

Policy-Guided Economic MPC for Multi-Energy Microgrids with Seasonal Thermal Energy Stor- age

G.D. Voogt

Master of Science Thesis

Policy-Guided Economic MPC for Multi-Energy Microgrids with Seasonal Thermal Energy Storage

MASTER OF SCIENCE THESIS

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

G.D. Voogt

June 30, 2026

Faculty of Mechanical Engineering (ME) · Delft University of Technology

Abstract

Microgrids combining electrical and seasonal thermal energy storage offer significant operational flexibility, but their management is complicated by a severe separation of time scales: electrical components evolve over minutes, while aquifer thermal energy storage (ATES) systems charge and discharge over seasonal cycles spanning weeks to months. Conventional Economic Model Predictive Control (EMPC) cannot account for this long-term seasonal objective within a computationally tractable prediction horizon, and existing approaches close this gap through heuristic terminal costs or surrogate-derived terminal constraints, both of which require careful tuning and offer limited theoretical guarantees. This thesis investigates Policy-Guided Model Predictive Control (PG-MPC) as an alternative, in which a guiding policy embedded within the short-horizon MPC problem encodes the long-term seasonal objective, rather than relying on a precomputed reference trajectory. A combined heat and power (CHP) microgrid model, incorporating a heat pump, auxiliary boiler and chiller, battery storage, and an ATES system, is formulated using the Mixed-Logical Dynamical (MLD) framework to capture the system's discrete operational modes and nonlinear heat pump characteristics. On this shared model, a baseline EMPC controller using surrogate-derived terminal ingredients is compared against a PG-MPC controller using a simple sinusoidal dead-band guiding policy. Both controllers are benchmarked in simulation under imperfect disturbance forecasts across four scenarios spanning a range of climatological and electricity-pricing conditions, and both satisfy the end-of-year thermal balance constraint in every case. Operational cost performance is strongly scenario-dependent: PG-MPC achieves cost reductions under high price volatility and above average cooling demand, and performance comparable economically but slightly worse in terms of comfort constraint satisfaction across most scenarios. The improvement offered by PG-MPC is therefore best understood as a gain in adaptivity and robustness to realised weather and price conditions rather than a uniform cost reduction: it removes the dependence on a precomputed annual trajectory while preserving guaranteed recursive feasibility, and remains competitive with EMPC even with a deliberately simple guiding policy. The framework thus offers a tractable, theoretically grounded route to embedding seasonal storage objectives in receding-horizon microgrid control, with realised performance governed by the quality of the guiding policy and the calibration of the terminal ingredients.

Table of Contents

1	Introduction	1
2	Microgrid modeling and Control	5
2-1	Electrical Modeling	5
2-1-1	Sources and Loads	5
2-1-2	Batteries	6
2-1-3	Power Grid Interconnection	7
2-2	Thermal Modeling	8
2-2-1	Heat Pump	8
2-2-2	Auxiliary Thermal Units	9
2-2-3	Aquifer Thermal Energy Storage	9
2-2-4	Thermal Masses	11
2-3	Microgrid Control	12
2-3-1	Control Hierarchy	12
2-3-2	Energy Management	14
3	System Description	15
3-1	System Overview	15
3-2	Electrical Subsystem	16
3-3	Thermal Subsystem	16
3-4	Disturbance Signals	17
3-5	State-Space Representation	19
3-6	System Parameters	20

4	Economic Model Predictive Control for Microgrids	23
4-1	Model Predictive Control Fundamentals	23
4-2	Economic Model Predictive Control	25
4-3	Mixed-Logical Dynamical Model Predictive Control	27
4-3-1	Discrete Actuators and Integer Programming	27
4-3-2	The MLD Framework	28
4-3-3	Binary Decisions as Linear Constraints	29
4-3-4	PWA Approximation of Nonlinear Functions	30
4-3-5	MLD-MPC	32
4-3-6	Stability of MLD-MPC	33
5	Time-Varying EMPC Controller Design	35
5-1	MLD Model	35
5-1-1	Battery	35
5-1-2	Grid Interconnection	36
5-1-3	Thermal System	37
5-1-4	Heat Pump	37
5-1-5	ATES	39
5-2	Cost Function Design	41
5-2-1	Full EMPC Optimization Problem	41
5-3	Terminal Constraint Set and Terminal Cost Construction	42
5-3-1	The Need for Terminal Ingredients	43
5-3-2	Surrogate Optimization Problem	44
5-3-3	Terminal Constraint Set	47
5-3-4	Terminal Cost	47
5-4	Forecast Signal Generation	48
6	Policy-Guided MPC Controller Design	51
6-1	Policy-Guided MPC Fundamentals	51
6-1-1	Introduction to PG-MPC	51
6-1-2	Theoretical Properties	52
6-2	Policy-Guided MPC Controller	55
6-2-1	Policy Design	55
6-2-2	Terminal Value Function Construction	56
6-2-3	Terminal Constraint Set Design	58
6-2-4	Full PG-MPC Problem	59

7 Case Study	61
7-1 Experimental Setup and Scenarios	61
7-1-1 Weather	61
7-1-2 Energy Prices	63
7-2 Performance Metrics	65
7-3 Simulation Results	66
7-3-1 Nominal System Trajectory	66
7-3-2 Comparative Analysis	68
7-4 Discussion	74
7-4-1 Closed-Loop Cost	74
7-4-2 Thermal Balance Constraint Satisfaction	76
7-4-3 Controller behavior and Policy Deviation	76
7-4-4 Policy Design	77
7-4-5 Scenario Robustness	78
7-4-6 Limitations	78
8 Conclusion and Future Work	81
8-1 Answering the Research Questions	81
8-2 Recommendation for Future Work	83
Bibliography	85
Glossary	93
List of Acronyms	93

List of Figures

2-1	The seasonal mismatch between thermal energy availability and heating demand (a), and the application of thermal storage in seasonal load shifting of thermal supply and demand.	9
2-2	Working principle of a two-well ATES system. Figure adapted from [37].	10
2-3	Schematic representation of the energy flows in the ATES system in both heating and cooling modes.	11
2-4	The hierarchical control architecture for microgrids. Figure adapted from [44]. . .	13
3-1	Overview of the system that will be used for the rest of this thesis.	16
3-2	The thermal dynamics of the reservoir, floor and building, visualized. Figure adapted from [19].	18
3-3	Representative two day samples of the load profile for each season, visualizing the mathematical description of the profile shown in Eq. (3-10).	19
4-1	Two time steps illustrating the working principle of MPC.	24
4-2	Difference between continuous input set (a), integer constrained input set (b), and binary constrained input set (c). Figure adapted from [66].	28
4-3	PWA approximation of a function of one variable, illustrating the relationship between partition fineness and corresponding approximation error.	31
4-4	Illustration of the 2D PWA approximation scheme with upper triangle active on the left and lower triangle active on the right. Figure adapted from [72].	32
5-1	Two-dimensional PWA approximation of the heat pump model and the resulting absolute approximation error.	39
5-2	ATES temperatures (a) and energy balance (b) when controlling the system using EMPC without terminal ingredients.	43

5-3	The control architecture used for the EMPC controller.	44
5-4	The disturbance signals used when solving the surrogate optimization problem. .	46
5-5	Solution to the surrogate optimization problem, with (a) showing the ATEs well temperatures and (b) showing the imbalance state as well as cumulative heating and cooling done by the heat pump.	49
5-6	Comparison of forecasted disturbance signals given to MPC and the real disturbance signals.	50
6-1	An example of two time steps using the PG-MPC controller.	53
6-2	Two examples of construction of terminal cost that satisfies Eq. (6-4), with left based on the global Lipschitz constant, and right based on the directional Lipschitz constants L^+ and L^- relative to x_f . Notation: $(x - x_f)^+ = \max(x - x_f, 0)$, $(x - x_f)^- = \max(-(x - x_f), 0)$	54
6-3	A block diagram showing the PG-MPC control architecture.	55
6-4	Representation of dead-band policy behavior over time at different values of the imbalance state b	56
6-5	The approximate value function surface along with its time-varying directional Lipschitz constants relative to the trajectory $\tilde{b}(t)$	59
7-1	The monthly average temperature across all four scenarios.	63
7-2	Cumulative solar irradiation across the four scenarios, highlighting inter-scenario variability in cooling demand (a) and PV generation potential (b).	63
7-3	The monthly average electricity price across all four scenarios, showing the temporal trends of the statistics listed in Table 7-2	64
7-4	The price duration curves for all four scenarios. The 2018 and TMY scenarios showing stable electricity prices compared to the more volatile 2021 and 2025 scenarios.	65
7-5	Nominal system states and inputs using the time-varying EMPC baseline controller using the TMY disturbance signals.	67
7-6	Nominal system states and inputs using the PGM-MPC controller using the TMY disturbance signals.	67
7-7	Cumulative operational cost and cumulative constraint violation penalties for both the EMPC and PG-MPC controllers for the TMY scenario.	68
7-8	Hot and cold well temperatures, as well as thermal imbalance trajectories for both EMPC and PGMPC for the TMY scenario.	69
7-9	Building temperature for both EMPC (left) and PGMPC (right) over the operational year for the TMY scenario.	69
7-10	Cumulative operational cost and cumulative constraint violation penalties for both the EMPC and PG-MPC controllers for the 2021 (Mild) scenario.	70

7-11 Hot and cold well temperatures, as well as thermal imbalance trajectories for both EMPC and PGMPC for the 2021 (Mild) scenario.	71
7-12 Building temperature for both EMPC (left) and PGMPC (right) over the operational year for the 2021 (Mild) scenario.	71
7-13 Cumulative operational cost and cumulative constraint violation penalties for both the EMPC and PG-MPC controllers for the 2025 (Moderate) scenario.	72
7-14 Hot and cold well temperatures, as well as thermal imbalance trajectories for both EMPC and PGMPC for the 2025 (Moderate) scenario.	72
7-15 Building temperature for both EMPC (left) and PGMPC (right) over the operational year for the 2025 (Moderate) scenario.	73
7-16 Cumulative operational cost and cumulative constraint violation penalties for both the EMPC and PG-MPC controllers for the 2018 (Extreme) scenario.	73
7-17 Hot and cold well temperatures, as well as thermal imbalance trajectories for both EMPC and PG-MPC for the 2018 (Extreme) scenario.	74
7-18 Building temperature for both EMPC (left) and PGMPC (right) over the operational year for the 2018 (Extreme) scenario.	74

Chapter 1

Introduction

The rapid growth of renewable generation in the Netherlands over the last years, along with increasing electrification of heat and transport, is placing increasing strain on the power grid [1], [2]. In particular, the replacement of conventional synchronous generators with intermittent renewable sources such as photovoltaic (PV) and wind reduces system inertia and, consequently, grid stability [3]. At the same time, the rapid increase in distributed generation and electricity demand has substantially increased grid congestion [4]. The scale of this challenge is illustrated by the Association of Energy Network Operators in the Netherlands [5], which shows that, as of 2026, large parts of the Netherlands face severe shortages in transport capacity, resulting in waiting lists for new grid connections. These developments highlight the need for new approaches that can manage generation and demand more locally while reducing the burden on the transmission and distribution network.

Microgrids and energy hubs have emerged as promising concepts to address this challenge by coordinating distributed energy resources at the local level [6], [7]. A microgrid provides localized control over distributed generation, loads, and storage technologies, enabling higher levels of self-consumption while also providing ancillary services to the main grid [8], [9]. By balancing local generation and demand, microgrids can reduce power exchanged with the main power network, thereby alleviating grid congestion and limiting the need for costly grid reinforcements. At the same time, their ability to coordinate flexible resources improves the integration of intermittent renewable generation, contributing to a more resilient and efficient power system.

Lasseter [6] originally envisioned microgrids as integrated systems combining both electrical and thermal energy. While much of the subsequent microgrid literature focused primarily on the electrical domain [10], recent work, particularly on energy hubs, has renewed interest in multi-energy systems through the integration of combined heat and power (CHP) units, Battery Energy Storage (BES), Thermal Energy Storage (TES), and electrically driven heating technologies [11]. By explicitly modeling the thermal subsystem, these systems can exploit the coupling between electricity and heat, providing additional operational flexibility. This flexibility enables more effective coordination of local generation, storage, and demand

in response to disturbances such as variations in solar irradiance, wind speed, and ambient temperature, further reducing reliance on the upstream grid.

While integrating electrical and thermal energy systems increases operational flexibility, it also substantially increases the complexity of operating the microgrid. To fully exploit the available flexibility, advanced energy management systems are required to coordinate distributed generation, controllable loads, and multiple storage technologies while accounting for varying demand, renewable generation, and operational constraints. This challenge becomes even more pronounced in microgrids equipped with hybrid energy storage systems, particularly when seasonal thermal energy storage is employed. The large separation between the time scales of electrical and thermal storage introduces operational objectives that extend well beyond the prediction horizons typically used for real-time control [12]. Consequently, the design of effective control algorithms has become a central research topic in the operation of modern microgrids [13], [14].

There is a substantial body of literature addressing the microgrid energy management system (EMS) problem, with optimization- and learning-based approaches receiving considerable attention [13], [15]. Learning-based methods, such as reinforcement learning (RL), have been investigated as an alternative to model-based control, but their reliance on extensive training data and the lack of explicit constraint guarantees limit their applicability to many practical energy management problems. Among optimization-based approaches, model predictive control (MPC) has become one of the most widely adopted techniques due to its ability to optimize system operation over a receding prediction horizon while explicitly accounting for system dynamics and operational constraints. Numerous studies have demonstrated its effectiveness for energy management in microgrids under uncertain renewable generation and electricity pricing [16]–[19].

Despite these advantages, the finite prediction horizon fundamentally limits the ability of conventional MPC to account for long-term operational objectives. Requirements such as maintaining seasonal thermal energy storage levels or enforcing annual energy balances extend far beyond the typical EMS optimization horizon of several hours [20, p. 11]. These objectives cannot be incorporated directly into the optimization problem without sacrificing computational tractability. Existing approaches typically rely on heuristic terminal costs [21] or surrogate constraints that encourage behavior similar to the desired long-term objective [22]. However, both approaches require careful tuning and provide only indirect control over long-term system behavior.

This thesis investigates the application of Policy-Guided MPC (PG-MPC) [23], [24] to the energy management of CHP microgrids with hybrid electrical and thermal energy storage. PG-MPC addresses the limitations of conventional MPC by embedding a guiding policy within the short-horizon optimization problem. This policy captures operational objectives that extend beyond the prediction horizon, enabling the controller to account for long-term requirements, such as seasonal thermal storage constraints, while retaining the computational tractability of receding-horizon optimization. To address this objective, the following main research question is considered:

How can Policy-Guided Model Predictive Control improve the energy management of CHP microgrids with seasonal thermal energy storage compared with conventional Economic Model Predictive Control?

This question is addressed through the following subquestions:

1. How can Policy-Guided MPC be formulated for a CHP microgrid with hybrid electrical and thermal energy storage?
2. How does the long-term operational performance of Policy-Guided MPC compare with a conventional EMPC controller?
3. What are the advantages and limitations of Policy-Guided MPC for practical microgrid energy management?

The scope of this thesis will be limited to presenting a proof-of-concept of applying the PG-MPC algorithm to the energy management of a simulated CHP microgrid. The aim is assessing the suitability of PG-MPC compared to conventional economic model predictive control (EMPC) controller design and providing insights into the performance, advantages, and disadvantages of applying PG-MPC.

The rest of this thesis will proceed as follows. Chapter 2 introduces the relevant background on both the electrical and thermal component models that will be used for simulating the microgrid system. Particular attention is given to the modeling of thermal energy storage in the context of microgrid systems. A short overview of the relevant background on microgrid control is also provided. Chapter 3 details the concrete microgrid model that will be used as the case study for controller design. Chapter 4 outlines the relevant background on MPC, EMPC, and the mixed logical dynamics (MLD) framework that is required for the controller design. Chapters 5 and 6 detail the design of the baseline EMPC and the PG-MPC controllers. The PG-MPC chapter also provides a comprehensive overview of the relevant background of the PG-MPC algorithm. Chapter 7 provides a detailed description of the scenarios and metrics used to benchmark the controllers, after which the simulation results are presented and discussed. Finally, Chapter 8 draws conclusions, provides explicit answering of the research questions, and sets out recommendations for future research.

Microgrid modeling and Control

This chapter will describe the mathematical modeling of key microgrid components. Each model captures the important operational characteristics, constraints, and efficiencies relevant to microgrid energy management, providing a comprehensive overview of component modeling for EMS applications. The components of a CHP microgrid can broadly be categorized as electrical and thermal components. Electrical components will be discussed in Section 2-1, and thermal components will be discussed in Section 2-2. Aspects of microgrid control are discussed in Section 2-3.

2-1 Electrical Modeling

The electrical subsystem of a CHP microgrid consists of generation, storage, and load components connected to a common electrical bus. This section presents the models used to describe each of these components, as well as the power balance constraint that governs their interaction.

2-1-1 Sources and Loads

Many components in microgrid energy management can be considered sources or loads, and do not have any dynamic models associated with them. Their power output or consumption is either uncontrollable, in the case of renewable energy sources (RES) generation and passive loads, or directly specified as a control input.

PV generation is a form of uncontrollable generation in which the power output is driven by the solar irradiance disturbance signal. The generated power depends on environmental conditions including panel tilt, cell temperature, and solar angle of incidence [20, pp. 50–52]. For the purpose of this thesis, the primary interest lies in the qualitative shape of the generation profile rather than absolute accuracy. These secondary effects are therefore

omitted, and PV power generation is estimated from the solar irradiance using a lumped efficiency model

$$P_{pv}(t) = \eta_{pv} A_{pv} s(t), \quad (2-1)$$

where η_{pv} is the lumped panel efficiency, A_{pv} the total panel area in m^2 , and $s(t)$ the solar irradiance in W/m^2 .

The electrical load is treated as an uncontrollable disturbance signal, representing the aggregate power demand of the building excluding heating and cooling, which are explicitly modelled through the thermal subsystem. This load reflects the combined consumption of appliances, lighting, and other plug loads, and is taken as exogenously given rather than derived from a first-principles model, consistent with standard practice in microgrid EMS studies.

2-1-2 Batteries

Batteries are among the most widely adopted storage solutions in microgrids due to their versatility and fast deployment characteristics [25]. They allow excess electrical energy to be stored during periods of low demand or high RES generation, and deployed when needed, enabling temporal decoupling of generation and consumption.

The state of charge of the battery can be modelled as a simple integrator of the power input [16], [20, p. 60]. Batteries are subject to self-discharge, in which the state of charge slowly declines over time in the absence of applied currents. Self-discharge depends on factors such as battery chemistry, charging current, and state of charge [26, p. 440]. For lithium-ion batteries, the self-discharge rate is typically small compared to active charging and discharging currents [27], and the self-discharge rate is commonly approximated as a small constant [16]. The resulting continuous-time dynamical model for the battery state of charge is given by

$$\frac{dx_b(t)}{dt} = \frac{1}{C_{\max}} (-\eta(P_b(t))P_b(t) - x_{sd}), \quad (2-2)$$

where $x_b(t) \in [0, 1]$ is the state of charge, $C_{\max}$ is the maximum battery capacity, $P_b(t)$ is the power supplied to the battery, and x_{sd} captures the self-discharge rate. The charging and discharging processes are subject to different efficiency losses [28], which are captured through a piecewise-constant efficiency model [16]

$$\eta(P_b) = \begin{cases} \eta_c, & P_b \leq 0, \\ \frac{1}{\eta_d}, & P_b > 0, \end{cases} \quad (2-3)$$

where η_c is the charging efficiency and η_d is the discharging efficiency.

When minimizing operational costs in microgrids, it is important to consider the degradation associated with the components, as poorly chosen operating points can lead to excessive degradation. Battery degradation depends on various factors, including temperature, charge and discharge currents, and depth of discharge [29], [30]. In the context of microgrid EMSs, the charge and discharge currents and depth of discharge are the most important factors to consider. The state of charge can easily be constrained to remain within acceptable bounds. The charge and discharge currents are directly related to the battery power output, which is

a quantity controlled by the EMS. In many applications, the degradation cost of the battery as a result of current drawn is characterized by [16], [17]

$$c_{\text{bat,deg}}(t) = \text{OM}_b |P_b(t)|, \quad (2-4)$$

where OM_b is the battery's operational and maintenance cost, with OM_b often chosen to reflect the ratio between capital expenditure and the number of battery life cycles [17], [31].

A way to characterize and compare batteries, and energy storage systems more generally, is its time constant. The time constant is defined as the ratio between maximum capacity to maximum discharge rate:

$$\tau = \frac{C_{\text{max}}}{P_{\text{max}}}, \quad (2-5)$$

where C_{max} is the maximum storage capacity in kWh and P_{max} is the maximum discharge rate in kW. The time constant τ is then effectively the minimum time in which the energy storage system can be fully charged or discharged. Large time constants correspond to energy storage systems with large capacity relative to the discharge rate, whereas smaller time constant systems can charge and discharge nearly instantaneously relative to its capacity.

2-1-3 Power Grid Interconnection

A key feature of grid-connected microgrids is their ability to exchange power bidirectionally with the utility grid, acting as a *prosumer* that consumes or supplies power depending on the local balance between generation and demand [32]. In many microgrid applications, the microgrid is not modelled as a balancing responsible party (BRP) (i.e., an entity contractually responsible for maintaining power balance within a defined portfolio), and the grid power is instead a free variable bounded only by the maximum power draw from the utility grid [16], [17]. This simplifies the energy management problem considerably: the grid acts as a large power buffer, absorbing surplus generation and covering shortfalls instantaneously, so the real-time power balance is always guaranteed and the EMS need only optimize the dispatch of local assets.

The grid interconnection is subject to a maximum import and export capacity $P_{g,\text{max}}$, reflecting physical limits of the connection. The cost of exchanging power with the grid is determined by the time-varying electricity price $c_{\text{el}}(t)$, expressed in €/kWh. Since electricity prices can vary significantly over the course of a day, particularly under dynamic pricing schemes such as day-ahead market pricing, the grid interconnection introduces a direct coupling between the EMS cost function and the electricity market. Importing power from the grid incurs a cost proportional to $c_{\text{el}}(t)$, while exporting surplus power may generate revenue, depending on the prevailing market price and the applicable feed-in tariff [32]. The instantaneous cost of grid exchange is therefore modelled as

$$C_g(t) = c_{\text{el}}(t)P_g(t), \quad (2-6)$$

where $P_g(t) > 0$ and $P_g(t) < 0$ correspond to power import and export respectively. For the single-bus microgrid considered in this thesis, power balance on the electrical bus must be maintained at all times. Since the grid interconnection is unconstrained in direction, it provides the degree of freedom needed to satisfy this balance for any feasible combination of the remaining power flows. This balance is enforced as an algebraic equality constraint and is discussed in the context of the full system model in Section 3-2.

2-2 Thermal Modeling

The thermal subsystem of a CHP microgrid encompasses heat generation, thermal storage, and building heat distribution. This section presents the models used to describe each of these components, beginning with the heat pump that couples the electrical and thermal subsystems, followed by auxiliary thermal units, Aquifer Thermal Energy Storage (ATES) seasonal storage, and the generic thermal mass model used to describe hydronic and building components.

2-2-1 Heat Pump

A heat pump is a thermodynamic device that transfers heat from a low-temperature source to a high-temperature sink by consuming mechanical work. In the context of CHP microgrids, heat pumps provide the primary coupling between the electrical and thermal subsystems. The heat pump is modelled using a Coefficient of Performance (COP) based approach, in which the thermal output is related to the electrical input through a performance factor derived from the ideal Carnot cycle [33, Sec. 11-1–11-2].

The Carnot COP represents the theoretical upper bound on heat pump performance for a given pair of source and sink temperatures, and is given by

$$\text{COP}_{\text{Carnot}}(T_{\text{sink}}, T_{\text{source}}) = \frac{T_{\text{sink}}}{T_{\text{sink}} - T_{\text{source}}}, \quad (2-7)$$

where T_{sink} and T_{source} denote the absolute temperatures of the sink and source respectively. Real heat pumps operate below this theoretical limit due to irreversibilities in the cycle. This is captured through a second-law efficiency $\eta_{\text{hp}} \in (0, 1)$, yielding the effective COP

$$\text{COP}(T_{\text{sink}}, T_{\text{source}}) = \eta_{\text{hp}} \cdot \text{COP}_{\text{Carnot}}(T_{\text{sink}}, T_{\text{source}}). \quad (2-8)$$

While T_{sink} and T_{source} are here written as generic arguments to define the COP relationship, in practice they correspond to time-varying temperatures and are written as $T_{\text{sink}}(t)$ and T_{source} from this point forward. The thermal power delivered by the heat pump is related to the electrical input power $\dot{W}$ through the COP as

$$\dot{Q}_{\text{out}}(t) = \dot{W}(t) \cdot \text{COP}(T_{\text{sink}}(t), T_{\text{source}}(t)). \quad (2-9)$$

Conservation of energy requires that the heat pump also satisfies the first law of thermodynamics

$$\dot{Q}_{\text{out}}(t) = \dot{Q}_{\text{in}}(t) + \dot{W}(t), \quad (2-10)$$

where $\dot{Q}_{\text{in}}$ is the heat extracted from the source.

Heat pumps can be operated in both heating mode and cooling mode. In cooling mode, the direction of heat transfer is reversed: heat is extracted from a low temperature source at T_{source} , and rejected into a high temperature sink T_{sink} . The COP expressions in Eqs. (2-7) and (2-8) remain valid, but the useful output is now the heat extracted from the source rather than the heat delivered to the sink, reducing the COP. In the case that the source temperature exceeds that of the sink temperature, a heat pump can be used for the purpose of active cooling, increasing the capacity for heat transfer compared to passive cooling via a heat exchanger.

2-2-2 Auxiliary Thermal Units

Heat pumps are typically sized below the peak thermal load for economic reasons, as sizing for peak demand would result in significant overcapacity during normal operation. Auxiliary thermal units, such as electric boilers and chillers, are therefore installed to supplement the heat pump during periods of peak heating or cooling demand. Additionally, at large temperature lifts the COP given by Eq. (2-8) decreases significantly, reducing the efficiency of the heat pump. In such conditions, direct generation from an auxiliary unit may be preferable. Both the auxiliary boiler and chiller are modelled as simple COP devices, in which the heat output is proportional to the electrical input

$$\dot{Q}_b(t) = \eta_b \cdot \dot{W}_b(t), \quad (2-11)$$

$$\dot{Q}_c(t) = \eta_{ch} \cdot \dot{W}_c(t). \quad (2-12)$$

Where η_b and η_c are the constant COP values of the boiler and chiller respectively, and $\dot{W}_b$ and $\dot{W}_c$ are their electrical power inputs.

2-2-3 Aquifer Thermal Energy Storage

There is a fundamental seasonal mismatch between thermal energy availability and heating demand. In winter, heating demand is high while solar availability is low, and in summer the situation is reversed. This mismatch is illustrated in Fig. 2-1a, which shows the heating availability and heating demand for a typical year in the Netherlands. The non-overlapping regions motivate the use of seasonal thermal storage to shift summer heat surplus to winter heating demand, as illustrated in Fig. 2-1b.

Among seasonal underground thermal storage technologies, ATEs is the most widely deployed in the Netherlands [34], [35]. An ATEs system stores thermal energy as sensible heat in a pair of wells: a warm well and a cold well [36]. During winter, warm water is extracted from the warm well for heating, and the cooled return water is injected into the cold well. During summer, the process is reversed: cold water is extracted for cooling, and the warmed return water is stored in the warm well. This working principle is also shown in Fig. 2-2.




Fig. 2-1: The seasonal mismatch between thermal energy availability and heating demand (a), and the application of thermal storage in seasonal load shifting of thermal supply and demand.




Fig. 2-2: Working principle of a two-well ATEs system. Figure adapted from [37].

The present model follows the framework of Rostampour and Keviczky [22], in which each well is characterized by its volume of stored water and specific enthalpy content. In the present work, this model is simplified using a constant control volume assumption, whereby extracted groundwater is assumed to be immediately replenished at the ambient groundwater temperature. This is consistent with standard modeling practice for low-temperature ATEs systems, in which equal injection and extraction volumes per seasonal cycle result in negligible net volumetric change per well [36]. Furthermore, for systems operating below 25 °C, density variations due to thermal expansion are negligible [38], making the constant volume assumption appropriate for the system-level simulation context considered here.

