UPTEC STS 26017
Examensarbete 30 hp
Juni 2026

Intelligent Heating Control

Digital Twin based Economic Model Predictive Control for Heating in a Large Office Building using a Reversible Heat Pump

Max Bergström & Matviy Khotyaintsev

Civilingejörsprogrammet i System i Teknik och Samhälle

UPPSALA
UNIVERSITET

Intelligent Heating Control

Matviy Khotyaintsev & Max Bergström

Abstract

As wholesale electricity prices on the Nord Pool day-ahead market become increasingly volatile, shifting a commercial building's energy consumption to cheaper hours offers significant economic benefits. However, conventional reactive controllers, such as static weather-compensation curves, lack the predictive foresight required to utilize a building's heavy structural mass as a thermal battery for load shifting. This thesis investigates the potential of implementing a Model Predictive Control (MPC) strategy for a 40 000 m² commercial office building in Stockholm. The facility features a heavy concrete frame and is equipped with a reversible heat pump plant connected to a borehole thermal energy storage field. To accurately represent the building's thermal dynamics, four grey-box Resistance-Capacitance (RC) model structures were evaluated. Both a single-state lumped model and a three-state model demonstrated high predictive stability, maintaining a 24-hour prediction error below 0.4°C.

The predictive controller minimizes a quadratic cost function based on hourly electricity spot prices, utilizing soft constraints to keep indoor temperatures comfortably between 21°C and 23°C. The proposed MPC and the legacy controller were evaluated over the 2025 calendar year using identical inputs within a thermodynamic sandbox model, ensuring that differences in energy performance reflect true algorithmic advantages.

By proactively pre-heating the structural mass during cheap nighttime hours and curbing compressor usage during expensive morning price peaks, the MPC effectively decoupled comfort delivery from peak energy costs. Under conservative operational adjustments, the MPC achieved a 39% reduction in compressor energy and a 62% decrease in overall operating costs compared to the legacy system. These substantial savings were realized without deteriorating the indoor thermal climate, demonstrating the potential in MPC as a highly viable framework for demand-side management in commercial buildings.

Teknisk-naturvetenskapliga fakulteten

Uppsala universitet, Utgivningsort Uppsala

Handledare: Carl Asp, Emma Ulfsparre Ämnesgranskare: Steffi Knorn

Examinator: Elísabet Andrésdóttir

Populärvetenskaplig sammanfattning

Varje timme på dygnet har elpriset ett nytt värde. Mitt i natten när efterfrågan är låg kan priset sjunka till en bråkdel av vad det kostar under morgontimmarna när kontor, fabriker och tunnelbanor startar igång. För ett vanligt hushåll spelar den här variationen kanske inte så stor roll, men för en stor kontorsbyggnad kan skillnaden mellan att värma byggnaden på rätt timme jämfört med fel timme innebära tiotusentals kronor per år.

Det är precis det problemet som det här examensarbetet handlar om. Vi har studerat en stor kontorsbyggnad i Stockholm med en uppvärmd yta på ungefär 40 000 kvadratmeter. Byggnaden värms och kyls av ett modernt värmepumpsystem som är kopplat till ett borrhålslager i berget under fastigheten, där överskottsvärme och kyla kan lagras och hämtas tillbaka när det behövs. Det är ett tekniskt avancerat system med stor potential att spara pengar, förutsatt att det styrs på ett smart sätt.

Problemet är att styrningen inte alltid är så smart. Den styrenhet som idag kontrollerar systemet följer en fast kurva som enbart tittar på utomhustemperaturen och bestämmer hur varmt vattnet i värmesystemet ska vara. Det spelar ingen roll vad elen kostar just nu, och det spelar ingen roll om byggnaden redan är uppvärmd sedan natten. Systemet reagerar bara på nuet och har ingen förmåga att planera framåt. Konsekvensen är att kompressorerna ofta kör för fullt under morgontimmarna, precis när elpriset är som högst.

