

Economic Model Predictive Control for HVAC Systems

Under day-ahead electricity prices and a monthly peak demand charge

Max Wervers

Master of Science Thesis



Economic Model Predictive Control for HVAC Systems

**Under day-ahead electricity prices and a monthly peak demand
charge**

MASTER OF SCIENCE THESIS

For the degree of Master of Science in Systems and Control at Delft
University of Technology

Max Wervers

June 30, 2025

Abstract

Congestion on the electricity grid is a growing issue in the Netherlands. To reduce stress on the grid, electricity providers are introducing new pricing structures. This thesis considers a pricing model that combines day ahead electricity prices with an added monthly peak demand charge. The challenge is that classical heating, ventilation, and air-conditioning (HVAC) control strategies are not responsive to these pricing mechanisms, leading to increased operational costs.

This project focuses on a heat pump-based floor heating HVAC system designed to maintain thermal comfort in a small office environment. Economic model predictive control (EMPC) can leverage the building's thermal mass to exploit variations in time-of-use (TOU) pricing. However, the monthly peak demand charge complicates the control problem, as its long-term effect is difficult to capture with the relatively short prediction horizons typical of MPC. To address this limited horizon, a method is proposed that artificially extends the prediction horizon to better account for long-term cost associated with the peak demand charge.

This artificially extended horizon is enabled by leveraging historical data for both the disturbance and electricity prices, allowing it to be approximated offline as a terminal cost function for the MPC. The proposed method, incorporating this artificial horizon extension into the MPC formulation, is compared to alternative strategies that address the peak demand charge by scaling it down within the cost function based on the length of the prediction horizon.

All methods are evaluated using information on electricity prices and disturbances that would realistically be available in practice. Results show that the MPC approach incorporating a terminal cost function economically outperforms methods that apply aggressive down scaling of the peak demand charge. However, similar economic performance is observed when compared to a more conservative scaling approach. Highlighting how conservatism with respect the peak demand is beneficial when operating under uncertainty.

Table of Contents

Acknowledgements	v
1 Introduction	1
1-1 Electricity Grid Congestion in the Netherlands	1
1-2 Novel Electricity Pricing Models	2
1-2-1 Day-Ahead Pricing	3
1-2-2 Peak Demand Charges	3
1-3 HVAC Systems	4
1-3-1 Heat Pump-Based HVAC Model	5
1-3-2 General Problem Statement	7
2 Background	9
2-1 Model Predictive Control	9
2-1-1 Economic Model Predictive Control	10
2-1-2 Consequence of Uncertainty Within Model Predictive Control	11
2-1-3 Formulating MPC with a Limited Horizon Under Monthly Peak Demand Charges	12
2-2 Proposed Prediction Horizon-Extension Strategy	15
2-2-1 Intuitive Motivation	15
2-2-2 Research Questions	16
3 Terminal Cost Approximation for Artificially Extended Horizon MPC	17
3-1 Artificially Extended Horizon MPC Formulation	17
3-1-1 Forecasting Limitations and Available Data	17
3-1-2 Artificially Extended Horizon Optimization Problem	20
3-1-3 Problem Decomposition and Reformulation	21
3-2 Terminal Cost construction	24
3-2-1 Perfect foresight Terminal Cost Approximation	25
3-2-2 Integration of Terminal Cost into MPC Algorithm	30

4	MPC Under Uncertainty	35
4-1	Terminal Cost Approximation and Implementation Under Forecast Uncertainty . . .	35
4-2	Alternative Short-Horizon MPC Methods	38
4-3	Performance Evaluation and Results	39
5	Conclusion and Discussion	51
5-1	Recommendations for Future Work	53
A	LP formulations	55
A-1	Linear MPC formulations	55
A-1-1	Equation (3-21)	55
A-1-2	Equation (3-28)	57
A-1-3	Equation (3-29)	58
A-1-4	Equation (3-30)	58
A-1-5	Equation (4-1)	59
A-1-6	Equation (4-2)	59
A-1-7	Equation (4-3)	60
A-1-8	Equation (4-5)	60
A-1-9	Equation (4-6)	61
A-2	PLP formulation	62
B	Simulation results	63
B-1	Modeling and Other Constants	63
B-1-1	Historical vs Real Temperatures 2023	64
B-2	Perfect Foresight Terminal Cost Functions	65
B-3	Terminal Cost Functions under Uncertainty	71
B-4	Monthly Distribution of Bill Cost Simulations in 2023	77
	Bibliography	85
	Glossary	89
	List of Acronyms	89
	List of Symbols	89

Acknowledgements

I would like to thank my supervisor, Koty McAllister, for introducing me to the topic of HVAC control under peak demand charges, for the insightful discussions, and for guiding me through the process of writing this MSc thesis.

Delft, University of Technology
June 30, 2025

Max Wervers

Chapter 1

Introduction

Heating, ventilation, and air-conditioning (HVAC) systems are major energy consumers in building climate management. Buildings account for up to 40% of total energy use in both the European Union and the United States [1], with HVAC systems contributing a significant portion of this consumption [2]. The energy demand of HVAC systems is dependent on the outside weather and also shows cyclic behavior driven by daily occupancy patterns, such as the typical 9-to-5 office hours when buildings are the most populated. This cyclic nature of energy use adds stress to electricity grids.

As grid loads continue to grow, electricity providers are shifting away from flat-rate pricing rates, adopting measures like time-of-use (TOU) pricing and peak demand charges to better manage energy demand. TOU electricity prices vary throughout the day, with the incentive for users to shift electricity consumption to off-peak periods. Furthermore, peak demand charges impose additional fees based on the highest power consumption throughout the billing period. These emerging pricing structures have challenged the classic HVAC control strategies for cost effective operation.

Classical HVAC control methods, such as on/off control and proportional-integral-derivative (PID) control, are commonly used due to their simplicity and effectiveness in regulating towards a desirable system state. However, these classical methods do not have the capacity to anticipate external disturbances and adapt to TOU rates and peak demand charges, resulting in a substantial increase in operating cost. Yet, the available thermal energy storage (TES) supplied by the mass of the building allows advanced control strategies like model predictive control (MPC) to exploit the variability associated with these TOU rates and peak demand charges.

1-1 Electricity Grid Congestion in the Netherlands

Congestion on the electricity grid poses a significant challenge in the current energy landscape of the Netherlands. Figure 1-1 displays a grid capacity map provided by TenneT [3], showing available capacity on the medium- to low-voltage grids managed by regional grid providers

responsible for delivering electricity to end users. The map highlights areas with a waiting list for grid connections in red and orange, indicating that grid congestion is currently a widespread issue across many parts of the Netherlands.




Figure 1-1: Grid capacity map of the Netherlands as of 19 May 2025, showing available capacity on medium- and low-voltage networks. Red and orange areas indicate regions with a waiting list for new grid connections. Source: TenneT.[3]

Grid congestion can largely be attributed to the increasing use of renewable energy sources. As fossil fuels are phased out, solar and wind energy are becoming more prominent. However, solar and wind energy are currently transmitted via the electricity grid, contributing to congestion. Moreover, grid congestion limits the capacity for integrating additional solar and wind energy, making it a critical bottleneck in the transition toward a renewable energy future.

Electrification also plays a role in worsening grid congestion, with the rise of electric vehicles being a prominent example. This trend extends to HVAC systems as well. According to the International Energy Agency (IEA), global heat pump capacity is projected to triple between 2021 and 2030 [4]. HVAC systems, particularly in the non-residential sector, often operate cyclically, with peak demand occurring during daytime or working hours. This cyclic usage increases peak loads on electricity grids, which must be sized to handle these surges. As a result, managing peak loads has become a critical aspect of modern grid operation.

1-2 Novel Electricity Pricing Models

Renewable energy and electrification are not the only contributors to grid congestion. Consumer behavior also plays a significant role. Consumers often follow patterns of electricity usage that align with peak demand hours. Electricity providers are aware of this and, given the increasing concern around grid congestion, are introducing new types of electricity pricing structures. These pricing models aim to incentivize consumers to reduce their electricity consumption during peak hours, thereby helping to alleviate grid congestion.

1-2-1 Day-Ahead Pricing

One of these new pricing structures for electricity is the dynamic energy contract. The dynamic energy contract became relevant at the start of 2023 and has since taken on a market share of 6% and 9% for residential and small businesses, respectively [5]. The dynamic energy contract is a form of TOU pricing and bases the cost of electricity on the so-called day-ahead prices. These day-ahead prices are time-varying, where each hour of the day is assigned a price. The prices are determined one day in advance and get published daily at 13:00 CET [6]. This thesis project uses the actual day-ahead prices in the Netherlands which are obtained through the ENTSO-e Transparency Platform [7]. Figure 1-2a presents some historical examples illustrating how day-ahead prices can vary throughout the day.




Figure 1-2: (a) Examples of day-ahead electricity prices in the Netherlands, showing hourly price variations across several days. (b) Distribution of day-ahead prices per hour over a one-month period, illustrating typical variability throughout the day. Data obtained from the ENTSO-e Transparency Platform [7]

As shown in Figure 1-2a, the hourly price structure often exhibits a recurring pattern: prices tend to peak during traditional high-demand periods and drop during off-peak hours. However, this pattern is not always guaranteed. The box plots in Figure 1-2b, generated from a month's worth of price data, further illustrate that while hourly price variability exists, the general trend aligns with peak-hour pricing being the most expensive.

This pattern makes dynamic energy contracts appealing to both providers and consumers. Consumers are able to shift their electrical loads to off-peak hours when electricity is cheaper. Meanwhile, providers benefit from reduced operating costs due to the lower strain on electricity grids during peak periods.

1-2-2 Peak Demand Charges

Another emerging pricing structure is the so-called peak demand charge. This charge adds an extra cost to the electricity bill based on the consumer's highest demand during the billing period. Typically, peak demand is calculated by averaging the power usage (kW) over a 15-minute interval. Figure 1-3 illustrates how the peak demand charge is applied. The gray area




Figure 1-3: Illustration of a peak demand charge. The grey area represents total energy consumption over time, while the red marker indicates the highest level of power demand during the billing period, which is used to calculate the peak demand charge.

represents the cost associated with total energy consumption, while the red marker indicates the peak demand which will be charged.

Currently, such peak demand charges are not implemented in the Netherlands. However, their introduction is plausible, especially since Belgium recently adopted a monthly peak demand charge aimed at reducing peak demands and therefore reducing stress on the electricity grid.

In this thesis project, a pricing model that combines a dynamic energy contract based on day-ahead prices with a monthly peak demand charge is considered. The peak demand charge is set at €5 per kW, which corresponds to the typical magnitude used in Belgium. One could argue that a pricing structure of day-ahead prices and peak demand charges would incentivize consumers to adopt a smoother energy consumption profile, avoiding high peaks to minimize peak demand charges, and shifting usage to off-peak hours to benefit from lower dynamic rates. However, this assumes that control systems, such as those used in HVAC, are responsive to these pricing signals. In reality, traditional HVAC controls are unaware of these pricing mechanisms, potentially leading to increased operational costs.

1-3 HVAC Systems

Most parts of our lives are spent indoors, where maintaining a comfortable indoor climate, insulated from external weather conditions, is essential. This comfortable climate is achieved through HVAC systems, which play a central role in regulating indoor environments. Figure 1-4 illustrates a simple schematic representation of an HVAC system in a Dutch climate. HVAC systems respond to various disturbances, such as fluctuations in outdoor weather conditions and changes in occupancy rates, to maintain the desired indoor climate.

In the Netherlands, for small to medium-sized buildings, heating is typically the most common method of achieving a comfortable indoor climate. Heating being the only objective allows for relatively simple HVAC systems. However, as buildings increase in size or entire districts, such as campuses, are considered, HVAC systems become significantly more complex and operate on a much larger scale. These large-scale HVAC systems are generally divided into two primary subsystems: the *airside* and the *waterside*.




Figure 1-4: Simplified schematic of an HVAC system in a Dutch climate

Patel et al. [8] conducted a case study on the HVAC system at Stanford University's campus. The airside subsystem, which includes building zones and air handling units (AHUs), is responsible for regulating indoor climates. Meanwhile, the waterside subsystem comprises pumps, boilers, chillers, cooling towers, heat recovery chillers, and hot and cold water storage systems. This subsystem ensures that the required load requested by the airside subsystem is met effectively. Similarly, Fathollahzadeh et al. [9] examined part of the HVAC system at the Colorado School of Mines. In their study, the airside subsystem consisted of multiple buildings equipped with AHUs, while the waterside subsystem included pumps, cooling towers, and chillers.

1-3-1 Heat Pump-Based HVAC Model

For this thesis project, a ground source heat pump supplying thermal energy to a floor heating system is considered. The system dynamics are adapted from the research done by Halvgaard et al. [10]. The building is modeled as a small office space consisting of a single air zone. The heat pump draws electrical power from the grid, which results in heat transfer to the water circulated through the floor heating system. This heated water raises the floor temperature, which in turn increases the air temperature inside the zone. The indoor air temperature is affected by the ambient temperature, which is treated as an external disturbance.

To model the building thermodynamics, the following assumptions are made:

- The air temperature in the zone is spatially uniform (well-mixed air).
- No mechanical ventilation or air exchange is present.
- Humidity effects are neglected (air is treated as dry).
- The ambient (outside) temperature is the only disturbance.

Figure 1-5 provides a schematic overview of the system components and their interactions, including the variables present in the system.

The thermodynamics of the system can be described using an energy balance equation in which the evolution of a zone temperature is determined by the net rate of heat transfer with neighboring zones, as shown in Equation (1-1).




Figure 1-5: Schematic representation of the ground source heat pump system supplying a floor heating system in a single-zone office building

$$C_{p,i}\dot{T}_i = \dot{Q}_{i,in} - \dot{Q}_{i,out} \quad (1-1)$$

Using Equation (1-1) for zone temperatures T_r , T_f , T_w the thermodynamics in Equations (1-2) are obtained.

$$C_{p,r}\dot{T}_r = \dot{Q}_{fr} - \dot{Q}_{ra} \quad (1-2a)$$

$$C_{p,f}\dot{T}_f = \dot{Q}_{wf} - \dot{Q}_{fr} \quad (1-2b)$$

$$C_{p,w}\dot{T}_w = \dot{Q}_c - \dot{Q}_{wf} \quad (1-2c)$$

The heat transfer rate between zones is modeled in Equations (1-3) as the product of the heat transfer coefficient $K_{i,j}$ and the temperature difference between adjacent zones. These heat transfer coefficients and other modeling constants can be found in Appendix B-1.

$$\dot{Q}_{ra} = K_{ra}(T_r - T_a) \quad (1-3a)$$

$$\dot{Q}_{fr} = K_{fr}(T_f - T_r) \quad (1-3b)$$

$$\dot{Q}_{wf} = K_{wf}(T_w - T_f) \quad (1-3c)$$

To model the heat transfer rate to the water from the heat pump's vapor-compression refrigeration cycle, it is assumed that the ground source temperature remains constant. Under this assumption, the relationship between the heat transfer rate and the electrical power supplied to the compressor is given by Equation (1-4).

$$\dot{Q}_c = COP \cdot W_c \quad (1-4)$$

It is assumed that the coefficient of performance (COP) remains constant and that the heat pump is only used for heating. Substituting Equations (1-3) and (1-4) into Equation (1-2) yields the system thermodynamics presented in Equation (1-5).

$$C_{p,r}\dot{T}_r = K_{fr}(T_f - T_r) - K_{ra}(T_r - T_a) \quad (1-5a)$$

$$C_{p,f}\dot{T}_f = K_{wf}(T_w - T_f) - K_{fr}(T_f - T_r) \quad (1-5b)$$

$$C_{p,w}\dot{T}_w = COP \cdot W_c - K_{wf}(T_w - T_f) \quad (1-5c)$$

Based on the system dynamics, a continuous linear state space model can be formulated, as shown in Equation (1-6). For the states (x) the inside air temperature (T_r), floor temperature (T_f), and water temperature (T_w) are used. The input (u) is given by the power supplied to the compressor inside the heat pump (W_c), while the disturbance input (d) is represented by the outside temperature (T_a)

$$\begin{aligned} \dot{x} &= Ax + Bu + B^d d \\ x &= \begin{bmatrix} T_r \\ T_f \\ T_w \end{bmatrix}, \quad u = W_c, \quad d = T_a \end{aligned} \quad (1-6)$$

Using the defined states and inputs, Equations (1-5) can be rewritten in state-space form, resulting in the system matrices shown in Equation (1-7).

$$A = \begin{bmatrix} -\frac{K_{fr}+K_{ra}}{C_{p,r}} & \frac{K_{fr}}{C_{p,r}} & 0 \\ \frac{K_{fr}}{C_{p,f}} & -\frac{K_{wf}+K_{fr}}{C_{p,f}} & \frac{K_{wf}}{C_{p,f}} \\ 0 & \frac{K_{wf}}{C_{p,w}} & -\frac{K_{wf}}{C_{p,w}} \end{bmatrix} \quad B = \begin{bmatrix} 0 \\ 0 \\ \frac{COP}{C_{p,w}} \end{bmatrix} \quad B_d = \begin{bmatrix} \frac{K_{ra}}{C_{p,r}} \\ 0 \\ 0 \end{bmatrix} \quad (1-7)$$

The continuous-time linear state-space model presented in Equation (1-6) is discretized using zero-order hold with a sampling time of $t_s = 15$ minutes. Which results in the discrete-time system in Equation (1-8)

$$x(k+1) = A_d x(k) + B_d u(k) + B_d^d d(k) \quad (1-8)$$

The sampling time t_s is selected to ensure an accurate representation of the building's thermodynamics while also aligning with the standard 15-minute interval used in the calculation of peak demand charges.