The dynamics of each well are governed by an energy balance derived from the first law of thermodynamics, accounting for the enthalpy carried by injected and extracted water and for thermal losses to the surrounding aquifer. The operation of the system in both modes is shown schematically in Fig. 2-3. The thermal loss terms in heating and cooling mode are

$$\begin{aligned}\dot{Q}_{h,l}(t) &= k_{\text{ATES}} \left(\frac{h_h(t)}{c_{p,w}} - T_{\text{gw}} \right), \\ \dot{Q}_{c,l}(t) &= k_{\text{ATES}} \left(T_{\text{gw}} - \frac{h_c(t)}{c_{p,w}} \right).\end{aligned}$$

The dynamic model for the change in enthalpy of the wells in heating mode is then given as

$$V_h \rho_h \frac{dh_h(t)}{dt} = \dot{m}(t)(h_h(t) - c_{p,w} T_{\text{gw}}) - \dot{Q}_{h,l}(t), \quad (2-13)$$

$$V_c \rho_c \frac{dh_c(t)}{dt} = \dot{m}(t)(h_{\text{HX,out}}^H(t) - h_c(t)) - \dot{Q}_{c,l}(t), \quad (2-14)$$

and in cooling mode, with the flow direction reversed, as

$$V_h \rho_h \frac{dh_h(t)}{dt} = \dot{m}(t)(h_{\text{HX,out}}^C(t) - h_h(t)) - \dot{Q}_{h,l}(t), \quad (2-15)$$

$$V_c \rho_c \frac{dh_c(t)}{dt} = \dot{m}(t)(c_{p,w} T_{\text{gw}} - h_c(t)) - \dot{Q}_{c,l}(t), \quad (2-16)$$




Fig. 2-3: Schematic representation of the energy flows in the ATES system in both heating and cooling modes.

where $h_h(t)$ and $h_c(t)$ are the specific enthalpies of the warm and cold wells, $h_{\text{HX,out}}^H$ and $h_{\text{HX,out}}^C$ are the heat exchanger output enthalpies in heating and cooling mode respectively, T_{gw} is the ambient groundwater temperature, $c_{p,w}$ is the isobaric heat capacity of water, and k_{ATES} is the thermal conductance of the wells to the surrounding aquifer. Under the constant volume assumption, the well volumes V_h , V_c are treated as fixed parameters. The heat exchanger output enthalpies $h_{\text{HX,out}}^H$ and $h_{\text{HX,out}}^C$ are determined by the heat pump model presented in Section 2-2-1. Together, Eqs. (2-13) to (2-16) and the heat pump model fully describe the enthalpy dynamics of the ATES wells across both operating modes.

An important operational constraint in ATES systems is that the total energy extracted from and injected into the ATES system should be balanced across an operational year [39], [40]. This leads to a balance constraint that applies over the operational year. In this thesis, the balance constraint is formulated at the system level, requiring that the net thermal energy exchange with the ATES over the operational year is zero:

$$\int_{t_0}^T \left(\dot{Q}_{\text{out}}^H(t) - \dot{Q}_{\text{in}}^C(t) \right) dt = 0. \quad (2-17)$$

This differs from the well-level formulation of the constraint used by Rostampour and Keviczky [22], in which the energy balance is enforced over each ATES well pair. Notably, the definition in Eq. (2-17) and the one used by Rostampour and Keviczky [22] coincide in the case that only a single ATES doublet is considered, which is the case for this thesis.

The constraint in Eq. (2-17) fundamentally upper-bounds the amount of heating that can be extracted by the amount of cooling that is injected. In the Netherlands, there is significant mismatch between annual heating and cooling demand [41], with heating demand far exceeding cooling demand. Were the two roughly equal, satisfaction of this constraint would be nearly automatic, as the cooling injected over the course of a year would naturally replenish what is extracted for heating. The large imbalance between the two means, however, that the constraint is non-trivial to satisfy and must be actively managed by the EMS.

2-2-4 Thermal Masses

Many components in thermal energy systems store and release heat over time due to their finite heat capacity. Such components are referred to as thermal masses, and are commonly

modelled under the lumped capacitance assumption, which treats the temperature of the component as spatially uniform [42, p. 29]. This assumption is appropriate when the internal thermal resistance of the component is small relative to the resistance at its boundary, which holds for well-mixed fluid volumes and building elements with high thermal conductivity. Examples relevant to this thesis include water volumes in hydronic circuits and the building structure itself.

Under the lumped capacitance assumption and assuming no mass accumulation, the energy balance for a thermal mass with specific enthalpy h_m follows from the first law of thermodynamics, and is given by

$$\rho V \frac{dh_m(t)}{dt} = \dot{Q}_{\text{in}}(t) - \dot{Q}_{\text{out}}(t) + \dot{m}(t)(h_{\text{in}}(t) - h_{\text{out}}(t)), \quad (2-18)$$

where $\dot{Q}_{\text{in}}$ and $\dot{Q}_{\text{out}}$ are the net heat flows into and out of the element, and the final term accounts for energy transport by mass flow, with h_{in} and h_{out} the specific enthalpies of fluid entering and leaving the element. The isobaric heat capacity c_p varies weakly with temperature for incompressible substances and is treated as constant throughout.

Heat exchange between a thermal mass and its surroundings through conduction, convection, and radiation is commonly aggregated into a single heat flow using Newton's law of cooling [42, p. 18]

$$\dot{Q}(t) = (UA) \cdot (T_2(t) - T_1(t)), \quad (2-19)$$

where (UA) is the lumped thermal conductance, T_1 is the surrounding temperature, and T_2 is the body temperature. Solar radiative heat gain is modelled separately as

$$\dot{Q}_{\text{rad}}(t) = A_{\text{eff}} s(t), \quad (2-20)$$

where A_{eff} is the effective surface area for solar heat absorption. The models in Eqs. (2-18) to (2-20) form the generic building blocks for the reservoir, floor, and building thermal dynamics developed in Section 3-3.

2-3 Microgrid Control

Microgrids integrate a diverse set of generation, storage, and load components, each operating under different physical constraints and timescales. Coordinating these components to achieve reliable and economical operation requires a structured control framework. This section introduces the hierarchical control architecture commonly adopted in microgrid literature, identifies the control layer addressed in this thesis, and provides a high-level overview of energy management approaches.

2-3-1 Control Hierarchy

Microgrid control spans multiple timescales, ranging from microsecond-level power electronics regulation to hour-level operational planning. A widely adopted framework for organizing these tasks is the three-layer hierarchy [20, pp. 7–11], [14], [43], illustrated in Fig. 2-4.




Fig. 2-4: The hierarchical control architecture for microgrids. Figure adapted from [44].

The primary control layer operates at the shortest timescale and is responsible for local stability of voltage and frequency. In grid-connected mode, the microgrid frequency is dictated by the utility grid and P/Q control is applied. In islanded mode, the microgrid must independently maintain voltage and frequency through V/f control [14]. Primary control can leave steady-state offsets in voltage and frequency, which are corrected by the secondary control layer. Secondary control also ensures real-time power balance through load sharing between components. The tertiary control layer operates at the longest timescale and is responsible for economic and operational planning of the microgrid. This layer generates setpoints for the secondary and primary layers, accounting for forecasts of renewable generation, load, and energy prices over a planning horizon. This layer is commonly referred to as the EMS.

This thesis is concerned exclusively with the tertiary control layer. Several modeling assumptions follow from this focus. The EMS operates on timescales of minutes to hours, which are orders of magnitude slower than the transient dynamics of the electrical components. The long-term behavior of the system is therefore largely decoupled from these fast dynamics, and a steady-state assumption for all microgrid components can be adopted without significant loss of accuracy [20, p. 85], [16]. Furthermore, the system operates in grid-connected mode, meaning that the utility grid maintains the microgrid frequency within acceptable bounds. Voltage stability and power quality are similarly handled by the lower control layers, and need not be addressed by the EMS directly. It is therefore assumed that the primary and secondary controllers track the setpoints issued by the EMS with sufficient accuracy, so that their internal dynamics can be abstracted away at the energy management level [14], [16]. Finally, a single-bus representation is adopted, which is appropriate for the building-level microgrid considered in this thesis given its small spatial extent, and is consistent with the steady-state modeling assumption above. This modeling approach is consistent with standard practice in the microgrid EMS literature [20, Sec. 4.3.1], [17], [45].

2-3-2 Energy Management

The EMS is responsible for coordinating the generation, storage, and controllable loads of the microgrid in order to ensure economic and reliable operation [20, p. 84], [14]. This problem can broadly be decomposed into two sub-problems: the unit commitment (UC) problem and the economic dispatch (ED) problem [46, Ch. 3–4]. The UC problem is a combinatorial scheduling problem that determines which generators should be online at each time step over the planning horizon, accounting for startup and shutdown costs and minimum up- and downtime constraints. Given a commitment decision, the ED problem then determines the optimal continuous power dispatch of all committed generators and storage systems to meet the load at minimum cost, subject to capacity and ramp-rate constraints [46, Ch. 3–4]. In traditional power systems, these problems are solved separately due to the intractability of solving them jointly at large scale [47]. In microgrids, however, an integrated approach is often adopted given the smaller problem size, allowing the UC and ED problems to be solved simultaneously within a single optimization framework [16].

A wide range of methods have been proposed for solving the microgrid EMS problem. Mathematical programming approaches, in particular mixed-integer linear program (MILP) and mixed-integer non-linear program (MINLP), are the most widely adopted due to their ability to explicitly handle operational constraints and binary decisions such as unit commitment [16], [48]. More recently, data-driven approaches such as RL have gained interest for their ability to implicitly learn long-term value functions without relying on accurate forecasts, though these methods generally lack the constraint handling and theoretical guarantees required for safety-critical applications [49], [50].

Among optimization-based approaches, MPC has become a prominent framework for microgrid EMS [13], [43]. Its receding horizon structure naturally incorporates updated forecasts of renewable generation, load, and energy prices into the control decision at each time step, providing inherent robustness to forecasting errors through feedback. The explicit constraint handling of MPC also makes it well-suited for safety-critical microgrid applications, where operational limits on storage, generation, and grid exchange must be respected at all times [16], [51]. The formulation and application of MPC to the microgrid EMS problem is discussed in detail in Chapters 5 and 6.

Chapter 3

System Description

This chapter describes the specific microgrid system that will be used throughout this thesis for the implementation and evaluation of the control algorithms. The system is instantiated from the modeling framework established in Chapter 2, with concrete component selections, parameter assignments, and disturbance signal definitions. The electrical and thermal subsystems are described in Sections 3-2 and 3-3 respectively, followed by a description of the disturbance signals in Section 3-4, and the full state-space representation of the combined system in Section 3-5.

3-1 System Overview

The system under consideration for this thesis is a building level CHP microgrid, consisting of an electrical subsystem and a thermal subsystem that are coupled through a heat pump and auxiliary thermal units. An overview of the system is shown in Fig. 3-1.

The electrical subsystem consists of a PV array, a battery, an electrical load, and an interconnection with the main grid. The thermal subsystem consists of an ATES system for seasonal energy storage, an auxiliary electric boiler, an auxiliary chiller, and a building thermal model including a reservoir, underfloor heating, and the building envelope itself. The heat pump serves as the primary coupling between the subsystems, as well as the interface between the thermal demand and the thermal storage system.

The system is driven by four exogenous disturbance signals: the electricity price, solar irradiance, electrical load, and ambient temperature. These signals are uncontrollable and must be forecasted over the MPC horizon in order to apply predictive control. A detailed description of each disturbance is given in Section 3-4.

The control objective will be minimizing the operational cost of the microgrid system over a receding horizon, subject to physical constraints of both subsystems and a comfort constraint on the building temperature. Both the electrical and thermal subsystems are described in more detail in Sections 3-2 and 3-3.




Fig. 3-1: Overview of the system that will be used for the rest of this thesis.

3-2 Electrical Subsystem

The electrical subsystem consists of a PV array, a battery, an electrical load, and an inter-connection with the main grid. The PV power output is treated as a disturbance signal, determined from the solar irradiance via Eq. (2-1). The electrical load is similarly treated as a disturbance, modelled as a typical usage profile consisting of a sustained base load with periodic peaks in the morning and evening. The battery is the only dynamical component in the electrical subsystem, with its state of charge governed by Eq. (2-2).

For microgrid operation, power balance on the electrical bus must be maintained at all times. This is expressed as the algebraic constraint requiring the sum of all power flows at the bus to be zero. For the system shown in Fig. 3-1, this yields

$$P_g(t) - P_b(t) + P_{pv}(t) - P_l(t) - \dot{W}^H(t) - \dot{W}^C(t) - \dot{W}_b(t) - \dot{W}_c(t) = 0, \quad (3-1)$$

where P_g is the power imported from or exported to the grid, P_b is the battery power, P_{pv} is the PV generation, P_l is the electrical load, $\dot{W}_b$ and $\dot{W}_c$ are boiler and chiller work, and $\dot{W}^H$ and $\dot{W}^C$ are the electrical work inputs to the heat pump in heating and cooling mode, respectively.

3-3 Thermal Subsystem

The thermal subsystem consists of an ATES system for seasonal thermal storage, an auxiliary electric boiler, an auxiliary chiller, and a building thermal circuit. The heat pump is the primary coupling between the electrical and thermal subsystems, transferring heat between

the ATES and the reservoir. The auxiliary boiler and chiller are modelled as simple COP devices, providing supplementary heating and cooling capacity when the heat pump alone is insufficient to meet the thermal demand. The ATES dynamics are governed by Eqs. (2-13) to (2-16).

The thermal circuit consists of a reservoir, an underfloor heating circuit, and the building air. Heat is supplied to or extracted from the reservoir by the heat pump, boiler, and chiller. The reservoir in turn exchanges heat with the floor, which exchanges heat with the building interior, which finally exchanges heat with the ambient environment. This chain of components and heat flows is also shown in Fig. 3-2. Following the lumped thermal mass model and applying Newton's law of cooling at each interface, the relevant heat flows are written as

$$\dot{Q}_{rf}(t) = (UA)_{rf}(T_r(t) - T_f(t)), \quad (3-2)$$

$$\dot{Q}_{fb}(t) = (UA)_{fb}(T_f(t) - T_b(t)), \quad (3-3)$$

$$\dot{Q}_{bo}(t) = (UA)_{bo}(T_b(t) - T_{\text{amb}}(t)). \quad (3-4)$$

where $(UA)_{rf}$, $(UA)_{fb}$, and $(UA)_{bo}$ denote the thermal conductances at the reservoir-floor, floor-building, and building-ambient interfaces respectively. Applying Eq. (2-18) to each thermal mass in the circuit, the temperature dynamics of the reservoir, floor, and building are then given by

$$\begin{cases} m_r c_{p,w} \dot{T}_r(t) = \dot{Q}_{\text{out}}^H(t) - \dot{Q}_{\text{in}}^C(t) + \dot{Q}_b(t) - \dot{Q}_c(t) - \dot{Q}_{rf}(t), \\ m_f c_{p,f} \dot{T}_f(t) = \dot{Q}_{rf}(t) - \dot{Q}_{fb}(t), \\ m_{b,\text{eff}} c_{p,a} \dot{T}_b(t) = \dot{Q}_{fb}(t) - \dot{Q}_{bo}(t) + \dot{Q}_{\text{rad}}(t). \end{cases} \quad (3-5)$$

where m_r and m_f are the masses of water in the reservoir and floor circuit respectively, $c_{p,w}$ and $c_{p,f}$ are the corresponding isobaric heat capacities, and $m_{b,\text{eff}}$ is the effective thermal mass of the building, obtained from the specific thermal mass per unit floor area as defined in EN ISO 13786 [52]. This norm provides a standardized method for characterizing the dynamic thermal properties of building components, making it an appropriate basis for determining the effective thermal mass in a lumped building model of this form. The terms $\dot{Q}_{\text{out}}^H$ and $\dot{Q}_{\text{in}}^C$ denote the heat delivered to and extracted from the reservoir by the heat pump in heating and cooling mode respectively, $\dot{Q}_b$ and $\dot{Q}_c$ are the heat contributions of the auxiliary boiler and chiller, and $\dot{Q}_{\text{rad}}$ is the radiative solar heat gain of the building, modelled as described in Section 2-2.

3-4 Disturbance Signals

There are many signals in microgrid applications that are important for economic operation, but cannot be modeled in terms of a dynamic model explicitly. These signals appear in the form of disturbances, which are uncontrollable inputs to the system. The following vector of disturbance signals are present in the microgrid shown in Fig. 3-1:

$$d(t) = \begin{bmatrix} c_{\text{el}}(t) \\ s(t) \\ P_l(t) \\ T_{\text{amb}}(t) \end{bmatrix}, \quad (3-6)$$




Fig. 3-2: The thermal dynamics of the reservoir, floor and building, visualized. Figure adapted from [19].

where $c_{el}(t)$ is the electricity price in €/kWh, $s(t)$ the solar irradiance in W/m^2 , $P_l(t)$ is the electrical load in kW, and $T_{amb}(t)$ is the ambient temperature in K.

The solar irradiance and ambient temperature are obtained from the Royal Netherlands Meteorological Institute (KNMI) Environmental Data Retrieval (EDR) Application Programming Interface (API) [53], and Dutch electricity pricing data is obtained through the European Network of Transmission System Operators for Electricity (ENTSO-E) transparency platform API [1]. The electrical load is generated synthetically using a parametric profile consisting of a seasonally modulated base load and two Gaussian peaks representing morning and evening consumption patterns. Crucially, this profile does not include power consumed for heating and cooling, as this is explicitly captured by the heat pump terms in Eq. (3-1).

The seasonal modulation factors for the base load and the two peaks are defined in terms of the day of the year $n \in \{1, \dots, 365\}$ as

$$\phi_b(n) = 1 + 0.20 \cos\left(\frac{2\pi(n-1)}{365}\right), \quad (3-7)$$

$$\phi_m(n) = 1 + 0.15 \cos\left(\frac{2\pi(n-1)}{365}\right), \quad (3-8)$$

$$\phi_e(n) = 1 + 0.10 \cos\left(\frac{2\pi(n-1)}{365}\right), \quad (3-9)$$

reflecting higher electricity consumption in winter months. The load profile is then defined using a Gaussian kernel. Defining $\mathcal{G}(\tau; \mu, \sigma) = \exp(-(\tau - \mu)^2 / 2\sigma^2)$, the load profile is written as

$$P_l(t) = \bar{P}_b \phi_b(n) + \bar{P}_m \phi_m(n) \mathcal{G}(\tau(t); 7.5, 1.5) + \bar{P}_e \phi_e(n) \mathcal{G}(\tau(t); 19.5, 2.0) + w(t), \quad (3-10)$$

where $w(t)$ is additive zero-mean white noise, $\tau(t) \in [0, 24)$ denotes the time of day in hours corresponding to time t , $\bar{P}_b$, $\bar{P}_m$, and $\bar{P}_e$ are the nominal base load, morning peak, and evening




Fig. 3-3: Representative two day samples of the load profile for each season, visualizing the mathematical description of the profile shown in Eq. (3-10).

peak amplitudes respectively. The morning peak is centred at 07:30 with a standard deviation of 1.5 h, and the evening peak is centred at 19:30 with a standard deviation of 2.0 h. The load profile is also visualized in Fig. 3-3 to illustrate how the profile develops across the four seasons. The nominal parameter values used in this thesis are listed in Table 3-1.

3-5 State-Space Representation

Collecting the dynamics derived in Sections 3-2 and 3-3, the system state, control input, and disturbance vectors are defined as

$$x(t) = \begin{bmatrix} x_b(t) \\ h_h(t) \\ h_c(t) \\ h_r(t) \\ h_f(t) \\ h_b(t) \end{bmatrix}, \quad u(t) = \begin{bmatrix} P_b(t) \\ \dot{W}^H(t) \\ \dot{W}^C(t) \\ \dot{Q}_b(t) \\ \dot{Q}_c(t) \\ P_g(t) \end{bmatrix}, \quad d(t) = \begin{bmatrix} c_{el}(t) \\ s(t) \\ P_l(t) \\ T_{amb}(t) \end{bmatrix}, \quad (3-11)$$

where x_b is the battery state of charge, h_h and h_c are the specific enthalpies of the hot and cold ATES wells, h_r , h_f , and h_b are the specific enthalpies of the reservoir, floor circuit, and building respectively.

Since the ATES dynamics differ between heating and cooling mode, the system admits two distinct operating modes. The state-space model in heating mode, denoted $\dot{x}(t) =$

$f^H(x(t), u(t), d(t))$, is given by

$$\begin{bmatrix} \dot{x}_b(t) \\ \dot{h}_h(t) \\ \dot{h}_c(t) \\ \dot{h}_r(t) \\ \dot{h}_f(t) \\ \dot{h}_b(t) \end{bmatrix} = \begin{bmatrix} \frac{1}{C_{\max}} (-\eta(P_b)P_b(t) - x_{sd}) \\ \frac{1}{V_h \rho_h} \left(\dot{m}(t)(h_h(t) - c_{p,w}T_{gw}) - k_{ATES} \left(\frac{h_h(t)}{c_{p,w}} - T_{gw} \right) \right) \\ \frac{1}{V_c \rho_c} \left(\dot{m}(t)(h_{HX,out}^H(t) - h_c(t)) - k_{ATES} \left(T_{gw} - \frac{h_c(t)}{c_{p,w}} \right) \right) \\ \frac{1}{m_r c_{p,w}} \left(\dot{Q}_{out}^H(t) + \dot{Q}_b(t) - \dot{Q}_c(t) - \dot{Q}_{rf}(t) \right) \\ \frac{1}{m_f c_{p,f}} \left(\dot{Q}_{rf}(t) - \dot{Q}_{fb}(t) \right) \\ \frac{1}{m_{b,eff} c_{p,a}} \left(\dot{Q}_{fb}(t) - \dot{Q}_{bo}(t) + \dot{Q}_{rad}(t) \right) \end{bmatrix}, \quad (3-12)$$

and in cooling mode, denoted $\dot{x}(t) = f^C(x(t), u(t), d(t))$, as

$$\begin{bmatrix} \dot{x}_b(t) \\ \dot{h}_h(t) \\ \dot{h}_c(t) \\ \dot{h}_r(t) \\ \dot{h}_f(t) \\ \dot{h}_b(t) \end{bmatrix} = \begin{bmatrix} \frac{1}{C_{\max}} (-\eta(P_b)P_b(t) - x_{sd}) \\ \frac{1}{V_h \rho_h} \left(\dot{m}(t)(h_{HX,out}^C(t) - h_h(t)) - k_{ATES} \left(\frac{h_h(t)}{c_{p,w}} - T_{gw} \right) \right) \\ \frac{1}{V_c \rho_c} \left(\dot{m}(t)(c_{p,w}T_{gw} - h_c(t)) - k_{ATES} \left(T_{gw} - \frac{h_c(t)}{c_{p,w}} \right) \right) \\ \frac{1}{m_r c_{p,w}} \left(\dot{Q}_{in}^C(t) + \dot{Q}_b(t) - \dot{Q}_c(t) - \dot{Q}_{rf}(t) \right) \\ \frac{1}{m_f c_{p,f}} \left(\dot{Q}_{rf}(t) - \dot{Q}_{fb}(t) \right) \\ \frac{1}{m_{b,eff} c_{p,a}} \left(\dot{Q}_{fb}(t) - \dot{Q}_{bo}(t) + \dot{Q}_{rad}(t) \right) \end{bmatrix}. \quad (3-13)$$

Note that the electrical, floor, and building dynamics are identical across both modes; only the ATES well dynamics and the reservoir heat input differ. The combined system can therefore be written compactly as

$$\dot{x}(t) = f(x(t), u(t), d(t)) = \begin{cases} f^H(x(t), u(t), d(t)), & \text{heating mode,} \\ f^C(x(t), u(t), d(t)), & \text{cooling mode,} \end{cases} \quad (3-14)$$

which is a nonlinear, hybrid state-space model. For simulation purposes, Eq. (3-14) is discretized using the RK4 method, yielding the discrete-time model

$$x(t+1) = \begin{cases} F^H(x(t), u(t), d(t)), & \text{heating mode,} \\ F^C(x(t), u(t), d(t)), & \text{cooling mode,} \end{cases} \quad (3-15)$$

where t denotes the discrete time index and F^H and F^C are the RK4 integrators applied to f^H and f^C over one sampling interval Δt .

3-6 System Parameters

The system parameters used throughout this thesis are listed in Table 3-1. The building is modelled as a four-story office building with a 50 m $\times$ 50 m floor plate, giving a total floor area of 10,000 m² and a total volume of 40,000 m³. The effective thermal mass of the building is derived from a specific thermal capacity of 150 kJ/m²K per floor, as per EN ISO 13786 [52]. The ATES wells are sized at 5,000 m³ each. While smaller than typical Dutch ATES systems [36], the volumes are sufficient for illustrating the algorithmic behavior of the EMS studied in this thesis. The heat pump second-law efficiency is set to $\eta = 0.5$, consistent with values reported for residential building-scale heat pump systems [54].

Table 3-1: System parameters used throughout this thesis.

Component	Parameter	Symbol	Value
Building	Floor area	A_{floor}	10,000 m ²
	Effective thermal mass	$c_{m,\text{eff}}$	150 kJ/m ² K
	Comfort band (lower)	$T_{b,\text{min}}$	20 °C
	Comfort band (upper)	$T_{b,\text{max}}$	21 °C
Building envelope	Wall area	A_{wall}	3,200 m ²
	Wall conductance	$(UA)_{bo}$	1.6 kW/K
	Effective solar area	A_{eff}	32 m ²
Floor	Thickness	δ_f	0.1 m
	Density	ρ_f	2400 kg/m ³
	Heat capacity	$c_{p,f}$	0.88 kJ/kgK
	Floor-reservoir conductance	$(UA)_{rf}$	80 kW/K
	Floor-building conductance	$(UA)_{fb}$	70 kW/K
Reservoir	Volume	V_r	25 m ³
	Temperature range	$[T_{r,\text{min}}, T_{r,\text{max}}]$	[17, 40] °C
	Max. mass flow	$\dot{m}_{r,\text{max}}$	0.5 kg/s
ATES	Hot well volume	V_h	5,000 m ³
	Cold well volume	V_c	5,000 m ³
	Groundwater temperature	T_{gw}	12 °C
	Thermal conductivity	k_{ATES}	0.03 kW/K
	Hot well temp. range	$[T_{h,\text{min}}, T_{h,\text{max}}]$	[15, 23.5] °C
	Cold well temp. range	$[T_{c,\text{min}}, T_{c,\text{max}}]$	[5, 12] °C
	Max. mass flow	$\dot{m}_{\text{max}}$	0.4 kg/s
Heat pump	Second-law efficiency	η_{hp}	0.5 [–]
	Max. electrical input	$\dot{W}_{\text{max}}$	200 kW
Aux. boiler	Max. heat output	$\dot{Q}_{b,\text{max}}$	100 kW
	COP	η_b	0.9 [–]
Aux. chiller	Max. cooling output	$\dot{Q}_{c,\text{max}}$	100 kW
	COP	η_{ch}	0.9 [–]
Battery	Capacity	C_{max}	10 kWh
	Charging efficiency	η_c	0.9 [–]
	Discharging efficiency	η_d	0.8 [–]
	Max. power	$P_{b,\text{max}}$	100 kW
	SoC range	$[x_{b,\text{min}}, x_{b,\text{max}}]$	[0, 1] [–]
PV array	Panel area	A_{pv}	500 m ²
	Panel efficiency	η_{pv}	0.17 [–]
Electrical load	Base load	$\bar{P}_b$	10 kW

Economic Model Predictive Control for Microgrids

This chapter presents the theoretical framework underlying the control design developed in this thesis. The chapter begins with a review of MPC fundamentals in Section 4-1, covering the finite horizon optimal control problem (OCP), terminal ingredients, and constraint softening. Section 4-2 then introduces EMPC, in which the stage cost directly reflects economic operation rather than tracking a setpoint, and discusses the role of dissipativity in establishing stability. The fundamental challenge of applying EMPC to microgrid systems, namely the absence of a meaningful steady state, is identified and motivates the time-varying terminal ingredient framework. Section 4-3 then introduces the MLD framework, which provides a systematic approach to incorporating the discrete operating mode decisions of microgrid components into the MPC optimization problem, yielding a tractable MILP formulation. Together, these ingredients form the basis for the controller design following in Chapters 5 and 6.

