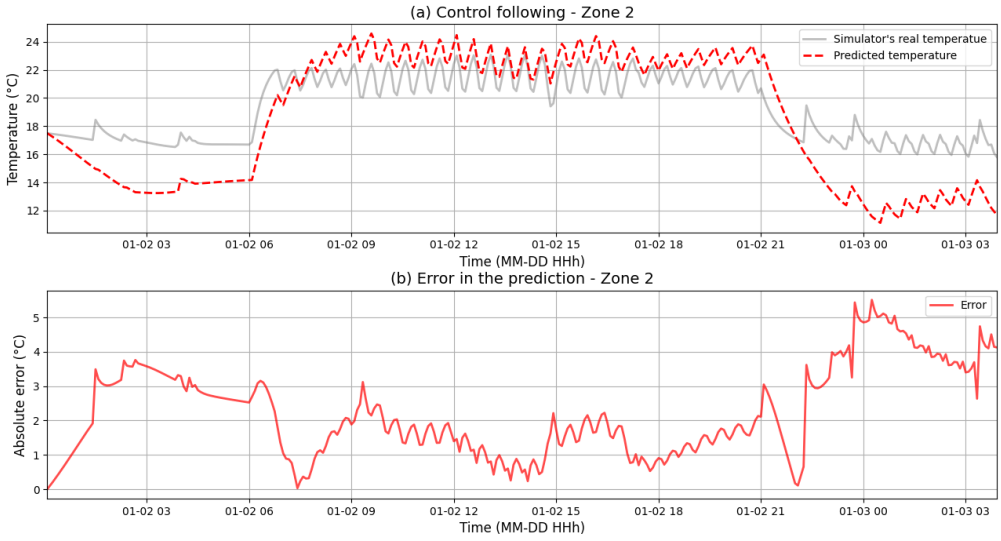

Figure B.2 Comparison between predicted and actual temperature evolution on January 3rd 2021 in zone 2 of the five-zone environment.

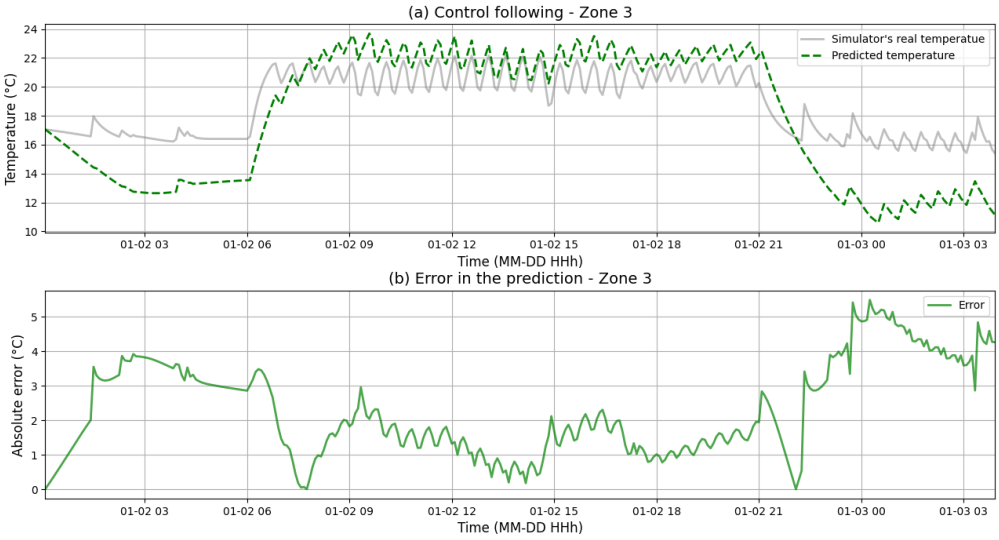

Figure B.3 Comparison between predicted and actual temperature evolution on January 3rd 2021 in zone 3 of the five-zone environment.