1-3-2 General Problem Statement

The purpose of the heat pump system is to maintain a comfortable indoor climate. While various metrics exist to evaluate indoor thermal comfort [11], this thesis considers only the room temperature T_r . Therefore, the system is required to keep the room temperature within the comfort bounds shown in Figure 1-6.




Figure 1-6: Indoor temperature comfort bounds used in this study. With stricter limits applied during working hours.

The comfort bounds are tightened during typical working hours, as the heat pump is assumed to be heating a small office space. Taking these comfort constraints into account leads to the general problem statement:

Economically controlling an HVAC system under day ahead electricity prices and monthly peak demand charges

Chapter 2

Background

This chapter discusses how the use of model predictive control (MPC) can reduce the operational costs associated with an electricity pricing structure that includes day-ahead electricity prices and monthly peak demand charges. However, incorporating such a pricing scheme into MPC introduces a new challenge: the monthly peak demand charge must be accounted for within the prediction horizon. Existing methods from the literature for addressing this issue are reviewed, and an alternative approach using historical data is proposed. The chapter concludes by presenting the research questions.

2-1 Model Predictive Control

Model predictive control (MPC) is a powerful control paradigm that uses a model of the system's dynamics to determine an optimal input sequence over a prediction horizon. This input sequence is determined through optimization of a selected quantity, also known as the objective function

This optimization is subject to the system dynamics and constraints on states and inputs. At each timestep, an optimization is performed, and typically the first input from the optimal control sequence is applied to the system. This process is repeated at the next timestep starting in the updated state. Figure 2-1 shows a diagram of how MPC interacts with the plant.

The model used to describe the system's dynamic behavior plays a crucial role in the success of MPC. A specific model choice presents a trade-off. One might strive to obtain a model that matches reality as closely as possible. However, due to the non-linear nature of building thermodynamics, this non-linear model would come at the cost of greater optimization complexity [12]. Therefore, the goal is to select a model with the least complexity required to achieve satisfactory results.

In a study conducted by Picard et al. [13], the complexity of modeling was explored for a six-room (zone) residential building. The building was initially modeled using a non-linear




Figure 2-1: Schematic diagram of MPC interacting with the plant (building).

approach in Modelica. The authors linearized this model to obtain a linear state-space representation and subsequently reduced the number of states to investigate the effect of model order on control outcome. They concluded that the controller model must at least be equipped with states for the independent air zones of the building. Furthermore, adding states to represent the walls that separate these zones helped filter the high-frequency fluctuations caused by neighboring zones.

In a comparative study, Zheng et al. [14] evaluated the controller performance obtained using different modeling approaches for a single-zone heating system. They compared a linear resistance-capacitance (RC) network with a artificial neural network (ANN) model that captures system nonlinearities. Both models were tested within an MPC framework to reduce operational costs and improve comfort. The study found that both approaches achieved comparable performance.

2-1-1 Economic Model Predictive Control

In the early days of MPC, its primary purpose was to track setpoints by using a quadratic cost function. However, the versatility of MPC extends beyond setpoint tracking when one changes the cost function to address other objectives. Rawlings et al. [15] introduced economic model predictive control (EMPC), which employs a cost function that measures economic performance. For HVAC systems, optimizing economic performance might involve minimizing energy consumption or the associated financial costs.

Since EMPC still has to provide a suitable living climate inside the building, one might ask themselves how EMPC is able to achieve cost reduction compared to the standard control methods. This capability arises from the building's inherent thermal mass, which slows thermodynamic processes and effectively acts as thermal energy storage (TES). Unlike traditional control methods, which do not capitalize on TES, EMPC can leverage it to shift energy loads away from high time-of-use (TOU) pricing periods and shave energy peaks when a significant peak demand charge is predicted.

Ma et al. [16] explored the application of EMPC in a simulated commercial building with

moderate thermal mass using EnergyPlus. The study considered both TOU prices and a daily peak demand charge, with a comfort range established by placing constraints on zone temperatures. Results demonstrated that the EMPC controller successfully shifted loads from on-peak pricing periods, achieving a cost saving of 28.52% compared to a baseline tracking controller. It also outperformed controllers employing pre-programmed pre-cooling strategies. In a subsequent study, the same authors implemented the EMPC algorithm in a physical building, further validating the controller's load-shifting capabilities [17]. However, both studies assumed perfect weather forecasts, neglecting the uncertainties associated with these forecasts, which may negatively impact controller performance.

2-1-2 Consequence of Uncertainty Within Model Predictive Control

Uncertainty can degrade the performance of MPC. One common source of uncertainty is a mismatch between the real system dynamics and the formulated dynamical model, often referred to as plant-model mismatch. In the context of buildings, this may arise, for instance, when heat transfer coefficients in the model deviate from actual values. Another source of uncertainty stems from external disturbances affecting the building envelope. While detailed building models may incorporate hundreds of such disturbances, it is often sufficient to include only the most influential ones in simplified models [18]. Dominant disturbances introducing uncertainty in MPC typically include ambient temperature, solar irradiance, and internal heat gains from occupants. However, as highlighted in the review article by Drgoňa et al. [19], there is no consensus in the literature on which disturbances are most critical, as this depends on factors such as building type and outdoor climate.

The risks associated with EMPC in the presence of uncertainty arise from its operational objectives. For example, balancing thermal comfort and cost reduction represents competing goals. When EMPC prioritizes financial cost as the primary optimization objective, the controller often pushes indoor temperatures toward the boundaries of acceptable constraints. In scenarios with uncertainty, such control actions may fail to maintain temperatures within the desired comfort range.




Figure 2-2: Demonstration of how forecast uncertainty grows over time.

One of the main disturbances when dealing with building thermodynamics is the ambient temperature. The predictive model uses forecasted values for the ambient temperature to determine the state trajectory. These forecasted values tend to become more uncertain the

further one looks into the future. Figure 2-2 shows an example of how this uncertainty can grow over time using a forecast. Because (forecast) uncertainty typically increases with the length of the prediction horizon, MPC controllers are often designed with limited prediction horizons to ensure that this uncertainty does not significantly degrade control performance. This presents a challenge: monthly peak demand charges must still be addressed within a limited prediction horizon.

2-1-3 Formulating MPC with a Limited Horizon Under Monthly Peak Demand Charges

The incorporation of a monthly demand charge within a MPC framework presents a challenge, as the prediction horizon typically does not span an entire month. We begin by defining the economic cost reflected in the monthly electricity bill, which accounts for day-ahead electricity prices and a monthly peak demand charge, as shown in Equation (2-1). Here, c_e and c_d represent the day-ahead and peak-demand rates, respectively, and u denotes energy usage. The summation spans the entire month, starting at $t = 0$ and ending at time point T .

$$\sum_{t=0}^{T-1} \left(c_e(t) \cdot u(t) \cdot \Delta \right) + c_d \cdot \max_{0 \leq t < T} u(t) \quad (2-1)$$

Cutsem et al. [20] conducted a comparative analysis of various economic model predictive control (EMPC) formulations for cooling buildings subject to TOU rates and monthly peak demand charges. These formulations aim to optimize the economic cost defined in Equation (2-1). The authors show, using different notation, that a common approach in the literature is the EMPC formulation presented in Equation (2-3). Since the system dynamics are formulated in linear state-space form, it is natural to formulate the MPC algorithm as a linear program (LP). However, the cost function in Equation (2-1) cannot be directly used in an LP due to the presence of the max operator. To address this, the max operator is replaced by the slack variable z .

Slack variable z must be combined with the constraint in Equation (2-2), this ensures that the optimization naturally drives z to be equal to the maximum energy usage within the prediction horizon.

$$z \geq u(k) \quad \forall k = t, \dots, t + N - 1 \quad (2-2)$$

Using a rolling horizon starting at time t with a prediction horizon of length N results in the MPC scheme shown in Equation (2-3):

$$\begin{aligned} \min_{u, z} \quad & \sum_{k=t}^{t+N-1} \left(c_e(k) \cdot u(k) \cdot \Delta \right) + c_d \cdot z \\ \text{s.t.} \quad & \text{system dynamics} \\ & \text{state \& input constraints} \\ & z \geq 0, \\ & z \geq u(k), \quad \forall k = t, \dots, t + N - 1 \end{aligned} \quad (2-3)$$

To demonstrate that the reformulation using the slack variable z is valid, we compare the corresponding objective functions. For simplicity, the objective functions are expressed solely in terms of the decision variables. Equations (2-4) and (2-5) present the objective functions for the original and the reformulated optimization problems, respectively. Both objective functions are subject to the same system dynamics, as well as identical state and input constraints.

$$J(\mathbf{u}) = \sum_{k=t}^{t+N-1} \left(c_e(k) \cdot u(k) \cdot \Delta \right) + c_d \cdot \max_k u(k) \quad (2-4)$$

$$\tilde{J}(\mathbf{u}, z) = \sum_{k=t}^{t+N-1} \left(c_e(k) \cdot u(k) \cdot \Delta \right) + c_d \cdot z \quad (2-5)$$

We define the optimization variables such that the objective functions attain their respective optimal values, denoted by J^* and $\tilde{J}^*$, as shown in Equation (2-6):

$$\begin{aligned} J^* &= J(\mathbf{u}^*) \\ \tilde{J}^* &= \tilde{J}(\hat{\mathbf{u}}, \hat{z}) \end{aligned} \quad (2-6)$$

We begin by substituting $\mathbf{u}^*$ for $\mathbf{u}$ in $\tilde{J}(\mathbf{u}, z)$. Since the slack variable z can be chosen freely as long as constraint (2-2) is satisfied, we set $z^* = \max_k u(k)$. Using Equations (2-4), (2-5), and (2-6), we conclude that:

$$\tilde{J}(\mathbf{u}^*, z^*) = J(\mathbf{u}^*) = J^* \quad (2-7)$$

From this, we observe that $\tilde{J}^* \leq \tilde{J}(\mathbf{u}^*, z^*)$, since $\tilde{J}^*$ is the minimum value of the reformulated objective. Therefore, by Equation (2-7), it follows that:

$$\tilde{J}^* \leq J^* \quad (2-8)$$

Moreover, we can analyze $\tilde{J}(\hat{\mathbf{u}}, \hat{z})$. In Equation (2-5), we observe that $c_d > 0$, which implies that the objective function is strictly increasing in z . Therefore, under the optimality condition and the constraint (2-2), it must hold that $\hat{z} = \max_k \hat{u}(k)$. Using Equations (2-4), (2-5), and (2-6), we conclude that:

$$\tilde{J}(\hat{\mathbf{u}}, \hat{z}) = J(\hat{\mathbf{u}}) = \tilde{J}^* \quad (2-9)$$

Since J^* is the minimal value of the original objective function, we have $J^* \leq J(\hat{\mathbf{u}})$. Therefore, by Equation (2-9), it follows that:

$$J^* \leq \tilde{J}^* \quad (2-10)$$

Combining the inequalities in Equations (2-8) and (2-10), we conclude that the optimal costs of both objective functions are equal. This establishes that the reformulated problem yields the same optimal value as the original, and thus the reformulation is valid.

However, formulation 2-3 does not account for peak demands that occurred earlier in the month, prior to the current prediction window N , and it weights the monthly peak demand charge against the TOU cost only within the prediction window. As a result, it tends to overestimate the contribution of the demand charge when the optimization is applied repeatedly throughout the billing period. Since the MPC algorithm operates with a moving horizon and solves an optimization problem at each time step, it is essential that each optimization is aware of the highest peak demand encountered so far in the month.

Keeping track of the highest peak demand so far is important because once a certain peak demand has been reached, additional usage up to that level does not incur any further demand charges. In other words, maintaining power consumption at or below the previous peak does not increase the demand cost, and the controller should take advantage of this to avoid unnecessary conservatism.

To address this limitation, Cutsem et al. propose two alternative formulations. The first introduces a threshold $u^{\text{thr}} = u^{\text{seen}}$, where u^{seen} denotes the highest peak demand observed in the current month up to the present time. This threshold is incorporated through the constraint $z \geq u(k) - u^{\text{thr}}$, effectively transforming the peak demand charge into an incremental form—penalizing only usage that exceeds the historical peak within the cost function. The second formulation extends this concept by defining the threshold as $u^{\text{thr}} = \max(u^{\text{seen}}, u^{\text{exp}})$, where u^{exp} represents the expected future peak demand. This approach enables the controller to consider both past and anticipated peak demands beyond the prediction horizon.

The results of the study indicate that the baseline formulation in Equation (2-3) yields the lowest monthly maximum peak demand. However, in terms of total monthly cost, both proposed formulations outperform the baseline. Notably, the formulation incorporating the expected future peak (u^{exp}) achieves slightly better performance than the one based solely on historical peaks (u^{seen}). Despite these findings, the study is limited in the sense that uncertainties are not considered, which would make it especially difficult to estimate the expected peak for the whole month.

Researchers Kumar et al. [21] took a different approach in formulating the cost function. Their reasoning was that avoiding peaks closer to the end of the month is more desirable, as peaks in the beginning of the month allow for peaks just as large during the month without added cost. To achieve this, a discount factor was introduced: $\sigma_t := \min\{(T-t)/N, 1\}$, where t and T represent the current and last time point of the month, respectively. Incorporating this idea into the cost function of Equation (2-3) involves replacing c_d with c_d/σ_t . However, the paper did not provide comparisons with other cost function formulations. One critique of this discounting approach is that summing the costs from each individual optimization over the horizon N does not equate to the actual financial cost associated with the electricity usage for the whole month. This raises the question of whether the financial cost is truly being optimized.

Alternatively, Risbeck et al. [22] proposed a cost function that accurately represents the true financial cost when summed over the entire month. The authors achieved this by introducing a telescoping series used to scale the peak demand charge. The telescoping series, denoted as $a(t)$, is defined as a sequence $a(0), a(1), \dots, a(T)$, where $a(0) = 0$ and $a(T) = 1$, with $t \in \{0, 1, \dots, T\}$ representing the current time step in the month, and T being the final time step. The resulting cost formulation is shown in Equation (2-11):

$$J(\mathbf{u}, z, t) = \sum_{k=t}^{t+N-1} \left(c_e(k) \cdot u(k) \cdot \Delta \right) + c_d(a(t+N)z - a(t)z^-) \quad (2-11)$$

Here, z^- denotes the previous peak. Moreover, the telescoping series $a(t)$ is not restricted to a specific shape. As long as the conditions $a(0) = 0$ and $a(T) = 1$ are satisfied, the cost function will accurately represent the total financial cost over the entire month.

The authors conducted a simulation on an energy storage system tasked with meeting a given energy demand by either supplying energy from storage or purchasing it at time-varying prices, or doing both simultaneously. A comparison was made among three methods: one employing a telescoping series with a linearly increasing weight, and two others that scaled the peak demand based on the ratio between the prediction horizon and the peak demand charge window. Among these, the method using a linearly increasing weight on the peak cost demonstrated the best performance. However, it should be noted that uncertainty, in electricity prices for example, was not taken into consideration. Additionally, the effects of using non-linear telescoping series to weight the cost were not explored, leaving room for further investigation.

2-2 Proposed Prediction Horizon-Extension Strategy

All approaches discussed so far rely on scaling the monthly peak demand to capture its impact within a short prediction horizon N . This horizon is simply limited, as the disturbance (ambient temperature) towards the end of the month remains unknown.

This thesis proposes an alternative strategy. Instead of restricting the optimization to this short horizon, the approach artificially extends it beyond N into a region of higher uncertainty regarding both the disturbance and day-ahead electricity prices. In this extended region, these unknown values are represented using historical averages. Since these averages remain static over time, the corresponding portion of the optimization can be pre-computed, thereby reducing computational complexity without compromising long-term planning capability.