4-1 Model Predictive Control Fundamentals

MPC has become a prominent framework for microgrid EMS [13], [43], owing to several properties that align well with the microgrid energy management problem. At each sampling instance, MPC solves a finite horizon OCP using the current system state and forecasts of the disturbance signals, such as renewable generation and electricity price, over the planning horizon [55]. Only the first element of the resulting optimal input sequence is applied, after which the horizon recedes and the problem is solved again with updated information. This principle is also visually illustrated in Fig. 4-1. The receding horizon structure introduces a feedback loop that provides inherent robustness to disturbances and model mismatch [56, Ch. 1], both of which are persistent challenges in microgrid applications [13]. Furthermore, operational constraints on storage, generation, and grid exchange are incorporated directly into the optimization problem, ensuring that the computed control actions are always physically admissible [51], [55]. Finally, MPC naturally handles the multi-input, multi-output




Fig. 4-1: Two time steps illustrating the working principle of MPC.

nature of microgrid systems, where multiple generation, storage, and load components must be coordinated simultaneously.

For MPC, dynamical systems of the discrete-time form

$$x^+ = f(x, u, d), \quad (4-1)$$

are considered, with $x \in \mathbb{X} \subseteq \mathbb{R}^n$ the state vector, $u \in \mathbb{U} \subseteq \mathbb{R}^m$ the input vector, and $d \in \mathbb{D} \subseteq \mathbb{R}^p$ the disturbance vector. The function $f : \mathbb{X} \times \mathbb{U} \times \mathbb{D} \rightarrow \mathbb{X}$ denotes state transition map describing the system dynamics. Based on this model, the finite horizon OCP solved at each sampling instant is formulated as follows. Adopting a slight abuse in notation in which time-indexed variables are written with subscripts (e.g., $x(k)$ as x_k), the finite horizon OCP solved at each sampling instance is stated as

$$\mathbb{P}_N(x_{\text{init}}, \mathbf{d}_t) : \begin{cases} \min_{\mathbf{u}} & \sum_{k=0}^{N-1} \ell(x_k, u_k, d_k) + V_f(x_N) \\ \text{s.t.} & x_{k+1} = f(x_k, u_k, d_k), \quad \forall k \in \mathbb{I}_{N-1}^+, \\ & x_0 = x_{\text{init}}, \\ & (x_k, u_k) \in \mathbb{Z}, \quad \forall k \in \mathbb{I}_{N-1}^+, \\ & x_N \in \mathbb{X}_f, \end{cases} \quad (4-2)$$

where $\mathbf{u} = \{u(k)\}_{k=0}^{N-1}$ is the sequence of control inputs to be optimized, $\mathbf{d}_t = \{d(t+k)\}_{k=0}^{N-1}$ is the forecast of disturbance signals over the horizon, and $\mathbb{Z} \subseteq \mathbb{X} \times \mathbb{U}$ is the set of feasible state-input pairs encoding the operational constraints of the system. The stage cost $\ell : \mathbb{X} \times \mathbb{U} \times \mathbb{D} \rightarrow \mathbb{R}$ captures the cost of operation at each step of the horizon. The optimal input sequence is denoted $\mathbf{u}^0 = \{u^0(k)\}_{k=0}^{N-1}$, and only the first element $u^0(0)$ is applied to the system. The horizon then recedes by one step and Eq. (4-2) is solved again at the next sampling instance with updated state and disturbance information.

The terms $V_f(\cdot)$ and $\mathbb{X}_f$ are referred to as the terminal cost and terminal constraint set respectively, and play a central role in establishing stability of the closed-loop system. Since

Eq. (4-2) is a finite horizon truncation of the underlying infinite horizon OCP, the finite horizon optimal value function $V_N^0(x)$ does not in general satisfy a Lyapunov decrease condition, and asymptotic stability of the closed-loop system is therefore not guaranteed without additional structure [57, Sec. 2.4]. The standard approach to recovering stability is to choose $V_f(\cdot)$ and $\mathbb{X}_f$ such that the following conditions hold [55][57, Theorem 2.19]: (i) $\mathbb{X}_f$ is positive invariant under a terminal control law $\kappa_f : \mathbb{X}_f \rightarrow \mathbb{U}$ and (ii) $V_f(\cdot)$ satisfies the Lyapunov decrease condition $V_f(f(x, \kappa_f(x))) - V_f(x) \leq -\ell(x, \kappa_f(x))$ for all $x \in \mathbb{X}_f$. Under these conditions, $V_N^0(x)$ can be shown to be a Lyapunov function for the closed-loop system, establishing asymptotic stability with respect to the desired steady state [55]. A standard choice of terminal cost is the optimal infinite horizon cost-to-go $V_\infty^0(x)$, since, by the Bellman optimality equation, this satisfies a global Lyapunov decrease property, thereby ensuring satisfaction of properties (i) and (ii) globally [57, p. 113]. In practice, $V_\infty^0(x)$ is generally not available in closed form, and an approximation $\tilde{V}_f(x) \approx V_\infty^0(x)$ is used instead. Since such approximations are typically only valid locally, the terminal constraint set $\mathbb{X}_f$ is introduced to restrict the terminal state to a region where the approximation is known to be accurate and the Lyapunov decrease condition can be verified [57, Sec. 2.4].

The joint state-input constraints in Eq. (4-2) are written as hard constraints, meaning that any violation renders the problem infeasible. In practice, when the system is subject to disturbances, hard constraints can cause the OCP to become infeasible, particularly when the disturbance realization is not perfectly anticipated by the forecast [57, p. 131]. It is therefore standard practice to relax constraints that are not safety-critical into soft constraints, which permit bounded violations at the cost of a penalty in the objective [57, p. 131]. For a constraint of the form $g(x, u) \leq 0$, the softened formulation introduces a slack variable $s \geq 0$ and replaces the hard constraint with

$$g(x, u) \leq s, \quad (4-3)$$

$$\tilde{\ell}(x, u, s) = \ell(x, u) + \rho s, \quad (4-4)$$

where $\tilde{\ell}$ is the modified stage cost and $\rho > 0$ is a penalty weight that governs the trade-off between constraint satisfaction and optimality [57], [58, p. 131]. The slack variable s becomes a decision variable in the optimization, and the optimizer will only permit a violation when doing so reduces the overall cost sufficiently. By choosing ρ sufficiently large, constraint violations can be made arbitrarily costly, recovering behavior close to the hard constraint formulation while maintaining feasibility.

4-2 Economic Model Predictive Control

EMPC generalizes the MPC framework by replacing the positive-definite tracking stage cost with an arbitrary economic stage cost $\ell(x, u)$ that directly reflects the operational cost of the system [59], [57, p. 153]. In the microgrid context, this cost typically includes electricity purchase costs, battery degradation, and auxiliary generation costs. The OCP formulation remains identical to Eq. (4-2), and soft constraints are applied for the same reasons discussed in Section 4-1.

The stability analysis of EMPC is more subtle than that of standard tracking MPC, since the stage cost is no longer positive definite with respect to any particular state, and $V_N^0(x)$

therefore cannot directly serve as a Lyapunov function candidate for the closed-loop system. The standard approach to establishing stability in EMPC is through the notion of dissipativity [59], [60]. The system is said to be dissipative with respect to the supply rate $s(x, u) = \ell(x, u) - \ell(x_s, u_s)$ if there exists a storage function $\lambda : \mathbb{X} \rightarrow \mathbb{R}$ such that

$$\lambda(f(x, u)) - \lambda(x) \leq s(x, u), \quad \forall (x, u) \in \mathbb{Z}, \quad (4-5)$$

where (x_s, u_s) is the optimal steady-state pair, defined as the state-input pair minimizing $\ell(x, u)$ subject to $x = f(x, u)$ and $(x, u) \in \mathbb{Z}$ [57, pp. 153–154]. Under this dissipativity condition, the rotated stage cost $\tilde{\ell}(x, u) = \ell(x, u) - \ell(x_s, u_s) + \lambda(x) - \lambda(f(x, u))$ is positive semi-definite, and the rotated value function $\tilde{V}_N^0(x) = V_N^0(x) + \lambda(x) - N\ell(x_s, u_s)$ can be shown to satisfy a Lyapunov decrease condition along the closed-loop trajectories. It follows that economic performance bounds can be established, showing that the EMPC asymptotic average cost is no worse than operating at steady state [59]. Notably, the stated result is weak in the sense that there may exist multiple trajectories satisfying this property, and hence convergence to x_s cannot be guaranteed. When the system is strictly dissipative the rotated stage cost $\tilde{\ell}(x, u)$ is positive definite rather than merely semi-definite. In this case it can be shown that the state x_s is an asymptotically stable equilibrium point for the closed-loop system $\mathcal{X}_N$ [59, Theorem. 2].

A fundamental challenge in applying the dissipativity framework to microgrid EMS is the absence of a meaningful steady state: the microgrid is driven by time-varying renewable generation, load, and price signals that admit no steady-state description, making Eq. (4-5) difficult to verify. The literature addresses this in two main ways. One is to add a regularization term penalizing deviation from a chosen steady state, ensuring dissipativity by design [18]. However, this steady state is a design parameter rather than a true economic optimum, and the regularization can dominate the stage cost, sacrificing economic performance for stability around a physically meaningless operating point. The other, more common in microgrid MPC, is to omit terminal ingredients entirely or replace them with heuristic end-of-horizon storage constraints [21], [45], which avoids the dissipativity problem but yields myopic behavior with no formal guarantees on recursive feasibility or disturbance robustness.

The framework of McAllister [58] establishes sufficient conditions on the terminal cost and terminal constraint set such that EMPC is robustly recursively feasible and inherently robust to small disturbances, without requiring dissipativity. Specifically, it is required that there exists a terminal control law $\kappa_f : \mathbb{X}_s \rightarrow \mathbb{U}$ and a continuous Lyapunov function $V_s : \mathbb{X} \rightarrow \mathbb{R}_{\geq 0}$ for the system $x^+ = f(x, \kappa_f(x))$ on a positive invariant set $\mathbb{X}_s \subseteq \mathbb{X}$ with respect to a steady state $x_s \in \mathbb{X}_s$. The terminal constraint set is defined as a sublevel set of V_s

$$\mathbb{X}_f = \{x \in \mathbb{X} \mid V_s(x) \leq \tau\}, \quad (4-6)$$

for some $\tau > 0$ chosen such that $\mathbb{X}_f \subseteq \mathbb{X}_s$, and the terminal cost $V_f : \mathbb{X} \rightarrow \mathbb{R}$ must satisfy the decrease condition

$$V_f(f(x, \kappa_f(x))) \leq V_f(x) - \ell(x, \kappa_f(x)) + \ell(x_s, u_s), \quad \forall x \in \mathbb{X}_f. \quad (4-7)$$

Under these conditions, McAllister [58, Theorem 10] establishes two key guarantees. First, the terminal constraint set $\mathbb{X}_f$ is robustly positive invariant and, consequently, the EMPC remains recursively feasible for all sufficiently small disturbances. Second, for any bounded

disturbance sequence, the degradation of nominal performance is bounded by the sum of the steady-state stage-cost and the magnitude of the disturbance signal.

A standard approach for constructing the terminal ingredients, outlined by McAllister [58] following Amrit et al. [61], proceeds by linearizing the system around the steady state (x_s, u_s) . A stabilizing terminal control law is obtained via discrete-time Linear Quadratic Regulator (LQR) control, and the associated Lyapunov function yields both the terminal set $\mathbb{X}_f$ as a sublevel set as in Eq. (4-6), chosen sufficiently small that all constraints are satisfied within it, and the terminal cost V_f , constructed to satisfy Eq. (4-7).

For time-varying systems, this construction must be carried out recursively backwards from a terminal time T . Starting from $V_f(T, x) = 0$ and $\mathbb{X}_f(T) \subseteq \mathbb{X}$ at $t = T$, the terminal ingredients at each preceding time step $t = T - 1, T - 2, \dots, 0$ are obtained by applying the linearized construction above to the system frozen at that time step, using the terminal cost from the subsequent step as the prior. This backward recursion propagates the value of future operation into the terminal cost at each time step, providing the EMS with an approximation of the long-term cost-to-go on a short horizon.

While this procedure is theoretically well-founded, it presents significant practical challenges in the microgrid context. The MLD system description makes implementation of an admissible terminal control law difficult. In particular, piecewise control laws need to be designed to account for the multitude of operating modes [62]. The construction of a terminal control law and terminal value function may also depend on time-varying disturbance signals, such as ambient temperature or solar irradiance, which are not available *a priori* in practical applications. Finally, for MLD systems, input constraints can be of the form $u \in [u_{\min}, u_{\max}] \cup \{0\}$, which is a disjoint, and hence non-convex, input constraints, further complicating the construction of admissible terminal control laws.

4-3 Mixed-Logical Dynamical Model Predictive Control

Many components in microgrid energy management involve discrete operating modes in addition to continuous dynamics. The heat pump switches between heating and cooling mode, the ATES system alternates between charging and discharging, and the battery has distinct efficiency characteristics depending on the direction of power flow. Incorporating these discrete decisions directly into the MPC framework requires a modeling formalism that can handle both continuous and discrete variables within a single optimization problem. The MLD framework provides exactly this, by encoding logical rules and discrete operating constraints as mixed-integer linear inequalities that can be solved to global optimality using standard branch-and-bound solvers. This section introduces the MLD framework, discusses the linearization of bilinear terms arising from discrete actuators, and presents the resulting mixed logical dynamical model predictive control (MLD-MPC) formulation. The stability properties of the resulting controller are then discussed.

4-3-1 Discrete Actuators and Integer Programming

A distinction can be drawn between systems which have both continuous and discrete dynamics, and systems with discrete actuators [63], [64]. A true hybrid system is typically




Fig. 4-2: Difference between continuous input set (a), integer constrained input set (b), and binary constrained input set (c). Figure adapted from [66].

formulated as [65]

$$\begin{cases} \dot{x} = f(x, u), & x \in \mathcal{C}, \\ x^+ = g(x, u), & x \in \mathcal{D}. \end{cases} \quad (4-8)$$

Here $f(\cdot)$ is the system flow map, and $g(\cdot)$ is the system jump map. The dynamics of this system are hybrid in the sense that states in the jump set $\mathcal{D}$ instantaneously jump in accordance to the jump map, whereas states in $\mathcal{C}$ behave according to the continuous dynamics. The systems considered in this thesis, are instead considered to be discrete-actuator systems. That is, systems which have both discrete and continuous inputs, but not necessarily continuous and discrete dynamics [66]. Note that the statement of the dynamics of these systems is, in principle, not different from that of regular dynamical systems. That is, the system dynamics can be stated as

$$\dot{x} = f(x, u), \quad (x, u) \in \mathbb{Z} \subseteq \mathbb{X} \times \mathbb{U}, \quad (4-9)$$

with the set $\mathbb{U}$ enforcing discreteness of 1 or more of the actuators. It should be noted that the term *hybrid system* is in the literature often used interchangeably with systems with discrete actuators [67], [68]. However, for the purpose of this thesis, a distinction will be drawn due to the difference in required stability theory as will be discussed in Section 4-3-6.

The presence of discrete actuators in the system results in $\mathbb{U}$ being written as a mixed-continuous discrete set rather than as a purely continuous set. This is visualized in figure Fig. 4-2, which shows different subsets of $\mathbb{R}^2$ defined using linear-, and mixed-integer linear constraints. Incorporating discrete input constraints into the MPC optimization problem state in Eq. (4-2) transforms the continuous optimization problem into a mixed-integer optimization problem. The resulting optimization problem is referred to as an MILP or MINLP depending on the nature of the cost function and constraints [69]. The MLD framework naturally gives rise to linear constraints [67], resulting in MILP problems given that the cost function is also modeled to be linear.

4-3-2 The MLD Framework

The MLD framework provides a structured approach to modeling dynamical systems that combine continuous linear dynamics with logical rules and operational constraints [67]. By

converting logical conditions into mixed-integer linear inequalities, the MLD framework allows discrete decisions such as the operational mode of a heat pump or the charging state of a battery to be incorporated directly into a linear system description. The resulting model takes the form [67]

$$\begin{cases} x(t+1) = Ax(t) + B_1u(t) + B_2\delta(t) + B_3z(t), \\ \text{s.t. } E_2\delta(t) + E_3z(t) \leq E_1u(t) + E_4x(t) + E_5, \end{cases} \quad (4-10)$$

where $x \in \mathbb{X}$ are the system states, $u \in \mathbb{U}$ the continuous inputs, $\delta \in \mathbb{B}^{n_\delta}$ the binary variables encoding discrete decisions, and $z \in \mathbb{R}^{n_z}$ continuous auxiliary variables introduced by the linearization of logical conditions. The matrices $E_1, \dots, E_5$ encode the mixed-integer linear constraints that enforce the logical rules of the system. While the MLD framework formally permits non-linear dynamics and constraints, the formulation is in practice restricted to linear dynamics to preserve tractability of the resulting optimization problem [67].

The logical conditions encoded in the MLD constraints arise from several sources in microgrid applications. Operational mode selections, such as whether the heat pump operates in heating or cooling mode, introduce binary variables that switch between different sets of continuous dynamics. Implication constraints of the form “if $\delta = 1$ then $g(x, u) \leq 0$ ” can be encoded as linear inequalities using the Big-M method, discussed in Section 4-3-3. Products of binary and continuous variables, which arise for example when a power flow is only active in a particular operating mode, can similarly be linearized [67]. Together, these tools allow the MLD framework to capture the combinatorial structure of the microgrid EMS problem within a single linear system description.

4-3-3 Binary Decisions as Linear Constraints

A common type of constraint encountered in microgrid applications is binary products between continuous variables and binary variables. Such bilinear terms arise naturally when a continuous quantity, such as a power flow, is only active in a particular discrete operating mode. These products take the form

$$z = x \cdot \delta, \quad (4-11)$$

where $x \in \mathbb{R}$ is a continuous variable, $\delta \in \mathbb{B}$ is a binary variable, and $z \in \mathbb{R}$ an auxiliary variable. Since Eq. (4-11) is non-linear, it cannot be incorporated directly into an MILP formulation. However, given that x has known lower- and upper-bounds m and M , the product can be linearized through the following four linear constraints [67]

$$\begin{cases} z \geq m\delta, \\ z \leq M\delta, \\ z \leq x - m(1 - \delta), \\ z \geq x - M(1 - \delta). \end{cases} \quad (4-12)$$

It is straightforward to verify that when $\delta = 1$, the constraints enforce $z = x$, and when $\delta = 0$, the constraints enforce $z = 0$, which is exactly the behavior described in Eq. (4-11). This linearization technique is commonly referred to as Big-M linearization.

The tightness of the Big-M bounds m and M has a direct impact on the computational tractability of the resulting MILP. During the Linear Program (LP) relaxation step of the

branch-and-bound algorithm used to solve MILP problems, the binary variables δ are treated as continuous variables in $[0, 1]$. The Big-M bounds determine the feasible region of this relaxation, and loose bounds result in a weak relaxation that poorly approximates the true integer feasible set. This increases the number of branch-and-bound nodes required to certify optimality, significantly increasing solver computation time [69]. It is therefore important that m and M are the tightest known lower- and upper-bound on x , which in the context of microgrids, typically follows from physical limits of the components, such as storage capacity or maximum power.

A common requirement is mutual exclusivity among binary decisions. For example, a system with distinct heating and cooling modes cannot occupy both simultaneously. Mutual exclusivity of binary variables δ_1 and δ_2 can be written as the constraint $\delta_1 \cdot \delta_2 = 0$, which is bi-linear. This bi-linear constraint can identically be expressed as the linear inequality

$$\delta_1 + \delta_2 \leq 1, \quad (4-13)$$

Many other logic rules admit similar linear encodings, but these will not be discussed here and the reader is referred to Bemporad and Morari [67] or Williams [70] for a comprehensive treatment.

Another useful transformation is the linearization of absolute values. Suppose the cost function contains terms of the form $|x|$, with $x \in \mathbb{R}$. These terms can not directly be expressed as linear constraints. Instead, auxiliary variable $s \geq 0$ is introduced, with linear constraints

$$\begin{cases} s \leq x, \\ s \geq -x, \end{cases} \quad (4-14)$$

imposed. The $|x|$ terms can then be replaced by s , since minimizing s subject to these constraints forces $s = |x|$ at optimality.

The absolute value $|x|$ can be linearized in two ways depending on the context. The formulation in Eq. (4-14) is appropriate when positive and negative deviation are penalized equally. If instead asymmetric penalties are desired (e.g., penalizing overshooting of a target more than undershooting) the signed deviation can be decomposed in components $s^+, s^- \geq 0$ and enforcing

$$\begin{cases} x = s^+ - s^-, \\ s = s^+ + s^-, \end{cases} \quad (4-15)$$

which again recovers $s = |x|$ at optimality. Distinct penalties $\gamma^+, \gamma^- \geq 0$ can be applied to s^+ and s^- separately in the objective, with the symmetric case being recovered for $\gamma^+ = \gamma^-$.

4-3-4 PWA Approximation of Nonlinear Functions

Many components in microgrid EMS admit non-linear models. Some of these, such as the directional battery efficiency, can be easily linearized. However, a heat pump COP, for example, is a non-linear function of the source and sink temperatures, and cannot be linearized through the methods of Section 4-3-3. Incorporating the non-linear model directly into an optimization problem would result in a MINLP, which is in general significantly harder to solve than an MILP [69]. A common strategy to preserve tractability of the MILP formulation




Fig. 4-3: PWA approximation of a function of one variable, illustrating the relationship between partition fineness and corresponding approximation error.

is to approximate non-linear functions using piece-wise affine (PWA) functions, which can be encoded as linear constraints in MLD format under the assumption that the function input domain is compact with bounded output [71]. As shown in Fig. 4-3, a finer partition yields a more accurate approximation at the cost of introducing additional binary variables and hence increasing the computational burden of the solver.

For a function of a single variable, a PWA approximation is constructed by partitioning the domain into intervals and fitting an affine function to each, with the approximation being exact at the interval endpoints. This is also shown visually in Fig. 4-3. The construction for functions of two variables, which is the case of interest in this thesis, follows a similar principle but requires additional care in handling the two-dimensional domain.

For a function of two variables $f : \mathcal{D} \rightarrow [f_{\min}, f_{\max}]$ defined on a compact rectangular domain $\mathcal{D} = [x_{\min}, x_{\max}] \times [y_{\min}, y_{\max}] \subset \mathbb{R}^2$, construction is more involved. Following Babayev [72], the domain $[x_{\min}, x_{\max}] \times [y_{\min}, y_{\max}]$ is partitioned into a rectangular grid of $n_x \times n_y$ cells. For each inner point (i, j) where $2 \leq i \leq n_x - 1$ and $2 \leq j \leq n_y - 1$, the upper left and lower right triangles are defined by the vertices $\{(i-1, j), (i, j+1), (i, j)\}$ and $\{(i, j-1), (i+1, j), (i, j)\}$ respectively. The PWA approximation is exact on each grid point (i, j) , and affine on the interior and edges of the triangles.

The approximation is implemented by associating a non-negative weight $\lambda_{ij} \geq 0$ with each grid point, constrained to form a convex combination over the grid, $\sum_{i=1}^{n_x} \sum_{j=1}^{n_y} \lambda_{ij} = 1$. The continuous variables x and y , and the approximated function value $f_{\text{PWA}}(x, y)$ are then expressed as the weighted sums of the grid point coordinates and function values

$$x = \sum_{i=1}^{n_x} \sum_{j=1}^{n_y} x_i \lambda_{ij}, \quad (4-16)$$

$$y = \sum_{i=1}^{n_x} \sum_{j=1}^{n_y} y_j \lambda_{ij}, \quad (4-17)$$

$$f_{\text{PWA}} = \sum_{i=1}^{n_x} \sum_{j=1}^{n_y} f_{ij} \lambda_{ij}, \quad (4-18)$$




Fig. 4-4: Illustration of the 2D PWA approximation scheme with upper triangle active on the left and lower triangle active on the right. Figure adapted from [72].

where $f_{ij} = f(x_i, y_j)$ are the function values evaluated at the vertices. Binary variables $\delta_{ij} \in \mathbb{B}$ and $\delta_j^i \in \mathbb{B}$ are introduced to enforce locality, ensuring that the active weights λ_{ij} are restricted to the corners of a single active triangle. Specifically, constraint

$$\lambda_{ij} \leq \delta_{i,j-1} + \delta_j^{i-1} + \delta_{i,j} + \delta_{j+1}^i + \delta_{i+1,j} + \delta_j^i, \quad (4-19)$$

together with

$$\sum_{i=1}^{n_x} \sum_{j=1}^{n_y} (\delta_{ij} + \delta_j^i) = 1, \quad (4-20)$$

ensures that exactly one triangle is active at a time and that only the three λ_{ij} values at the vertices of that triangle are permitted to be nonzero, with all others driven to zero by Eq. (4-19). Within the active triangle, the three nonzero weights form the barycentric coordinates of the point p with respect to the triangle vertices, and sum to one by Eq. (4-18). Each weight is proportional to the area of the sub-triangle formed by p and the opposite edge, so that vertices closer to p receive larger weights. The interpolated function value $\tilde{f}(x, y)$ in Eq. (4-18) is then a weighted average of the function values at the three active vertices, producing an affine interpolation of f over the active triangle that is exact at all grid points. The full set of constraints constitutes a valid MILP encoding of the PWA approximation, with the number of binary variables scaling as $\mathcal{O}(n_x \cdot n_y)$ [72]. This procedure is illustrated in Fig. 4-4, where the upper-left and lower-right triangles associated with grid point (i, j) are activated by $\delta_{i,j} = 1$ and $\delta_j^i = 1$ respectively.

4-3-5 MLD-MPC

The MLD-MPC framework combines the MLD system description of Section 4-3-2 with the MPC OCP of Eq. (4-2), yielding an optimization problem that simultaneously optimizes continuous power dispatch and discrete operational decisions over the planning horizon. Sub-

stituting the MLD dynamics into Eq. (4-2), the resulting OCP is stated as

$$\mathbb{P}_N^{\text{MLD}}(x_{\text{init}}, \mathbf{d}_t) : \begin{cases} \min_{\mathbf{u}, \boldsymbol{\delta}, \mathbf{z}} & \sum_{k=0}^{N-1} \ell(x_k, u_k, d_k, \delta_k, z_k) + V_f(x_N) \\ \text{s.t.} & x_{k+1} = Ax_k + B_1 u_k + B_2 \delta_k + B_3 z_k + B_4 d_k, \quad \forall k \in \mathbb{I}_{N-1}^+ \\ & E_2 \delta_k + E_3 z_k \leq E_1 u_k + E_4 x_k + E_6 d_k + E_5, \quad \forall k \in \mathbb{I}_{N-1}^+, \quad (4-21) \\ & x_0 = x_{\text{init}}, \\ & (x_k, u_k) \in \mathbb{Z}, \quad d_k \in \mathbb{D}, \quad z_k \in \mathbb{R}^p, \quad \delta_k \in \mathbb{B}^q, \quad \forall k \in \mathbb{I}_N^+, \\ & x_N \in \mathbb{X}_f \end{cases}$$

where $\boldsymbol{\delta} = \{\delta_k\}_{k=0}^{N-1}$ and $\mathbf{z} = \{z_k\}_{k=0}^{N-1}$ are the sequences of binary and auxiliary variables over the horizon. The second constraint encodes all MLD constraints, including operational constraints and Big-M linearizations from Section 4-3-3. When the stage cost $\ell(\cdot)$ is taken to be linear in its arguments, Eq. (4-21) reduces to an MILP, which can be solved to global optimality using standard branch-and-bound solvers [69]. In the microgrid context, a linear stage cost directly reflecting electricity purchase costs and auxiliary generation costs is a natural choice, as these are linear functions of the power flows. The MLD-MPC framework has been widely applied to microgrid EMS problems for exactly this reason, as it provides a tractable formulation that can handle both the continuous dispatch decisions and the discrete operating mode selections of the system within a unified optimization framework [16], [17], [31], [73].