Betong lagrar värme. Det är en enkel fysikalisk egenskap som visar sig vara ekonomiskt värdefull. En tung betongstomme kan hålla kvar värme under många timmar, vilket betyder att man kan värma upp byggnaden under natten när elen är billig och sedan låta den lagrade värmen räcka när priset stiger på morgonen. Byggnaden fungerar i princip som ett stort värmebatteri. Men för att utnyttja det här batteriet behöver styrningen veta vad elen kommer att kosta imorgon och hur temperaturen i byggnaden förväntas utvecklas under de kommande timmarna. Det kräver en annan typ av styrning än en enkel kurva.

Den typ av styrning vi har undersökt kallas "Model Predictive Control", förkortat MPC. Principen är att styrningen varje timme löser ett matematiskt optimeringsproblem: givet hur byggnaden ser ut termiskt just nu, hur vädret väntas bli och vad elen förväntas kosta de närmaste 24 timmarna, vad är den billigaste sekvensen av åtgärder som ändå håller inomhustemperaturen inom ett komfortabelt intervall? Svaret på det problemet avgör vilket vattentemperaturvärde som skickas till värmesystemet. Sedan löses problemet om på nytt nästa timme med uppdaterad information.

För att en sådan styrning ska fungera behövs en matematisk modell av hur byggnaden beter sig termiskt. Vi har identifierat en sådan modell från två år av mätdata från byggnadens styrsystem. Modellen är konstruerad så att den fångar byggnadens fysikaliska egenskaper utan att kräva detaljerade kunskaper om varje betongskiva och fönsters exakta värmeegenskaper. Det visade sig att redan den enklaste modellstrukturen vi testade klarade av att förutsäga inomhustemperaturen med en noggrannhet på under 0,4 grader Celsius för ett helt år av osedd testdata, vilket är tillräckligt bra för att styrningen ska kunna fatta välgrundade beslut.

En viktig del av projektet handlar om hur man jämför den nya styrningen mot den gamla på ett rättvist sätt. Det vanligt tillvägagångssätt i forskningen är att jämföra den simulerade MPC-styrningen mot verklig historisk energiförbrukning. Det låter rimligt men är i praktiken missvisande. En styrning som optimerar mot en matematisk modell av byggnaden kan komma att utnyttja de ställen där modellen avviker från verkligheten, och resultatet ser ut som besparing även när det inte är det. Vi har istället kört både den nya och den gamla styrningen i exakt samma simuleringsmiljö, med exakt samma förutsättningar. På så vis beror eventuella skillnader i utfall enbart på hur styrlogiken skiljer sig, inte på att en kontroller kör i en modell och den andra i verkligheten.

Resultaten visar att den prediktiva styrningen minskar kompressorernas driftkostnad med ungefär 62 procent jämfört med den befintliga styrningen, även efter att vi gjort konservativa antaganden om de månader där simuleringen är som mest osäker. Energiförbrukningen minskar med ungefär 39 procent. Det kanske mest intressanta resultatet är att kostnadsminskningen är betydligt större än energiminskningen. Det beror på att den nya styrningen inte i första hand sparar energi utan förflyttar förbrukningen i tid, från dyra timmar till billiga. En timmes el som förbrukas klockan tre på natten kostar en bråkdel av vad samma timme kostar vid åttatiden på morgonen.

Det är också värt att notera att det inte är den matematiska modellens komplexitet som avgör hur bra styrningen fungerar ekonomiskt. Vi testade två olika modellstrukturer, en enkel och en mer komplex, och båda gav nästan identiska ekonomiska resultat. Det är prissignalen och byggnadens totala termiska tröghet som driver besparingarna, inte hur detaljerad modellen är.

Slutsatsen är att den studerade byggnaden har en oanvänd flexibilitet som en prediktiv styrning kan ta tillvara. Besparingarna är robusta även under de mest försiktiga antagandena. Nästa steg är att testa styrningen på den riktiga byggnaden under en uppvärmningssäsong, för att se hur väl simuleringens löften håller när de möter verkligheten.