2-2-1 Intuitive Motivation

The motivation for this approach becomes clear when examining the economic cost function to be optimized, as shown in Equation (2-1). This cost comprises two main components: the cost associated with day-ahead electricity prices and the cost resulting from the monthly peak demand charge. The structure of these costs makes it such that the optimization can be seen as having two tasks: short-term decision-making and long-term planning. Short-term decisions are driven by fluctuating day-ahead electricity prices, while long-term planning is required to manage the peak demand charge, which depends on the maximum power drawn during the entire billing period.

By artificially extending the prediction horizon to cover the full month and using historical averages for electricity prices and disturbances in the uncertain region, the optimizer is able to estimate an optimal peak demand early in the billing period based on the available historical




Figure 2-3: Half-monthly temperature distribution (2011–2022) and average temperatures.

data. This estimated peak demand can then be used as a reference at the very beginning of the month, and can therefore be leveraged immediately for the short term decision making.

This approach raises an important question: can we trust the optimized peak demand prediction? Figure 2-3 plots the mean outdoor temperature from 2011 to 2022 together with the distribution for each half-month interval. Although noticeable deviations from the long-term average appear, the building’s large thermal mass means indoor temperatures rise or fall only slowly. The MPC leverages this inertia: when outdoor conditions stray from the seasonal norm, a brief cold spell for example, the MPC is able to shift heating in time, smoothing the load and keeping the peak demand close to the optimized target, rather than letting short-lived extremes dictate the peak demand. Consequently, peak demand is determined more by the predictable, slowly varying seasonal pattern than by isolated weather events. Therefore, the optimization using historical averages is expected to give a reliable indication of the peak demand.

2-2-2 Research Questions

This research project focuses on simulating the proposed method and demonstrating that the portion of the optimization involving historical averages for the disturbance and electricity prices can be treated as static. This allows a significant part of the problem to be encapsulated within a pre-computed approximated terminal cost function, reducing computational complexity while retaining economic effectiveness.

The research is guided by the following questions:

- How can a terminal cost function be obtained and approximated to support economic decision-making within the MPC framework?
- How does the performance of the MPC with a terminal cost function compare economically to approaches found in existing literature?

Terminal Cost Approximation for Artificially Extended Horizon MPC

The following chapter discusses how a model predictive control (MPC) formulation, using a short prediction horizon (several days), can be equipped with historical data to capture the peak demand charge at the end of the month. In particular, we show how the optimization problem with the artificially extended prediction horizon can be decomposed into two parts: one that covers the short-term decisions within the short prediction horizon, and a second part that can be pre-computed and encapsulated as a terminal cost. This terminal cost accounts for the remaining future impact beyond the short horizon, especially the influence of the peak demand charge. Furthermore, the chapter outlines how the terminal cost can be approximated efficiently and implemented into the MPC problem. By representing it as a convex piecewise-linear function, the terminal cost can be incorporated into a linear program (LP) framework.

3-1 Artificially Extended Horizon MPC Formulation

3-1-1 Forecasting Limitations and Available Data

The proposed method in 2-2 aims to improve the capture of peak demand charges by artificially extending the prediction horizon beyond the point at which forecast uncertainty becomes significant. It is important to clarify the information available within the horizon where forecasts remain reasonably accurate. The disturbance of interest, ambient temperature (T_a), can typically be predicted with good accuracy several days in advance using modern numerical weather prediction models [23].

To reflect the performance of such state-of-the-art models, the forecasted temperature used in this work includes an imposed error, as defined in Equation 3-1:

$$T_{a,forecast}(k+h|k) = T_{a,real}(k+h) + T_{a,error}(k+h|k) \quad (3-1)$$

Here, h represent the time steps into future with respect to k . The evolution of the forecast error is modeled as a discrete-time stochastic process, as shown in Equation 3-2:

$$\begin{aligned} T_{a,error}(k+h|k) &= T_{a,error}(k+h-1|k) + \mathcal{N}(0, \sigma^2) \\ T_{a,error}(k|k) &= 0 \end{aligned} \quad (3-2)$$

As a result, the forecasted ambient temperature follows a normal distribution centered around the true future temperature, with variance increasing linearly over time. This is formalized in Equation 3-3:

$$T_{a,forecast}(k+h|k) \sim \mathcal{N}(T_{a,real}(k+h), h\sigma^2) \quad (3-3)$$

Using the relation from Equation 3-3 the confidence intervals for various forecast lengths can be calculated using $CI \pm z_{\alpha/2}\sigma\sqrt{h}$. Table 3-1 shows these confidence intervals for $\sigma = 0.2$. It can be observed that the prediction remains relatively accurate for short forecast lengths. However, when forecasting over an entire month, the variation becomes significantly larger.

Confidence	$z_{\alpha/2}$	CI 48h	CI 72h	CI 96h	CI 30d
50%	0.674	± 0.93	± 1.14	± 1.32	± 3.62
68%	1.000	± 1.39	± 1.70	± 1.96	± 5.37
80%	1.282	± 1.78	± 2.18	± 2.51	± 6.88
90%	1.645	± 2.28	± 2.79	± 3.22	± 8.83
95%	1.960	± 2.72	± 3.33	± 3.84	± 10.52
99%	2.576	± 3.57	± 4.37	± 5.05	± 13.82
99.9%	3.291	± 4.56	± 5.59	± 6.45	± 17.66

Table 3-1: Confidence intervals across horizons (48h–30d) at various confidence levels.

In addition to ambient temperature, the day-ahead electricity price plays a critical role in the decision-making process of the MPC controller. These day-ahead prices are published every day at 13:00. This means that the amount of available information varies depending on the time of day. Let us denote the amount of available information at time point t by $M(t)$. Figure 3-1 illustrates how the availability of price information $M(t)$ changes throughout the day.




Figure 3-1: Availability of day-ahead electricity price information throughout the day, in hours. Prices are published daily at 13:00, resulting in varying amounts of known price data depending on the time of day.

This variation highlights that depending on the current time, different time-spans of future price information is available. To better utilize the thermal mass of the building and address the peak demand charge, the aim is to artificially extend the prediction horizon beyond the range of accurate ambient temperature and electricity price forecasts. To enable this extension, historical averages are used to approximate both the ambient temperature (T_a) and the day-ahead electricity prices (c_e) beyond the reliable forecasts.

Starting with the ambient temperature (T_a), the following prediction method is proposed:

$T_{a,pred}$: Predicted ambient temperatures based on the average temperature observed at the same hour, day, and month from the years 2011 to 2022.

These historical ambient temperatures are obtained from measurements done by the KNMI at De Bilt, The Netherlands [24]. The period 2011 till 2022 is selected to represent current climate conditions and allows for validation of the MPC algorithm for 2023 and onward.

For the prediction of day ahead electricity prices (c_e), the following prediction method is used:

$c_{e,pred}$: Predicted day-ahead electricity prices based on the average price at the same hour over the past 30 days.

Note that unlike the ambient temperature prediction, the predicted day-ahead electricity prices use a shorter historical window of 30 days. This choice reflects that electricity prices are subject to market behavior and the rapidly changing energy landscape. Additionally, the decision is made to only differentiate between different hours of the day and not different days of the week. Although electricity prices typically exhibit different behaviors during various

parts of the week, such as weekdays and weekends, no distinction is made in this prediction approach. Accounting for these differences would require a larger historical dataset to produce representative averages, exceeding 30 days. However, extending the historical window risks including data that no longer reflects current market dynamics. To balance temporal relevance with statistical reliability, the prediction uses a 30-day window without differentiating between weekdays and weekends.

3-1-2 Artificially Extended Horizon Optimization Problem

Equipped with this information, we now have the components to formulate an MPC algorithm with an artificially extended horizon that considers the entire month. The objective function in Equation (3-4a) includes the economic cost incurred from any time point $t \in \{0, 1, \dots, T-1\}$ until the end of the month T , associated with operating the HVAC system under day-ahead electricity prices and a monthly peak demand charge.

Within the cost function (3-4a), slack variables (ζ_1, ζ_2) are included and heavily penalized through a large weight w . These slack variables allow for constraint violations when absolutely necessary. This is important in cases where hard constraints cannot be satisfied, for example, when ambient temperatures become excessively high and maintaining comfort constraints is no longer feasible.

The MPC algorithm follows a structure similar to the formulation in Equation (2-3), with the constraint on z now incorporating the peak demand used so far, denoted as u_{peak}^- . The complete optimization problem is presented in Equation (3-4).

$$\min_{\mathbf{u}, \mathbf{x}, \boldsymbol{\zeta}, z} \sum_{k=t}^{T-1} \left(c_e(k) \cdot u(k) \cdot \Delta + w^\top \boldsymbol{\zeta}(k) \right) + c_d \cdot z \quad (3-4a)$$

$$\text{s.t. } x(k+1) = A_d x(k) + B_d u(k) + B_d^d d(k), \quad \forall k = t, \dots, T-1 \quad (3-4b)$$

$$T_{\min, r}(k) - \zeta_1(k) \leq T_r(k) \leq T_{\max, r}(k) + \zeta_1(k), \quad \forall k = t+1, \dots, T \quad (3-4c)$$

$$\zeta_1(k) \geq 0, \quad \forall k = t+1, \dots, T \quad (3-4d)$$

$$T_w(k) - \zeta_2(k) \leq T_{\max, w}(k), \quad \forall k = t+1, \dots, T \quad (3-4e)$$

$$\zeta_2(k) \geq 0, \quad \forall k = t+1, \dots, T \quad (3-4f)$$

$$u_{\min} \leq u(k) \leq u_{\max}, \quad \forall k = t, \dots, T-1 \quad (3-4g)$$

$$z \geq u(k) - u_{\text{peak}}^-, \quad \forall k = t, \dots, T-1 \quad (3-4h)$$

The optimization problem is subject to the system dynamics in (3-4b). The room temperature constraints (3-4c) ensure that thermal comfort is maintained, while the water temperature constraints (3-4e) ensure operation within a region where a constant coefficient of performance (COP) can be assumed for the heat pump.

Used within the objective function (3-4a) and the system dynamics (3-4b) are the values $c_e(k)$ and $d(k)$, respectively. These values are entries from the vectors $\mathbf{c}_e$ and $\mathbf{d}$, which are defined in Equation (3-5):

$$\mathbf{c}_e = \begin{bmatrix} \mathbf{c}_e^* \\ \hat{\mathbf{c}}_e \end{bmatrix} \in \mathbb{R}^{T-t-1}, \quad \mathbf{d} = \begin{bmatrix} \mathbf{d}^* \\ \hat{\mathbf{d}} \end{bmatrix} \in \mathbb{R}^{T-t-1} \quad (3-5)$$

In Equation (3-5) the known electricity prices are denoted by $\mathbf{c}_e^* \in \mathbb{R}^{M(t)}$. While the predicted electricity prices $c_{e,pred}$ are represented by $\hat{\mathbf{c}}_e \in \mathbb{R}^{T-t-M(t)-1}$. The disturbance vector consists of the reliable temperature forecast $T_{a,forecast}$, stored in $\mathbf{d}^* \in \mathbb{R}^N$, and the disturbance prediction $T_{a,pred}$ is represented as $\hat{\mathbf{d}} \in \mathbb{R}^{T-t-N-1}$. Note that N corresponds to both the length of the vector $\mathbf{d}^*$ and the prediction horizon.

Throughout this thesis, multiple MPC formulations will be presented. Across these different formulations, the constraints (3-4c), (3-4d), (3-4e), (3-4f), and (3-4g) remain unchanged. Therefore, these constraints are grouped and defined once in Equation (3-6), allowing us to refer to them concisely and improving overall readability.

$$\begin{aligned} T_{\min,r}(k) - \zeta_1(k) &\leq T_r(k) \leq T_{\max,r}(k) + \zeta_1(k), \\ \zeta_1(k) &\geq 0, \\ T_w(k) - \zeta_2(k) &\leq T_{\max,w}(k), \\ \zeta_2(k) &\geq 0, \\ u_{\min} &\leq u(k) \leq u_{\max}, \end{aligned} \quad (3-6)$$

3-1-3 Problem Decomposition and Reformulation

The MPC algorithm (3-4) results in a large LP problem, even when applied to our relatively simple system, which features a small state space and a single input. While real-time computation is feasible for the current setup, the computational burden may increase significantly as the state and input spaces grow for more complex systems. Moreover, to address the research questions of this thesis, system performance must be evaluated over multiple simulated months. Since the ambient temperature forecasts are generated stochastically (3-3), each month is simulated across several scenarios to assess the MPC performance. Using Gurobi [25], a single month of simulation under 30 scenarios currently requires approximately 35 hours of computation time. This is inefficient for extensive testing and would become more prohibitive with increased model complexity.

To mitigate this computational burden, the optimization problem is decomposed into two stages. We begin by reformulating the original problem in Equation (3-4) by augmenting the state vector $x(k)$ with the peak demand variable $z(k)$. In this formulation and from now on, the peak demand variable is treated as a dynamic state, where $z(k)$ represents the maximum demand observed up to time k . The resulting augmented state-space formulation is presented in Equation (3-7).

$$\tilde{x}(k) = \begin{bmatrix} x(k) \\ z(k) \end{bmatrix} \Rightarrow \tilde{x}(k+1) = \begin{bmatrix} A_d x(k) + B_d u(k) + B_d^d d(k) \\ \max(u(k), z(k)) \end{bmatrix} \quad (3-7)$$

Using this augmented state formulation, the total cost at the end of the month, J_T , starting at $t = 0$, can be expressed as a function of the initial augmented state $\tilde{x}(0)$, the input sequence

$\mathbf{u}$, and the disturbance and electricity price vectors, $\mathbf{d}$ and $\mathbf{c}_e$, respectively. The cost is defined in Equation (3-8):

$$J_T(\tilde{x}(0), \mathbf{u}, \mathbf{d}, \mathbf{c}_e) := \sum_{k=0}^{T-1} \ell(\tilde{x}(k), u(k), d(k), c_e(k)) + \ell_T(\tilde{x}(T)) \quad (3-8)$$

s.t. (3-7)

In Equation (3-8) the stage cost contains the economic cost related to the electricity prices as well as penalties associated with soft constraint violations. This stage cost is defined in Equation (3-9):

$$\ell(\tilde{x}(k), u(k), d(k), c_e(k)) = c_e(k) \cdot u(k) + w^\top p(x(k), k) \quad (3-9)$$

Within Equation 3-9, $p(x(k), k)$ is a penalty vector capturing the extent to which comfort and operational constraints are violated. It is defined in Equation (3-10):

$$p(x(k), k) = \begin{bmatrix} \max(0, T_r(k) - T_{\max,r}(k)) \\ \max(0, T_{\min,r}(k) - T_r(k)) \\ \max(0, T_w(k) - T_{\max,w}) \end{bmatrix} \quad (3-10)$$

This vector $p(x(k), k)$ quantifies the deviation of the room temperature T_r from its comfort bounds $[T_{\min,r}, T_{\max,r}]$, and the water temperature T_w from its operational limit $T_{\max,w}$.

The terminal cost term $\ell_T(\tilde{x}(T))$ in Equation (3-8) represents the peak demand charge and is defined in Equation (3-11):

$$\ell_T(\tilde{x}(T)) = c_d \cdot z(T) \quad (3-11)$$

$J_T(\tilde{x}(0), \mathbf{u}, \mathbf{d}, \mathbf{c}_e)$ depends on the disturbance vector $\mathbf{d}$. As shown in Equation (3-5), this vector is constructed out of two other vectors $\mathbf{d}^*$ and $\hat{\mathbf{d}}$. As discussed in Section 3-1-2 the forecasted disturbance, which represent the true disturbance with added uncertainty, is contained within $\mathbf{d}^*$. Whereas $\hat{\mathbf{d}}$ contains a predicted disturbance vector based on historical averages.