4-3-6 Stability of MLD-MPC

As discussed in Section 4-3-1, the class of systems considered in this thesis is not that of hybrid systems, but instead systems with discrete actuators. For this class of systems, all of the results for standard MPC also hold [66, Theorem 1]. Because of this, all the stability theory discussed in Sections 4-1 and 4-2 remains valid for systems with discrete actuators, and no new theory will need to be presented.

Time-Varying EMPC Controller Design

This chapter introduces the EMPC framework used for the operation of the CHP microgrid. The physical model described in Chapter 3 is first formulated in the MLD representation, allowing system dynamics, constraints, and logical decision variables to be incorporated into a single unified modeling framework, with particular attention given to the non-linear and discontinuous behavior present in the system, including the asymmetric battery charging efficiency, asymmetric import and export prices, and temperature dependence of the heat pump COP. Building on this MLD representation, an EMPC formulation is developed to minimize the operational costs of the microgrid while satisfying the long-term ATEs operational constraints; since the time-scale of seasonal storage is significantly longer than the permissible horizon length for on-line optimization, terminal ingredients are introduced, constructed from a surrogate optimization problem that uses a coarser time-step and simplified model to approximate the optimal solution and enforce satisfaction of the thermal balance by the end of the year. Finally, the generation of forecasting signals used in the EMPC controller is discussed.

5-1 MLD Model

The MLD framework is used to model the system for control. This section describes how the models presented in Chapter 3 are converted into MLD format, including the treatment of discrete operating modes, asymmetric efficiencies, and nonlinear heat pump dynamics.

5-1-1 Battery

The battery charging and discharging process is subject to asymmetric efficiency losses depending on the direction of power flow, as described by Eq. (2-3). To handle this asymmetry

within the MILP formulation, a binary variable $\delta_b(k)$ is introduced to indicate the charging state of the battery

$$\delta_b(k) = \begin{cases} 1, & \text{if } P_b(k) \geq 0, \quad (\text{charging}), \\ 0, & \text{if } P_b(k) < 0, \quad (\text{discharging}). \end{cases} \quad (5-1)$$

An auxiliary variable $z_b(k) = \delta_b(k)P_b(k)$ is then defined to represent the product of the binary and continuous power variables, which is linearized using the Big-M method of Section 4-3-3. An additional auxiliary variable $c_g(k)$ is also defined as [16]

$$c_g(k) = c_{\text{el}}^{\text{sell}}(k)P_g(k) + (c_{\text{el}}^{\text{buy}}(k) - c_{\text{el}}^{\text{sell}}(k))z_g(k), \quad (5-2)$$

to represent the cost incurred from importing and exporting electricity from the grid. In MLD format, the resulting constraints are stated as [16]

$$\begin{bmatrix} P_{b,\max} \\ -(P_{b,\max} + \varepsilon) \\ P_{b,\max} \\ P_{b,\max} \\ -P_{b,\max} \\ -P_{b,\max} \end{bmatrix} \delta_b(k) + \begin{bmatrix} 0 \\ 0 \\ 1 \\ -1 \\ 1 \\ -1 \end{bmatrix} z_b(k) \leq \begin{bmatrix} 1 \\ -1 \\ 1 \\ -1 \\ 0 \\ 0 \end{bmatrix} P_b(k) + \begin{bmatrix} P_{b,\max} \\ -\varepsilon \\ P_{b,\max} \\ P_{b,\max} \\ 0 \\ 0 \end{bmatrix}. \quad (5-3)$$

The battery dynamics are discretized using forward Euler integration. Under a zero-order hold assumption on the input, this approximation is exact, yielding the discrete-time model

$$x_b(k+1) = x_b(k) + \frac{\Delta t}{C_{\max}} \left(\left(\eta_c - \frac{1}{\eta_d} \right) z_b(k) + \frac{1}{\eta_d} P_b(k) \right). \quad (5-4)$$

where the auxiliary variable $z_b(k)$ correctly selects the charging efficiency η_c when the battery is charging and the discharging efficiency η_d when it is discharging, without introducing any non-linearity into the discrete-time model.

5-1-2 Grid Interconnection

The cost of exchanging power with the grid is asymmetric, as electricity is bought and sold at different prices $c_{\text{el}}^{\text{buy}}(k)$ and $c_{\text{el}}^{\text{sell}}(k)$ respectively. A binary variable $\delta_g(k)$ is introduced to indicate the direction of power flow

$$\delta_g(k) = \begin{cases} 1, & \text{if } P_g(k) \geq 0 \quad (\text{import}), \\ 0, & \text{if } P_g(k) < 0 \quad (\text{export}). \end{cases} \quad (5-5)$$

An auxiliary variable $z_g(k) = \delta_g(k)P_g(k)$ is defined to represent the import power, allowing the grid cost to be expressed as a linear function of $P_g(k)$ and $z_g(k)$. The resulting MLD constraints are time-varying as they depend on the electricity price at time k [16]

$$\begin{bmatrix} P_{g,\max} \\ -(P_{g,\max} + \varepsilon) \\ M_g \\ M_g \\ -M_g \\ -M_g \end{bmatrix} \delta_g(k) + \begin{bmatrix} 0 \\ 0 \\ 1 \\ -1 \\ 1 \\ -1 \end{bmatrix} c_g(k) \leq \begin{bmatrix} 1 \\ -1 \\ c_{\text{el}}^{\text{buy}}(k) \\ -c_{\text{el}}^{\text{buy}}(k) \\ c_{\text{el}}^{\text{sell}}(k) \\ -c_{\text{el}}^{\text{sell}}(k) \end{bmatrix} P_g(k) + \begin{bmatrix} P_{g,\max} \\ -\varepsilon \\ M_g \\ M_g \\ 0 \\ 0 \end{bmatrix}, \quad (5-6)$$

where $M_g = \max_k \{c_{\text{el}}^{\text{buy}}(k)\} \cdot P_{g,\max}$ is the Big-M bound for the grid cost auxiliary variable.

5-1-3 Thermal System

The reservoir, floor, and building dynamics do not require additional binary variables. However, the thermal states span a broad range of timescales. The reservoir responds quickly to heat pump inputs, while the building and floor have significantly larger thermal inertia. The resulting set of equations has large separation of timescales, making explicit discretization methods unstable. An implicit discretization is therefore used, in which the heat flow terms Eqs. (3-2) to (3-4) are evaluated at $t + 1$, yielding

$$h_r(k+1) = h_r(k) + \frac{\Delta t}{m_r} \left(\dot{Q}_{\text{out}}^H(k) - \dot{Q}_{\text{in}}^C(k) + \dot{Q}_b(k) - \dot{Q}_c(k) - \dot{Q}_{rf}(k+1) \right), \quad (5-7)$$

$$h_f(k+1) = h_f(k) + \frac{\Delta t}{m_f} \left(\dot{Q}_{rf}(k+1) - \dot{Q}_{fb}(k+1) \right), \quad (5-8)$$

$$h_b(k+1) = h_b(k) + \frac{\Delta t}{m_{b,\text{eff}}} \left(\dot{Q}_{fb}(k+1) - \dot{Q}_{bo}(k) + \dot{Q}_{\text{rad}}(k) \right). \quad (5-9)$$

Since the heat flow terms are linear in the states, the implicit scheme remains unconditionally stable regardless of the chosen time step Δt .

Enforcing this comfort band as hard constraints is problematic for reasons discussed in Section 4-1, so instead, slack variables $s_l \geq 0$ and $s_u \geq 0$ are introduced to relax the constraint. The constraints

$$\frac{1}{c_{p,a}} h_b(k) + s_l(k) \geq T_{b,\text{min}}, \quad (5-10)$$

$$\frac{1}{c_{p,a}} h_b(k) - s_u(k) \geq T_{b,\text{max}}, \quad (5-11)$$

are added so that s_l and s_u reflect the upper and lower violation of the comfort band respectively. s_l and s_u can then be incorporated into the cost function as linear terms to ensure that the slack is minimized and hence the constraint satisfied.

5-1-4 Heat Pump

The heat pump is subject to the first law of thermodynamics, which requires that energy is conserved across the heat pump cycle. Specifically, the heat delivered on the hot side must equal the sum of the electrical work input and the heat extracted from the cold side,

$$\dot{Q}_{\text{out}}^H(k) = \dot{W}^H(k) + \dot{Q}_{\text{in}}^H(k), \quad (5-12)$$

$$\dot{Q}_{\text{in}}^C(k) = \dot{Q}_{\text{out}}^C(k) + \dot{W}^C(k), \quad (5-13)$$

where $\dot{Q}_{\text{in}}^H$ and $\dot{Q}_{\text{in}}^C$ denote the heat extracted from the cold side in heating and cooling mode respectively. These constraints are linear and are incorporated directly into the MILP formulation.

The heat pump output is nonlinear in the system states, as the COP depends non-linearly on the temperature lift across the heat pump as well as the electrical input power $\dot{W}$. To incorporate this into the MILP formulation, a two-dimensional PWA approximation is applied as described in Section 4-3-4. As stated in Eq. (2-10), the heat pump output is in principle

a function of three variables: the reservoir temperature, the ATES well temperature, and $\dot{W}$. However, since the reservoir and ATES well temperatures operate over a narrow range relative to the temperature lift, the COP is approximated as a function of the temperature lift ΔT and $\dot{W}$ alone, with the hot-side temperature treated as fixed. The resulting heating and cooling outputs are then approximated as

$$\dot{Q}_{\text{out}}^H \approx f_{\text{PWA}}^H(\dot{W}^H, T_h - T_r), \quad (5-14)$$

$$\dot{Q}_{\text{out}}^C \approx f_{\text{PWA}}^C(\dot{W}^C, T_r - T_c), \quad (5-15)$$

where f_{PWA}^H and f_{PWA}^C are PWA approximations of $\text{COP}^H(\Delta T^H) \cdot \dot{W}^H$ and $\text{COP}^C(\Delta T^C) \cdot \dot{W}^C$ respectively, constructed over a grid in the $(\dot{W}, \Delta T)$ domain using the method of Section 4-3-4. Each approximation introduces the auxiliary variables λ_{ij} , δ_{ij} , and δ_j^i as defined in Eqs. (4-16) to (4-18), subject to the locality constraint Eq. (4-19) and the triangle selection constraint Eq. (4-20). The heat pump can operate in either heating or cooling mode, represented by binary variables $\delta^H(k) \in \mathbb{B}$ and $\delta^C(k) \in \mathbb{B}$ respectively. These binaries serve a dual purpose: they enforce mutual exclusivity of the two operating modes, and they activate or deactivate the PWA approximations for each mode. Additionally, the heat pump can only operate when the temperature lift is in a feasible region. Incorporating these elements in the PWA constraints from Section 4-3-4 results in

$$\Delta T^H(k) - \sum_{i=1}^n \sum_{j=1}^m \lambda_{ij}^H(k) \Delta T_i^H \leq M_{\text{max}}^H (1 - \delta^H(k)), \quad (5-16)$$

$$\Delta T^H(k) - \sum_{i=1}^n \sum_{j=1}^m \lambda_{ij}^H(k) \Delta T_i^H \geq -M_{\text{max}}^H (1 - \delta^H(k)), \quad (5-17)$$

$$\dot{W}^H(k) - \sum_{i=1}^n \sum_{j=1}^m \lambda_{ij}^H(k) \dot{W}_j^H \leq \dot{W}_{\text{max}} (1 - \delta^H(k)), \quad (5-18)$$

$$\dot{W}^H(k) - \sum_{i=1}^n \sum_{j=1}^m \lambda_{ij}^H(k) \dot{W}_j^H \geq -\dot{W}_{\text{max}} (1 - \delta^H(k)), \quad (5-19)$$

$$\dot{Q}_{\text{out}}^H(k) = \sum_{i=1}^n \sum_{j=1}^m \lambda_{ij}^H(k) f_{ij}^H, \quad (5-20)$$

$$\sum_{i=1}^n \sum_{j=1}^m \lambda_{ij}^H(k) = \delta^H(k), \quad (5-21)$$

$$\sum_{i=1}^n \sum_{j=1}^m \delta_j^{i,H}(k) + \delta_{ij}^H(k) = \delta^H(k), \quad (5-22)$$

where M_{max}^H is the maximum heat pump lift in heating mode and $\dot{W}_{\text{max}}$ the maximum heat pump work. These terms arise from a big-M linearization as the corresponding equality constraint should only be active when $\delta^H = 1$. The locality and triangle activation constraints in Eqs. (4-19) and (4-20) remain unchanged. An analogous set of constraints holds for the cooling mode, with $\delta^H(k)$, $\lambda_{ij}^H(k)$, $\Delta T^H(k)$, $\dot{W}^H(k)$, $\dot{Q}_{\text{out}}^H(k)$, $\dot{W}_j^H$, ΔT_i^H and f_{ij}^H replaced by their counterparts $\delta^C(k)$, $\lambda_{ij}^C(k)$, $\Delta T^C(k)$, $\dot{W}^C(k)$, $\dot{Q}_{\text{out}}^C(k)$, $\dot{W}_j^C$, ΔT_i^C and f_{ij}^C respectively, with constraints Eqs. (5-16) to (5-22) otherwise identically structured. The non-linear relation between heat-pump input work and output heat at various temperature lifts is shown in Fig. 5-1. The PWA approximation is generated by taking a coarse grid with $n = 3$ and $m = 2$. The




Fig. 5-1: Two-dimensional PWA approximation of the heat pump model and the resulting absolute approximation error.

coarse grid keeps the amount of additional variables bounded, at the expense of increased approximation error.

Some additional operational constraints on the heat pump and auxiliary thermal units are introduced. As the heat pump can only operate in one mode at a time, a mutual exclusivity constraint

$$\delta^H(k) + \delta^C(k) \leq 1, \quad (5-23)$$

is introduced to ensure consistency. The corresponding bounds on the heat pump electrical input are then

$$\dot{W}^H(k) \leq \dot{W}_{\max} \cdot \delta^H(k), \quad (5-24)$$

$$\dot{W}^C(k) \leq \dot{W}_{\max} \cdot \delta^C(k). \quad (5-25)$$

Finally, the auxiliary boiler and chiller should be constrained to not operate simultaneously, as this is considered thermodynamically wasteful. However, this would require introducing $2N$ additional binary variables as the auxiliary thermal unit power output can not be directly tied to the operating mode. This is because rebalancing actions for the ATEs system may lead to cooling output even when there is heating demand in the building, hence requiring auxiliary boiler heating to prevent the building temperature from leaving the comfort band. Instead, a constraint is introduced to ensure that the amount of thermodynamically wasteful operation is minimized

$$\frac{\dot{Q}_b(k)}{\dot{Q}_{b,\max}} + \frac{\dot{Q}_c(k)}{\dot{Q}_{c,\max}} \leq 1, \quad (5-26)$$

which ensures that the most simultaneous heating and cooling is limited to half of the max power output, with simultaneous operation becoming less likely as the boiler or chiller output increases beyond 50% of the maximum thermal output.

5-1-5 ATEs

The dynamics of the ATEs system described in Section 2-2-3 is bound by the same operational heating and cooling modes as the heat pump. To rewrite the dynamics of Eqs. (2-13) to (2-16) in a single unified model, the auxiliary variables $z_h^H(k) = \delta^H(k)h_h(k)$, $z_c^H(k) = \delta^H(k)h_c(k)$,

$z_h^C(k) = \delta^C(k)h_h(k)$, and $z_c^C(k) = \delta^C(k)h_c(k)$ are introduced, which can be linearized through the usage of big-M constraints. For brevity, the full list of the four sets big-M constraints are not explicitly stated here, but the reader is referred to Section 4-3-3 or Bemporad and Morari [67] for the procedure on how to construct these constraints. Additional auxiliary variables for the heat exchanger enthalpy output in heating and cooling mode are also introduced as a linear combination of existing variables

$$z_{\text{out}}^H(k) = z_h^H(k) - \frac{1}{\dot{m}}\dot{Q}_{\text{in}}^H(k), \quad (5-27)$$

$$z_{\text{out}}^C(k) = z_h^C(k) - \frac{1}{\dot{m}}\dot{Q}_{\text{out}}^C(k). \quad (5-28)$$

The resulting unified dynamical model is then written as

$$h_h(k+1) = h_h(k) + \frac{\Delta t}{m_h} \left(\dot{m} \left(-z_h^H(k) + c_w T_{gw} \delta^H(k) + z_{\text{out}}^C(k) - z_h^C(k) \right) - \dot{Q}_{h,1}(k) \right), \quad (5-29)$$

$$h_c(k+1) = h_c(k) + \frac{\Delta t}{m_c} \left(\dot{m} \left(z_{\text{out}}^H(k) - z_c^H(k) - z_c^C(k) + c_{p,w} T_{gw} \delta^C(k) \right) + \dot{Q}_{c,1}(k) \right), \quad (5-30)$$

which are 2 equations unifying the heating and cooling mode dynamic models from Section 2-2-3 into the MLD format.

Additional constraints are required on the heat exchanger output enthalpy in both operating modes. In heating mode, the heat exchanger output mixes with the cold well; if its specific enthalpy exceeds that of the cold well, the cold well is heated rather than charged, effectively discharging both wells simultaneously. Since no cost or value is directly associated with the stored energy in the wells, the controller would otherwise exploit these operating points freely. Constraints are therefore added to prevent this:

$$z_{\text{out}}^H(k) \geq h_{c,\min} \delta^H(k), \quad (5-31)$$

$$z_{\text{out}}^H(k) \leq z_c^H(k), \quad (5-32)$$

$$z_{\text{out}}^C(k) \geq h_{h,\max} \delta^C(k), \quad (5-33)$$

$$z_{\text{out}}^C(k) \leq z_h^C(k). \quad (5-34)$$

Note that no additional big-M constraints on z_{out}^H and z_{out}^C are needed, as the activation of these constraints is implicitly handled through the dependence on δ^H , δ^C , z_c^H , and z_h^C .

A key challenge in the ATES system is ensuring that the thermal balance is maintained over a full seasonal cycle. That is, the net heat injected into the system over a year should be zero as described in Eq. (2-17). Since this constraint spans a timescale far longer than the MPC horizon, it cannot be enforced directly as a hard constraint. Instead, an auxiliary state $b(k)$ is introduced to track the cumulative net heat injection into the ATES over the horizon

$$b(k+1) = b(k) + \Delta t \left(\dot{Q}_{\text{in}}^H(k) - \dot{Q}_{\text{out}}^C(k) \right), \quad (5-35)$$

where $\dot{Q}_{\text{in}}^H$ is the heat pump input heat, extracted from the ATES, and $\dot{Q}_{\text{out}}^C$ is the heat pump output heat, injected into the ATES. Note that the subscripts "in" and "out" refer to the heat pump perspective: $\dot{Q}_{\text{in}}^H$ denotes heat entering the heat pump (i.e., extracted from the ATES), while $\dot{Q}_{\text{out}}^C$ denotes heat leaving the heat pump (i.e., injected into the ATES).

5-2 Cost Function Design

The cost function in the EMPC problem is primarily designed to minimize operational costs. The operational costs of the microgrid can be divided into two main components: (i) the cost of importing and exporting electricity to the utility grid, and (ii) the degradation cost of the battery. The grid costs are captured in the variable $c_g(k)$ defined in Eq. (5-2) and can be directly minimized to minimize operational costs. The battery degradation cost from Eq. (2-4) is also incorporated, resulting in the economic stage cost

$$\ell_{\text{econ}}(c_g, P_b) = c_g + \text{OM}_b |P_b|. \quad (5-36)$$

The stage-cost also accounts for the constraint violation cost from the relaxed comfort band constraints. This cost term is given as

$$\ell_{\text{rel}}(s_l, s_u) = \beta_c (s_l + s_u), \quad (5-37)$$

where $\beta_c > 0$ is a constant picked arbitrarily large such that Eq. (5-37) behaves similarly to enforcing the comfort band as a hard constraint. The total stage cost is therefore the sum of the economic and relaxation terms

$$\ell(c_g, P_b, s_l, s_u) = \ell_{\text{econ}}(c_g, P_b) + \ell_{\text{rel}}(s_l, s_u), \quad (5-38)$$

which is minimized at each time step subject to the constraints of the EMPC problem stated in Eq. (5-44).

5-2-1 Full EMPC Optimization Problem

Before stating the full optimization problem, it is convenient to collect all the constraints from the preceding sections into compact notation. The binary and auxiliary variables can be collected in the vectors

$$\delta(k) = \begin{bmatrix} \delta_b(k) \\ \delta_g(k) \\ \delta^H(k) \\ \delta^C(k) \\ \delta_{\text{PWA}}^H \\ \delta_{\text{PWA}}^C \end{bmatrix}, \quad z(k) = \begin{bmatrix} z_b(k) \\ z_g(k) \\ c_g(k) \\ z_{\text{out}}^H(k) \\ z_h^H(k) \\ z_c^H(k) \\ z_{\text{out}}^C(k) \\ z_h^C(k) \\ z_c^C(k) \end{bmatrix}, \quad (5-39)$$

where δ_{PWA}^H and δ_{PWA}^C encompass all the binary variables introduced in heat pump PWA approximations. The system dynamics in Eqs. (5-4), (5-7) to (5-9), (5-29), (5-30) and (5-35) can then be collected into the generic map

$$x(k+1) = f(x(k), x(k+1), u(k), d(k), \delta(k), z(k)). \quad (5-40)$$

Recall that this mapping depends on both $x(k)$ and $x(k+1)$, as the thermal dynamics in Eqs. (5-7) to (5-9) are integrated via implicit Euler, whereas the rest of the system is integrated via forward Euler.

The auxiliary constraints in Eqs. (5-2), (5-3) and (5-6), the heat exchanger constraints in Eqs. (5-27), (5-28) and (5-31) to (5-34), the mutual exclusivity constraints in Eq. (5-26), and the unlabeled big-M constraints on z_h^H , z_c^H , z_h^C , and z_c^C can all be collected into the MLD format as

$$E_2\delta(k) + E_3z(k) \leq E_1u(k) + E_4x(k) + E_6d(k) + E_5 \quad (5-41)$$

in which $E_1, \dots, E_6$ are real matrices of appropriate dimension encoding all of the aforementioned constraints.

The heat pump PWA constraints stated in Eqs. (5-16) to (5-22) along with the corresponding versions of Eqs. (4-19) and (4-20), as well as the other heat pump constraint in Eqs. (5-12), (5-13), (5-24) and (5-25) can be collected into

$$\mathcal{C}^H(\lambda^H(k), \dot{Q}_{\text{out}}^H(k), \Delta T^H(k), \dot{W}^H(k), \delta^H(k)) \leq 0, \quad (5-42)$$

for heating mode, and the corresponding

$$\mathcal{C}^C(\lambda^C(k), \dot{Q}_{\text{out}}^C(k), \Delta T^C(k), \dot{W}^C(k), \delta^C(k)) \leq 0, \quad (5-43)$$

for cooling mode, with $\lambda^{(\cdot)} = (\lambda_{1,1}^{(\cdot)}, \lambda_{1,2}^{(\cdot)}, \dots, \lambda_{i,j}^{(\cdot)}, \dots, \lambda_{n,m}^{(\cdot)})^\top$.

Using the above definitions the EMPC optimization problem in full is then stated as

$$\mathbb{P}^E(t, x_{\text{init}}, \mathbf{d}_t) : \begin{cases} \min_{\mathbf{u}, \delta, \mathbf{z}} & \sum_{k=0}^{N-1} \ell(c_{g,k}, P_{b,k}, s_{l,k}, s_{u,k}) + V_f(t + N, x_N) \\ \text{s.t.} & x_{k+1} = f(x_k, x_{k+1}, u_k, d_k, \delta_k, z_k) \quad \forall k \in \mathbb{I}_{N-1}^+ \\ & E_2\delta_k + E_3z_k \leq E_1u_k + E_4x_k + E_6w_k + E_5, \quad \forall k \in \mathbb{I}_{N-1}^+, \\ & \mathcal{C}^H(\lambda_k^H, \dot{Q}_{\text{out},k}^H, \Delta T_k^H, \dot{W}_k^H, \delta_k^H) \leq 0, \quad \forall k \in \mathbb{I}_N^+, \\ & \mathcal{C}^C(\lambda_k^C, \dot{Q}_{\text{out},k}^C, \Delta T_k^C, \dot{W}_k^C, \delta_k^C) \leq 0, \quad \forall k \in \mathbb{I}_N^+, \\ & x_0 = x_{\text{init}}, \\ & x_N \in \mathbb{X}_f(N), \end{cases} \quad (5-44)$$

where the constraint with matrices $E_1, \dots, E_6$ incorporates all the constraints which are not explicitly referenced by label, such as the big-M constraints on z_h^H and z_c^C .

5-3 Terminal Constraint Set and Terminal Cost Construction

The EMPC formulation in Eq. (5-44) optimizes over a finite receding horizon and therefore cannot, on its own, account for the seasonal objective of the ATES subsystem, whose charging and discharging dynamics evolve over timescales far exceeding any practically implementable prediction horizon. The purpose of this section is to construct terminal ingredients that encode this long-term thermal balance requirement into the short-horizon optimization problem, ensuring that the annual ATES balance constraint is satisfied while retaining the tractability of receding-horizon control.




Fig. 5-2: ATES temperatures (a) and energy balance (b) when controlling the system using EMPC without terminal ingredients.

5-3-1 The Need for Terminal Ingredients

The EMPC formulation in Eq. (5-44) minimizes a finite-horizon economic objective over a receding window of length N . Without terminal ingredients, the controller accounts for no information beyond this horizon, resulting in myopic behavior with respect to economic consequences past the optimization horizon. This is particularly problematic for CHP microgrid EMS designs incorporating seasonal energy storage, where effective load shifting requires planning over timescales far exceeding the availability of meaningful electricity price and weather forecasts [74]. For ATES systems this limitation is especially pronounced: the ratio of stored thermal energy capacity to maximum charge/discharge power yields time constants of several weeks to months [39], well beyond any practically implementable prediction horizon for on-line optimization.

The myopic behavior is demonstrated in Fig. 5-2, which shows the ATES temperatures and cumulative energy balance when the EMPC controller is applied without terminal ingredients over an entire operational year. As Fig. 5-2b clearly shows, a substantial thermal imbalance remains by the end of the year, violating the ATES operational constraint in Eq. (2-17). This can be attributed to the controller consistently preferring heat pump operation over the auxiliary boiler on economic grounds. The warm well is depleted until empty, with the seasonal asymmetry in heating and cooling not providing sufficient heat injection into the ATES to cause balanced yearly operation. Thus, without any mechanism to account for satisfaction of the ATES operational constraint, balanced operation is not achieved. These results motivate the construction of appropriate terminal ingredients to ensure that the ATES balance constraint is satisfied.

Notably, no terminal constraint set is imposed on the battery state-of-charge. The time constant of battery storage is typically on the order of hours rather than weeks or months, and hence the relevant charging and discharging dynamics evolve well within the typical EMS planning horizon. Consequently, the controller can adequately capture the long-term impact of battery operation directly within the finite-horizon optimization problem, without requiring additional terminal ingredients. The terminal constraint set and terminal cost are therefore constructed specifically for the ATES subsystem only.




Fig. 5-3: The control architecture used for the EMPC controller.

5-3-2 Surrogate Optimization Problem

In order to construct the terminal ingredients for the time-varying EMPC controller, a surrogate optimization is introduced. This surrogate problem is based on a simplified version of Eq. (5-44), in which the time step is taken to be a full day rather than 15 minutes. The solution of the surrogate problem will then be used to inform the design of the terminal constraint set. The control architecture that will be used is also shown in Fig. 5-3. Notably, because of the lack of accurate year-ahead forecasts, this planning is done in an open-loop fashion, based on Typical Meteorological Year (TMY) data.