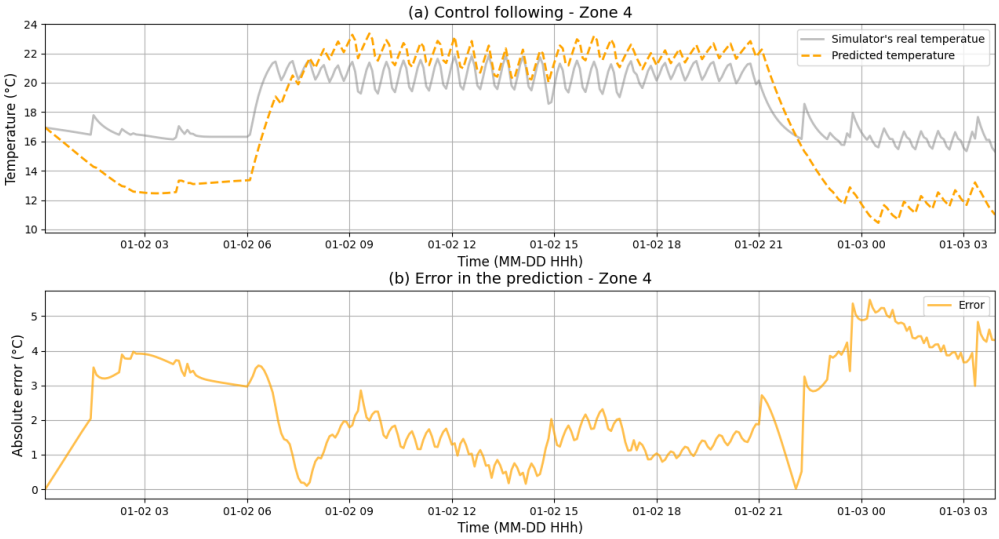

Figure B.4 Comparison between predicted and actual temperature evolution on January 3rd 2021 in zone 4 of the five-zone environment.

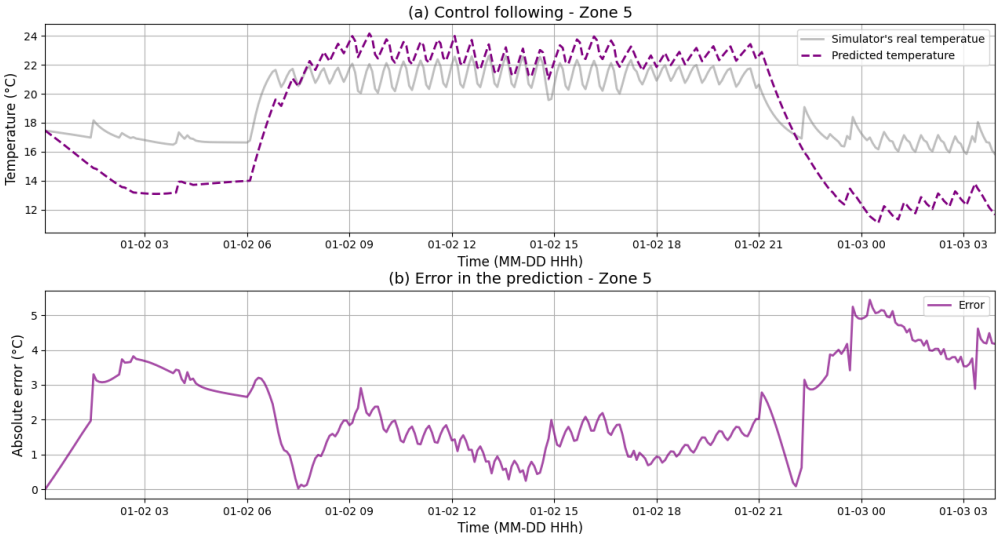

Figure B.5 Comparison between predicted and actual temperature evolution on January 3rd 2021 in zone 5 of the five-zone environment.