Using the two components of vector $\mathbf{d}$, the total cost $J_T(\tilde{x}(0), \mathbf{u}, \mathbf{d}, \mathbf{c}_e)$, again starting at $t = 0$, can be decomposed as shown in Equation (3-12)

$$J_T(\tilde{x}(0), \mathbf{u}, \mathbf{d}, \mathbf{c}_e) = J_{[0,N]}(\tilde{x}(0), \mathbf{u}, \mathbf{d}^*, \mathbf{c}_e^{(N)}) + J_{[N,T]}(\tilde{x}(N), \mathbf{u}, \hat{\mathbf{d}}, \hat{\mathbf{c}}_e) \quad (3-12)$$

With the vector $\mathbf{d}^* \in \mathbb{R}^N$ having length N the cost from $t = 0$ to $t = N$ can be formulated as a function of this vector $\mathbf{d}^*$, as shown in Equation (3-13)

$$J_{[0,N]}(x(0), \mathbf{u}, \mathbf{d}^*, \mathbf{c}_e^{(N)}) := \sum_{k=0}^{N-1} \ell(\tilde{x}(k), u(k), d^*(k), c_e^{(N)}(k)) \quad (3-13)$$

s.t. (3-7)

Since N may exceed the length M of the known electricity price vector, the partial electricity price vector $\mathbf{c}_e^{(N)} \in \mathbb{R}^N$ is defined as a concatenation of the known part $\mathbf{c}_e^* \in \mathbb{R}^M$ and the part $\hat{\mathbf{c}}_e \in \mathbb{R}^{T-M}$ based on historical averages, such that:

$$\mathbf{c}_e^{(N)} = [\mathbf{c}_e^*, \hat{\mathbf{c}}_e(1), \hat{\mathbf{c}}_e(2), \dots, \hat{\mathbf{c}}_e(N-M)] \in \mathbb{R}^N, \quad (3-14)$$

where the vector $\hat{\mathbf{c}}_e$ provides the remaining $N-M$ entries needed to evaluate the cost $J_{[0,N]}$. Moreover, if $M \geq N$ this would simply imply that $\mathbf{c}_e^{(N)} = \mathbf{c}_e^*$.

The cost for the rest of the billing period from $t = N$ till $t = T$ is represented by Equation (3-15):

$$J_{[N,T]}(\tilde{x}(N), \mathbf{u}, \hat{\mathbf{d}}, \hat{\mathbf{c}}_e) := \sum_{k=N}^{T-1} \ell(\tilde{x}(k), u(k), \hat{\mathbf{d}}(k), \hat{\mathbf{c}}_e(k)) + \ell_T(\tilde{x}(T)) \quad (3-15)$$

s.t. (3-7)

With the decomposed cost functions shown in Equations (3-13) and (3-15), the MPC optimization problem can be reformulated. Since MPC uses a moving horizon, we generalize the initial time to t , rather than fixing it at $t = 0$. The resulting formulation is shown in Equation (3-16):

$$\min_{u(t), \dots, u(N-1) \in \mathcal{U}} \left\{ J_{[0,N]}(x(0), \mathbf{u}, \mathbf{d}^*, \mathbf{c}_e^{(N)}) + \min_{u(N), \dots, u(T-1) \in \mathcal{U}} \left\{ J_{[N,T]}(\tilde{x}(N), \mathbf{u}, \hat{\mathbf{d}}, \hat{\mathbf{c}}_e) \right\} \right\} \quad (3-16)$$

Here, $\mathcal{U}$ denotes the set of admissible control inputs, as defined in Equation (3-17):

$$\mathcal{U} := \{u(k) \in \mathbb{R} \mid u_{\min} \leq u(k) \leq u_{\max}\} \quad (3-17)$$

The formulation in Equation (3-16) reveals the structure of a two-level optimization: an outer minimization over the trusted forecast window, and an inner minimization over the uncertain future. Where the inner problem excludes all information considered accurate ($\mathbf{c}_e^*$) or reliable ($\mathbf{d}^*$). Moreover, since $\hat{\mathbf{d}}$ and $\hat{\mathbf{c}}_e$ are based on historical averages they are a function only of time. Therefore the inner minimization becomes a strong candidate to be represented as a precomputed terminal cost function, as this terminal cost would be static with respect to $\hat{\mathbf{d}}$ and $\hat{\mathbf{c}}_e$.

This terminal cost function could therefore be represented as shown in Equation (3-18):

$$V_f(\tilde{x}(t+N), t+N) = \min_{u(N), \dots, u(T-1) \in \mathcal{U}} \left\{ J_{[N,T]}(\tilde{x}(N), \mathbf{u}, \hat{\mathbf{d}}, \hat{\mathbf{c}}_e) \right\} \quad (3-18)$$

Meaning that the MPC formulation shown in Equation 3-16 can be displayed as optimizing over a stage cost and terminal cost function, as shown in Equation 3-19

$$\begin{aligned}
& \min_{u(t), \dots, u(N-1) \in \mathcal{U}} \left\{ \sum_{k=t}^{t+N-1} \ell(\tilde{x}(k), u(k), d^*(k), c_e^{(N)}(k)) + V_f(x(t+N), t+N) \right\} \\
& \text{s.t. } \tilde{x}(k+1) = \begin{bmatrix} A_d x(k) + B_d u(k) + B_d^d d(k) \\ \max(u(k), z(k)) \end{bmatrix}, \quad \forall k \in [t, T-1]
\end{aligned} \tag{3-19}$$

However, solving Equation (3-19) with the exact precomputed terminal cost from Equation (3-18) is computationally intractable due to the curse of dimensionality [26]. Specifically, the terminal cost would need to be evaluated for every possible combination of the terminal state $\tilde{x}(t+N)$ and time $t+N$. To address this, the terminal cost function will be approximated, as discussed subsequently.

3-2 Terminal Cost construction

To further motivate this terminal cost approximation, we will briefly discuss AlphaZero, which also takes advantage of terminal cost approximation in the context of an even more complex sequential decision-making problem. Developed by DeepMind, AlphaZero [27] is an algorithm designed to master the games of chess, shogi, and Go. At the time of its release in 2017, it was arguably the strongest engine across these board games. Therefore, AlphaZero is widely regarded as one of the most impressive achievements in reinforcement learning at the time of its release in 2017.

Without going into the specific details of the algorithm, AlphaZero uses reinforcement learning through self-play to train a neural network that outputs both a policy and a value estimate for game positions. In principle, this extensively trained neural network can play the board games independently. However, the learned policy is not used directly during online play, as it is not sufficiently accurate due to approximation errors introduced during the training of the neural network.




Figure 3-2: AlphaZero's online decision-making architecture, combining a look-ahead search tree with an offline-trained position evaluator, taken from [28]

Instead the online player combines both the use of a look-ahead tree with the offline trained position evaluator, as shown in Figure 3-2. Therefore, Bertsekas [28] draws parallels with AlphaZero and the MPC architecture. Bertsekas describes the general MPC scheme by the use of an ℓ -step minimization, as shown in Equation (3-20). It can be seen that this ℓ -step minimization reflects the MPC formulation from earlier shown in Equation (3-19), with $\alpha = 1$.

$$\min_{u_k, u_{k+1}, \dots, u_{k+\ell-1}} \left\{ \sum_{m=0}^{\ell-1} \alpha^m g(x_{k+m}, u_{k+m}) + \alpha^\ell \tilde{J}(x_{k+\ell}) \right\} \quad (3-20)$$

Similar to the scheme in Equation (3-20), Alphazero also uses an optimization over a rolling horizon with a terminal cost. This time the horizon is represented by the look-ahead tree and the terminal cost by the approximated value function given by the neural network. And it is precisely that combination that lead to the success of AlphaZero, thereby quoting Bertsekas: *"It is presently believed that the principal contributor to their success is long lookahead, which uses an efficient on-line play algorithm that involves various forms of tree pruning. The off-line trained position evaluator and player have also contributed to success, although likely to a lesser extent."*

What we can draw from this anecdote is that an approximated value function (or terminal cost) does not necessarily have to represent a perfect offline player in order to achieve strong performance within an ℓ -step minimization scheme. Returning to the topic of this thesis, the idea was to artificially extend the prediction horizon to a length T in order to estimate a peak demand that can immediately be used. Thereby we ended at the MPC formulation in Equation (3-19) where most part of the prediction horizon is now captured within this terminal cost function that needs to be approximated.

Therefore, we must ask ourselves what information the terminal cost function should contain in order to effectively support short-term decision making. It is unlikely that the terminal cost evaluated at the precise final state of the prediction horizon N is critical for determining short-term control actions. Instead, what matters more is the magnitude of the peak demand that occurs beyond the prediction horizon and the cost associated with that peak. Consequently, the terminal cost function is designed to depend primarily on the peak demand itself.

3-2-1 Perfect foresight Terminal Cost Approximation

In order to obtain the terminal cost function we must solve the optimization in equation (3-18). Doing so is similar to the optimization shown in Equation (3-4). We begin by assuming that the disturbance vector $\mathbf{d}$ and electricity price vector $\mathbf{c}_e$ are known over the entire horizon and reflect the true, realized values, rather than forecasts or estimates. This assumption allows us to isolate and analyze the impact of the approximations we will introduce. The resulting optimization, which defines the terminal cost function, is shown in Equation (3-21).

$$\begin{aligned}
\min_{\mathbf{u}, \mathbf{x}, \zeta, z(T)} \quad & \sum_{k=t+N}^{T-1} \left(\check{c}_e(k) \cdot u(k) \cdot \Delta + w^\top \zeta(k) \right) + c_d \cdot z(T) \\
\text{s.t.} \quad & x(k+1) = A_d x(k) + B_d u(k) + B_d^d \check{d}(k), \quad \forall k = t, \dots, T-1 \\
& (3-6), \quad \forall k = t, \dots, T-1 \\
& z(T) \geq z(t+N) \quad \forall k = t, \dots, T-1 \\
& x(t+N) = \begin{bmatrix} 20 & 20 & 30 \end{bmatrix}^T
\end{aligned} \tag{3-21}$$

The differences between the artificially extended horizon optimization in Equation (3-4) and the terminal cost optimization in Equation (3-21) lie primarily in the treatment of the disturbance and electricity price vectors. In the terminal cost optimization, these are represented by $\check{d}$ and $\check{c}_e$, which contain the actual (true) values. Additionally, the constraint on the peak demand variable $z(T)$ is different. In Equation (3-21), the constraint $z(T) \geq z(t+N)$ is introduced, which informs the optimizer that at time point $t+N$ we already have obtained a peak demand of $z(t+N)$. By solving the optimization problem (3-21) over a range of $z(t+N)$ values, the terminal cost can be obtained as a function of $z(t+N)$.

To make the computation of the terminal cost function tractable the terminal cost is approximated as shown in Equation (3-22).

$$V_f(z(t+N), \tau) \approx V_f(\tilde{x}(t+N), t+N) \tag{3-22}$$

This approximation is achieved by considering the following aspects:

1. Initial condition: only a single initial condition $x(t+N)$ is considered. This initial condition is set to $x(t+N) = \begin{bmatrix} 20 & 20 & 30 \end{bmatrix}^T$ meaning that the room temperature T_r is set to be in the middle of the comfort bounds shown in Figure 1-6. It might be more accurate to set the initial condition closer to the bottom of the comfort bounds as the economic model predictive control (EMPC) is likely to push the room temperature state to the bottom of the comfort bounds. However, by setting the initial state close to the comfort bound one can run the risk that the optimizer has to start with a high peak to prevent comfort bound violations, therefore the safe approach is setting the initial condition in the middle of the comfort bounds.

2. Time points $t+N$: We consider time points $t+N$ that correspond with midnight at the end of each day within the billing period. Let $\tau \in \mathbb{N}$ index the days of the billing period, and define time points $t+N$ such that:

$$t+N = \tau \cdot \frac{24 \cdot 60}{t_s},$$

where $t_s = 15$ minutes, is the sample time. This choice for τ was initially made because of how the length ($M(t)$) of the known day-ahead electricity price information varies throughout the day. Figure 3-1 shows that at each time point t the known information about the electricity prices ends at some point τ . The initial idea was to match the prediction horizon N according

to size of M at each time point t . However this turned out to be a too short of a prediction horizon resulting in comfort bound violations. Yet the time points τ were still retained at midnight, since the comfort bounds are less restrictive during nighttime hours.

3. range of $z(t + N)$ values: The range of $z(t + N)$ values is defined as:

$$\mathcal{Z}_{t+N} = \{10, 20, \dots, 140\} \cup \{150, 175, \dots, 475\} \cup \{500, 550, \dots, 1000\}$$

covering the interval from 10 to 1000 with progressively larger step sizes.

This non-uniform spacing is chosen to provide finer resolution at lower values of $z(t+N)$, where the optimal $z(T)$ tends to vary more sensitively if little inputs are required, for example during the less cold months of the year. At higher values of $z(T)$, the region around the optimum is less sensitive to changes in $z(T)$, allowing for coarser resolution without significantly impacting performance.

The optimization in Equation (3-21) is made into a linear program and solved with the defined approximations. The linear program can be found in Appendix A. Figure 3-3 shows the calculated terminal cost for the whole month of January in 2023. With the vertical axis indicating the economic cost and the horizontal axes indicating the $z(t + N)$ values and the indexes τ corresponding to individual days of the month.




Figure 3-3: Perfect foresight terminal cost function January 2023.

Inspecting the terminal costs for a given τ , we observe that the economic cost initially remains constant across a range of $z(t + N)$ values. At a certain point, highlighted in black, the economic cost begins to increase. This behavior arises because the terminal cost is defined as a function of the minimum (initial) peak demand $z(t + N)$. The flat region of the curve implies

that, for all $z(t+N)$ in that range except the point in black, the optimizer selects a higher value of peak demand $z(T)$ as optimal, which is than actually used within the optimization.

The terminal cost shown in Figure 3-3 corresponds to a cold month. Similarly, the terminal cost can also be computed for a milder month. Figure 3-4 presents the terminal cost for March 2023. As seen in the figure, the economic cost for each individual day τ initially remains constant up to the optimal point, indicated in black. As expected, the optimal points now occur at lower values of $z(T)$, since less heating is required in March compared to January.




Figure 3-4: Perfect foresight terminal cost function March 2023.

Moving to an even milder month, June 2023, the terminal cost plot shown in Figure 3-5 displays an almost linear increase in cost for any given day τ . This behavior is expected, as the month of July typically requires little to no heating. As a result, the terminal cost begins to resemble a linear function driven primarily by the peak demand charge term $c_d \cdot z(T)$, with minimal contribution from day-ahead electricity prices. Terminal cost plots for all other months in 2023 are provided in Appendix B-2.

Interestingly, when inspecting the optimal peak demand values $z(T)$, highlighted in black, we observe that for $\tau = 0$, the first day in the month, the optimal $z(T)$ is at its highest. For subsequent τ values, the optimal $z(T)$ gradually decreases. This behavior is illustrated in Figure 3-6, which shows the optimal $z(T)$ values for each index τ during January 2023. This decreasing optimal value for $z(T)$ as the month progresses confirms that it is indeed the start of the month where you would want to set the peak demand.




Figure 3-5: Perfect foresight terminal cost function June 2023.




Figure 3-6: Optimal peak demand values $z(T)$ (highlighted in black) for each day index τ during January 2023.

3-2-2 Integration of Terminal Cost into MPC Algorithm

Using the approximated terminal cost function $V_f(z(t+N), \tau)$, we can now solve the optimization problem for the MPC algorithm as formulated in Equation (3-19). Expanding Equation (3-19) leads to the optimization problem given in Equation (3-23).

$$\min_{u, x, \zeta, z(t+N)} \sum_{k=t}^{t+N-1} \left(\check{c}_e(k) \cdot u(k) \cdot \Delta + w^\top \zeta(k) \right) + V_f(z(t+N), \tau) \quad (3-23a)$$

$$\text{s.t. } x(k+1) = A_d x(k) + B_d u(k) + B_d^d \check{d}(k) \quad \forall k = t, \dots, T-1 \quad (3-23b)$$

$$(3-6), \quad \forall k = t, \dots, T-1 \quad (3-23c)$$

$$z(t+N) \geq u(k) \quad \forall k = t, \dots, T-1 \quad (3-23d)$$

$$z(t+N) \geq z(t) \quad (3-23e)$$

Using the terminal cost function $V_f(z(t+N), \tau)$ within the objective function in Equation (3-23b) implies that $z(t+N)$ no longer represents an incremental peak demand charge. Therefore, the constraints in Equations (3-23d) and (3-23e) are now applied to properly account for the peak demand. Moreover, the inclusion of $V_f(z(t+N), \tau)$ within the objective function must fit in the LP framework. Which is possible since $V_f(z(t+N), \tau)$ is guaranteed to be a convex piecewise-affine function.

$V_f(z(t+N), \tau)$ is guaranteed to represent a convex piecewise-affine function because the optimization problem in Equation (3-21) can be reformulated as a parametric linear program (PLP), as shown in Equation (3-24). Here, θ denotes the vector of parameters and contains a single non-zero entry corresponding to $z(t+N)$. The reformulation from Equation (3-21) to Equation (3-24) is provided in Appendix A-2.

$$\begin{aligned} \min_x \quad & c^T x \\ \text{s.t.} \quad & Gx \leq w + F\theta \end{aligned} \quad (3-24)$$

Solving this PLP has the property that the obtained terminal cost function is convex and piecewise affine [29]. This means that the terminal cost function can be represented by the maximum of affine functions, as shown in Equation (3-25):

$$V_f(z(t+N), \tau) = \max_{i \in \mathcal{I}_\tau} (a_i z(t+N) + b_i) \quad (3-25)$$

For each τ , there exists a set of affine functions that collectively represent the terminal cost as a function of $z(t+N)$. Figure 3-7 illustrates this for the case of January with $\tau = 1$, where the sampled terminal costs are shown as red points. The dashed lines indicate the affine segments connecting these points, and the blue line represents the maximum of this set of affine functions, resulting in the convex piecewise-affine terminal cost function.