To accurately represent building heating and cooling demand over the course of an entire day, the concepts of Heating Degree Days (HDD) and Cooling Degree Days (CDD) are applied [75]. HDD and CDD convert hourly ambient temperature profiles into daily thermal energy demand by aggregating the difference between the ambient temperature and a reference base temperature, thereby quantifying the cumulative heating or cooling load required to maintain a comfortable indoor temperature. For a given day, HDD and CDD are computed as

$$\text{HDD} = \frac{1}{24} \sum_{h=1}^{24} \max(T_{\text{base}} - T_d(h), 0), \quad (5-45)$$

$$\text{CDD} = \frac{1}{24} \sum_{h=1}^{24} \max(T_{\text{sol-air}}(h) - T_{\text{base}}, 0), \quad (5-46)$$

where $T_d(h)$ is the ambient temperature at hour h , T_{base} is a base temperature determined by the building characteristics, and $T_{\text{sol-air}}(h)$ is the sol-air temperature, which augments the ambient temperature with a solar irradiance correction to account for solar heat gains through the building envelope:

$$T_{\text{sol-air}}(h) = T_d(h) + \frac{\alpha}{h_o} s(h), \quad (5-47)$$

where $s(h)$ is the solar irradiance at hour h in W/m^2 , α is the solar absorptivity of the building's exterior surface, and h_o is the external convective heat transfer coefficient. The

ratio α/h_o characterizes the thermal interaction between the building envelope and incident solar radiation [76]. The solar irradiance correction is only applied to CDD, as solar heat gains act to reduce heating demand rather than increase it, and are therefore neglected in the HDD formulation to obtain a more conservative estimate of the heating demand.

All of the detailed thermal dynamics described in Eq. (3-5), and by extension the variable heat pump COP model, will be ignored in the surrogate problem. Instead, these will be replaced with a constraint on the daily thermal energy balance consisting of the heating, cooling, and thermal demand. This is added as a soft constraint on the imbalance state $\Delta Q(t)$, defined as

$$\Delta Q(t) = Q_{\text{dem}}(t) + (\dot{Q}_b(t) - \dot{Q}_c(t) + \dot{Q}_{\text{out}}^H(t) - \dot{Q}_{\text{in}}^C(t))\Delta t. \quad (5-48)$$

The thermal demand $Q_{\text{dem}}(t)$ is positive when there is a heating demand and negative when there is a cooling demand. The absolute value of $\Delta Q(t)$ is penalized, where the absolute value can be implemented via linearization of absolute values as described in Section 4-3-2. The cost term then becomes

$$\ell_{\text{imb}}^s(\Delta Q) = \beta_Q |\Delta Q|. \quad (5-49)$$

Where β_Q is a constant penalty term chosen arbitrarily large so that the relaxed constraint is behaviorally identical to a hard constraint. As the ambient temperature profile cannot be accurately forecasted one year ahead, TMY data is used to generate a thermal demand profile. EnergyPlus Weather (EPW) files provide hourly meteorological data, including dry-bulb temperature, and solar radiation, in a standardised format originally developed for building energy simulation [77], making them well-suited for generating representative annual thermal demand profiles. The EPW file used corresponds to De Bilt, Netherlands (WMO station 062600), with records spanning 1949–2023 [78].

Because of the coarse time step, the electrical subsystem will be ignored. The time constant of an electrical battery, defined in Eq. (2-5), is typically on the order of one to a few hours. Since this time constant is much shorter than the time step of 24 hours, the battery dynamics are ignored. It is assumed that all electricity will be bought from the grid against the historical daily average electricity price. The resulting stage cost only consists of the electrical costs of buying electricity to run the heating and cooling units to satisfy the thermal demand:

$$\ell_{\text{cost}}^s(t, u) = \tilde{c}_{\text{el}}(t) \left(\frac{\Delta t}{\eta_{\text{ch}}} \dot{Q}_c + \frac{\Delta t}{\eta_b} \dot{Q}_b + \frac{\Delta t}{\text{COP}^H} \dot{Q}_{\text{out}}^H + \frac{\Delta t}{\text{COP}^C} \dot{Q}_{\text{out}}^C \right), \quad (5-50)$$

where $\tilde{c}_{\text{el}}(t)$ is the historical daily average electricity price in the Netherlands computed based on data from the ENTSO-E API [1]. The disturbance signals for the electrical price and thermal demand profile used when solving the surrogate problem are shown in Fig. 5-4.

Since the surrogate problem optimizes over a full operational year, the balance constraint from Eq. (2-17) can be added explicitly in form

$$\sum_{t=0}^{365} (\dot{Q}_{\text{in}}^H(t) - \dot{Q}_{\text{out}}^C(t))\Delta t = 0. \quad (5-51)$$

In practice, enforcing this as a strict equality is both numerically undesirable and physically unrealistic. A 24 hour time step introduces discretization errors that make exact annual balance unattainable, so a tolerance Δb_{max} is introduced such that

$$\left| \sum_{t=0}^{365} (\dot{Q}_{\text{in}}^H(t) - \dot{Q}_{\text{out}}^C(t)) \Delta t \right| \leq \Delta b_{\text{max}}, \quad (5-52)$$




Fig. 5-4: The disturbance signals used when solving the surrogate optimization problem.

where the absolute value can again be linearized to obtain a constraint compatible with the MILP formulation.

The reduced set of states, inputs and disturbances for the surrogate problem are then

$$x^s(k) = \begin{bmatrix} h_h(k) \\ h_c(k) \\ b(k) \end{bmatrix}, \quad u^s(k) = \begin{bmatrix} \dot{W}^H(k) \\ \dot{W}^C(k) \\ \dot{Q}_b(k) \\ \dot{Q}_c(k) \end{bmatrix}, \quad d^s(k) = \begin{bmatrix} \tilde{c}_{el}(k) \\ Q_{dem}(k) \end{bmatrix}, \quad (5-53)$$

and the reduced set of binary and auxiliary variables are

$$\delta^s(k) = \begin{bmatrix} \delta^H(k) \\ \delta^C(k) \end{bmatrix}, \quad z^s(k) = \begin{bmatrix} z_{out}^H(k) \\ z_h^H(k) \\ z_c^H(k) \\ z_{out}^C(k) \\ z_h^C(k) \\ z_c^C(k) \end{bmatrix}. \quad (5-54)$$

These can then be used to express the constraints of the surrogate problem compactly. The heat exchanger constraints in Eqs. (5-24), (5-27), (5-31) and (5-32) as well as their cooling mode counterparts, the mutual exclusivity constraints in Eqs. (5-23) and (5-26), and the heat pump first law constraints in Eqs. (5-12) and (5-13) are collected in the standard MLD constraint representation

$$E_2^s \delta^s(k) + E_3^s z^s(k) \leq E_1^s u^s(k) + E_4^s x^s(k) + E_6^s d^s(k) + E_5^s. \quad (5-55)$$

The system dynamics from Eqs. (5-29), (5-30) and (5-35) can be represented compactly as

$$x^s(k+1) = f^s(x^s(k), u^s(k), d^s(k), \delta^s(k), z^s(k)). \quad (5-56)$$

Combining the above elements, the surrogate optimization problem is formulated as an MILP over a horizon of 365 days. The problem takes as inputs the initial system state x_{init}^s and

sequence of disturbance vectors $\mathbf{d}^s$, and optimizes the daily dispatch of all thermal units over the full year:

$$\mathbb{P}^{\text{sur}}(x_{\text{init}}^s, \mathbf{d}^s) : \begin{cases} \min_{\mathbf{u}^s, \delta^s, \mathbf{z}^s} & \sum_{k=0}^{365} \ell_{\text{imb}}^s(\Delta Q_k) + \ell_{\text{cost}}^s(k, u_k) \\ \text{s.t.} & x_{k+1}^s = f^s(x_k^s, u_k^s, d_k^s, \delta_k^s, z_k^s), \quad \forall k \in \mathbb{I}_{364}^+ \\ & E_2^s \delta_k^s + E_3^s z_k^s \leq E_1^s u_k^s + E_4^s x_k^s + E_6^s d_k^s + E_5^s, \quad \forall k \in \mathbb{I}_{365}^+, \\ & \left| \sum_{k=0}^{365} (\dot{Q}_{\text{in},k}^H - \dot{Q}_{\text{out},k}^C) \Delta t \right| \leq \Delta b_{\text{max}}, \\ & x_0^s = x_{\text{init}}^s. \end{cases} \quad (5-57)$$

The decision variables are the continuous inputs $\mathbf{u}^s$, comprising the heat pump power, boiler, and chiller outputs, the binary variables δ^s encoding the discrete operating mode selections, and the auxiliary variables $\mathbf{z}^s$ introduced by auxiliary constraints. The solution to $\mathbb{P}^{\text{sur}}$ yields an optimal annual operating trajectory under TMY conditions denoted $b^*(t) = \arg \min \mathbb{P}^{\text{sur}}(x, \mathbf{d}^s)$. This trajectory is subsequently used to construct the terminal constraint set for the EMPC controller, which is defined as a neighbourhood of the optimal imbalance trajectory. Notably, this trajectory is planned in an open-loop fashion, meaning the optimization problem in Eq. (5-57) is solved once *a priori*, rather than being dynamically replanned on-line.

5-3-3 Terminal Constraint Set

Solving Eq. (5-57) using historical average daily electricity price and a thermal demand profile based on TMY data results in optimal trajectory $b^*(t)$, shown in Fig. 5-5. This trajectory is interpolated and used to construct the terminal trajectory for the EMPC problem. Theoretically, this should be added as a terminal equality constraint, as the aim is to satisfy the thermal balance by the end of the year exactly. However, since the surrogate solves a simplified version of the problem, $b^*(t)$ is not guaranteed to be reachable by the EMPC controller at time t . To ensure feasibility, a relaxation is introduced as a tube around the target state:

$$\mathbb{X}_f(t) = \{b \in \mathbb{R} \mid b^*(t) - \varepsilon \leq b \leq b^*(t) + \varepsilon\}, \quad (5-58)$$

where $\varepsilon \geq 0$ is a tuning parameter controlling the tube width. As discussed in Section 4-2, it is desirable to establish recursive feasibility of the EMPC problem. However, since the state $b^*(t)$ is not guaranteed to be reachable, the theory presented in Section 4-2 and McAllister [58] cannot be applied directly, and recursive feasibility cannot be established via the usual method of showing positive invariance of $\mathbb{X}_f$. As a practical alternative, ε is tuned by trial-and-error methods to find values that ensure Eq. (5-44) does not go infeasible during the simulation.

5-3-4 Terminal Cost

Ideally the terminal cost V_f satisfies a Lyapunov decrease property on $\mathbb{X}_f$ under the terminal control law, as this ensures that the problem remains feasible and has bounded degradation in

the nominal economic performance when subject to bounded disturbances [58]. As discussed in Section 4-2, constructing a terminal control law and terminal cost-to-go approximation satisfying Eq. (4-7) is not straightforward. As mentioned in Section 5-3-3, $b^*(t)$ is derived from a simplified surrogate problem and is not guaranteed to be reachable by the EMPC controller, and hence, the robustness argument from McAllister [58] does not apply in this case. This also applies to the dissipativity results discussed in Section 4-2

In view of these difficulties, a pragmatic alternative is adopted, following approaches used in the economic MPC literature for energy systems with storage states [45], [79]. The terminal cost is chosen as a linear penalty on deviation of the terminal ATEs imbalance from the center of the terminal constraint set:

$$V_f(t, b) = \gamma |b - b^*(t)|. \quad (5-59)$$

Since γ is chosen sufficiently large relative to the stage cost, it serves as an upper bound on the marginal value of stored energy, ensuring the terminal cost dominates any remaining cost-to-go. Furthermore, given the linear dynamics and stage cost, the Dynamic Programming (DP) principle suggests the true cost-to-go is piecewise linear in b , lending further support to the choice Eq. (5-59).

Without this terminal cost, the economic objective provides no incentive for the controller to track $b^*(t)$ closely within the terminal constraint set $\mathbb{X}_f(t)$: operating near the boundary of the set may be economically cheap over the finite horizon, yet a terminal state at the boundary of $\mathbb{X}_f(t)$ can easily lead to model mismatch or disturbances rendering $\mathbb{X}_f(t+1)$ unreachable after disturbance realisation at timestep t . The linear penalty Eq. (5-59) counteracts this by imposing a cost that increases with deviation from $b^*(t)$. When γ is chosen sufficiently large, this penalty dominates the economic stage cost near the boundary of $\mathbb{X}_f(t)$, biasing the terminal state toward the interior of the terminal constraint set and thereby increasing the robustness to disturbances and model mismatch.

While this terminal cost does not satisfy the Lyapunov decrease condition required for formal guarantees on robust recursive feasibility and economic performance degradation, it provides a practical mechanism for discouraging boundary operation without requiring the construction of a terminal control law. The weight γ is treated as a tuning parameter alongside ε , and is selected such that the closed-loop simulation remains feasible and the annual ATEs balance constraint is satisfied over the full evaluation horizon.

5-4 Forecast Signal Generation

Acquisition of disturbance signals was discussed in Section 3-4. Practically speaking, the MPC controller will not have perfect forecasting information available. In real-time applications, forecasts of disturbance signals are generated based on current and past weather measurements and climate data [80]. The design and implementation of real-time forecasting is considered outside the scope of this work. Instead, the focus of this work lies on the controller design and its robustness to forecast imperfections. Providing the MPC controller with exact, perfect disturbance forecasts would constitute an unrealistic benchmark that does not reflect real-world conditions. Therefore, a deliberate mismatch is introduced by supplying the




Fig. 5-5: Solution to the surrogate optimization problem, with (a) showing the ATES well temperatures and (b) showing the imbalance state as well as cumulative heating and cooling done by the heat pump.

controller with transformed versions of the true disturbance signals, so that the signals used by MPC for planning differ from those actually experienced by the system.

In contrast to the weather-dependent disturbance signals, electricity prices do not require forecasting in the context of the Dutch energy market. Day-ahead electricity prices for the Netherlands are determined through the EPEX SPOT auction and published by 12:40 each day for all hours of the following day [81], making the full price profile available at least 11 hours, and up to 35 hours, in advance. The price signal used by MPC can therefore be treated as known exactly over the planning horizon, and no mismatch is introduced for this disturbance.

Of the disturbance signals that do require forecasting, the electrical load forecast is arguably the least realistic, as it is identical to the true load profile, but with no additive noise applied. However, the implications of this simplification are limited. In a grid-connected microgrid, any imbalance between supply and demand that arises from a load forecast error is absorbed instantaneously by the grid interconnection, which acts as a slack resource. Consequently, load forecast accuracy has a comparatively minor effect on controller performance in this setting, and the idealization introduced here is not expected to significantly affect the validity of the results.

The forecast mismatch for ambient temperature is introduced by applying a moving average filter with window size W to the true temperature signal, which attenuates high-frequency fluctuations while preserving the overall diurnal trend. For solar irradiance, the forecast is constructed by applying a Gaussian filter ($\sigma = 5$) to smooth the signal, followed by a time shift of $k = 4$ samples, which introduces a phase offset intended to reflect the imperfect temporal alignment typical of real-world irradiance forecasts. A notable limitation of this approach is that the combination of smoothing and phase shift causes the forecast to systematically lag behind the true signal, resulting in a one-sided mismatch: irradiance is consistently underestimated, and overestimation does not occur. This is reflected in the negative Mean Bias Error (MBE) of -6.1 W/m^2 . In practice, real-world irradiance forecasts exhibit errors in both directions, and this asymmetry represents a limitation of the chosen signal transformation approach.




Fig. 5-6: Comparison of forecasted disturbance signals given to MPC and the real disturbance signals.

The resulting forecast errors are reported in Table 5-1 and are based on the same scenario data used for the optimization problem in Section 5-3. Representative samples are also shown in Fig. 5-6, showing this mismatch between the real signals and the artificially created forecast signals. For temperature, the Mean Absolute Error (MAE) is 1.15 °C and the Root Mean Squared Error (RMSE) is 1.51 °C, with an Normalized Root Mean Squared Error (nRMSE) of 14.8% relative to the mean true temperature and a near-zero MBE of -0.001 °C, indicating no systematic bias. These values are consistent with short-term outdoor temperature forecast errors reported in the literature, where standard deviations of 1.2 °C to 2.2 °C are reported for horizons ranging from 1-hour to 6-hours [82], [83]. For irradiance, computed over daytime periods only to exclude the trivially accurate nighttime zeros, the MAE is 58.7 W/m² and the RMSE is 70.9 W/m², yielding an nRMSE of 29.4% relative to mean daytime irradiance, with a slight negative MBE of -6.1 W/m². The irradiance nRMSE is consistent with values reported in the literature for short-term Global Horizontal Irradiance (GHI) forecasting, where nRMSE values in the range of 27%–37% are commonly observed across forecast horizons of 30 to 180 minutes [84], suggesting that the mismatch introduced here is representative of realistic forecast uncertainty, despite being derived from signal transformations rather than a dedicated forecasting model.

Table 5-1: Summary of forecast error metrics for ambient temperature and solar irradiance

Signal	MAE	RMSE	nRMSE	MBE
Temperature	1.15 °C	1.51 °C	14.8%	-0.001 °C
Irradiance (daytime)	58.7 W/m ²	70.9 W/m ²	29.4%	-6.1 W/m ²

Policy-Guided MPC Controller Design

This chapter presents the design of a PG-MPC controller for the microgrid presented in Chapter 3. Conventional EMPC formulations, discussed in Section 4-2, characterize performance relative to a steady state, which is ill-suited to the seasonal dynamics of the ATES system. Additionally, as discussed in Sections 5-3-3 and 5-3-4, constructing appropriate terminal ingredients satisfying the required properties is difficult for the system under study. PG-MPC addresses this by replacing the surrogate based terminal ingredients with ones derived from a guiding policy, allowing performance to be characterized relative to a guiding policy instead. The theoretical foundations of PG-MPC are first reviewed, covering recursive feasibility and closed-loop performance guarantees, as well as the design of the terminal cost and terminal constraint set. Building on this, a bang-bang dead-band guiding policy is designed to capture the seasonal charging and discharging behavior of the ATES. The policy value function is then approximated via Monte-Carlo rollouts and fitted using an artificial neural network (ANN), from which time-varying Lipschitz constants are extracted to construct the terminal cost.

6-1 Policy-Guided MPC Fundamentals

This section will outline and discuss the important theory and results in PG-MPC that will be used to justify the PG-MPC controller design. The algorithm will be described in detail, as well as the theoretical properties that can be guaranteed. Design procedures for the terminal constraint set and terminal value function will also be discussed. The theory presented in Sections 6-1-1 and 6-1-2 is based on the work by Msaad and McAllister [23].

6-1-1 Introduction to PG-MPC

Conventional EMPC stability theory, discussed in Section 4-2, characterizes stability and performance relative to a steady state (x_s, u_s) . This is problematic in EMS applications, where a steady state may be suboptimal compared to other trajectories. PG-MPC addresses this by constructing the terminal cost and terminal constraint set based on a guiding policy

$\pi(\cdot)$, so that performance is characterized relative to the guiding policy rather than a steady-state solution.

The guiding policy must be admissible, but is otherwise unrestricted and can be obtained in many ways: from domain knowledge (e.g., Rule Based Control (RBC)), data-driven methods (e.g., RL or approximate dynamic programming (ADP)), or existing controllers such as a PID-like policy. Notably, the guiding policy is not required to be optimal, nor even an approximation of an optimal policy. However, one practical consideration is that the guiding policy should be fast to evaluate relative to the MPC solver, as it is evaluated at every time step.

At the first time step, the guiding policy is rolled out over the full MPC horizon N to determine the terminal constraint set center x_f . Specifically, $x_f = \phi_\pi(N; x(0))$, where $\phi_\pi(k; x)$ denotes the state reached after applying policy $\pi(\cdot)$ for k steps starting from the initial state x . The MPC problem from Eq. (4-2) is then solved with $\mathbb{X}_f(x_f)$ some compact set centered at x_f , yielding a terminal state $x^0(N; x, x_f) \in \mathbb{X}_f(x_f)$. At the next time step, the terminal constraint set center is updated by rolling out the policy for a single step from the previous MPC terminal state $x^0(N; x, x_f)$. The next terminal constraint set center is thus not dependent on the previous terminal constraint set center, but instead based on the previous terminal state:

$$x_f^+ = \pi(x^0(N; x, x_f)). \quad (6-1)$$

This procedure is also shown in Fig. 6-1, illustrating the propagation of the terminal constraint set over time relative to both the policy and MPC trajectories. Because MPC is free to select any $x \in \mathbb{X}_f(x_f)$, and the next center propagates from $x^0(N; x, x_f)$ rather than x_f , the center can drift over time relative to the pure policy rollout trajectory. The extent of this drift at each step is governed by the size of $\mathbb{X}_f(x_f)$, which bounds the per-step deviation, and by the terminal cost $V_f(x, x_f)$, which penalizes off-policy terminal states.

Two limiting cases recover on-policy behavior at the terminal time. First, if $\mathbb{X}_f(x_f)$ is a singleton, i.e., the terminal constraint $x(N) = x_f$ is enforced exactly, the optimizer has no freedom and always acts on-policy. Second, if $V_f(x, x_f)$ is sufficiently large such that no cost savings over the horizon can offset the cost of deviating from x_f , the optimizer will again choose $x(N) = x_f$ regardless of the constraint set size. In the microgrid EMS context, this second condition means that no amount of load shifting or input allocation over the horizon justifies acting off-policy at the end of the horizon.

6-1-2 Theoretical Properties

The construction of terminal ingredients in PG-MPC differs fundamentally from conventional EMPC. In conventional EMPC, discussed in Section 4-2, the terminal set is rendered positive invariant by the terminal control law, under which the terminal cost also satisfies a Lyapunov decrease property. Because these arguments do not apply to PG-MPC, new theory is required to establish guarantees on recursive feasibility and closed-loop performance relative to the guiding policy.

In PG-MPC the terminal constraint set is typically defined to be an n -sphere or box around the terminal state x_f , but in principle can be any compact set centered at x_f . Explicitly, the terminal constraint set can be given as

$$\mathbb{X}_f(x_f) = \{x \in \mathbb{X} \mid |x - x_f| \leq \tau\}, \quad (6-2)$$




Fig. 6-1: An example of two time steps using the PG-MPC controller.

where $\|\cdot\|$ denotes any p -norm. Remarkably, the terminal constraint set in the context of PG-MPC is not required to be positive invariant with respect to a terminal control law. Despite this, recursive feasibility of the PG-MPC optimization problem can be established. Under the standard MPC assumptions [57, Assumption. 2.2-2.3] and admissibility of the policy, it can be shown that the PG-MPC optimization problem is recursively feasible [23, Lemma. 1]. Intuitively, this property holds because the terminal constraint set center x_f^+ is propagated from the MPC terminal state $x^0(N; x, x_f) \in \mathbb{X}_f(x_f)$ via the admissible guiding policy, ensuring that a feasible solution always exists at the next time step by following the policy over the remainder of the horizon.

Performance guarantees of the PG-MPC algorithm can be established relative to the guiding policy. Consider the policy value function $J_\pi(x)$ defined as the infinite horizon cost of the policy

$$J_\pi(x) = \sum_{k=0}^{\infty} \ell(\phi_\pi(k, x), \pi(\phi_\pi(k; x))). \quad (6-3)$$

Under the assumption that $J_\pi(x) < \infty$, and provided that the terminal cost satisfies [23, Lemma. 2]

$$V_f(x, x_f) \geq J_\pi(x) - J_\pi(x_f), \quad (6-4)$$

along with $V_f(x_f, x_f) = 0$, it can be shown that the PG-MPC closed-loop transient performance satisfies

$$\sum_{k=0}^{T-1} \ell(x(k), u(k)) \leq J_\pi(x), \quad (6-5)$$

for all $T \geq 1$ [23, Theorem. 1]. Here, $x(k)$ and $u(k)$ denote the state and input at time k resulting from the PG-MPC control law in the closed loop. Hence, Eq. (6-5) states that the PG-MPC closed-loop cost is upper-bounded by the cost of the policy rollout. The key condition for this result is that the terminal cost $V_f(x)$ serves as an over-approximation of the cost increase resulting from acting off-policy, as expressed in Eq. (6-4).

The case that x is scalar will be considered moving forward, but it should be kept in mind that the stated theory generalizes to $x \in \mathbb{R}^n$. In the case that $J_\pi(x)$ is Lipschitz continuous, a sufficient condition for satisfying Eq. (6-4) globally is to choose the terminal cost as

$$V_f(x) = \rho|x - x_f|, \quad (6-6)$$




Fig. 6-2: Two examples of construction of terminal cost that satisfies Eq. (6-4), with left based on the global Lipschitz constant, and right based on the directional Lipschitz constants L^+ and L^- relative to x_f . Notation: $(x - x_f)^+ = \max(x - x_f, 0)$, $(x - x_f)^- = \max(-(x - x_f), 0)$.

with $\rho \geq L$, where L is the Lipschitz constant of $J_\pi(x)$. This suffices for Eq. (6-4) since $J_\pi(x_f)$ is constant in x , and Lipschitz constants are invariant under the addition of constants. Note that, as illustrated in Fig. 6-2, the Lipschitz constant might be a conservative approximation in particular directions of the state-space. In the context of microgrid energy management this might occur when the costs associated with having more or less charge in energy storage systems are asymmetric. To construct a terminal cost that still satisfies Eq. (6-4) the Lipschitz constants can be defined in the negative and positive direction relative to x_f . This is also shown in the right panel of Fig. 6-2. In this case, the terminal cost is given as the piecewise-linear function

$$V_f(x, x_f) = \begin{cases} L^+(x - x_f), & x \geq x_f, \\ L^-(x_f - x), & x < x_f. \end{cases} \quad (6-7)$$

Here, L^- represents the worst-case marginal cost increase in the negative direction relative to x_f , while L^+ does the same in the positive direction. Because the Lipschitz constants represent the worst-case cost increase, it directly follows that Eq. (6-7) over-approximates the true marginal cost increase $J_\pi(x) - J_\pi(x_f)$ in either direction and therefore satisfies Eq. (6-4). Thus, Eq. (6-7) can be used as a terminal cost in the PG-MPC algorithm, providing a less conservative alternative to Eq. (6-6).

As an alternative, the terminal cost can be constructed on-line as a local linearization of $J_\pi(x)$ centered at x_f , yielding a structure similar to Eq. (6-7). The advantage of this approach is that it exploits the local properties of $J_\pi(x)$, leading to less conservative performance compared to the Lipschitz bound. In particular, the Lipschitz constant based terminal cost by definition associates an increase in cost with all direction in the state-space, whereas local linearization can account for directions in the state-space having a zero (or negative) gradient. However, care should be taken as the role of the terminal constraint set changes in when using this approach. Rather than acting purely as a design parameter governing the allowable drift from the on-policy trajectory, it must be designed carefully to ensure that terminal states are constrained to a region where Eq. (6-4) is known to be satisfied.




Fig. 6-3: A block diagram showing the PG-MPC control architecture.

6-2 Policy-Guided MPC Controller

The controller design discussed in Chapter 5, along with the optimization problem stated in Eq. (5-44) remain identical for the PG-MPC controller. The key difference relative to the EMPC controller is the design of the terminal ingredients. The design of these ingredients in the context of PG-MPC is discussed in Section 6-1. This section proceeds with the design of the guiding policy and corresponding terminal ingredients for the CHP microgrid system described in Chapter 3, the overall architecture of which is shown in Fig. 6-3.

6-2-1 Policy Design

Fundamentally, the guiding policy will be designed to achieve the same task as the terminal constraints for the EMPC controller as discussed in Section 5-3-1. The guiding policy should be designed to ensure that the thermal balance is satisfied by the end of the year. There are a multitude of possible approaches for devising a policy which manages the ATES system, including methods such as RL and ADP to construct approximations of the optimal control policy. An example of this is given by Msaad et al. [24] in a greenhouse climate control setting, who apply an RL-guided MPC scheme that follows the same structure as PG-MPC, including the use of the learned policy for terminal constraint set construction. The results indicate performance gains relative to MPC at shorter horizons despite the authors remarking that the RL policy underperforms relative to the other controller evaluated. However, for the purpose of this thesis, a deliberately simple RBC will be used as this is both straightforward to design and implement, easy to interpret, and sufficient in showing the key advantages of the PG-MPC framework.