Acknowledgements

We are deeply grateful to our supervisors at Vasakronan AB, Emma Ulfsparre and Carl Asp, for their invaluable support and guidance throughout this project. We would also like to thank Vasakronan AB for the opportunity to conduct our research within their organization and for providing access to the data and resources that made this work possible. Our sincere thanks go to our subject reviewer, Steffi Knorn, for her constructive feedback, for recognizing the potential in our work, and for motivating us to take the extra step toward publishing it as a conference paper. Finally, we would like to express our appreciation to Elísabet Andrésdóttir for her never-ending dedication throughout our time in the STS program and for always beeing a source of inspiration.

Max Bergström and Matviy Khotyaintsev, Uppsala, June 2026

Contents

Part I

Framing

Chapter 1

Introduction

The electricity grid that serves Swedish cities operates at a different cadence from the one that shaped the energy regulations still governing commercial buildings. As wind and solar generation has grown, wholesale prices on the NordPool day-ahead market have become structurally volatile (Halvgaard et al. 2012). Within a single day the SE3 spot price for Stockholm often becomes cheap overnight and expensive during morning peaks. The economic question for a building operator is therefore no longer only how much electricity a building consumes annually but also when it is drawn from the grid, as concentrating consumption during the cheapest hours captures significant cost reductions (Drgoňa et al. 2020).

Large commercial buildings with active thermal plant are better positioned to exploit this than almost any other building type. A heat pump connected to a thermal storage medium, whether building mass, buffer tanks, or a borehole field, is a physical mechanism for shifting electricity demand in time. By running the heating more aggressively during cheap hours and less during expensive ones, the same delivered comfort can be purchased at a fraction of the cost. The enabling condition is thermal inertia: the heavier the structure, the further into the future a unit of stored energy remains useful, and the wider the window over which a predictive controller can arbitrage the price signal (Reynders et al. 2013).

The large commercial office building in Stockholm that is studied in this work is owned and managed by Vasakronan AB, Sweden's largest commercial property company. This property combines all of the requisite structural and mechanical conditions in a single installation, featuring 40 000 m^2 of heated area on a heavy concrete frame, served by a reversible heat pump plant connected to a borehole thermal energy storage (BTES) field. The heavy structure stores thermal energy over many hours, and the plant can both heat and cool simultaneously. This research directly aligns with Vasakronan's strategic climate roadmap targeting a climate-neutral value chain by 2030.

But the controller currently governing this plant has no mechanism to exploit any of it. The supply water temperature that drives every compressor in the building follows a fixed outdoor compensation curve, unchanged since commissioning. It carries no information about tomorrow's electricity prices, today's solar forecast, or the thermal energy already stored in the building's structure. This thesis asks what a predictive controller would recover from that missed opportunity, and whether the savings figures that the building MPC literature typically reports would hold up for the building in question.

1.1 Research aim and questions

The primary aim of this thesis is to investigate the techno-economic potential of implementing an Model Predictive Control (MPC) strategy for a reversible heat pump plant with borehole thermal energy storage in a large commercial office building. Specifically, the study seeks to evaluate how grey-box thermal modeling and day-ahead spot price optimization can reduce operational costs compared to the limitations of conventional reactive control strategies and how the energy consumption is affected.

To fulfill this aim, the research is guided by the following three questions:

1. How can the thermal dynamics of a large, heavy-mass office building be adequately identified and represented using grey-box thermal modeling to serve as a reliable prediction model for control purposes?

2. To what extent can the implementation of an MPC strategy reduce the operational energy costs of the building compared to the existing rule-based heating curve and how is the overall thermal energy consumption affected?

3. How does the MPC affect the indoor thermal comfort?

1.2 Scope and Limitations

This study focuses on a single large office building in Stockholm, where the supply water temperature setpoint serves as the sole manipulated variable. Zone-level controls, ventilation systems, and auxiliary plant equipment are left unchanged and fall outside the scope of the work.

Neither controller is deployed on the physical plant. Instead, both the predictive controller and the existing weather-compensation curve are evaluated within the same simulation environment, using identical inputs and the same building model. This ensures that any difference in performance reflects the control logic itself, not the gap between the model and the real building.