Obtaining the coefficients a_i and b_i for every τ is done by selecting each pair of adjacent samples of the terminal cost, $(z_i(t+N), V_f(z_i(t+N), \tau))$ and $(z_{i+1}(t+N), V_f(z_{i+1}(t+N), \tau))$, and computing a_i and b_i as in (3-26).




Figure 3-7: Convex piecewise-affine representation of the terminal cost for January 2023 with $\tau = 1$.

$$a_i = \frac{V(z_{i+1}(t+N), \tau) - V(z_i(t+N), \tau)}{z_{i+1}(t+N) - z_i(t+N)}, \quad (3-26)$$

$$b_i = V(z_i(t+N), \tau) - a_i z_i(t+N)$$

Similar to the peak demand charge, the max operator over a set of affine functions cannot be directly represented in a standard LP. To address this, a new optimization variable v is introduced. By pairing v with the constraints defined in Equation (3-27), the terminal cost can be incorporated within the LP.

$$v \geq a_i z(t+N) + b_i, \quad \forall i \in \mathcal{I}_\tau. \quad (3-27)$$

Introducing the optimization variable v along with its associated constraint into the MPC algorithm results in the optimization problem shown in Equation (3-28). In this formulation, the terminal cost function in the objective (Equation (3-28a)) is replaced by the variable v .

$$\min_{\mathbf{u}, \mathbf{x}, \boldsymbol{\zeta}, v, z(t+N)} \sum_{k=t}^{t+N-1} \left(\check{c}_e(k) \cdot u(k) \cdot \Delta + w^\top \zeta(k) \right) + v \quad (3-28a)$$

$$\text{s.t. } x(k+1) = A_d x(k) + B_d u(k) + B_d^d \check{d}(k), \quad \forall k = t, \dots, T-1 \quad (3-28b)$$

$$(3-6), \quad \forall k = t, \dots, T-1 \quad (3-28c)$$

$$v \geq a_i z(t+N) + b_i, \quad \forall i \in \mathcal{I}_\tau \quad (3-28d)$$

$$z(t+N) \geq u(k), \quad \forall k = t, \dots, T-1 \quad (3-28e)$$

$$z(t+N) \geq z(t) \quad (3-28f)$$

Note that we are still using the vectors $\check{\mathbf{c}}_e$ and $\check{\mathbf{d}}$, which represent the true values of the electricity price and disturbance, respectively. These vectors are used to evaluate how the

MPC formulation in Equation (3-28) compares to a naive implementation of the demand charge, in terms of convergence toward the perfect-foresight optimal cost. In order to make this comparison, we first define how the perfect-foresight optimal cost is obtained.

This perfect-foresight optimal cost is obtained by solving the open-loop optimization problem defined in Equation (3-29), where a single optimization is performed over the entire month, starting at $t = 0$ and continuing up to T .

$$\min_{u, x, \zeta, z(T)} \sum_{k=0}^{T-1} \left(\check{c}_e(k) \cdot u(k) \cdot \Delta + w^\top \zeta(k) \right) + c_d \cdot z(T) \quad (3-29a)$$

$$\text{s.t. } x(k+1) = A_d x(k) + B_d u(k) + B_d^d \check{d}(k), \quad \forall k = 0, \dots, T-1 \quad (3-29b)$$

$$(3-6), \quad \forall k = 0, \dots, T-1 \quad (3-29c)$$

$$z(T) \geq u(k), \quad \forall k = 0, \dots, T-1 \quad (3-29d)$$

$$(3-29e)$$

Moreover, the naive implementation of the peak demand charge within the cost function is represented by the MPC scheme formulated in Equation 3-30, which now adopts a closed-loop form with a prediction horizon N .

$$\min_{u, x, \zeta, z(t+N)} \sum_{k=t}^{t+N-1} \left(\check{c}_e(k) \cdot u(k) \cdot \Delta + w^\top \zeta(k) \right) + c_d \cdot z(t+N) \quad (3-30a)$$

$$\text{s.t. } x(k+1) = A_d x(k) + B_d u(k) + B_d^d \check{d}(k) \quad \forall k = t, \dots, t+N-1 \quad (3-30b)$$

$$(3-6), \quad \forall k = t, \dots, t+N-1 \quad (3-30c)$$

$$z(t+N) \geq u(k) - z(t) \quad \forall k = t, \dots, t+N-1 \quad (3-30d)$$

$$z(t+N) \geq 0 \quad (3-30e)$$

Both the terminal cost (3-28) and naive (3-30) implementation of the peak demand charge within the MPC formulation use prediction horizon N . This prediction horizon is defined as shown in Equation (3-31):

$$N(t) = M(t) + N_{\text{ext}} \quad (3-31)$$

Since the terminal cost function $V_f(z(t+N), \tau)$ is defined only at discrete time points τ , the prediction horizon is constructed such that it always ends on a τ -aligned time step. As previously discussed, $M(t)$ is guaranteed to end on such a time step. Therefore, if the sampling rate would be one hour, selecting $N_{\text{ext}} \in \{24, 48, 72, \dots\}$ ensures that the total prediction horizon $N(t)$, as defined in Equation (3-31), also aligns with a valid τ . Although this alignment is not required for the MPC scheme using the naive peak demand charge implementation in Equation (3-30), we adopt it to allow a consistent comparison with the terminal cost-based formulation in Equation (3-28).

Both MPC schemes (3-28) and (3-30), as well as the open-loop optimization (3-29), are initialized at the initial state $x(0) = [20 \ 20 \ 30]^\top$. The simulation covers the month of January 2023. Figure 3-8 shows the bill values computed using (2-1) for both MPC schemes, plotted against N_{ext} , expressed in hours. The horizontal dashed line indicates the optimal cost obtained from the open-loop optimization (3-29).




Figure 3-8: Bill values for the two MPC schemes ((3-28),(3-30)) plotted against N_{ext} (in hours), with the optimal open-loop cost shown as a horizontal dashed line.

Figure 3-8 demonstrates that for $N_{\text{ext}} = 120$ hours, the bill value of MPC scheme (3-28) converges toward the optimal open-loop bill value, whereas MPC scheme (3-30) does not. Admittedly, this result is to be expected, as the terminal cost can be interpreted as artificially extending N to cover the entire month. Nevertheless, the convergence toward the optimal cost suggests that the terminal cost approximation is effective, and beyond $N_{\text{ext}} = 120$ hours the terminal temperature state has no significant impact on the MPC solution.

MPC Under Uncertainty

So far the simulated model predictive control (MPC) algorithm used a terminal cost based on perfect knowledge of the future disturbance and electricity price. In reality we do not have perfect knowledge. Therefore, in this chapter the MPC algorithm will be equipped with only the disturbance and electricity price information that would be available in an online setting. Moreover, the MPC algorithm with terminal cost will be compared to previously proposed MPC formulations from the literature.

4-1 Terminal Cost Approximation and Implementation Under Forecast Uncertainty

To start, the terminal cost $V_f(z(t+N), \tau)$ needs to be recalculated, now using vectors $\hat{\mathbf{d}}$ and $\hat{\mathbf{c}}_e$ as defined in Equation (3-5). This results in solving the problem in Equation (4-1) while using the same approximations as shown in Section 3-2-1

$$\min_{\mathbf{u}, \mathbf{x}, \zeta, z(T)} \sum_{k=t+N}^{T-1} \left(\hat{\mathbf{c}}_e(k) \cdot \mathbf{u}(k) \cdot \Delta + w^\top \zeta(k) \right) + c_d \cdot z(T) \quad (4-1a)$$

$$\text{s.t. } \mathbf{x}(k+1) = \mathbf{A}_d \mathbf{x}(k) + \mathbf{B}_d \mathbf{u}(k) + \mathbf{B}_d^d \hat{\mathbf{d}}(k), \quad \forall k = t, \dots, T-1 \quad (4-1b)$$

$$(3-6), \quad \forall k = t, \dots, T-1 \quad (4-1c)$$

$$z(T) \geq z(t+N), \quad \forall k = t, \dots, T-1 \quad (4-1d)$$

Solving Equation (4-1) for the month of January 2023 results in the terminal cost function shown in Figure 4-1.

Comparing the terminal cost shown in Figure 4-1 with the perfect-foresight terminal cost in Figure 3-3, we observe that the transitions between indices τ take on a flatter profile. This effect can be attributed to the construction of $\hat{\mathbf{d}}$ and $\hat{\mathbf{c}}_e$, which are based on historical averages




Figure 4-1: Historical average based terminal cost function January 2023.




Figure 4-2: Historical average based terminal cost function March 2023.

rather than actual values. Nevertheless, the pattern of optimal peak demands highlighted in black remains consistent, with the highest values occurring at the beginning of the month.

Similar to Section 3-2-1, the terminal cost functions for the milder month of March and the warmer month of June are also examined. Figure 4-2 shows the terminal cost for March 2023, which like January exhibits a flatter profile due to the use of historical averages, but also reflects a lower optimal peak demand compared to the colder month. As expected, the hotter month of June 2023, shown in Figure 4-3, displays an almost linear profile, similar to the perfect-foresight terminal cost in Figure 3-5. Terminal cost plots for all other months in 2023 are provided in Appendix B-3.




Figure 4-3: Historical average based terminal cost function June 2023.

Implementing the terminal cost into the MPC algorithm is again done using the maximum of affine functions as explained in Section 3-2-2. The resulting algorithm will be referred to as MPC-V, which is shown in Equation (4-2).

MPC-V:

$$\min_{\mathbf{u}, \mathbf{x}, \boldsymbol{\zeta}, v, z(t+N)} \sum_{k=t}^{t+N-1} \left(c_e^{(N)}(k) \cdot u(k) \cdot \Delta + w^\top \boldsymbol{\zeta}(k) \right) + v \quad (4-2a)$$

$$\text{s.t. } x(k+1) = A_d x(k) + B_d u(k) + B_d^d d^*(k), \quad \forall k = t, \dots, T-1 \quad (4-2b)$$

$$(3-6), \quad \forall k = t, \dots, T-1 \quad (4-2c)$$

$$v \geq a_i z(t+N) + b_i, \quad \forall i \in \mathcal{I}_\tau \quad (4-2d)$$

$$z(t+N) \geq u(k), \quad \forall k = t, \dots, T-1 \quad (4-2e)$$

$$z(t+N) \geq z(t) \quad (4-2f)$$

The MPC-V algorithm in Equation (4-2) now utilizes the electricity price vector $\mathbf{c}^{(N)}$, which contains the accurate price information as defined in Equation (3-14). In addition, the reliable forecast $\mathbf{d}^*$ is also incorporated into the formulation.

4-2 Alternative Short-Horizon MPC Methods

Having formulated the MPC-V algorithm in Equation (4-2), which uses a terminal cost approximation to effectively extend the prediction horizon, we are now in a position to compare this approach with alternative formulations introduced in Section 2-1-3.

The first formulation is based on a time-dependent scaling factor $\sigma_t := \min\{(T - t)/N, 1\}$, as proposed by Kumar et al. [21], which increases the weight on the peak demand as the end of the billing period approaches. Incorporating this scaling factor results in the formulation presented in Equation (4-3), referred to as MPC-1.

MPC-1:

$$\min_{\mathbf{u}, \mathbf{x}, \boldsymbol{\zeta}, z(t+N)} \sum_{k=t}^{t+N-1} \left(c_e^{(N)}(k) \cdot u(k) \cdot \Delta + w^\top \boldsymbol{\zeta}(k) \right) + \frac{c_d}{\sigma_t} \cdot z(t+N) \quad (4-3a)$$

$$\text{s.t. } x(k+1) = A_d x(k) + B_d u(k) + B_d^d d^*(k) \quad \forall k = t, \dots, t+N-1 \quad (4-3b)$$

$$(3-6), \quad \forall k = t, \dots, t+N-1 \quad (4-3c)$$

$$z(t+N) \geq u(k) - z(t) \quad \forall k = t, \dots, t+N-1 \quad (4-3d)$$

$$z(t+N) \geq 0 \quad (4-3e)$$

The second formulation is based on the approach proposed by Risbeck et al. [22], which introduces a telescoping series to scale the peak demand charge over time. The key idea is to distribute the peak demand cost across the month in such a way that the cumulative sum of these scaled charges equals the true financial cost of the peak demand charge. The telescoping series used to scale the peak demand charge is defined in Equation (4-4).

$$a(t) = \left(\frac{t}{T} \right)^2 \quad (4-4)$$

A quadratic telescoping series is used to promote the use of peak demand earlier in the month. Incorporating this scaling approach results in the formulation referred to as MPC-2, which is presented in Equation (4-5).

MPC-2

$$\min_{\mathbf{u}, \mathbf{x}, \boldsymbol{\zeta}, z(t+N)} \sum_{k=t}^{t+N-1} \left(c_e^{(N)}(k) \cdot u(k) \cdot \Delta + w^\top \boldsymbol{\zeta}(k) \right) + a(t+N) \cdot c_d \cdot z(t+N) \quad (4-5a)$$

$$\text{s.t. } x(k+1) = A_d x(k) + B_d u(k) + B_d^d d^*(k) \quad \forall k = t, \dots, t+N-1 \quad (4-5b)$$

$$(3-6), \quad \forall k = t, \dots, t+N-1 \quad (4-5c)$$

$$z(t+N) \geq u(k) - z(t) \quad \forall k = t, \dots, t+N-1 \quad (4-5d)$$

$$z(T+N) \geq 0 \quad (4-5e)$$

The third and final formulation applies a simple scaling of the peak demand charge based on the ratio between the prediction horizon and the length of the billing period, given by $\frac{N}{T}$. This approach is referred to as MPC-3, and is shown in Equation (4-6).

MPC-3:

$$\min_{u, x, \zeta, z(t+N)} \sum_{k=t}^{t+N-1} \left(c_e^{(N)}(k) \cdot u(k) \cdot \Delta + w^\top \zeta(k) \right) + \frac{N}{T} \cdot c_d \cdot z(t+N) \quad (4-6a)$$

$$\text{s.t. } x(k+1) = A_d x(k) + B_d u(k) + B_d^d d^*(k) \quad \forall k = t, \dots, t+N-1 \quad (4-6b)$$

$$(3-6), \quad \forall k = t, \dots, t+N-1 \quad (4-6c)$$

$$z(t+N) \geq u(k) - z(t) \quad \forall k = t, \dots, t+N-1 \quad (4-6d)$$

$$z(t+N) \geq 0 \quad (4-6e)$$

Each MPC method utilizes the same time-varying prediction horizon N , as defined in Equation (3-31). Furthermore, at the start of each month ($t = 0$), the MPC is initialized at the states $x(0) = [20 \ 20 \ 30]^\top$. In practice, the initial state at the beginning of each month would be determined by the final state of the previous month. Depending on the season, this initial state would likely vary, with low values for (T_r, T_f) paired with a high value for T_w , or vice versa. For simplicity only a single initial state is considered. While $T_w = 30$ may represent a relatively high initial value during warmer months, it is not expected to impact the performance of the MPC methods relative to each other.

4-3 Performance Evaluation and Results

Each MPC algorithm is simulated over the entire year of 2023 for various prediction horizon lengths, given by $N_{\text{ext}} \in \{48, 72, 96\}$. Since d^* is determined stochastically (see Equation (3-3)), each algorithm is run 30 times per month. The resulting monthly bill values (Equation (2-1)) are averaged and presented in Figure 4-4.




Figure 4-4: Average monthly electricity bills for each MPC algorithm across different prediction horizon lengths ($N_{\text{ext}} \in \{48, 72, 96\}$) over the year 2023.

Figure 4-4 is divided into groups of columns. The single column on the far left represents the bill value from the perfect-foresight open-loop optimization (Equation (3-29)), which serves as a theoretical lower bound on the achievable cost under perfect information. The remaining groups correspond to the different MPC methods, where each column within a group represents a different prediction horizon length N_{ext} .

The average bill values indicate that MPC-V and MPC-1 perform similarly in terms of cost, while MPC-2 and MPC-3 yield significantly higher bills. The performance gap between MPC-1 and MPC-2/MPC-3 may be attributed to the magnitude of the scaling factors employed in these methods. Specifically, MPC-2 and MPC-3 use scaling factors smaller than 1, effectively reducing the weight of the peak demand charge in the cost function. In contrast, MPC-1 applies a scaling factor that remains at 1 for most of the billing period and only increases towards the end.