Inspecting the solution to the surrogate optimization problem in Fig. 5-5 reveals there is a clear seasonal trend in the charging and discharging of the thermal storage. Ideally, the policy design should reflect this seasonal trend, motivating the usage of reference tracking policy that




Fig. 6-4: Representation of dead-band policy behavior over time at different values of the imbalance state b .

attempts to track a sinusoidal trajectory on the imbalance. The ATEs dynamics have both an upper- and lower-bound to the amount of heat that can be extracted. Because of this, the simplest possible policy design is a bang-bang dead-band controller, with the dead-band being around a time-varying trajectory $\tilde{b}(t)$. The policy is then a function $\pi : \mathcal{T} \times \mathbb{R} \rightarrow \{\dot{Q}_{\text{in,max}}^H, 0, \dot{Q}_{\text{out,max}}^C\}$, such that the policy always pushes the imbalance state b back to the dead-band with maximum control effort. The dead-band will be defined as $\tilde{b}(t) \pm \Delta$, with $\Delta \geq 0$ the dead-band width. The reference signal $\tilde{b}(t)$ is taken as a sinusoid

$$\tilde{b}(t) = A \sin\left(\frac{2\pi}{N_{\text{year}}}t\right), \quad (6-8)$$

where A is the design amplitude and N_{year} is the number of time steps in a year. The policy can then be written as:

$$\pi(t, b) = \begin{cases} \dot{Q}_{\text{max}}^H, & \text{if } b < \tilde{b}(t) - \Delta, \\ \dot{Q}_{\text{max}}^C, & \text{if } b > \tilde{b}(t) + \Delta, \\ 0, & \text{otherwise.} \end{cases} \quad (6-9)$$

A schematic representation of the policy behavior over an operational year is shown in Fig. 6-4, illustrating how the (t, b) space is partitioned into regions relative to the dead-band. Outside the dead-band, the policy applies maximum effort to drive b back towards it: maximum cooling above and maximum heating below.

6-2-2 Terminal Value Function Construction

As discussed in Section 6-1-2, the terminal value function V_f should be chosen such that it is an over-approximation of $J_\pi(x) - J_\pi(x_f)$, which is the additional cost incurred from acting off-policy. An approximation of J_π can be constructed via Monte-Carlo rollouts of policy $\pi(t, b)$ for a random sample of initial conditions $\{(t_0^{(i)}, b_0^{(i)})\}_{i=0}^m$, yielding a dataset $\mathcal{D} = \{(t_0^{(i)}, b_0^{(i)}, J_\pi(t_0^{(i)}, b_0^{(i)}))\}_{i=0}^m$ of initial conditions and their associated rollout costs. The rollout cost is determined by a stage cost $\ell(t, \dot{W}^H, \dot{W}^C)$, capturing instantaneous operational

costs, and a terminal cost $\ell_T(b)$, capturing the cost of remaining imbalance on the final day of the year. The stage cost captures the electrical costs of injecting and extracting energy from the ATES system as

$$\ell(t, \dot{W}^H, \dot{W}^C) = \tilde{c}_{\text{el}}(t)(\dot{W}^H + \dot{W}^C)\Delta t, \quad (6-10)$$

where $\dot{W}^H$ and $\dot{W}^C$ are obtained from the policy actions $\dot{Q}_{\text{max}}^H$ and $\dot{Q}_{\text{max}}^C$ via the constant COP approximation, and $\tilde{c}_{\text{el}}(t)$ is taken as the hourly average electricity price over the period 2020–2025.

The terminal cost $\ell_T(b)$ reflects the economic cost of ending the year with a non-zero imbalance. Any remaining imbalance must be resolved by supplying thermal energy from an auxiliary source. Here, the boiler is assumed to supply the heating required to discharge remaining cooling imbalance, and the chiller is assumed to supply the cooling required to discharge remaining heating imbalance. The associated cost consists of two parts: (i) the electrical cost of running the heat pump to rebalance the ATES, and (ii) the additional cost of operating the boiler or chiller to provide the required thermal energy. Note that there is also an opportunity cost arising from the ATES being unavailable for heating during January while rebalancing, which is not explicitly accounted for. Instead, the maximum observed electricity price is used as a conservative proxy for the true cost of imbalance, giving

$$\ell_T(b) = \max_t(\tilde{c}_{\text{el}}(t)) \left(b^- \left(\frac{\text{COP}_H}{\text{COP}_H - 1} + \frac{1}{\eta_{ch}} \right) + b^+ \left(\frac{\text{COP}_C}{\text{COP}_C - 1} + \frac{1}{\eta_b} \right) \right), \quad (6-11)$$

where $b^+ = \max(b(T), 0)$ and $b^- = \max(-b(T), 0)$ denote the remaining heating and cooling imbalance respectively. The $\frac{\text{COP}}{\text{COP}-1}$ terms convert imbalance to electrical energy via the ATES heat pump, while $\frac{1}{\eta_b}$ and $\frac{1}{\eta_{ch}}$ capture the additional electrical cost of operating the boiler and chiller to supply the rebalancing thermal energy. The empirical policy value function is then given by the sum of the stage and terminal costs accumulated along the policy rollout:

$$J_\pi(t_0, b_0) = \sum_{k=t_0}^{T-1} \ell(k, \dot{W}^H(k), \dot{W}^C(k)) + \ell_T(b(T)). \quad (6-12)$$

Note that $J_\pi(t_0, b_0)$ is finite by construction, and hence satisfies the property $J_\pi(x) < \infty$, as the policy rollout is defined over a finite horizon.

Value function approximation has been studied extensively, with methods ranging from radial basis functions and spline interpolation to Gaussian processes and linear basis function architectures [85, Ch. 7]. Here, the approximation is formulated as a supervised learning problem, training an ANN with parameters θ to map initial conditions to an approximation of their associated empirical cost such that $\tilde{J}_\pi^\theta(t, b) \approx J_\pi(t, b)$. The resulting loss function is then

$$\mathcal{L}(\theta, \mathcal{D}) = \frac{1}{N_s} \sum_{i=1}^{N_s} \left(\tilde{J}_\pi^\theta(t_0^{(i)}, b_0^{(i)}) - J_\pi(t_0^{(i)}, b_0^{(i)}) \right)^2, \quad (6-13)$$

which is minimized using the ADAM optimizer [86]. A neural network is chosen primarily because the Lipschitz constants of this surface will be needed for construction of the time-varying terminal cost. The time-varying Lipschitz constants can be found as the upper-bound on the gradient of the value function at fixed t :

$$L(t) = \max_b \left| \frac{\partial \tilde{J}_\pi^\theta(t, b)}{\partial b} \right|_t, \quad (6-14)$$

which is easily obtained efficiently and exactly via automatic differentiation [87].

Since the heat pump COP differs between heating and cooling modes, it is desirable to incorporate this asymmetry into the terminal cost structure. Several approaches exist for achieving this. One option is to embed the ANN directly into the MILP, which is possible when using a multilayer perceptron with ReLU activations [88], [89]. However, this approach is computationally intractable even for small networks, as each neuron introduces a binary variable and a big-M constraint [90]. A second option is to replace the ReLU activations with tanh and construct an online local linearization of $\tilde{J}_\pi^\theta(t, b)$. While this better captures local properties of the value function, it requires careful tuning of the terminal constraint set to ensure that Eq. (6-4) remains satisfied within this set.

The simplest approach, and the one adopted here, is to use directional Lipschitz constants relative to the reference trajectory $\tilde{b}(t)$. Define

$$L^+(t) = \max_{\tilde{b}(t) \leq b \leq b_{\max}} \left. \frac{\partial \tilde{J}_\pi^\theta(t, b)}{\partial b} \right|_t, \quad L^-(t) = \max_{b_{\min} \leq b \leq \tilde{b}(t)} \left. \frac{\partial \tilde{J}_\pi^\theta(t, b)}{\partial b} \right|_t, \quad (6-15)$$

where $L^+(t)$ and $L^-(t)$ are the Lipschitz constants associated with deviating from $\tilde{b}(t)$ in the positive (excess heating) and negative (excess cooling) directions, respectively. These constants have a natural economic interpretation: $L^+(t)$ and $L^-(t)$ represent the worst-case marginal cost of increasing the imbalance in either direction at time t relative to the reference trajectory. This approach conveys more local information about the value function than a single global Lipschitz constant, while preserving the PG-MPC guarantees. A further practical advantage is that $L^+(t)$ and $L^-(t)$ can be computed off-line, avoiding any additional online computation.

Using the directional Lipschitz constants, the terminal cost can be constructed following Eq. (6-7), resulting in

$$V_f(b, b_f) = \begin{cases} L^+(b - b_f), & b \geq b_f, \\ L^-(b_f - b), & b < b_f, \end{cases} \quad (6-16)$$

which captures the economic cost associated with deviating from the on-policy terminal state b_f in either direction.

The approximate value function surface and the resulting time-varying directional Lipschitz constants are shown in Fig. 6-5. For interpretability, the surface displayed is the shifted value function $\tilde{J}_\pi^\theta(t, b) - \tilde{J}_\pi^\theta(t, \phi_\pi(t; \tilde{b}(t)))$, which subtracts the on-policy cost at each time step. Since $\tilde{J}_\pi^\theta(t, \tilde{b}(t))$ is constant in b , this shift does not affect the Lipschitz constants and therefore leaves the terminal cost construction unchanged. The shifted surface is zero along the reference trajectory $\tilde{b}(t)$ by construction, making deviations from the on-policy behavior directly visible.

6-2-3 Terminal Constraint Set Design

The terminal constraint set $\mathbb{X}_f(b_f)$ is defined as a box constraint around the terminal constraint set center x_f :

$$\mathbb{X}_f(b_f) = \{b \in \mathbb{R} \mid |b - b_f| \leq \Delta_f\}, \quad (6-17)$$




Fig. 6-5: The approximate value function surface along with its time-varying directional Lipschitz constants relative to the trajectory $\tilde{b}(t)$.

where $\Delta_f \geq 0$ is the half-width of the terminal constraint set. As discussed in Section 6-1-2, $\mathbb{X}_f$ is not required to be positive invariant, and can be treated as a pure design parameter limiting the allowable drift of the terminal state away from the on-policy trajectory. A larger Δ_f grants the PG-MPC controller more freedom to deviate from the guiding policy at the end of the horizon, potentially improving short-horizon performance at the cost of acting less on-policy. Conversely, a smaller Δ_f more closely constrains the optimizer to follow the guiding policy, resulting in overall slower drift from the on-policy trajectory. It should also be noted that Δ_f interacts directly with the terminal cost V_f : a large terminal cost relative to Δ_f will naturally discourage the optimizer from exploiting the full extent of the terminal constraint set, while a small terminal cost may allow excessive drift.

6-2-4 Full PG-MPC Problem

The additional constraints introduced for modeling the piecewise terminal cost using the procedure outlined in Section 4-3-3 are collected into the full set of MLD constraints given in Eq. (5-41). All other constraints discussed in Section 5-2-1 remain identical for the PG-MPC controller. The full statement of the PG-MPC problem is identical to Eq. (5-44), where in the PG-MPC case, V_f and $\mathbb{X}_f$ correspond to Eqs. (6-6) and (6-17) respectively.

Chapter 7

Case Study

This chapter presents the experimental setup, simulation results, and discussion for the evaluation of the PG-MPC controller against the EMPC baseline. The chapter is structured as follows. Section 7-1 describes the experimental setup, including the weather scenarios and electricity price profiles used to evaluate controller performance. Section 7-2 defines the performance metrics used for the comparative analysis. Section 7-3 presents the simulation results across all four scenarios, and Section 7-4 discusses the results in terms of operational cost, constraint satisfaction, and controller behavior, concluding with a reflection on the role of policy design and the robustness of the PG-MPC framework.

7-1 Experimental Setup and Scenarios

This section describes the simulation scenarios and performance metrics used to evaluate both controllers. The scenarios are constructed to span a representative range of climatological and electricity market conditions relevant to the Netherlands, allowing the sensitivity of both controllers to year-specific variability to be assessed.

7-1-1 Weather

To evaluate the performance of the EMS under a range of climatological conditions representative of the Netherlands, four annual simulation scenarios are considered. The scenarios are constructed using weather data retrieved from the KNMI EDR API [53]. The scenarios are selected to span a range of summer thermal conditions relevant to the building cooling loads and photovoltaic generation profiles.

Although the scenarios are evaluated over a full operational year, the summer profile is the main driver for inter-scenario variability. The thermal balance of the ATES system requires that energy extracted when heating in winter and injected during cooling in summer must remain balanced. In the Netherlands, this annual balance is asymmetric. Heating demand in

the winter is reliably larger than cooling demand in summer. The summer cooling introduces uncertainty as the magnitude of cooling demand varies substantially from year to year. As the heating from thermal storage is cheap compared to using the auxiliary boiler, the operational costs are mainly impacted by the cooling demand profile during summer, as this is the main driver for available heating during the winter season. The four scenarios are therefore selected to systematically span varying summer conditions observed in the Netherlands.

The reference scenario uses a TMY sourced from EPW data [78]. This scenario represents statistically average conditions based on measurements from De Bilt in the Netherlands during the period 1949–2023 and serves as the baseline against which all other scenarios will be compared. The EPW is used to characterise nominal controller behavior. Additionally, the EPW data will be used to generate thermal demand profiles for the baseline EMPC controller. As forecasting the temperature profile a year ahead is not feasible, it is standard practice to use TMY data for assessing building thermal demand [91].

The weather data from the year 2021 is used to represent a year with a near-average to mild summer. With a mean temperature close to the long term, this scenario produces moderate cooling demand and serves as a test case which approaches TMY conditions while still capturing realistic, year-specific variability in solar irradiance and thermal load profiles.

The 2025 scenario represents a moderately warm year, with two official heatwaves recorded and above-average temperatures across multiple months, providing an intermediate scenario between the average and extreme cases. Given projections of increasing heat extremes in the Netherlands across all emission scenarios [92], [93], conditions such as those observed in 2025 are likely to become increasingly representative of future Dutch summers, making it a relevant scenario for evaluating controller performance.

Finally, the 2018 scenario is used to represent a year with a particularly hot summer. The summer of 2018 is the warmest meteorological summer on the record in the Netherlands, resulting in sustained elevated cooling loads. This scenario is used to assess robustness under high thermal demand.

Table 7-1 shows a summary of the scenarios discussed above. The difference in scenarios is also shown in Figs. 7-2a and 7-2b, showing the amount of cumulative CDD and the cumulative solar irradiance. It is worth noting that, as presented in Fig. 7-2a, 2025 started as hot as 2018, particularly in the months April up until July. However, this heat was not sustained beyond July, justifying the categorization of 2025 as a moderate scenario rather than an extreme scenario. These trends can also be observed in Fig. 7-1, showing the 2018 and 2025 scenarios reach monthly average temperatures higher than 10 °C by April.

Table 7-1: Overview of the discussed assessment scenarios to be used to evaluate controller performance.

Scenario	Year	Summer Character	Role	Seasonal Mismatch
TMY	-	Climatological average	Nominal baseline	329 MWh
Mild	2021	Near-average	Low cooling demand	322 MWh
Moderate	2025	Moderately hot, two heat waves	Moderate to high cooling demand	252 MWh
Extreme	2018	Hottest summer on record	High cooling demand	260 MWh




Fig. 7-1: The monthly average temperature across all four scenarios.




Fig. 7-2: Cumulative solar irradiation across the four scenarios, highlighting inter-scenario variability in cooling demand (a) and PV generation potential (b).

7-1-2 Energy Prices

Another important aspect distinguishing the scenarios is the electricity price profiles. As discussed in Section 3-4, the prices are obtained from historical spot market data for the Netherlands. The TMY scenario uses a price profile constructed as the hourly average over the preceding five years. Table 7-2 summarises the key statistics by scenario, Fig. 7-4 shows the corresponding price duration curves, and Fig. 7-3 shows the monthly average electricity prices.

The TMY scenario shows the highest mean price but a low coefficient of variation, with limited mean daily spread and few negative price hours. The relatively stable character of this profile is a direct consequence of multi-year averaging, which smooths out the year-specific volatility present in the historical scenarios.

The 2021 scenario has a high mean price and the highest coefficient of variation across all scenarios, reflecting the energy market turbulence driven by rising gas prices during this period [94]. Fig. 7-3 shows that the electricity mean monthly electricity price remained constant up to and including June. After this period, both the mean and standard deviation of the monthly electricity price start to increase. Despite the volatility, negative price hours remain limited.




Fig. 7-3: The monthly average electricity price across all four scenarios, showing the temporal trends of the statistics listed in Table 7-2

Renewable penetration, in particular wind, was comparatively lower in 2021 [2], and as wind power is identified as the primary driver of negative electricity prices [95], this is consistent with renewable curtailment not yet being a dominant market feature at the time.

The 2025 price profile is characterised by a moderate mean price and coefficient of variation, yet exhibits by far the highest mean daily spread and number of negative price hours across all scenarios. This divergence highlights that Coefficient of Variation (CV) alone does not fully capture the structure of price volatility relevant to EMSs. The 2025 profile features large intraday price swings driven by periods of near-zero or negative midday prices, followed by elevated evening prices. This pattern is consistent with the substantial growth in renewable generation capacity in the Netherlands between 2019 and 2025, rising from 19% to over 50% [2], which increasingly suppresses midday spot prices while leaving peak demand periods unaffected. This further reinforces 2025 as a scenario representative of expected future operating conditions.

The 2018 scenario has the lowest mean price and lowest volatility across all metrics. With a limited share of renewable generation, dispatchable gas-fired plant set the marginal price for the majority of hours, resulting in a price profile that closely mirrors gas price dynamics, which is a well-documented relationship in liberalised European power markets [94].

Table 7-2: Summary of electricity price profile statistics for each scenario

Scenario	Year	Mean Price [€/kWh]	CV [%]	Mean Daily Spread [€/kWh]	Negative Price Hours [h]
TMY	-	0.109	38.1	0.096	13.0
Mild	2021	0.101	73.3	0.108	94.75
Moderate	2025	0.087	56.7	0.146	521.75
Extreme	2018	0.052	33.5	0.064	7.0




Fig. 7-4: The price duration curves for all four scenarios. The 2018 and TMY scenarios showing stable electricity prices compared to the more volatile 2021 and 2025 scenarios.

7-2 Performance Metrics

The performance of both controllers is evaluated across three main metrics. The first is operational cost, capturing the total economic cost incurred over the simulation year. The second is temperature comfort constraint satisfaction, measuring the extent to which the building zone temperature remains within the prescribed comfort band. The third is thermal balance constraint satisfaction, assessing whether the ATES system achieves the required end-of-year heat balance. These three metrics are evaluated jointly rather than in isolation, since improvements in one may come at the expense of another.

The operational costs capture the total economic cost incurred over the simulation year. This comprises two components: the net cost of exchanging power with the electricity grid, determined by a time-varying import and export cost, and the battery degradation cost, which accumulates as a function of charging and discharging power over the year. Together, these represent the primary economic objective that both controllers are designed to minimize. Operational costs are reported as cumulative costs over the simulation year, allowing the relative performance of both controllers to be assessed not only at year-end, but also in terms of the timing of cost accumulation throughout the year, which can reveal differences in control strategy that a single end-of-year figure would otherwise obscure.

Comfort constraint satisfaction is assessed by examining the cumulative constraint violation penalty incurred over the simulation year. A penalty is incurred at each timestep at which the building zone temperature falls outside the comfort band of 20–21 °C, proportional to the magnitude of the violation. Reporting the cumulative violation penalty alongside the operational cost allows the cost savings of one controller relative to the other to be assessed in the context of whether those savings are achieved at the expense of comfort constraint satisfaction, rather than from a genuine improvement in operational efficiency.

Thermal balance constraint satisfaction is assessed by examining the cumulative heat injected and extracted from the ATES system over the full operational year. A thermally balanced system requires that the balance between extracted and injected heat sums to zero by the end of the year, as expressed in Eq. (5-52). Rather than reporting only the end-of-year

balance, the full cumulative trajectory is examined for each scenario, as the path by which both controllers achieve the final balance is informative of the underlying control strategy and its interaction with the seasonal demand profile.

7-3 Simulation Results

This section presents the simulation results for both the EMPC and PG-MPC controllers across the four scenarios defined in Section 7-1. The results are first presented for the TMY scenario as a representative overview of nominal system behavior, before proceeding to the comparative analysis across all scenarios. The comparative analysis presents the metrics discussed in Section 7-2 for both the EMPC controller designed in Chapter 5 and the PG-MPC controller designed in Chapter 6, allowing the two controllers to be compared directly under identical disturbance conditions.

7-3-1 Nominal System Trajectory

This subsection presents the full system state and input trajectories of both the EMPC baseline controller and the PG-MPC controller under the TMY scenario, serving as a representative overview of nominal system behavior. The TMY scenario is chosen for this purpose as it represents statistically average climatological conditions. These results are not meant as a performance index, but rather to establish that the simulated system exhibits physically plausible behavior for both controllers before proceeding with the comparative analysis. This step is a necessary precursor to the comparative analysis, since any quantitative comparison would be of limited value if the underlying system behavior were not first shown to be physically sound.

Figs. 7-5a and 7-5b shows the evolution of all relevant system states and control inputs over the simulation year. Particularly, the temperatures, battery state of charge, thermal inputs, and electrical inputs are shown. The expected seasonal structure of the Dutch climate is clearly visible in the charging and discharging of the ATES wells. This seasonal trend can be seen in Fig. 7-5a, with the well temperatures decreasing in winter at the start and end of the year, and increasing during the summer in the middle of the year. The battery state of charge follows a diurnal pattern in response to the available PV generation. The observed patterns indicate that the system simulation behaves as expected and is physically consistent.

Figs. 7-6a and 7-6b show the corresponding system trajectories for the PG-MPC controller under the same TMY scenario. The overall system behavior is qualitatively similar, with the ATES well temperatures showing a similar seasonal trajectory, and the battery showing a similar diurnal cycle. This is as expected, as both controllers are operating using the same disturbance signals and the main difference between the two controllers is in the design of the terminal ingredients, which only affects the ATES charging and discharging behavior. The PG-MPC controller can therefore also be considered physically consistent, and both controllers are taken forward for the purpose of comparative analysis in the next subsection. Any quantitative differences observed in that analysis can therefore be attributed primarily to the terminal ingredient design, rather than to differences in the underlying system response.




Fig. 7-5: Nominal system states and inputs using the time-varying EMPC baseline controller using the TMY disturbance signals.




Fig. 7-6: Nominal system states and inputs using the PGM-MPC controller using the TMY disturbance signals.




Fig. 7-7: Cumulative operational cost and cumulative constraint violation penalties for both the EMPC and PG-MPC controllers for the TMY scenario.

7-3-2 Comparative Analysis

Section 7-3-1 shows the baseline simulation results for an entire year. This section will analyze comparatively the performance of both controllers under the metrics defined in Section 7-2 for each scenario.

Scenario - TMY

Fig. 7-7a shows the cumulative operational cost for both controllers under the TMY scenario. The PG-MPC controller shows a different trajectory compared to the EMPC controller, with the PG-MPC controller exhibiting lower operational costs throughout the year, but seeing a sharp increase at the start of November. The PG-MPC controller achieves a final cost of €18,580, compared to EMPC achieving a final cost of €17,945. Hence, the operational cost of the PG-MPC controller is higher by 3.54% compared to the baseline EMPC controller for this scenario. Fig. 7-7b shows that similar initial penalties for both EMPC and PG-MPC at the start of the year due to initialization outside of the comfort band. PG-MPC incurs more violation penalties across the operational year compared to EMPC. Particularly, from March forward with substantially more constraint violation penalty accumulated in September.

The trajectories for the well temperatures and thermal imbalance are shown in Fig. 7-8. Both controllers end the year with no remaining imbalance. The EMPC controller perfectly tracks the surrogate solution, as the surrogate solution for this scenario and the trajectory used for the EMPC controller are the same. The PG-MPC shows significant deviation from the surrogate trajectory throughout the entire year. Particularly early in the year, the PG-MPC controller closely follows the policy trajectory, but from March forward the PG-MPC controller shows significant deviation from the policy trajectory. The well temperature trajectories reflect the expected seasonal trajectories for both controllers throughout January up until November. From the November through to the end of December, the PG-MPC shows deviation from the expected discharging of the hot well and charging of the cold well during the winter months. This is also supported by the imbalance trajectory indicating that the PG-MPC controller is injecting thermal energy into the ATEs (i.e., cooling the reservoir)




Fig. 7-8: Hot and cold well temperatures, as well as thermal imbalance trajectories for both EMPC and PGMPC for the TMY scenario.




Fig. 7-9: Building temperature for both EMPC (left) and PGMPC (right) over the operational year for the TMY scenario.

during the months of November and December.

The increase in the amount of energy extracted from the building when cooling can also be clearly observed in the temperature profile in Fig. 7-9, showing that the PG-MPC controller tends to stay around 20 °C, compared to the EMPC controller observing spikes up to almost 21 °C in the cooling months of May to September. Neither controllers show significant violations of the comfort band of 20-21 °C except for at the very start of the year when the initial building temperature is initialized outside of this comfort band.

Scenario - 2021 (Mild)

The 2021 scenario represents a mild summer with near-average cooling demand and a moderate amount of negative price hours. Fig. 7-10a shows the cumulative operational cost for both controllers. Compared to the TMY scenario, this scenario shows similar operational costs. However, in this case, the PG-MPC method shows a lower operational cost of €14,869 compared to the EMPC controller which has an operational cost that is 8.7% higher at €16,164. In particular, Fig. 7-10a shows that the main cost savings of PG-MPC relative to EMPC happen during the periods January through March and October through December. Fig. 7-10b indicates that both algorithms incur similar penalties throughout most of the year for slight




Fig. 7-10: Cumulative operational cost and cumulative constraint violation penalties for both the EMPC and PG-MPC controllers for the 2021 (Mild) scenario.

constraint violations of the comfort band, with large spikes happening at the end of the year due to EMPC violating the comfort band around mid-to-late December shown in Fig. 7-12.

The well temperature and imbalance trajectories for the 2021 scenario are shown in Fig. 7-11. The results for the EMPC controller indicate that the terminal constraint set trajectory based on the TMY solution to the surrogate problem is similar in both shape and magnitude to the solution of the surrogate problem for the 2021 scenario, with only slight deviations happening around June and December. The PG-MPC controller shows a substantially different trajectory compared to the EMPC controller, with more aggressive extraction of heat from the ATES system during the January to March period. The PG-MPC controller also starts cooling slightly earlier in the year at around April, compared to the EMPC controller, which starts around mid to late May. The earlier start of cooling coincides with PG-MPC overtaking EMPC in terms of operational cost in the early March to late April period. Overall, the PG-MPC injects more heat into the ATES system during summer, allowing for deployment of more heating in the November and December winter period. Near the end of the year, the PG-MPC shows significant deviation from the policy trajectory, particularly from July forward, the magnitude of the imbalance is higher by about 10 MWh compared to the policy trajectory.

In terms of well temperatures, both controllers leave the thermal storage at nearly identical temperatures on the final day of the year. Additionally, early cooling by the PG-MPC controller is observed, with cooling starting slightly earlier compared to the EMPC controller. The PG-MPC controller shows larger temperature ranges in both the hot and cold wells, indicating deeper charging and discharging cycles of the ATES system.

The building temperatures in Fig. 7-12 indicate slight comfort band constraint violation for the EMPC controller at around mid December. Again, higher temperatures in the building are observed near the summer period for the EMPC controller compared to the PG-MPC, which tends to stick closer to the lower bound of 20 °C comparatively.




Fig. 7-11: Hot and cold well temperatures, as well as thermal imbalance trajectories for both EMPC and PGMPC for the 2021 (Mild) scenario.




Fig. 7-12: Building temperature for both EMPC (left) and PGMPC (right) over the operational year for the 2021 (Mild) scenario.