Moreover, Figure 4-4 shows that increasing the prediction horizon does not necessarily lead to a lower average bill, particularly for the better-performing methods MPC-V and MPC-1. The non-decreasing trend in average bill values highlights how the growing uncertainty in $\mathbf{d}^*$ and $\mathbf{c}_e^{(N)}$ becomes a limiting factor in further reducing costs. This increasing uncertainty is particularly significant due to the slow dynamics of the building, as incorrect forecasts of $\mathbf{d}^*$ and $\mathbf{c}_e^{(N)}$ can steer the system state in the wrong direction, potentially requiring a large corrective action to maintain comfort. Such corrective actions may result in high peak demands, which would lead to elevated peak demand charges, further increasing the overall cost.

Besides the average bill cost, it is also of interest to examine the distribution of simulation results within each month. Figure 4-5 presents the results for the cold month of January 2023. The data is grouped by prediction horizon, and within each group, every MPC method is represented by a box plot.




Figure 4-5: Distribution of total bill cost simulations for January 2023, grouped by prediction horizon and MPC method.

Figure 4-5 shows that for January 2023, MPC-V outperforms MPC-1. However, this result does not necessarily generalize across all months. In many cases, MPC-V performs only slightly and not significantly better than MPC-1, as illustrated by the results for March and June 2023, shown in Figures 4-6 and 4-7, respectively.

With respect to increasing the prediction horizon and consequently the uncertainty in $\mathbf{d}^*$ and $\mathbf{c}_e^{(N)}$, Figures 4-5, 4-6, and 4-7 indicate that the different MPC methods exhibit similar behavior as uncertainty grows. The complete set of box plots for the remaining months of 2023 is provided in Appendix B-4.

Interestingly, during the warmer months June through September 2023, the average bill values across all methods are relatively similar, with the exception of July. In this month, methods MPC-2 and MPC-3 perform significantly worse. To investigate this behavior in more detail, a single simulation for each method with $N_{\text{ext}} = 48$ is shown in Figures 4-8, 4-9, 4-10, and 4-11.

Each of Figures 4-8, 4-9, 4-10, and 4-11 contains two subplots. The top subplot displays the evolution of the system states over time, with the comfort bounds indicated by dashed lines. Additionally, the real disturbance $\tilde{\mathbf{d}}$, representing the ambient temperature T_a , is shown in this plot. The bottom subplot presents the control input u alongside the corresponding day-ahead electricity prices.




Figure 4-6: Distribution of total bill cost simulations for March 2023, grouped by prediction horizon and MPC method.




Figure 4-7: Distribution of total bill cost simulations for June 2023, grouped by prediction horizon and MPC method.



Figure 4-8: Simulation results for MPC-V during July 2023 with $N_{\text{ext}} = 48$.




Figure 4-9: Simulation results for MPC-1 during July 2023 with $N_{\text{ext}} = 48$.




Figure 4-10: Simulation results for MPC-2 during July 2023 with $N_{\text{ext}} = 48$.




Figure 4-11: Simulation results for MPC-3 during July 2023 with $N_{\text{ext}} = 48$.

At the beginning of the month there is a period during which the electricity price becomes strongly negative. While MPC-V and MPC-1 do not respond to this price signal, MPC-2 and MPC-3 do. This behavior occurs because the cost associated with the peak demand charge is substantially scaled down at that point in time for MPC-2 and MPC-3. As a result, these methods interpret the negative price as an opportunity to profit by taking on a large peak demand. However, in reality the economic cost related to the peak demand charge does not allow for profit to be made while using such a large peak demand.

This example raises the question of whether blindly combining day-ahead prices with peak demand charges is truly the most effective strategy for grid operators to reduce stress on the grid. A negative price signal typically indicates a surplus of energy that needs to be absorbed. However, a peak demand charge can discourage consumption at these times, thereby limiting the grid's ability to offload excess energy. This highlights an interesting research direction not explored in this thesis: what electricity pricing structures are most optimal for grid operation?

Another interesting observation is that the control inputs generated by MPC-V and MPC-1 are identical, as shown in Figures 4-8 and 4-9. This behavior can be explained by examining the terminal cost plot for July, provided in Appendix B-3, Figure B-20. It reveals that the terminal cost function for July is linear with proportional constant c_d . As a result, the terminal cost in the MPC-V objective simplifies to $V_f(z(t+N), \tau) = c_d \cdot z(t+N)$, which during the period when the MPC controllers were active, is equivalent to the objective function used by MPC-1.

However, as shown in Figures 4-1 and 4-2, the colder months do not represent a linear terminal cost function $V_f(z, \tau)$. Despite this, the resulting average costs for MPC-V and MPC-1 remain closely matched. To investigate this behavior, we analyze a single simulation for February 2023 with a prediction horizon of $N_{\text{ext}} = 48$ h for both methods. The corresponding results are shown in Figures 4-12 and 4-13.

Figure 4-12 shows the simulation using MPC-V. It can be observed that the controller immediately sets the peak demand, which reaches a value of 375 W. The resulting bill for this simulation is 10.41 Euro. The simulation using MPC-1 is shown in Figure 4-13, where the controller gradually builds up the peak demand, eventually reaching the peak demand of 270 W, with a resulting bill of 10.58 Euro.

Interestingly, the controllers adopt different strategies yet result in similar monthly bills. A comparison of their control inputs with the perfect-foresight solution, formulated in Equation (3-29), is shown in Figure 4-14, which shows the theoretically optimal control trajectory. MPC-V demonstrates its ability to immediately leverage the peak demand, similar to how the perfect-foresight controller immediately uses its peak demand. In contrast, MPC-1 stands out in its ability to obtain a peak demand that is close to the optimal peak demand given by the perfect foresight controller. However, attaining this optimal peak demand does not mean that MPC-1 starts to act like the perfect foresight controller as MPC-1 is still subject to uncertainty in $\mathbf{b}^*$ and $\mathbf{c}^{(N)}$.

Because February 2023 was relatively warm compared to the historical average (Appendix B-1-1), the estimated peak demand in MPC-V turned out to be higher than the optimal peak demand determined by the perfect-foresight simulation. However, since MPC-1 required approximately 10 days to approach the optimal peak demand, it lost out on potential cost savings during that initial period. Consequently, MPC-1 achieved a lower peak demand than




Figure 4-12: Simulation results for MPC-V during February 2023 with $N_{\text{ext}} = 48$.




Figure 4-13: Simulation results for MPC-1 during February 2023 with $N_{\text{ext}} = 48$.

MPC-V, which helped offset the higher time-of-use (TOU) charges MPC-1 incurred, leading to a final bill value in a similar range for both methods.




Figure 4-14: Simulation results for perfect-foresight controller during February 2023

Another month worth considering is April 2023, which experienced ambient temperatures below the historical averages, as shown in Appendix B-1-1. Moreover, April also exhibited frequent negative pricing signals. A single simulation with $N_{\text{ext}} = 48$ h was conducted for both methods, MPC-V and MPC-1, with the results shown in Figure 4-15 and Figure 4-16, respectively.

MPC-1 starts with a lower peak demand than MPC-V, but after a few days, both methods attain peak demand levels that are quite similar. After April 10th, more negative price signals begin to appear. It can be observed that MPC-V increases its peak demand more than MPC-1. This occurs because, at that point in time, the increase in peak demand incurs a lower cost in $V_f(z(t+N), \tau)$ compared to $c_d \cdot z(t+N)$.

However, when analyzing the attained bill values and corresponding peak demands, we observe that MPC-V results in a bill of 5.52 Euro with a peak demand of 562 W, while MPC-1 achieves a lower bill of 5.21 Euro with a peak demand of 247 W. For comparison, the peak demand attained by the perfect-foresight controller was 439 W. Although MPC-V reached a peak demand closer to that of the perfect-foresight controller, its total bill ended up being slightly higher than that of MPC-1.

The fact that the bill for MPC-V is higher is noteworthy, as one might expect that a higher peak demand, closer to the optimal value, would better leverage the TOU prices and result in a lower bill. The reason this lower bill was not achieved could be attributed to the uncertainty




Figure 4-15: Simulation results for MPC-V during April 2023 with $N_{\text{ext}} = 48$.




Figure 4-16: Simulation results for MPC-1 during April 2023 with $N_{\text{ext}} = 48$.

present in $\mathbf{b}^*$ and $\mathbf{c}^{(N)}$. In particular, 'costly wrong' decisions resulting from this uncertainty are likely to be amplified when operating at higher peak demand levels.

Assessing MPC-2 and MPC-3 for April 2023 under the same conditions results in a bill of 7.03 Euro with a peak demand of 1351 W for MPC-2, while MPC-3 attains a bill of 8.24 Euro with a peak demand of 1333 W. Figure 4-17 and Figure 4-18 show the simulation results for MPC-2 and MPC-3, respectively.




Figure 4-17: Simulation results for MPC-2 during April 2023 with $N_{\text{ext}} = 48$.

Figure 4-17 shows how MPC-2, similar to its behavior in July 2023, overreacts to a negative price signal, resulting in a large peak demand. Moreover, MPC-3, as shown in Figure 4-18, does not overreact to the same price signal, but instead responds later in the month when an even more negative price signal appears. This indicates that both methods are susceptible to negative price signals, resulting in high peak demands.




Figure 4-18: Simulation results for MPC-3 during April 2023 with $N_{\text{ext}} = 48$.

Conclusion and Discussion

This thesis project explored the use of model predictive control (MPC) to economically control an HVAC system under day-ahead electricity prices combined with a monthly peak demand charge. In doing so, it proposed a approach for solving the optimization problem in a way that captures the full billing period. The research was guided by the following questions:

- How can a terminal cost function be obtained and approximated to support economic decision-making within the MPC framework?

To answer this question, the MPC optimization over the full billing period was first formulated. It was shown how the unknown disturbance and electricity price data could be substituted with averages derived from historical data. Subsequently, the artificially extended horizon optimization was decomposed into an outer minimization over the trusted forecast window and an inner minimization over the uncertain future.

As a result, the inner minimization could be approximated offline by a terminal cost function. To make the computation of this terminal cost tractable, it was necessary to approximate it. The first approximation was designed to ensure that the terminal cost would depend primarily on the peak demand. This was achieved by evaluating the terminal cost using a single initial state at a limited number of time points within the billing period.

The terminal cost function was initially obtained using perfect foresight, allowing for comparison with a perfect-foresight open-loop optimization. To implement the approximated terminal cost in the MPC formulation, we made use of the fact that the optimization used to derive the terminal cost corresponds to a parametric linear program (PLP), with the peak demand as a parameter. By sampling this optimization over a range of peak demand values, a convex piecewise affine function could be constructed. This function, represented as the maximum over a set of affine functions, could then be incorporated into the MPC formulation.

After integrating the terminal cost into the MPC, it was shown that with an increasing prediction horizon, the controller's performance quickly converged to the optimal open-loop cost. Which indicated that the terminal cost approximation was effective while given perfect information on the electricity price and ambient temperature.

- How does the performance of the MPC with a terminal cost function compare economically to approaches found in existing literature?

To evaluate the performance of the MPC with a terminal cost under realistic conditions, the terminal cost was recalculated without perfect foresight, using only information that would actually be available in practice, such as historical averages. A comparison was then made between the terminal-cost-based formulation (MPC-V) and alternative approaches that apply a scaling factor to the peak demand charge (MPC-1, MPC-2, and MPC-3). Each method was simulated for every month of 2023, with 30 simulations conducted per month to account for forecast uncertainty.

Based on the average cost and the distribution of individual simulations, MPC-2 and MPC-3 were found to perform the worst. This is likely due to the scaling factors used in these methods, which are less than 1 and therefore reduce the weight of the peak demand charge in the cost function. As a result, the controllers tend to overreact to low or negative price signals, leading to high peak demand charges. In contrast, the better-performing methods, MPC-V and MPC-1, show similar performance, both in terms of average cost and variability across simulations. Interestingly, despite their similar outcomes, the control strategies of MPC-V and MPC-1 was different in the cooler months: while MPC-V is more inclined to immediately set its peak demand, MPC-1 gradually built up to its peak over time.

A comparison of the control inputs from MPC-V and MPC-1 with the perfect-foresight open-loop inputs reveals where each method loses out on potential cost savings. During February 2023 where the ambient temperature is hotter than the historical average MPC-V adopts a strategy that mirrors the optimal controller's immediate use of peak demand. However, the initial peak it selects is higher than the truly optimal value, leading to an elevated monthly bill. Conversely, MPC-1 follows a different trajectory: it builds up to its peak demand gradually. Although the eventual peak demand is close to the optimal value, the prolonged build up means that the controller loses out on cost-saving opportunities early in the billing period.

However, the strategies employed during February by MPC-V and MPC-1 do not necessarily translate over to all months, as demonstrated by the results for April 2023. This month was colder than the historical average and featured a high amount of negative price signals. As a result, the optimal peak demand under perfect foresight was relatively high. Due to its conservative nature, MPC-1 achieved a peak demand further from the optimal value compared to peak demand achieved by MPC-V. Nevertheless, even though MPC-V attained a peak demand closer to the perfect-foresight benchmark, the average cost did not show a significant improvement over MPC-1. This outcome may be attributed to inaccuracies in the estimates of $\mathbf{b}^*$ and $\mathbf{c}^{(N)}$, which can lead to suboptimal control actions. These inaccuracies are likely to be amplified when operating at higher peak demand levels, diminishing results.

The main takeaway is that, under uncertainty, it is generally best to adopt a conservative approach with respect to the peak demand charge. This is demonstrated by MPC-V and MPC-1, both of which follow conservative strategies, in contrast to MPC-2 and MPC-3. However, conservatism with respect to the peak demand charge walks a fine line: on one hand, staying too far below the optimal peak demand may limit the amount of leverage one has on time-of-use (TOU) pricing; on the other hand, allowing the peak demand to become too high can amplify the impact of suboptimal decisions resulting from inaccuracies in $\mathbf{b}^*$ and $\mathbf{c}^{(N)}$. Given that MPC-1 and MPC-V perform similarly in terms of the final electricity bill,

and that MPC-1 is the most conservative in handling the peak demand charge, this raises the question of whether a slightly more conservative version of MPC-V—though not as restrictive as MPC-1, could result in gaining an edge over MPC-1 in terms of economic performance.

5-1 Recommendations for Future Work

A slightly more conservative variant of MPC-V could be achieved by modifying the historical prediction of ambient temperature, $\hat{\mathbf{d}}$. Rather than using the historical average, one could instead adopt, for example, the upper quartile of historical temperature data. This adjustment would lead the terminal cost function to favor a lower peak demand, thereby inducing a more conservative control strategy.

Another aspect worth exploring is how the amount of thermal mass in the building affects simulation results. Thermal mass is beneficial in that it enables load shifting to take advantage of TOU pricing. However, increasing thermal mass also slows down the building's thermal dynamics, thereby requiring a longer prediction horizon to capture all relevant effects. A longer horizon, in turn, introduces greater uncertainty in $\mathbf{b}^*$ and $\mathbf{c}^{(N)}$. Therefore, investigating the optimal amount of thermal mass for economic performance could be an interesting direction for future research.

Another important aspect related to the building model is how the performance of the MPC is affected when not all states are observable. For instance, this project assumed access to the floor temperature state. However, in practice, such a state may not be directly measurable, necessitating some form of state estimation to model the dynamics.

An additional direction worth exploring is how different values of the peak demand charge coefficient, c_d , influence economic performance. This project considered a peak demand charge of 5 €/kW. A valuable extension would be to compare the methods under varying peak demand charge levels to assess their performance relative to different peak demand charges.

Moreover, regarding the peak demand charge, this project sampled the terminal cost function using values from $\mathcal{Z}_{t+N}$. However, the terminal cost plots indicate that the values at the end of $\mathcal{Z}_{t+N}$ begin to exhibit a linear increase in the terminal cost. This suggests that, once the terminal cost starts increasing linearly, the remaining values in $\mathcal{Z}_{t+N}$ could be discarded, thereby reducing the computational burden associated with evaluating the terminal cost.

Appendix A

LP formulations

A-1 Linear MPC formulations

The MPC formulations are rewritten to the form shown in Equation (A-1) such that the optimizations can be solved using Gurobi

$$\min_{\boldsymbol{\eta}} f^{\top} \boldsymbol{\eta} \quad \text{such that} \quad \begin{cases} E_{eq} \cdot \boldsymbol{\eta} = e_{eq}, \\ E_{in} \cdot \boldsymbol{\eta} \leq e_{in}, \\ lb \leq \boldsymbol{\eta} \leq ub. \end{cases} \quad (\text{A-1})$$

Within the formulations n refers to the dimension of the state vector and m refers to the dimension of the input vector.

A-1-1 Equation (3-21)

$$\boldsymbol{\eta} = \begin{bmatrix} \boldsymbol{u} \\ \boldsymbol{x} \\ \boldsymbol{\zeta} \\ \boldsymbol{z} \end{bmatrix} = \begin{bmatrix} u(0) \\ u(1) \\ \vdots \\ u(N-1) \\ x(1) \\ x(2) \\ \vdots \\ x(N) \\ \zeta(0) \\ \zeta(1) \\ \vdots \\ \zeta(N-1) \\ z(T) \end{bmatrix} \in \mathbb{R}^{N(m+n+2)+1}, \quad f = \begin{bmatrix} \check{c}_e(0) \cdot dt \\ \check{c}_e(1) \cdot dt \\ \vdots \\ \check{c}_e(N-1) \cdot dt \\ 0 \\ 0 \\ \vdots \\ 0 \\ w \\ w \\ \vdots \\ w \\ c_d \end{bmatrix} \in \mathbb{R}^{N(m+n+2)+1} \quad (\text{A-2})$$

$$E_{eq} = \begin{bmatrix} -B_d & 0 & 0 & \cdots & 0 & I_n & 0 & 0 & \cdots & 0 & 0 & \cdots & 0 \\ 0 & -B_d & 0 & \cdots & 0 & -A_d & I_n & 0 & \cdots & 0 & 0 & \cdots & 0 \\ 0 & 0 & -B_d & \cdots & 0 & 0 & -A_d & I_n & \cdots & 0 & 0 & \cdots & 0 \\ \vdots & \vdots & \vdots & \ddots & \vdots & \vdots & \vdots & \vdots & \ddots & \vdots & \vdots & \cdots & 0 \\ 0 & 0 & 0 & \cdots & -B_d & 0 & 0 & 0 & \cdots & -A_d & I_n & \cdots & 0 \end{bmatrix} \in \mathbb{R}^{Nn \times N(m+n+2)+1} \quad (\text{A-3})$$

$$e_{eq} = \begin{bmatrix} B_d^d \check{d}(0) + A_d x(0) \\ B_d^d \check{d}(1) \\ B_d^d \check{d}(2) \\ \vdots \\ B_d^d \check{d}(N-1) \end{bmatrix} \in \mathbb{R}^{Nn} \quad (\text{A-4})$$

$$E_{in} = \begin{bmatrix} 0 & -I_{x,1} & -I_{\zeta,1} & 0 \\ 0 & I_{x,1} & -I_{\zeta,1} & 0 \\ 0 & I_{x,3} & -I_{\zeta,2} & 0 \\ 0 & 0 & 0 & -1 \end{bmatrix} \in \mathbb{R}^{3N+1 \times N(m+n+2)+1}, \quad e_{in} = \begin{bmatrix} -T_{\min,r} \\ T_{\max,r} \\ T_{\max,w} \\ -z(t+N) \end{bmatrix} \in \mathbb{R}^{3N+1} \quad (\text{A-5})$$

Where $T_{\min,r}$, $T_{\max,r}$, $T_{\max,w}$ are vectors in $\mathbb{R}^N$. Moreover, $I_{x,i}$ and $I_{\zeta,i}$ represent selection matrices:

$$I_{x,1} = \begin{bmatrix} 1 & 0 & 0 & 0 & \cdots & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 & \cdots & 0 & 0 & 0 & 0 \\ \vdots & \vdots & \vdots & \vdots & \ddots & \vdots & \vdots & \vdots & \vdots \\ 0 & 0 & 0 & 0 & \cdots & 0 & 1 & 0 & 0 \end{bmatrix} \in \mathbb{R}^{N \times Nn}, \quad I_{\zeta,1} = \begin{bmatrix} 1 & 0 & 0 & \cdots & 0 & 0 & 0 & 0 \\ 0 & 0 & 1 & \cdots & 0 & 0 & 0 & 0 \\ \vdots & \vdots & \vdots & \vdots & \ddots & \vdots & \vdots & \vdots \\ 0 & 0 & 0 & 0 & \cdots & 0 & 1 & 0 \end{bmatrix} \in \mathbb{R}^{N \times 2N} \quad (\text{A-6})$$

$$I_{x,3} = \begin{bmatrix} 0 & 0 & 1 & 0 & 0 & 0 & \cdots & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 1 & \cdots & 0 & 0 \\ \vdots & \vdots & \vdots & \vdots & \vdots & \vdots & \ddots & \vdots & \vdots \\ 0 & 0 & 0 & 0 & 0 & 0 & \cdots & 0 & 1 \end{bmatrix} \in \mathbb{R}^{N \times Nn}, \quad I_{\zeta,2} = \begin{bmatrix} 0 & 1 & 0 & 0 & \cdots & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 & \cdots & 0 & 0 & 0 \\ \vdots & \vdots & \vdots & \vdots & \ddots & \vdots & \vdots & \vdots \\ 0 & 0 & 0 & 0 & \cdots & 0 & 0 & 1 \end{bmatrix} \in \mathbb{R}^{N \times 2N} \quad (\text{A-7})$$

$$lb = \begin{bmatrix} u_{\min} \\ -\infty \\ 0 \\ 0 \end{bmatrix} \in \mathbb{R}^{N(m+n+2)+1}, \quad ub = \begin{bmatrix} u_{\max} \\ \infty \\ \infty \\ \infty \end{bmatrix} \in \mathbb{R}^{N(m+n+2)+1} \quad (\text{A-8})$$

A-1-2 Equation (3-28)

$$\boldsymbol{\eta} = \begin{bmatrix} \mathbf{u} \\ \mathbf{x} \\ \boldsymbol{\zeta} \\ v \\ z \end{bmatrix} = \begin{bmatrix} u(0) \\ u(1) \\ \vdots \\ u(N-1) \\ x(1) \\ x(2) \\ \vdots \\ x(N) \\ \zeta(0) \\ \zeta(1) \\ \vdots \\ \zeta(N-1) \\ v \\ z(t+N) \end{bmatrix} \in \mathbb{R}^{N(m+n+2)+2}, \quad \mathbf{f} = \begin{bmatrix} \check{c}_e(0) \cdot dt \\ \check{c}_e(1) \cdot dt \\ \vdots \\ \check{c}_e(N-1) \cdot dt \\ 0 \\ 0 \\ \vdots \\ 0 \\ w \\ w \\ \vdots \\ w \\ 1 \\ 0 \end{bmatrix} \in \mathbb{R}^{N(m+n+2)+2} \quad (\text{A-9})$$

$$E_{eq} = \begin{bmatrix} -B_d & 0 & 0 & \cdots & 0 & I_n & 0 & 0 & \cdots & 0 & 0 & \cdots & 0 \\ 0 & -B_d & 0 & \cdots & 0 & -A_d & I_n & 0 & \cdots & 0 & 0 & \cdots & 0 \\ 0 & 0 & -B_d & \cdots & 0 & 0 & -A_d & I_n & \cdots & 0 & 0 & \cdots & 0 \\ \vdots & \vdots & \vdots & \ddots & \vdots & \vdots & \vdots & \vdots & \ddots & \vdots & \vdots & \cdots & 0 \\ 0 & 0 & 0 & \cdots & -B_d & 0 & 0 & 0 & \cdots & -A_d & I_n & \cdots & 0 \end{bmatrix} \in \mathbb{R}^{Nn \times N(m+n+2)+2} \quad (\text{A-10})$$

$$e_{eq} = (\text{A-29}) \quad (\text{A-11})$$

$$E_{in} = \begin{bmatrix} 0 & -I_{x,1} & -I_{\zeta,1} & 0 & 0 & 0 \\ 0 & I_{x,1} & -I_{\zeta,1} & 0 & 0 & 0 \\ 0 & I_{x,3} & -I_{\zeta,2} & 0 & 0 & 0 \\ I & 0 & 0 & 0 & 0 & -\mathbf{1} \\ 0 & 0 & 0 & 0 & 0 & -1 \\ 0 & 0 & 0 & 0 & -\mathbf{1} & \mathbf{a} \end{bmatrix} \in \mathbb{R}^{4N+I+1 \times N(m+n+2)+2}, \quad e_{in} = \begin{bmatrix} -T_{\min,r} \\ T_{\max,r} \\ T_{\max,w} \\ 0 \\ -z(t) \\ -\mathbf{b} \end{bmatrix} \in \mathbb{R}^{4N+I+1} \quad (\text{A-12})$$

Where $T_{\min,r}$, $T_{\max,r}$, $T_{\max,w}$ are vectors in $\mathbb{R}^N$. Moreover, $I_{x,i}$ and $I_{\zeta,i}$ represent selection matrices:

$$I_{x,1} = (\text{A-6}) \quad I_{\zeta,1} = (\text{A-6}) \quad (\text{A-13})$$

$$I_{x,3} = (\text{A-7}) \quad I_{\zeta,2} = (\text{A-7}) \quad (\text{A-14})$$

And $\mathbf{a}$, $\mathbf{b}$ are given by:

$$\mathbf{a} = \begin{bmatrix} a_1 \\ a_2 \\ \vdots \\ a_{\mathcal{I}} \in \mathbb{R}^{\mathcal{I}} \end{bmatrix}, \quad \mathbf{b} = \begin{bmatrix} b_1 \\ b_2 \\ \vdots \\ b_{\mathcal{I}} \end{bmatrix} \in \mathbb{R}^{\mathcal{I}} \quad (\text{A-15})$$

$$lb = \begin{bmatrix} u_{min} \\ -\infty \\ 0 \\ -\infty \\ 0 \end{bmatrix} \in \mathbb{R}^{N(m+n+2)+2}, \quad ub = \begin{bmatrix} u_{max} \\ \infty \\ \infty \\ \infty \\ \infty \end{bmatrix} \in \mathbb{R}^{N(m+n+2)+2} \quad (\text{A-16})$$

A-1-3 Equation (3-29)

Same as (A-1-1) except for:

$$E_{in} = \begin{bmatrix} 0 & -I_{x,1} & -I_{\zeta,1} & 0 \\ 0 & I_{x,1} & -I_{\zeta,1} & 0 \\ 0 & I_{x,3} & -I_{\zeta,2} & 0 \\ I & 0 & 0 & -1 \end{bmatrix} \in \mathbb{R}^{4N \times N(m+n+2)+1}, \quad e_{in} = \begin{bmatrix} -T_{min,r} \\ T_{max,r} \\ T_{max,w} \\ 0 \end{bmatrix} \in \mathbb{R}^{3N+1} \quad (\text{A-17})$$

A-1-4 Equation (3-30)

Same as (A-1-1) except for:

$$\boldsymbol{\eta} = \begin{bmatrix} \mathbf{u} \\ \mathbf{x} \\ \boldsymbol{\zeta} \\ z \end{bmatrix} = \begin{bmatrix} u(0) \\ u(1) \\ \vdots \\ u(N-1) \\ x(1) \\ x(2) \\ \vdots \\ x(N) \\ \zeta(0) \\ \zeta(1) \\ \vdots \\ \zeta(N-1) \\ z(t+N) \end{bmatrix} \in \mathbb{R}^{N(m+n+2)+1} \quad (\text{A-18})$$

$$E_{in} = \begin{bmatrix} 0 & -I_{x,1} & -I_{\zeta,1} & 0 \\ 0 & I_{x,1} & -I_{\zeta,1} & 0 \\ 0 & I_{x,3} & -I_{\zeta,2} & 0 \\ I & 0 & 0 & -1 \end{bmatrix} \in \mathbb{R}^{4N \times N(m+n+2)+1}, \quad e_{in} = \begin{bmatrix} -T_{min,r} \\ T_{max,r} \\ T_{max,w} \\ z(t) \end{bmatrix} \in \mathbb{R}^{3N+1} \quad (\text{A-19})$$

A-1-5 Equation (4-1)

Same as (A-1-1) except for:

$$f = \begin{bmatrix} \hat{c}_e(0) \cdot dt \\ \hat{c}_e(1) \cdot dt \\ \vdots \\ \hat{c}_e(N-1) \cdot dt \\ 0 \\ 0 \\ \vdots \\ 0 \\ w \\ w \\ \vdots \\ w \\ 1 \\ 0 \end{bmatrix} \in \mathbb{R}^{N(m+n+2)+2} \quad (\text{A-20})$$

$$e_{eq} = \begin{bmatrix} B_d^d \hat{d}(0) + A_d x(0) \\ B_d^d \hat{d}(1) \\ B_d^d \hat{d}(2) \\ \vdots \\ B_d^d \hat{d}(N-1) \end{bmatrix} \in \mathbb{R}^{Nn} \quad (\text{A-21})$$

A-1-6 Equation (4-2)

Same as (A-1-2) except for:

$$f = \begin{bmatrix} c_e^{(N)}(0) \cdot dt \\ c_e^{(N)}(1) \cdot dt \\ \vdots \\ c_e^{(N)}(N-1) \cdot dt \\ 0 \\ 0 \\ \vdots \\ 0 \\ w \\ w \\ \vdots \\ w \\ 1 \\ 0 \end{bmatrix} \in \mathbb{R}^{N(m+n+2)+2} \quad (\text{A-22})$$

$$e_{eq} = \begin{bmatrix} B_d^d d^*(0) + A_d x(0) \\ B_d^d d^*(1) \\ B_d^d d^*(2) \\ \vdots \\ B_d^d d^*(N-1) \end{bmatrix} \in \mathbb{R}^{Nn} \quad (\text{A-23})$$

A-1-7 Equation (4-3)

Same as (A-1-1) except for:

$$\eta = \begin{bmatrix} u \\ x \\ \zeta \\ z \end{bmatrix} = \begin{bmatrix} u(0) \\ u(1) \\ \vdots \\ u(N-1) \\ x(1) \\ x(2) \\ \vdots \\ x(N) \\ \zeta(0) \\ \zeta(1) \\ \vdots \\ \zeta(N-1) \\ z(t+N) \end{bmatrix} \in \mathbb{R}^{N(m+n+2)+1}, \quad f = \begin{bmatrix} c_e^{(N)}(0) \cdot dt \\ c_e^{(N)}(1) \cdot dt \\ \vdots \\ c_e^{(N)}(N-1) \cdot dt \\ 0 \\ 0 \\ \vdots \\ 0 \\ w \\ w \\ \vdots \\ w \\ \frac{c_d}{\sigma_t} \end{bmatrix} \in \mathbb{R}^{N(m+n+2)+1} \quad (\text{A-24})$$

$$e_{eq} = \begin{bmatrix} B_d^d d^*(0) + A_d x(0) \\ B_d^d d^*(1) \\ B_d^d d^*(2) \\ \vdots \\ B_d^d d^*(N-1) \end{bmatrix} \in \mathbb{R}^{Nn} \quad (\text{A-25})$$

A-1-8 Equation (4-5)

Same as (A-1-1) except for:

$$\boldsymbol{\eta} = \begin{bmatrix} \boldsymbol{u} \\ \boldsymbol{x} \\ \boldsymbol{\zeta} \\ z \end{bmatrix} = \begin{bmatrix} u(0) \\ u(1) \\ \vdots \\ u(N-1) \\ x(1) \\ x(2) \\ \vdots \\ x(N) \\ \zeta(0) \\ \zeta(1) \\ \vdots \\ \zeta(N-1) \\ z(t+N) \end{bmatrix} \in \mathbb{R}^{N(m+n+2)+1}, \quad \boldsymbol{f} = \begin{bmatrix} c_e^{(N)}(0) \cdot dt \\ c_e^{(N)}(1) \cdot dt \\ \vdots \\ c_e^{(N)}(N-1) \cdot dt \\ 0 \\ 0 \\ \vdots \\ 0 \\ w \\ w \\ \vdots \\ w \\ a(t+N) \cdot c_d \end{bmatrix} \in \mathbb{R}^{N(m+n+2)+1} \quad (\text{A-26})$$

$$\boldsymbol{e}_{eq} = \begin{bmatrix} B_d^d d^*(0) + A_d x(0) \\ B_d^d d^*(1) \\ B_d^d d^*(2) \\ \vdots \\ B_d^d d^*(N-1) \end{bmatrix} \in \mathbb{R}^{Nn} \quad (\text{A-27})$$

A-1-9 Equation (4-6)