Scenario - 2025 (Moderate)

The 2025 scenario introduces an early summer with 2 recorded heat-waves and significantly more negative price hours. Together, these conditions create a year that departs substantially from the climatological average underlying the TMY scenario, making it a useful test of how each controller adapts to atypical seasonal and pricing conditions. Fig. 7-13a shows the cumulative operational cost, showing lower operational cost of €7,917 for PG-MPC compared to €12,035, amounting to a 34.2% decrease in operational cost for the PG-MPC algorithm. Fig. 7-13b again shows similar accumulation of constraint violation over the year for both algorithms, with PG-MPC ending slightly higher compared to EMPC. Notably, the additional cost incurred from constraint violation by PG-MPC is less than the economic cost savings relative to EMPC, indicating that PG-MPC has a lower total cost in this scenario. This combination of lower operational cost alongside marginally higher constraint violation is consistent with the trend observed across the other scenarios, in which PG-MPC achieves its cost advantage in part by operating closer to the boundary of the comfort band.

The surrogate solution in Fig. 7-14 indicates that for this scenario, there is significant deviation between the TMY trajectory that EMPC tracks and the solution to the surrogate problem for the 2025 scenario. In particular, the cooling period starts earlier than expected




Fig. 7-13: Cumulative operational cost and cumulative constraint violation penalties for both the EMPC and PG-MPC controllers for the 2025 (Moderate) scenario.




Fig. 7-14: Hot and cold well temperatures, as well as thermal imbalance trajectories for both EMPC and PG-MPC for the 2025 (Moderate) scenario.

in around late April for the 2025 scenario, compared to the mid March in the TMY trajectory. There is also significantly more cooling demand than expected in the months May through September. PG-MPC shows a trajectory that closely matches the surrogate solution in terms of magnitude early in the year, but there is a temporal mismatch throughout the year between the start of the heating and cooling seasons. Additionally, PG-MPC shows a higher amount of cooling deployed compared to the surrogate solution, particularly in September. The Well temperatures at the final day of the year are similar across both methods again, but PG-MPC shows a very minor deviation, with both the hot and cold wells being slightly lower temperature compared to the EMPC solution.

The building temperatures in Fig. 7-15 reflect the higher overall temperatures during the summer period, with both the EMPC and PG-MPC controllers frequently operating near the upper-bound of the comfort band. Again, the only constraint violations throughout the year are due to minor ambient temperature disturbance signal driving the building temperature above or below the comfort band. The PG-MPC also shows higher temperatures in November, coinciding with the higher amount of ATEs heating deployment relative to the EMPC controller during the same period.




Fig. 7-15: Building temperature for both EMPC (left) and PGMPC (right) over the operational year for the 2025 (Moderate) scenario.




Fig. 7-16: Cumulative operational cost and cumulative constraint violation penalties for both the EMPC and PG-MPC controllers for the 2018 (Extreme) scenario.

Scenario - 2018 (Extreme)

The 2018 scenario represents the most extreme scenario in terms of temperatures, with this scenario having the hottest recorded summer in the Netherlands. However, in terms of electricity pricing, this scenario is the least volatile, with the lowest coefficient of variation, mean daily spread, and negative price hours. In the 2018 scenario, Fig. 7-16a shows that both controllers present similar operational costs across the year. The EMPC controller has an operational cost of €6,924 compared to the very slightly higher cost of €6,953 for the PG-MPC controller. The constraint violation again indicates that by the end of the year, PG-MPC has incurred slightly higher constraint violation penalties compared to EMPC.

The surrogate solution in Fig. 7-17 suggest that the magnitude of cooling demand for this year is larger than the TMY predicts. However, the trajectory planned on TMY data compared shows similar magnitude and timing of heating and cooling compared to the surrogate solution. The PG-MPC controller in this case closely follows the sinusoidal trajectory. The well temperatures, as well as the heating and cooling curves again show that PG-MPC has deeper charging and discharging of the thermal storage system.




Fig. 7-17: Hot and cold well temperatures, as well as thermal imbalance trajectories for both EMPC and PG-MPC for the 2018 (Extreme) scenario.




Fig. 7-18: Building temperature for both EMPC (left) and PGMPC (right) over the operational year for the 2018 (Extreme) scenario.

The hot summer temperatures of the 2018 scenario are clearly observed in Fig. 7-18, showing frequent operation near the upper-bound of the comfort band. The PG-MPC instead deploys more cooling as shown in Fig. 7-17, and tends to operate closer to the lower-bound instead.

7-4 Discussion

This section discusses the results presented in Section 7-3 and the analysis across the four scenarios. The discussion is structured around the performance metrics and constraints discussed in Section 7-2: operational cost, comfort constraint satisfaction, and thermal balance constraint satisfaction. This is followed by a discussion of the observed controller behavior and its implications for the design of the PG-MPC controller.

7-4-1 Closed-Loop Cost

Across the four scenarios, the PG-MPC and EMPC controllers show comparable operational cost, with the exception of the 2025 (Moderate) scenario, in which PG-MPC achieves a

substantially lower operational cost. The PG-MPC controller also shows slightly lower operational cost in the 2021 scenario. In the TMY and 2018 scenarios, the PG-MPC controller incurs marginally higher operational costs, though the difference in the 2018 scenario is negligible at less than 0.40%, and both controllers can be considered to perform equivalently in terms of economic cost in these scenarios. It is worth noting that the TMY scenario represents the most favourable conditions for the EMPC controller, as the time-varying terminal constraint set is derived directly from the TMY surrogate solution, giving EMPC an *a priori* informational advantage that would not be available in practice. Despite this, the PG-MPC controller achieves competitive operational cost, albeit via a substantially different imbalance trajectory.

The large performance gap between the two controllers in the 2025 scenario can be explained by a couple of factors. First, the surrogate solution for the 2025 scenario shows significant deviation from the EMPC terminal constraint trajectory from the start of January. This leads to unmet heating demand which needs to be satisfied by deploying boiler heat, as the ATES heating deployment is limited by the terminal constraint. 2025 also had a particularly early summer, as shown in Fig. 7-2a, resulting in EMPC forcing deployment of heating around early May even though there was no heating demand. The TMY trajectory also significantly underestimates the 2025 cooling demand by nearly 10 MWh, resulting in under-utilization of the ATES system. The PG-MPC controller better adapts to the particularly hot and early summer, leading to more aggressive cooling, which is also reflected in the building temperature trajectories shown in Fig. 7-15, where PG-MPC operates at lower average temperatures during these months. This aggressive cooling between May and September allows for more ATES heating to be deployed in November, leading to substantially lower economic operational cost during this period as seen in Fig. 7-13a. It should be noted that Figs. 7-13a and 7-13b together indicate that the total closed-loop cost of the full year is lower for the PG-MPC controller compared to the EMPC controller.

While the 2021 scenario also saw substantial price volatility, the difference in operational cost between the two controllers is smaller. This can be attributed to the TMY surrogate solution already being a close approximation of the optimal solution for 2021, which is supported by the similar cumulative CDD and solar irradiance shown in Figs. 7-2a and 7-2b. In the period July to September, the PG-MPC controller shows significant deviation from the policy trajectory, as the lower-than-expected cooling demand leads it to consistently choose off-policy actions whose stage-cost reduction outweighs the additional terminal cost incurred. Despite the lower cooling demand, PG-MPC still injects substantially more heat into the ATES system during May to September compared to EMPC, leaving more ATES heating available for deployment later in the year and reducing costs in the October to December period.

In terms of comfort constraint satisfaction, all controllers show a large initial spike corresponding to initialization outside of the comfort band. Across all scenarios, PG-MPC incurs slightly higher cumulative constraint violation cost compared to EMPC. Inspection of the building temperature trajectories reveals that PG-MPC consistently operates closer to the bounds of the comfort band, which increases exposure to disturbances and forecast mismatch that can push the temperature state outside the comfort band. It should therefore be noted that the operational cost savings of PG-MPC in the 2021 scenario, and its competitive performance in the TMY and 2018 scenarios, are partly achieved by exploiting this margin: PG-MPC implicitly trades a modest increase in constraint violations for reduced energy cost by consistently operating near the boundary of the comfort band. The increased operation

near particularly the bottom of the comfort band during summer in the TMY, 2025, and 2021 scenarios, can be attributed to the guiding policy pushing for more cooling compared to the EMPC terminal trajectory during this period. Hence, the building temperature operates closer to the lower-bound of the comfort band as a direct consequence of the increased cooling.

7-4-2 Thermal Balance Constraint Satisfaction

Both controllers satisfy the thermal energy balance constraint by the final date of the year in all scenarios. Across the three scenarios based on historical weather data, that is the 2021, 2025, and 2018 scenarios, the PG-MPC controller follows the expected seasonal pattern, with net heat injection during the cooling months of May through September and net heat extraction during the heating months of January through April and October through December. The notable exception is the TMY scenario, in which the PG-MPC controller exhibits substantial deviation from this behavior. As shown in Fig. 7-8, the imbalance trajectory diverges significantly from the policy trajectory from March onward. Due to the low cooling demand in this scenario, the cumulative imbalance stays strictly positive throughout the entire year, leaving no extraction available for the November and December period. With the system at positive imbalance at the start of November, the terminal cost penalty is substantial, as illustrated in Fig. 6-5, driving the controller toward active cooling to restore the thermal balance by year-end. This rebalancing behavior is a desirable property of the algorithm, as it guarantees thermal balance constraint satisfaction. However, as a trade-off, it leads to non-ideal operation in December, where the boiler must be relied upon not only to satisfy building thermal demand, but also to compensate for the additional load imposed by rebalancing the thermal storage, resulting in elevated operational costs. The underlying cause of this divergence is not fully understood: the 2021 scenario exhibits similarly low cooling demand yet does not produce the same behavior, suggesting the difference is not driven by cooling demand alone, but more likely by some interaction between the realized price trajectory and the dead-band policy's response to it.

7-4-3 Controller behavior and Policy Deviation

The PG-MPC shows minimal deviation compared to the sinusoidal policy trajectory in the 2018 (Extreme) scenario. This can be attributed to interaction between the lower mean price, and lower volatility of electricity prices in 2018 and the design of the terminal cost penalty used in the PG-MPC controller. The Lipschitz constant used in construction of the terminal cost represents the worst case marginal cost increase from acting off-policy based on the historically average electricity price in the last five years. As discussed in Section 7-1-2 and partially shown in Table 7-2, electricity price volatility, as well as mean price, saw a substantial increase post-2020, leading to the historical average in the last five year being a significant over-estimation of the true cost of acting off-policy for the 2018 scenario. Due to the large over-estimation, and low amount of volatility to exploit in the electricity pricing, there is little room left for PG-MPC to find solutions which save sufficient cost over the horizon to offset the terminal cost incurred from acting off-policy, and as a result, PG-MPC will always pick either the next on-policy terminal state, or states very close to the next on-policy terminal state.

The TMY scenario, shown in Fig. 7-8, saw substantial deviation from the on-policy trajectory, particularly from March forward. The deviation from policy around March can partially be explained by Fig. 7-2a showing that cooling demand only starts around mid-to-late April, whereas the policy trajectory is pushing to cool starting from late March to early April. This pushes the controller to cool in the period mid May through to the start of September. However, due to the low cooling demand, building temperature gets pushed to the lower-bound of the comfort band, with large costs associated with the controller cooling more beyond this point. Fig. 7-3 shows an increase in both the mean and variance of the electricity prices during the months August and September, further compounding the effect and causing the controller to act off-policy to save on costs further in this period. Fig. 7-7a also suggest this, as at this point in time, the cost gap between the controllers is largest. These short term savings by PG-MPC ultimately lead to an increase in operational cost by the end of the year, as the thermal storage is left in a bad state by November. As a result, the PG-MPC supplies cooling to the building which needs to be paid for and compensated for with extra boiler heat during the November to December period in order to meet the thermal balance constraint by the end of the year.

7-4-4 Policy Design

A key consideration in interpreting the PG-MPC simulation results is the role of the policy design in shaping performance outcomes. Many observed limitations, particularly the mistiming between the heating and cooling season transitions, are not inherent weaknesses of the PG-MPC framework itself, but rather artifacts of the design of the underlying policy.

The policy employed in this work is intentionally kept simple to serve as a baseline. By construction, the sinusoidal trajectory is suboptimal for any particular year, as it does not reflect the true climate conditions. The framework relies on the PG-MPC controller finding solutions that locally save cost by drifting away from the policy trajectory, resulting in cost savings and constraint satisfaction over the operational year. This is most strongly observed in the TMY and 2021 scenarios, in which the PG-MPC trajectories significantly drift away from the policy trajectories. If the electricity prices are low relative to the terminal cost constructed, such as in the 2018 scenario, the PG-MPC controller reduces to reference tracking of the policy trajectory.

The degree to which the policy limits performance depends directly on how well the policy is designed, providing a natural direction for improving the economic performance. The only practical requirements on the policy are that it produces feasible inputs and states, and that it is fast to evaluate. Beyond these, a wide range of alternatives could be substituted without any modification to the underlying controller structure. For instance, an RL or ADP-derived policy could directly address long-horizon cost reduction; a more sophisticated RBC could better account for building-specific complexities and adapt to particular climate years rather than relying on a one-size-fits-all trajectory; and a PI-controller could provide more contextually appropriate control actions by naturally scaling its response to the deviation from the reference. These are but a few examples, and the results presented here establish a baseline from which further policy improvements are expected to yield meaningful performance gains.

7-4-5 Scenario Robustness

The results across all four scenarios demonstrates that the PG-MPC controller is robust to the scenario variability in the sense that the thermal balance constraint and building temperature comfort constraints are satisfied in all cases, with the controller achieving competitive or superior operational cost. However, as discussed in Section 7-4-3 the 2018 scenario and TMY scenario reveal that the performance of the PG-MPC controller is sensitive to the interaction between electricity price and terminal ingredient design.

The PG-MPC controller performed best relative to the EMPC controller under conditions of above average cooling demand and high price volatility. The 2025 scenario exemplifies this, in which price volatility, particularly the mean daily spread, was highest. These conditions allow the PG-MPC controller to exploit price volatility that the EMPC, constrained to its TMY trajectory, cannot. The 2021 scenarios also shows moderate improvements relative to the EMPC baseline controller.

The PG-MPC performs comparable or slightly worse relative to the EMPC controller in the scenarios for which the TMY solution provides a reasonable approximation of the surrogate solution for that particular year. This is most clearly shown in the 2018 and 2021 scenarios. The EMPC controller also outperforms the PG-MPC controller in the TMY scenario, but this cannot be considered a fair comparison, given that the EMPC controller has an informational advantage in that it tracks an approximation of the optimal solution.

7-4-6 Limitations

Several limitations of this study should be acknowledged. The results presented in this chapter are based only on simulation. The performance of both controllers under real-operating conditions may differ due to increased model mismatch, sensor noise, and communication latency.

The ATES used in this study is also subject to simplifying assumptions. As discussed in Section 2-2-3, the ATES model assumes a constant volume ATES well model. While there is justification for such assumptions, modeling the ATES as variable volume allows for more accurately capturing the amount of energy stored in the ATES wells. Additionally, the constant volume assumption fixes the thermal storage capacity of each well as a function of the parameters V_h and V_c , whereas in a variable volume model, the effective storage capacity evolves in response to the history of heat injected and extracted. However, as the storage capacity constraints are never observed to be binding across any of the four scenarios considered in this study, this simplification is not expected to have a material impact on the presented results.

The building model described in Chapter 3 assumes a single-zone simplification, in which the entire building is represented by a single thermal mass and a single temperature state. In practice, a multi-zone building model would capture spatial temperature gradients and room-level dynamics that may influence the comfort constraint satisfaction results. However, for the purpose of evaluating EMS at the system level, the single-zone approximation is considered appropriate, as the primary interest lies in the building's aggregate thermal demand, rather than controlling the intra-building temperature distribution.

The policy used in the PG-MPC controller is based on a dead-band controller tracking a sinusoidal reference signal, representing an intentionally simple baseline choice. As discussed

in Sections 7-4-3 and 7-4-4, many of the observed limitations in controller performance can be attributed to the simple policy design rather than to fundamental limitations of the PG-MPC framework itself. More sophisticated policy designs, whether through more complex rule-based logic, RL methods, or fuzzy logic, could better capture the relationship between seasonal storage state, price volatility, and optimal control actions, potentially yielding substantial performance improvements at the cost of increased policy design complexity.

The Lipschitz constants are constructed based on an approximation of the policy value function. This approximation is constructed based on the electricity profile used for the TMY scenario. Demonstrably, this terminal cost was unsuited for the 2018 scenario, in which the low volatility and low mean of electricity prices resulted in PG-MPC effectively tracking the sinusoidal reference signal for the entire year, with little to no deviations. The TMY scenario suggests that during the summer months, the cost of deviating in the direction of positive imbalance was also not sufficient to prevent active rebalancing by the end of the year. However, the thermal balance constraint was satisfied by the end of the year.

All PG-MPC and EMPC simulations were carried out with a prediction horizon of four hours (16 steps), kept short to limit computational time over year-long simulations. Extending the horizon length is expected to improve performance for both controllers, as a longer horizon provides greater opportunity for cost savings through battery load shifting and more informed ATEs dispatch decisions. For the PG-MPC controller in particular, a longer horizon allows the controller to identify and exploit off-policy opportunities that are not visible over a four-hour window. However, the benefits of a longer horizon are not unconditional. A horizon that consistently surfaces short-term cost savings from off-policy actions may progressively drive the system toward imbalance states from which recovery is costly, particularly if the terminal cost penalty is insufficiently weighted relative to the stage cost savings accumulated over the extended horizon. The interaction between horizon length and terminal cost design therefore warrants careful consideration.

Finally, the four scenarios considered in this thesis, while spanning a wide range of climatological and electricity pricing conditions, do not constitute an exhaustive evaluation. None of the scenarios considered represent the combination of a particularly cold winter coupled with a mild summer. In terms of electricity pricing, a scenario considering a moderate or extreme summer, combined with high price volatility is also not considered.

Conclusion and Future Work

This thesis investigated the application of PG-MPC to the energy management of CHP microgrids equipped with hybrid electrical and seasonal thermal energy storage. The central difficulty in such systems is the separation of time scales: the seasonal charging and discharging of an ATES system imposes operational objectives that extend far beyond the prediction horizons that remain computationally tractable for MPC. Conventional EMPC addresses this through either heuristic terminal costs or surrogate-derived terminal constraint sets, both of which require careful tuning and design, with theoretical guarantees being difficult to provide. PG-MPC instead embeds a guiding policy within the short-horizon optimisation problem, allowing long-term objectives such as the end-of-year thermal balance to be encoded through this policy. A baseline EMPC controller and a PG-MPC controller were designed for an identical microgrid model and benchmarked across four scenarios spanning a wide range of climatological and electricity-pricing conditions, with both controllers operating under imperfect predictions of the disturbance signals. The remainder of this chapter answers the research questions posed in Chapter 1 and sets out recommendations for future work.

8-1 Answering the Research Questions

1. *How can Policy-Guided MPC be formulated for a CHP microgrid with hybrid electrical and thermal energy storage?*

The first subquestion was addressed in Chapters 5 and 6. The microgrid model described in Chapter 3 was modelled in the MLD framework so that the discrete actuation of the heat pump, boiler, and chiller, together with the PWA of the nonlinear heat pump, could be represented as linear constraints. On this shared model, the PG-MPC controller was obtained by replacing the EMPC terminal ingredients rather than by altering the underlying optimisation problem. A deliberately simple bang-bang dead-band guiding policy was designed to track a sinusoidal imbalance reference that reflects the seasonal charging and discharging pattern observed in the solution of a surrogate optimization problem. The terminal value function

was constructed as a Lipschitz-based linear over-approximation of the policy value function, estimated from Monte-Carlo rollouts, and the terminal constraint set was defined as a box around the policy-propagated terminal state. This construction inherits the theoretical properties of the PG-MPC framework: under the standard MPC assumptions and admissibility of the policy, the PG-MPC optimization problem is recursively feasible, and closed-loop cost is upper-bounded by the policy cost. The simulation results indicate that in all scenarios, the PG-MPC controller manages to achieve the primary objective of satisfying the yearly thermal balance constraint. The formulation therefore demonstrates that a seasonal thermal-balance objective can be encoded directly through a guiding policy, without imposing a precomputed year-long terminal trajectory on the controller.

2. How does the long-term operational performance of PG-MPC compare with a conventional EMPC controller?

The second subquestion was addressed through the case study of Chapter 7. Both controllers satisfied the end-of-year thermal balance constraint in every scenario, confirming that the seasonal objective is reliably enforced by both terminal-ingredient designs. In terms of operational cost, performance was strongly scenario-dependent. PG-MPC achieved substantially lower cost in the 2025 (Moderate) scenario (€7,917 against €12,035, a reduction of 34.2%) and a clear improvement in the 2021 (Mild) scenario in both operational cost (€14,869 against €16,164, a reduction of 8.7%) as well as constraint satisfaction. In the 2018 (Extreme) scenario the two controllers were effectively equivalent, with PG-MPC marginally higher by less than 0.4%. In the TMY scenario PG-MPC was higher by 3.5%; however, this comparison favours EMPC by construction, since its terminal constraint set is derived directly from the TMY surrogate solution and so has an a priori informational advantage that would not be available in practice. The largest gains arose under conditions of above-average cooling demand and high price volatility, where PG-MPC showed deeper charging and discharging of the thermal storage compared to EMPC, which was constrained to its fixed TMY trajectory. Across all scenarios, the cost savings are partly obtained by operating closer to the comfort-band boundary, and PG-MPC consequently incurred higher cumulative comfort-violation penalties. Overall, PG-MPC matched or outperformed the conventional EMPC controller in the realistic, non-idealised scenarios while never relying on a precomputed optimal trajectory.

3. What are the advantages and limitations of Policy-Guided MPC for practical microgrid energy management?

The principal advantage of PG-MPC is that it can encode the long-term seasonal objective through a guiding policy rather than a precomputed reference trajectory, freeing the controller to adapt to the realized weather and price conditions of a given year while retaining recursive feasibility and a performance guarantee relative to that policy. This adaptivity is the direct source of its cost advantage in the volatile scenarios. A second advantage is modularity: the only practical requirements on the policy are that it produces feasible states and inputs and is fast to evaluate, so the guiding policy can be replaced, for instance by an RL, ADP, rule-based, or PI policy, without any change to the controller structure.

The limitations fall into two categories. The first concerns the design of the terminal ingredients. The terminal cost is built from a Lipschitz constant estimated on the TMY price profile,

which substantially over-estimates the cost of acting off-policy under the low-volatility, low-mean prices of the 2018 scenario. In this case PG-MPC effectively collapsed to tracking the sinusoidal reference. Conversely, in the TMY scenario the penalty for deviating in the positive-imbalance direction was insufficient to prevent the controller from not cooling enough during the summer months, forcing costly end-of-year rebalancing through the auxiliary chiller. The interaction between price conditions and terminal-ingredient design is therefore the dominant sensitivity of the method. The second category concerns the guiding policy itself: the intentionally simple sinusoidal dead-band policy is sub-optimal for any specific year, and several observed shortcomings, most notably the mistiming of the heating-to-cooling season transitions, are artefacts of this policy rather than of the framework. Finally, the study is based entirely on simulation with a simplified single-zone building, a constant-volume ATES model, and a short four-hour prediction horizon, so the quantitative results should be read as a proof of concept rather than as predictions of real-world performance.

How can Policy-Guided Economic Model Predictive Control improve the energy management of CHP microgrids with seasonal thermal energy storage compared with conventional Economic Model Predictive Control?

The results show that PG-MPC can improve microgrid energy management by replacing a fixed, surrogate-derived terminal trajectory of conventional EMPC with a guiding policy. The PG-MPC framework allows for preserving guaranteed recursive feasibility, whereas the designed EMPC controller cannot due to the terminal constraint trajectory potentially being unreachable. Using a guiding policy along with properly designed terminal cost allows the controller to respond to the actual conditions of a given operating year, yielding comparable operational cost in the realistic scenarios and substantial savings when conditions deviate from the climatological average and when price volatility is high. The improvement is thus best understood as a gain in robustness and adaptivity rather than a uniform reduction in cost: PG-MPC removes the dependence on an accurate precomputed annual trajectory, and even with a deliberately naive guiding policy it remains competitive with the EMPC controller. The framework therefore offers a tractable and theoretically grounded route to embedding seasonal storage objectives in receding-horizon microgrid control, with the realized benefit governed by the quality of the guiding policy and the calibration of the terminal ingredients.

8-2 Recommendation for Future Work

Based on the finding in this thesis, several directions for future research naturally follow.

The most immediate is the design of the guiding policy. Because the framework is agnostic to the policy, replacing the sinusoidal dead-band controller with more capable alternatives, such as richer rule-based controllers, fuzzy-logic based controllers, or a learned RL policy. RL in particular provides an interesting direction as this provides a comprehensive set of tools to both construct a policy that accounts for long-horizon behavior, as well as the accompanying value function, providing an excellent surrogate for the cost-to-go. This is expected to yield meaningful performance gains, while reusing the existing controller.

A second direction would be more calibration of the PG-MPC terminal ingredients. The Lipschitz cost based on averaged electricity prices over the last five years provided both an

estimate which was too conservative in the 2018 scenario, and one which was not conservative enough in the TMY scenario. Constructing terminal cost based on a local linearization, or adapting the Lipschitz constant based on the electricity price regime observed during operation could reduce conservatism and better balance stage-cost savings against the terminal penalty. The coupling between prediction horizon length and terminal cost weighting also warrants further study, since longer horizons provide more opportunities to save operational costs in the short term, the interaction between the horizon length and terminal cost design should be carefully considered. A controller repeatedly surfacing short-term savings may lead to more off-policy actions, drifting further from the on-policy trajectory, and potentially recovering some of the myopic behavior that PG-MPC attempts to address.

A third direction concerns the system model. Lifting the simplifying assumptions of the present study: a variable-volume, variable mass flow rate ATES model, a multi-zone building representation, and longer prediction horizons would all narrow the gap between the simulated and the achievable real-world performance.

Finally, a broader and experimental evaluation is recommended. The four scenarios, while diverse, omit important combinations. Notably a cold winter with a mild summer, and an extreme summer coupled with high price volatility were not considered. Additionally, the entire study remains in simulation. Extending the scenario set and, in due course, validating PG-MPC on a hardware or hardware-in-the-loop setup under model mismatch, sensor noise, and communication latency would be the decisive test of its practical value.

Bibliography

- [1] ENTSO-E, *Power Statistics*, European Network of Transmission System Operators for Electricity, 2025. <https://www.entsoe.eu/data/power-stats/>.
- [2] Centraal Bureau voor de Statistiek (CBS), *Electricity balance sheet; supply and consumption*, 2026. <https://www.cbs.nl/en-gb/figures/detail/84575ENG>.
- [3] A. Ulbig, T. S. Borsche, and G. Andersson, “Impact of Low Rotational Inertia on Power System Stability and Operation,” *IFAC Proceedings Volumes*, 19th IFAC World Congress, vol. 47, no. 3, pp. 7290–7297, Jan. 1, 2014, ISSN: 1474-6670. DOI: 10.3182/20140824-6-ZA-1003.02615.
- [4] K. K. Mehmood, R. A. Arruda Moura, A. Van Der Molen, R. Tonkoski, P. Van Der Wielen, and P. H. Nguyen, “Efficient Congestion Management: Optimizing Non-Firm Capacities for Maximized Grid Utilization,” in *2024 IEEE International Conference on Environment and Electrical Engineering and 2024 IEEE Industrial and Commercial Power Systems Europe (EEEIC / I&CPS Europe)*, Jun. 2024, pp. 1–6. DOI: 10.1109/EEEIC/ICPSEurope61470.2024.10751051.
- [5] Association of Energy Network Operators in the Netherlands. “Capaciteitskaart elektriciteitsnet | Totaal - Afname.” (2026), <https://data.partnersinenergie.nl/capaciteitskaart/totaal/afname>.
- [6] R. Lasseter, “MicroGrids,” in *2002 IEEE Power Engineering Society Winter Meeting. Conference Proceedings (Cat. No.02CH37309)*, vol. 1, Jan. 2002, 305–308 vol.1. DOI: 10.1109/PESW.2002.985003.
- [7] M. Geidl, G. Koepfel, P. Favre-Perrod, B. Klockl, G. Andersson, and K. Frohlich, “Energy hubs for the future,” *IEEE Power and Energy Magazine*, vol. 5, no. 1, pp. 24–30, Jan. 2007, ISSN: 1558-4216. DOI: 10.1109/MPAE.2007.264850.
- [8] R. Lasseter, A. Akhil, C. Marnay, et al. “Integration of distributed energy resources. The CERTS Microgrid Concept,” Apr. 1, 2002. <https://escholarship.org/uc/item/9w88z7z1>.

- [9] J. L. Martínez-Ramos, A. Marano-Marcolini, F. P. García-López, et al. "Provision of Ancillary Services by a Smart Microgrid: An OPF Approach," in *2018 International Conference on Smart Energy Systems and Technologies (SEST)*, Sep. 2018, pp. 1–6. DOI: 10.1109/SEST.2018.8495883.
- [10] M. Uddin, H. Mo, D. Dong, S. Elsayah, J. Zhu, and J. M. Guerrero, "Microgrids: A review, outstanding issues and future trends," *Energy Strategy Reviews*, vol. 49, p. 101 127, Sep. 1, 2023, ISSN: 2211-467X. DOI: 10.1016/j.esr.2023.101127.
- [11] Y.-G. Son, B.-C. Oh, M. A. Acquah, R. Fan, D.-M. Kim, and S.-Y. Kim, "Multi Energy System With an Associated Energy Hub: A Review," *IEEE Access*, vol. 9, pp. 127 753–127 766, 2021, ISSN: 2169-3536. DOI: 10.1109/ACCESS.2021.3108142.
- [12] X. Lin and R. Zamora, "Controls of hybrid energy storage systems in microgrids: Critical review, case study and future trends," *Journal of Energy Storage*, vol. 47, p. 103 884, Mar. 1, 2022, ISSN: 2352-152X. DOI: 10.1016/j.est.2021.103884.
- [13] J. Hu, Y. Shan, Y. Yang, et al. "Economic Model Predictive Control for Microgrid Optimization: A Review," *IEEE Transactions on Smart Grid*, vol. 15, no. 1, pp. 472–484, Jan. 2024, ISSN: 1949-3061. DOI: 10.1109/TSG.2023.3266253.
- [14] D. E. Olivares, A. Mehrizi-Sani, A. H. Etemadi, et al. "Trends in Microgrid Control," *IEEE Transactions on Smart Grid*, vol. 5, no. 4, pp. 1905–1919, Jul. 2014, ISSN: 1949-3061. DOI: 10.1109/TSG.2013.2295514.
- [15] A. Joshi, S. Capezza, A. Alhaji, and M.-Y. Chow, "Survey on AI and Machine Learning Techniques for Microgrid Energy Management Systems," *IEEE/CAA Journal of Automatica Sinica*, vol. 10, no. 7, pp. 1513–1529, Jul. 2023, ISSN: 2329-9274. DOI: 10.1109/JAS.2023.123657.
- [16] A. Parisio, E. Rikos, and L. Glielmo, "A Model Predictive Control Approach to Microgrid Operation Optimization," *IEEE Transactions on Control Systems Technology*, vol. 22, no. 5, pp. 1813–1827, Sep. 2014, ISSN: 1558-0865. DOI: 10.1109/TCST.2013.2295737.
- [17] F. Garcia-Torres and C. Bordons, "Optimal Economical Schedule of Hydrogen-Based Microgrids With Hybrid Storage Using Model Predictive Control," *IEEE Transactions on Industrial Electronics*, vol. 62, no. 8, pp. 5195–5207, Aug. 2015, ISSN: 1557-9948. DOI: 10.1109/TIE.2015.2412524.
- [18] J. Patiño, A. Márquez, and J. Espinosa, "An economic MPC approach for a microgrid energy management system," in *2014 IEEE PES Transmission & Distribution Conference and Exposition - Latin America (PES T&D-LA)*, Sep. 2014, pp. 1–6. DOI: 10.1109/TDC-LA.2014.6955203.
- [19] R. Halvgaard, N. K. Poulsen, H. Madsen, and J. B. Jørgensen, "Economic Model Predictive Control for building climate control in a Smart Grid," in *2012 IEEE PES Innovative Smart Grid Technologies (ISGT)*, Jan. 2012, pp. 1–6. DOI: 10.1109/ISGT.2012.6175631.
- [20] C. Bordons, F. Garcia-Torres, and M. A. Ridao, *Model Predictive Control of Microgrids* (Advances in Industrial Control). Cham: Springer International Publishing, 2020, ISBN: 978-3-030-24569-6 978-3-030-24570-2. DOI: 10.1007/978-3-030-24570-2.

- [21] H. A. Schlüter, D. Boiroux, N. K. Poulsen, H. Madsen, and J. B. Jørgensen, "Economic Model Predictive Control for Energy Systems in Smart Homes," in *2019 IEEE Conference on Control Technology and Applications (CCTA)*, Aug. 2019, pp. 598–604. DOI: 10.1109/CCTA.2019.8920663.
- [22] V. Rostampour and T. Keviczky, "Energy Management for Building Climate Comfort in Uncertain Smart Thermal Grids with Aquifer Thermal Energy Storage*," *IFAC-PapersOnLine*, 20th IFAC World Congress, vol. 50, no. 1, pp. 13 156–13 163, Jul. 1, 2017, ISSN: 2405-8963. DOI: 10.1016/j.ifacol.2017.08.2170.
- [23] S. Msaad and R. D. McAllister, "Economic Model Predictive Control with Policy-Guided Terminal Ingredients," Unpublished Manuscript, Jul. 2026.
- [24] S. Msaad, M. Harraway, and R. D. McAllister, "RL-Guided MPC for Autonomous Greenhouse Control," *IFAC-PapersOnLine*, 8th IFAC Conference on Sensing, Control and Automation Technologies for Agriculture AGRICONTROL 2025, vol. 59, no. 23, pp. 449–454, Jan. 1, 2025, ISSN: 2405-8963. DOI: 10.1016/j.ifacol.2025.11.829.
- [25] M. Faisal, M. A. Hannan, P. J. Ker, A. Hussain, M. B. Mansor, and F. Blaabjerg, "Review of Energy Storage System Technologies in Microgrid Applications: Issues and Challenges," *IEEE Access*, vol. 6, pp. 35 143–35 164, 2018, ISSN: 2169-3536. DOI: 10.1109/ACCESS.2018.2841407.
- [26] P. T. Moseley and J. Garche, *Electrochemical Energy Storage for Renewable Sources and Grid Balancing*. Amsterdam: Elsevier, 2015, ISBN: 978-0-444-62616-5.
- [27] A. Zimmerman, "Self-discharge losses in lithium-ion cells," *IEEE Aerospace and Electronic Systems Magazine*, vol. 19, no. 2, pp. 19–24, Feb. 2004, ISSN: 1557-959X. DOI: 10.1109/MAES.2004.1269687.
- [28] E. M. Krieger and C. B. Arnold, "Effects of undercharge and internal loss on the rate dependence of battery charge storage efficiency," *Journal of Power Sources*, vol. 210, pp. 286–291, Jul. 15, 2012, ISSN: 0378-7753. DOI: 10.1016/j.jpowsour.2012.03.029.
- [29] P. Arora, R. E. White, and M. Doyle, "Capacity Fade Mechanisms and Side Reactions in Lithium-Ion Batteries," *Journal of The Electrochemical Society*, vol. 145, no. 10, p. 3647, Oct. 1, 1998, ISSN: 1945-7111. DOI: 10.1149/1.1838857.
- [30] R. L. Hartmann, "An Aging Model for Lithium-Ion Cells," University of Akron, 2008. https://etd.ohiolink.edu/acprod/odb_etd/etd/r/1501/10?clear=10&p10_accession_num=akron1226887071.
- [31] J. Tobajas, F. Garcia-Torres, P. Roncero-Sánchez, J. Vázquez, L. Bellatreche, and E. Nieto, "Resilience-oriented schedule of microgrids with hybrid energy storage system using model predictive control," *Applied Energy*, vol. 306, p. 118 092, Jan. 15, 2022, ISSN: 0306-2619. DOI: 10.1016/j.apenergy.2021.118092.
- [32] Y. Parag and B. K. Sovacool, "Electricity market design for the prosumer era," *Nature Energy*, vol. 1, no. 4, p. 16 032, Mar. 21, 2016, ISSN: 2058-7546. DOI: 10.1038/nenergy.2016.32.
- [33] Y. A. Çengel, M. A. Boles, and M. Kanoğlu, *Thermodynamics: An Engineering Approach*, Ninth edition. New York: McGraw Hill Education, 2019, 982 pp., ISBN: 978-1-260-09268-4.

- [34] P. Fleuchaus, B. Godschalk, I. Stober, and P. Blum, "Worldwide application of aquifer thermal energy storage – A review," *Renewable and Sustainable Energy Reviews*, vol. 94, pp. 861–876, Oct. 1, 2018, ISSN: 1364-0321. DOI: 10.1016/j.rser.2018.06.057.
- [35] T. Yang, W. Liu, G. J. Kramer, and Q. Sun, "Seasonal thermal energy storage: A techno-economic literature review," *Renewable and Sustainable Energy Reviews*, vol. 139, p. 110 732, Apr. 1, 2021, ISSN: 1364-0321. DOI: 10.1016/j.rser.2021.110732.
- [36] M. Bloemendal and N. Hartog, "Analysis of the impact of storage conditions on the thermal recovery efficiency of low-temperature ATES systems," *Geothermics*, vol. 71, pp. 306–319, Jan. 1, 2018, ISSN: 0375-6505. DOI: 10.1016/j.geothermics.2017.10.009.
- [37] M. Bloemendal, M. Jaxa-Rozen, and T. Olsthoorn, "Methods for planning of ATES systems," *Applied Energy*, vol. 216, pp. 534–557, Apr. 15, 2018, ISSN: 0306-2619. DOI: 10.1016/j.apenergy.2018.02.068.
- [38] R. J. Caljé, "Future use of aquifer thermal energy storage below the historic centre of amsterdam," M.S. thesis, TU Delft, Department of Water Management, Delft, 2010.
- [39] M. Bloemendal, T. Olsthoorn, and F. Boons, "How to achieve optimal and sustainable use of the subsurface for Aquifer Thermal Energy Storage," *Energy Policy*, vol. 66, pp. 104–114, Mar. 1, 2014, ISSN: 0301-4215. DOI: 10.1016/j.enpol.2013.11.034.
- [40] P. Fleuchaus, S. Schüppler, B. Godschalk, G. Bakema, and P. Blum, "Performance analysis of Aquifer Thermal Energy Storage (ATES)," *Renewable Energy*, vol. 146, pp. 1536–1548, Feb. 1, 2020, ISSN: 0960-1481. DOI: 10.1016/j.renene.2019.07.030.
- [41] J. de Wind, J. Alpízar-Castillo, J. Visser, and L. Ramírez-Elizondo, "An in-depth analysis of residential E-cooling demand in The Netherlands," *Case Studies in Thermal Engineering*, vol. 73, p. 106 469, Sep. 1, 2025, ISSN: 2214-157X. DOI: 10.1016/j.csite.2025.106469.
- [42] A. F. Mills and C. F. M. Coimbra, *Basic Heat and Mass Transfer*, Third edition. San Diego, California: Prentice Hall, 2015, ISBN: 978-0-9963053-0-3.
- [43] J. Hu, Y. Shan, J. M. Guerrero, A. Ioinovici, K. W. Chan, and J. Rodriguez, "Model predictive control of microgrids - An overview," *Renewable and Sustainable Energy Reviews*, vol. 136, p. 110 422, Feb. 1, 2021, ISSN: 1364-0321. DOI: 10.1016/j.rser.2020.110422.
- [44] R. Bhadani, S. Banik, H. Tu, S. Lukic, and G. Karsai. "On Simulation of Power Systems and Microgrid Components with SystemC-AMS." arXiv: 2407.06217 [eess.SY]. (Jul. 4, 2024), pre-published.
- [45] J. Vasilj, S. Gros, D. Jakus, and M. Zanon, "Day-Ahead Scheduling and Real-Time Economic MPC of CHP Unit in Microgrid With Smart Buildings," *IEEE Transactions on Smart Grid*, vol. 10, no. 2, pp. 1992–2001, Mar. 2019, ISSN: 1949-3061. DOI: 10.1109/TSG.2017.2785500.
- [46] A. J. Wood, B. F. Wollenberg, and G. B. Sheblé, *Power Generation, Operation, and Control*, Third edition. Hoboken, New Jersey: Wiley-Interscience, 2014, 632 pp., ISBN: 978-0-471-79055-6.

- [47] R. Sioshansi, R. O'Neill, and S. S. Oren, "Economic Consequences of Alternative Solution Methods for Centralized Unit Commitment in Day-Ahead Electricity Markets," *IEEE Transactions on Power Systems*, vol. 23, no. 2, pp. 344–352, May 2008, ISSN: 1558-0679. DOI: 10.1109/TPWRS.2008.919246.
- [48] A. Bischi, L. Taccari, E. Martelli, et al. "A detailed MILP optimization model for combined cooling, heat and power system operation planning," *Energy*, vol. 74, pp. 12–26, Sep. 1, 2014, ISSN: 0360-5442. DOI: 10.1016/j.energy.2014.02.042.
- [49] H. Shuai, J. Fang, X. Ai, J. Wen, and H. He, "Optimal Real-Time Operation Strategy for Microgrid: An ADP-Based Stochastic Nonlinear Optimization Approach," *IEEE Transactions on Sustainable Energy*, vol. 10, no. 2, pp. 931–942, Apr. 2019, ISSN: 1949-3037. DOI: 10.1109/TSTE.2018.2855039.
- [50] C. Hu, Z. Cai, Y. Zhang, R. Yan, Y. Cai, and B. Cen, "A soft actor-critic deep reinforcement learning method for multi-timescale coordinated operation of microgrids," *Protection and Control of Modern Power Systems*, vol. 7, no. 1, p. 29, Aug. 5, 2022, ISSN: 2367-0983. DOI: 10.1186/s41601-022-00252-z.
- [51] A. Parisio, E. Rikos, G. Tzamalīs, and L. Glielmo, "Use of model predictive control for experimental microgrid optimization," *Applied Energy*, vol. 115, pp. 37–46, Feb. 15, 2014, ISSN: 0306-2619. DOI: 10.1016/j.apenergy.2013.10.027.
- [52] *EN ISO 13786:2017 — Thermal performance of building components — Dynamic thermal characteristics — Calculation methods*, Geneva, 2017. <https://www.iso.org/standard/65711.html>, Published.
- [53] Royal Netherlands Meteorological Institute (KNMI), *Hourly, validated and automated in-situ ground-based meteorological observations in the Netherlands*, version 1.0, KNMI Data Platform, May 15, 2026. <https://datapatform.knmi.nl/dataset/hourly-in-situ-meteorological-observations-validated-1-0>.
- [54] I. Staffell, D. Brett, N. Brandon, and A. Hawkes, "A review of domestic heat pumps," *Energy & Environmental Science*, vol. 5, pp. 9291–9306, Oct. 2012. DOI: 10.1039/C2EE22653G.
- [55] D. Q. Mayne, J. B. Rawlings, C. V. Rao, and P. O. M. Scokaert, "Constrained model predictive control: Stability and optimality," *Automatica*, vol. 36, no. 6, pp. 789–814, Jun. 1, 2000, ISSN: 0005-1098. DOI: 10.1016/S0005-1098(99)00214-9.
- [56] K. J. Åström and R. M. Murray, *Feedback Systems: An Introduction for Scientists and Engineers*, Second edition. Princeton Oxford: Princeton University Press, 2021, 505 pp., ISBN: 978-0-691-19398-4.
- [57] J. B. Rawlings, D. Q. Mayne, and M. Diehl, *Model Predictive Control: Theory, Computation, and Design*, 2nd edition. Santa Barbara: Nob Hill Publishing, LLC, 2024, 623 pp., ISBN: 978-0-9759377-8-5.
- [58] R. D. McAllister, "Inherently Robust Economic Model Predictive Control Without Dissipativity," *IFAC-PapersOnLine*, 8th IFAC Conference on Nonlinear Model Predictive Control NMPC 2024, vol. 58, no. 18, pp. 188–193, Jan. 1, 2024, ISSN: 2405-8963. DOI: 10.1016/j.ifacol.2024.09.029.

- [59] D. Angeli, R. Amrit, and J. B. Rawlings, "On Average Performance and Stability of Economic Model Predictive Control," *IEEE Transactions on Automatic Control*, vol. 57, no. 7, pp. 1615–1626, Jul. 2012, ISSN: 1558-2523. DOI: 10.1109/TAC.2011.2179349.
- [60] M. Diehl, R. Amrit, and J. B. Rawlings, "A Lyapunov Function for Economic Optimizing Model Predictive Control," *IEEE Transactions on Automatic Control*, vol. 56, no. 3, pp. 703–707, Mar. 2011, ISSN: 1558-2523. DOI: 10.1109/TAC.2010.2101291.
- [61] R. Amrit, J. B. Rawlings, and D. Angeli, "Economic optimization using model predictive control with a terminal cost," *Annual Reviews in Control*, vol. 35, no. 2, pp. 178–186, Dec. 1, 2011, ISSN: 1367-5788. DOI: 10.1016/j.arcontrol.2011.10.011.
- [62] D. Mignone, G. Ferrari-Trecate, and M. Morari, "Stability and stabilization of piecewise affine and hybrid systems: An LMI approach," in *Proceedings of the 39th IEEE Conference on Decision and Control (Cat. No.00CH37187)*, vol. 1, Dec. 2000, 504–509 vol.1. DOI: 10.1109/CDC.2000.912814.
- [63] M. J. Risbeck, "Mixed-integer model predictive control with applications to building energy systems," The University of Wisconsin-Madison, 2018.
- [64] M. Branicky, V. Borkar, and S. Mitter, "A unified framework for hybrid control: Model and optimal control theory," *IEEE Transactions on Automatic Control*, vol. 43, no. 1, pp. 31–45, Jan. 1998, ISSN: 1558-2523. DOI: 10.1109/9.654885.
- [65] R. Goebel, R. G. Sanfelice, and A. R. Teel, "Hybrid dynamical systems," *IEEE Control Systems Magazine*, vol. 29, no. 2, pp. 28–93, Apr. 2009, ISSN: 1941-000X. DOI: 10.1109/MCS.2008.931718.
- [66] J. B. Rawlings and M. J. Risbeck, "Model predictive control with discrete actuators: Theory and application," *Automatica*, vol. 78, pp. 258–265, Apr. 1, 2017, ISSN: 0005-1098. DOI: 10.1016/j.automatica.2016.12.024.
- [67] A. Bemporad and M. Morari, "Control of systems integrating logic, dynamics, and constraints," *Automatica*, vol. 35, no. 3, pp. 407–427, Mar. 1, 1999, ISSN: 0005-1098. DOI: 10.1016/S0005-1098(98)00178-2.
- [68] E. F. Camacho, D. R. Ramirez, D. Limon, D. Muñoz de la Peña, and T. Alamo, "Model predictive control techniques for hybrid systems," *Annual Reviews in Control*, vol. 34, no. 1, pp. 21–31, Apr. 1, 2010, ISSN: 1367-5788. DOI: 10.1016/j.arcontrol.2010.02.002.
- [69] P. Belotti, C. Kirches, S. Leyffer, J. Linderoth, J. Luedtke, and A. Mahajan, "Mixed-integer nonlinear optimization," *Acta Numerica*, vol. 22, pp. 1–131, May 2013, ISSN: 0962-4929, 1474-0508. DOI: 10.1017/S0962492913000032.
- [70] H. P. Williams, *Model Building in Mathematical Programming*, 5. ed. Chichester, West Sussex: Wiley, 2013, 411 pp., ISBN: 978-1-118-44333-0.
- [71] W. P. M. H. Heemels, B. De Schutter, and A. Bemporad, "Equivalence of hybrid dynamical models," *Automatica*, vol. 37, no. 7, pp. 1085–1091, Jul. 1, 2001, ISSN: 0005-1098. DOI: 10.1016/S0005-1098(01)00059-0.
- [72] D. A. Babayev, "Piece-wise linear approximation of functions of two variables," *Journal of Heuristics*, vol. 2, no. 4, pp. 313–320, Mar. 1, 1997, ISSN: 1572-9397. DOI: 10.1007/BF00132502.

- [73] H. Li, A. T. Eseye, J. Zhang, and D. Zheng, “Optimal Energy Management for Industrial Microgrids with High-Penetration Renewables,” *Protection and Control of Modern Power Systems*, vol. 2, no. 2, pp. 1–14, Apr. 2017, ISSN: 2367-0983. DOI: 10.1186/s41601-017-0040-6.
- [74] E. Prat, P. Pinson, R. M. Lusby, R. Plougonven, J. Badosa, and P. Drobinski, “Optimal Operation of a Building with Electricity-Heat Networks and Seasonal Storage,” in *2024 IEEE PES Innovative Smart Grid Technologies Europe (ISGT EUROPE)*, Oct. 2024, pp. 1–5. DOI: 10.1109/ISGTEUROPE62998.2024.10863412.
- [75] M. De Rosa, V. Bianco, F. Scarpa, and L. A. Tagliafico, “Heating and cooling building energy demand evaluation; a simplified model and a modified degree days approach,” *Applied Energy*, vol. 128, pp. 217–229, Sep. 1, 2014, ISSN: 0306-2619. DOI: 10.1016/j.apenergy.2014.04.067.
- [76] *ISO 52016-1:2017 Energy performance of buildings — Energy needs for heating and cooling, internal temperatures and sensible and latent heat loads*, Geneva, 2017. <https://www.iso.org/standard/65696.html>, Published.
- [77] D. Crawley, L. Lawrie, C. Pedersen, et al. “ENERGYPLUS, a new-generation building energy simulation program,” ser. Proceedings of Building Simulation ’99, vol. 1, Sep. 1999.
- [78] L. K. Lawrie and D. B. Crawley, *Development of global typical meteorological years (TMYx)*, 2026. <https://climate.onebuilding.org>.
- [79] W. C. Clarke, C. Manzie, and M. J. Brear, “Hierarchical economic MPC for systems with storage states,” *Automatica*, vol. 94, pp. 138–150, Aug. 1, 2018, ISSN: 0005-1098. DOI: 10.1016/j.automatica.2018.04.012.
- [80] D. P. e Silva, J. L. F. Salles, J. F. Fardin, et al. “Measured and forecasted weather and power dataset for management of an island and grid-connected microgrid,” *Data in Brief*, vol. 39, p. 107513, Dec. 1, 2021, ISSN: 2352-3409. DOI: 10.1016/j.dib.2021.107513.
- [81] next-kraftwerke. “What is day-ahead trading?” (), <https://www.next-kraftwerke.be/en/knowledge-hub/day-ahead-trading/>.
- [82] S. Seal, B. Boulet, and V. R. Dehkordi, “Centralized model predictive control strategy for thermal comfort and residential energy management,” *Energy*, vol. 212, p. 118456, Dec. 1, 2020, ISSN: 0360-5442. DOI: 10.1016/j.energy.2020.118456.
- [83] J. A. Candanedo, J.-M. Hardy, É. Saloux, R. Platon, V. Raissi-Dehkordi, and A. Côté, “Preliminary assessment of a weather forecast tool for building operation,” 2018.
- [84] L. Yang, X. Gao, J. Hua, and L. Wang, “Intra-day global horizontal irradiance forecast using FY-4A clear sky index,” *Sustainable Energy Technologies and Assessments*, vol. 50, p. 101816, Mar. 1, 2022, ISSN: 2213-1388. DOI: 10.1016/j.seta.2021.101816.
- [85] W. B. Powell, *Approximate Dynamic Programming: Solving the Curses of Dimensionality*. John Wiley & Sons, 2007, vol. 703.
- [86] D. P. Kingma and J. Ba. “Adam: A Method for Stochastic Optimization.” arXiv: 1412.6980 [cs.LG]. (Jan. 30, 2017), pre-published.

- [87] A. G. Baydin, B. A. Pearlmutter, A. A. Radul, and J. M. Siskind, “Automatic differentiation in machine learning: A survey,” *Journal of Machine Learning Research*, vol. 18, no. 153, pp. 1–43, 2018. <http://jmlr.org/papers/v18/17-468.html>.
- [88] B. Grimstad and H. Andersson, “ReLU networks as surrogate models in mixed-integer linear programs,” *Computers & Chemical Engineering*, vol. 131, p. 106580, Dec. 5, 2019, ISSN: 0098-1354. DOI: 10.1016/j.compchemeng.2019.106580.
- [89] R. Anderson, J. Huchette, W. Ma, C. Tjandraatmadja, and J. P. Vielma, “Strong mixed-integer programming formulations for trained neural networks,” *Mathematical Programming*, vol. 183, no. 1, pp. 3–39, Sep. 1, 2020, ISSN: 1436-4646. DOI: 10.1007/s10107-020-01474-5.
- [90] M. Fischetti and J. Jo, “Deep neural networks and mixed integer linear optimization,” *Constraints*, vol. 23, no. 3, pp. 296–309, Jul. 1, 2018, ISSN: 1572-9354. DOI: 10.1007/s10601-018-9285-6.
- [91] *ISO 15927-1:2003 — Hygrothermal performance of buildings — Calculation and presentation of climatic data*, Geneva, Nov. 2003. <https://www.iso.org/standard/29559.html>, Published.
- [92] K. van der Wiel, J. Beersma, H. van den Brink, et al. “KNMI’23 Climate Scenarios for the Netherlands: Storyline Scenarios of Regional Climate Change,” *Earth’s Future*, vol. 12, no. 2, e2023EF003983, 2024, ISSN: 2328-4277. DOI: 10.1029/2023EF003983.
- [93] Royal Netherlands Meteorological Institute (KNMI), “KNMI’23 climate scenarios for the netherlands,” Royal Netherlands Meteorological Institute (KNMI), De Bilt, Netherlands, KNMI-Publicatie 23-03, 2023. https://cdn.knmi.nl/system/ckeditor/attachment_files/data/000/000/358/original/KNMI23_climate_scenarios_user_report.pdf.
- [94] B. Zakeri, I. Staffell, P. E. Dodds, et al. “The role of natural gas in setting electricity prices in Europe,” *Energy Reports*, vol. 10, pp. 2778–2792, Nov. 2023, ISSN: 2352-4847. DOI: 10.1016/j.egyr.2023.09.069.
- [95] N. Rafizadeh, “The economics of negative price phenomenon in renewable-integrated electricity markets,” *Energy Economics*, vol. 153, p. 109086, Jan. 2026, ISSN: 01409883. DOI: 10.1016/j.eneco.2025.109086.

Glossary

List of Acronyms

ADP	approximate dynamic programming
ANN	artificial neural network
API	Application Programming Interface
ATES	Aquifer Thermal Energy Storage
BES	Battery Energy Storage
BRP	balancing responsible party
CDD	Cooling Degree Days
CHP	combined heat and power
COP	Coefficient of Performance
DP	Dynamic Programming
EDR	Environmental Data Retrieval
ED	economic dispatch
EMPC	economic model predictive control
EMS	energy management system
ENTSO-E	European Network of Transmission System Operators for Electricity
EPW	EnergyPlus Weather
GHI	Global Horizontal Irradiance
HDD	Heating Degree Days
KNMI	Royal Netherlands Meteorological Institute
LP	Linear Program
LQR	Linear Quadratic Regulator
MAE	Mean Absolute Error
MBE	Mean Bias Error

MILP	mixed-integer linear program
MINLP	mixed-integer non-linear program
MLD-MPC	mixed logical dynamical model predictive control
MLD	mixed logical dynamics
MPC	model predictive control
OCP	optimal control problem
PG-MPC	Policy-Guided MPC
PV	photovoltaic
PWA	piece-wise affine
RBC	Rule Based Control
RES	renewable energy sources
RL	reinforcement learning
RMSE	Root Mean Squared Error
TES	Thermal Energy Storage
TMY	Typical Meteorological Year
UC	unit commitment
nRMSE	Normalized Root Mean Squared Error
CV	Coefficient of Variation

Nomenclature

CV Coefficient of variation of the electricity price