Same as (A-1-1) except for:

$$\boldsymbol{\eta} = \begin{bmatrix} \boldsymbol{u} \\ \boldsymbol{x} \\ \boldsymbol{\zeta} \\ z \end{bmatrix} = \begin{bmatrix} u(0) \\ u(1) \\ \vdots \\ u(N-1) \\ x(1) \\ x(2) \\ \vdots \\ x(N) \\ \zeta(0) \\ \zeta(1) \\ \vdots \\ \zeta(N-1) \\ z(t+N) \end{bmatrix} \in \mathbb{R}^{N(m+n+2)+1}, \quad \boldsymbol{f} = \begin{bmatrix} c_e^{(N)}(0) \cdot dt \\ c_e^{(N)}(1) \cdot dt \\ \vdots \\ c_e^{(N)}(N-1) \cdot dt \\ 0 \\ 0 \\ \vdots \\ 0 \\ w \\ w \\ \vdots \\ w \\ \frac{N}{T} \cdot c_d \end{bmatrix} \in \mathbb{R}^{N(m+n+2)+1} \quad (\text{A-28})$$

$$e_{eq} = \begin{bmatrix} B_d^d d^*(0) + A_d x(0) \\ B_d^d d^*(1) \\ B_d^d d^*(2) \\ \vdots \\ B_d^d d^*(N-1) \end{bmatrix} \in \mathbb{R}^{Nn} \quad (\text{A-29})$$

A-2 PLP formulation

Using the matrices and vectors defined in A-1-1 we can reformulate the problem as a parametric linear program (PLP):

$$\begin{aligned} \min_{\eta} \quad & f^T \eta \\ \text{s.t.} \quad & G\eta \leq w + F\theta \end{aligned} \quad (\text{A-30})$$

Defining the components as:

$$G_1 = E_{in}, \quad G_2 = \begin{bmatrix} E_{eq} \\ -E_{eq} \end{bmatrix}, \quad G_3 = \begin{bmatrix} -I \\ I \end{bmatrix} \quad (\text{A-31})$$

$$w_1 = \begin{bmatrix} -T_{\min,r} \\ T_{\max,r} \\ T_{\max,w} \\ 0 \end{bmatrix}, \quad w_2 = \begin{bmatrix} e_{eq} \\ -e_{eq} \end{bmatrix}, \quad w_3 = \begin{bmatrix} -lb \\ ub \end{bmatrix} \quad (\text{A-32})$$

Where $T_{\min,r}$, $T_{\max,r}$, $T_{\max,w}$ are vectors in $\mathbb{R}^N$.

$$G = \begin{bmatrix} G_1 \\ G_2 \\ G_3 \end{bmatrix} \in \mathbb{R}^{2Nn+3N+1+2(N(m+n+2)+1) \times N(n+N+2)+1} \quad (\text{A-33})$$

$$w = \begin{bmatrix} w_1 \\ w_2 \\ w_3 \end{bmatrix} \in \mathbb{R}^{2Nn+3N+1+2(N(m+n+2)+1)} \quad (\text{A-34})$$

$$F = \begin{bmatrix} \mathbf{0}_{3N \times 1} \\ -1 \\ \mathbf{0}_{2Nn \times 1} \\ \mathbf{0}_{2(N(m+n+2)+1) \times 1} \end{bmatrix} \in \mathbb{R}^{2Nn+3N+1+2(N(m+n+2)+1)} \quad (\text{A-35})$$

$$\theta = z(t + N) \quad (\text{A-36})$$

Appendix B

Simulation results

Code used for this thesis project is available in: <https://github.com/MaxWervers/MScThesis>

B-1 Modeling and Other Constants

$K_{ra} = 28 \text{ kJ}/(^{\circ}\text{C} \cdot \text{h})$	Heat transfer coefficient: room to ambient
$K_{fr} = 624 \text{ kJ}/(^{\circ}\text{C} \cdot \text{h})$	Heat transfer coefficient: floor to room
$K_{wf} = 28 \text{ kJ}/(^{\circ}\text{C} \cdot \text{h})$	Heat transfer coefficient: water to floor
$C_{p,r} = 810 \text{ kJ}/^{\circ}\text{C}$	Thermal capacity of the room zone
$C_{p,f} = 3315 \text{ kJ}/^{\circ}\text{C}$	Thermal capacity of the floor zone
$C_{p,w} = 836 \text{ kJ}/^{\circ}\text{C}$	Thermal capacity of the water zone
$COP = 3$	Coefficient of performance of the heat pump
$t_s = 15 \text{ min}$	Sampling time (time step for discrete modeling)
$\Delta = 0.25$	Time-of-use (TOU) cost scaling factor
$c_d = 5 \text{ €/kW}$	Peak demand charge
$u_{\min} = 0 \text{ kW}$	Lower bound on control input (power)
$u_{\max} = 10 \text{ kW}$	Upper bound on control input (power)
$T_{\min,r} = \text{see 1-6}$	Minimum room temperature
$T_{\max,r} = \text{see 1-6}$	Maximum room temperature
$T_{\max,w} = 45 \text{ }^{\circ}\text{C}$	Upper bound on water temperature
$w = 10,000$	Soft constraint penalty (same for each ζ_1, ζ_2)

B-1-1 Historical vs Real Temperatures 2023

Figure B-1 shows the real temperatures of 2023 and the temperatures based on historical values. A moving average over 3 days is used to make the plot more readable.




Figure B-1: Historical $T_{a,pred}$ vs real ambient temperatures T_a in 2023

B-2 Perfect Foresight Terminal Cost Functions




Figure B-2: Perfect foresight terminal cost function January 2023.




Figure B-3: Perfect foresight terminal cost function February 2023.




Figure B-4: Perfect foresight terminal cost function March 2023.




Figure B-5: Perfect foresight terminal cost function April 2023.




Figure B-6: Perfect foresight terminal cost function May 2023.




Figure B-7: Perfect foresight terminal cost function June 2023.




Figure B-8: Perfect foresight terminal cost function July 2023.




Figure B-9: Perfect foresight terminal cost function August 2023.




Figure B-10: Perfect foresight terminal cost function September 2023.




Figure B-11: Perfect foresight terminal cost function October 2023.




Figure B-12: Perfect foresight terminal cost function November 2023.




Figure B-13: Perfect foresight terminal cost function December 2023.

B-3 Terminal Cost Functions under Uncertainty




Figure B-14: Historical average based terminal cost function January 2023.




Figure B-15: Historical average based terminal cost function February 2023.




Figure B-16: Historical average based terminal cost function March 2023.




Figure B-17: Historical average based terminal cost function April 2023.




Figure B-18: Historical average based terminal cost function May 2023.




Figure B-19: Historical average based terminal cost function June 2023.




Figure B-20: Historical average based terminal cost function July 2023.




Figure B-21: Historical average based terminal cost function August 2023.




Figure B-22: Historical average based terminal cost function September 2023.




Figure B-23: Historical average based terminal cost function October 2023.




Figure B-24: Historical average based terminal cost function November 2023.




Figure B-25: Historical average based terminal cost function December 2023.

B-4 Monthly Distribution of Bill Cost Simulations in 2023

This section presents the distribution of total bill cost simulations for each month in 2023. The data is grouped by prediction horizon, and each box plot represents the results for a specific model predictive control (MPC) method.




Figure B-26: Distribution of total bill cost simulations for January 2023, grouped by prediction horizon and MPC method.




Figure B-27: Distribution of total bill cost simulations for February 2023, grouped by prediction horizon and MPC method.




Figure B-28: Distribution of total bill cost simulations for March 2023, grouped by prediction horizon and MPC method.




Figure B-29: Distribution of total bill cost simulations for April 2023, grouped by prediction horizon and MPC method.




Figure B-30: Distribution of total bill cost simulations for May 2023, grouped by prediction horizon and MPC method.




Figure B-31: Distribution of total bill cost simulations for June 2023, grouped by prediction horizon and MPC method.




Figure B-32: Distribution of total bill cost simulations for July 2023, grouped by prediction horizon and MPC method.




Figure B-33: Distribution of total bill cost simulations for August 2023, grouped by prediction horizon and MPC method.




Figure B-34: Distribution of total bill cost simulations for September 2023, grouped by prediction horizon and MPC method.




Figure B-35: Distribution of total bill cost simulations for October 2023, grouped by prediction horizon and MPC method.




Figure B-36: Distribution of total bill cost simulations for November 2023, grouped by prediction horizon and MPC method.




Figure B-37: Distribution of total bill cost simulations for December 2023, grouped by prediction horizon and MPC method.

Bibliography

- [1] Xiaodong Cao, Xilei Dai, and Junjie Liu. Building energy-consumption status worldwide and the state-of-the-art technologies for zero-energy buildings during the past decade. *Energy and Buildings*, 128:198–213, 2016.
- [2] IEA International Energy Agency and International Partnership for Energy Efficiency Cooperation. Building energy performance metrics. Technical report, IEA Publications, 2015.
- [3] TenneT. Interactive grid capacity map. <https://www.tennet.eu/nl-en/grid-capacity-map>, n.d. Accessed: 19-5-2025.
- [4] IEA International Energy Agency. The future of heat pumps. Technical report, IEA Publications, 2022.
- [5] Autoriteit Consument & Markt. Monitor Consumentenmarkt Energie: Aantal contracten per contracttype. <https://public.tableau.com/views/MonitorConsumentenmarktEnergie/Aantalcontractenpercontracttype>, 2025. Tableau Public dashboard, accessed June 18, 2025.
- [6] Nord Pool. Day-Ahead Market. <https://www.nordpoolgroup.com/en/the-power-market/Day-ahead-market/>, 2025. Accessed June 18, 2025.
- [7] ENTSO-e Transparency Platform. Day-ahead energy prices. <https://newtransparency.entsoe.eu/market/energyPrices>, 2025. Accessed: 2025-06-19.
- [8] Nishith R. Patel, Michael J. Risbeck, James B. Rawlings, Christos T. Maravelias, Michael J. Wenzel, and Robert D. Turney. A case study of economic optimization of hvac systems based on the stanford university campus airside and waterside systems. *International High Performance Buildings Conference*, 2018.
- [9] Mohammad Hassan Fathollahzadeh and Paulo Cesar Tabares-Velasco. Integrated framework for optimization of air- and water-side hvac systems to minimize electric utility cost of existing commercial districts. *Energy and Buildings*, 273:112328, 2022.

- [10] Rasmus Halvgaard, Niels Kjølstad Poulsen, Henrik Madsen, and John Bagterp Jørgensen. Economic model predictive control for building climate control in a smart grid. In *2012 IEEE PES innovative smart grid technologies (ISGT)*, pages 1–6. IEEE, 2012.
- [11] Salvatore Carlucci and Lorenzo Pagliano. A review of indices for the long-term evaluation of the general thermal comfort conditions in buildings. *Energy and Buildings*, 53:194–205, 2012.
- [12] Sébastien Gros, Mario Zanon, Rien Quirynen, Alberto Bemporad, and Moritz Diehl. From linear to nonlinear mpc: bridging the gap via the real-time iteration. *International Journal of Control*, 93(1):62–80, 2020.
- [13] Damien Picard, Ján Drgoňa, Michal Kvasnica, and Lieve Helsen. Impact of the controller model complexity on model predictive control performance for buildings. *Energy and Buildings*, 152:739–751, 2017.
- [14] Wanfu Zheng, Dan Wang, and Zhe Wang. Economic model predictive control for building hvac system: A comparative analysis of model-based and data-driven approaches using the boptest framework. *Applied Energy*, 374:123969, 2024.
- [15] James B. Rawlings and Rishi Amrit. *Optimizing Process Economic Performance Using Model Predictive Control*, pages 119–138. Springer Berlin Heidelberg, Berlin, Heidelberg, 2009.
- [16] Jingran Ma, Joe Qin, Timothy Salsbury, and Peng Xu. Demand reduction in building energy systems based on economic model predictive control. *Chemical Engineering Science*, 67(1):92–100, 2012. Dynamics, Control and Optimization of Energy Systems.
- [17] Jingran Ma, S. Joe Qin, and Timothy Salsbury. Application of economic mpc to the energy and demand minimization of a commercial building. *Journal of Process Control*, 24(8):1282–1291, 2014. Economic nonlinear model predictive control.
- [18] Ján Drgoňa, Damien Picard, Michal Kvasnica, and Lieve Helsen. Approximate model predictive building control via machine learning. *Applied Energy*, 218:199–216, 2018.
- [19] Ján Drgoňa, Javier Arroyo, Iago Cupeiro Figueroa, David Blum, Krzysztof Arendt, Donghun Kim, Enric Perarnau Ollé, Juraj Oravec, Michael Wetter, Draguna L. Vrabie, and Lieve Helsen. All you need to know about model predictive control for buildings. *Annual Reviews in Control*, 50:190–232, 2020.
- [20] Olivier Van Cutsem, Maher Kayal, David Blum, and Marco Pritoni. Comparison of mpc formulations for building control under commercial time-of-use tariffs. In *2019 IEEE Milan PowerTech*, pages 1–6, 2019.
- [21] Ranjeet Kumar, Michael J. Wenzel, Mohammad N. ElBsat, Michael J. Risbeck, Kirk H. Drees, and Victor M. Zavala. Stochastic model predictive control for central hvac plants. *Journal of Process Control*, 90:1–17, 2020.
- [22] Michael J. Risbeck and James B. Rawlings. Economic model predictive control for time-varying cost and peak demand charge optimization. *IEEE Transactions on Automatic Control*, 65(7):2957–2968, 2020.

-
- [23] Peter Bauer, Alan Thorpe, and Gilbert Brunet. The quiet revolution of numerical weather prediction. *Nature*, 525(7567):47–55, 2015.
 - [24] Royal Netherlands Meteorological Institute (KNMI). Klimatologie: Uurgegevens, n.d. Accessed: 2025-01-27.
 - [25] Gurobi Optimization, LLC. *Gurobi Optimizer Reference Manual*, 2024. <https://www.gurobi.com>.
 - [26] Richard Bellman. *Dynamic Programming*. Princeton University Press, Princeton, NJ, USA, 1 edition, 1957.
 - [27] David Silver, Thomas Hubert, Julian Schrittwieser, Ioannis Antonoglou, Matthew Lai, Arthur Guez, Marc Lanctot, Laurent Sifre, Dhharshan Kumaran, Thore Graepel, et al. A general reinforcement learning algorithm that masters chess, shogi, and go through self-play. *Science*, 362(6419):1140–1144, 2018.
 - [28] Dimitri P Bertsekas. Model predictive control and reinforcement learning: A unified framework based on dynamic programming. *IFAC-PapersOnLine*, 58(18):363–383, 2024.
 - [29] Francesco Borrelli, Alberto Bemporad, and Manfred Morari. Geometric algorithm for multiparametric linear programming. *Journal of optimization theory and applications*, 118:515–540, 2003.

Glossary

List of Acronyms

HVAC	heating, ventilation, and air-conditioning
AHUs	air handling units
PID	proportional-integral-derivative
MPC	model predictive control
EMPC	economic model predictive control
LP	linear program
COP	coefficient of performance
TOU	time-of-use
LP	linear program
PLP	parametric linear program
TES	thermal energy storage
RC	resistance-capacitance
ANN	artificial neural network

List of Symbols

Symbol	Description	Unit
T_r	Room temperature	$^{\circ}\text{C}$
T_f	Floor temperature	$^{\circ}\text{C}$
T_w	Water temperature	$^{\circ}\text{C}$
T_a	Ambient temperature	$^{\circ}\text{C}$
$\dot{Q}_{ra}$	Heat transfer rate: room to ambient	W
$\dot{Q}_{fr}$	Heat transfer rate: floor to room	W
$\dot{Q}_{wf}$	Heat transfer rate: water to floor	W
$\dot{Q}_c$	Heat supplied by the heat pump	W
W_c	Compressor electrical power input	W
u	Control input (equal to W_c)	W
d	Disturbance input (equal to T_a)	$^{\circ}\text{C}$
$\mathbf{c}_e^*$	Vector of known electricity prices	€/Wh
$\hat{\mathbf{c}}_e$	Vector of predicted electricity prices	€/Wh
$\check{\mathbf{c}}_e$	Vector of true electricity prices	€/Wh
$\mathbf{d}^*$	Forecast ambient temperature (reliable)	$^{\circ}\text{C}$
$\hat{\mathbf{d}}$	Predicted ambient temperature vector	$^{\circ}\text{C}$
$\check{\mathbf{d}}$	True ambient temperature vector	$^{\circ}\text{C}$
ζ_1	Slack variable for room temperature constraint	–
ζ_2	Slack variable for water temperature constraint	–
$z(\cdot)$	Peak demand variable	W
v	Terminal cost slack variable	€
$V_f(z(t+N), \tau)$	Approximated terminal cost function	€
t	Current time index within the month	–
T	Final time index of the month	–
M	Number of known electricity price values	–
N	Prediction horizon (length of forecast window)	–
τ	Index of midnight (once per day)	–