The evaluation covers a single calendar year, 2025. Both energy consumption and operating cost are directly affected by the prevailing electricity prices, so the reported figures reflect the specific price and weather conditions of that year and should not be interpreted as multi-year averages.

Chapter 2

Background

This chapter describes the studied building and its energy plant. It covers the building's thermal properties, how the heat pump plant and borehole field operate together, and why the existing weather-compensation curve is insufficient as a control strategy for a building with this level of flexibility.

2.1 The Studied Building

This thesis focuses on a large commercial office building in Stockholm, Sweden, owned and managed by Vasakronan AB. The building, which has a total heated area of approximately 40000 m^2, features a massive concrete structure. From a thermodynamic perspective, its substantial physical properties place it in the highest thermal inertia category defined by ISO 52016-1 (ISO 2017). Because of its size, complex thermal dynamics, and modern energy systems, the building represents an ideal pilot asset for testing Vasakronan's digital-twin and smart-control initiatives. Through this research collaboration, direct access was established to the building's automated operating systems, bypassing the data-collection bottlenecks that typically hinder academic studies of advanced predictive controllers in real-world commercial properties.

The building's indoor climate is managed by a centralised reversible heat pump plant connected to a borehole thermal energy storage (BTES) field beneath the property. Depending on the season, the plant extracts heat from or dumps heat into the bedrock. Heating and cooling are distributed throughout the building via a common hydronic loop, where the supply water temperature serves as the primary energy carrier. Adjusting this temperature is the single most influential lever available for controlling the building's overall energy balance, and it is the variable the MPC will later be tasked with optimising.

2.2 Limitations of the Existing Control Strategy

The building's heat pump plant is currently governed by a static weather-compensation curve. This logic measures the outdoor temperature and derives a supply water setpoint from a fixed curve, with one setpoint for each outdoor temperature. The approach is designed to be reactive as it responds only to the present outdoor conditions and carries no information about tomorrow's electricity prices, expected solar gains, or the thermal energy already stored in the building's structure. Consequently, the system often operates the compressors at full capacity during morning hours, when both occupancy driven internal heat gains and Nord Pool prices are at their daily peak.

The thermal inertia of the concrete structure offers a natural solution to this inefficiency. Because the building retains heat over many hours, it is physically possible to pre-heat the structure during cheap overnight periods and reduce compressor output during expensive ones, without compromising indoor comfort. The BTES field extends this capability further, allowing thermal energy to be stored over longer time horizons than the building mass alone permits. However, neither of these storage mechanisms is accessible to a controller that cannot anticipate future conditions.

2.3 Towards Model Predictive Control

To enable this proactive operation, a more advanced framework is required: Model Predictive Control (MPC). By utilizing a mathematical model of the building's thermal dynamics combined with forecasts for weather and spot prices, an MPC algorithm can solve an optimization problem to determine the most cost-effective heating and cooling strategy for the coming 24 hours. At each hour, the controller solves an optimisation problem over a 24 hour planning horizon, applies only the first action, and then re-solves at the next hour with updated measurements and forecasts. This receding-horizon principle allows the controller to continuously correct for forecast errors and unexpected disturbances without committing to a fixed plan.

A critical aspect of an MPC is its objective function, which must carefully balance cost savings against indoor climate quality. The controller does not simply turn off the heating during expensive hours. It is instead provided with strict comfort boundaries, for instance maintaining the indoor room temperature between 21°C and 23°C. The algorithm explores the future thermal trajectory to find the absolute minimum operating cost while mathematically guaranteeing that the indoor climate never violates these predefined comfort limits.

2.4 The Role of System Identification

The success of load shifting and predictive control relies entirely on the accuracy of the underlying mathematical model. If the model misjudges the building's thermal inertia, the controller might pre-heat the structure too little, leading to comfort violations, or too much, wasting energy. Before the MPC can be deployed, the thermodynamic behaviour of the building must therefore be mathematically identified and validated against historical data.

This is not straightforward for the studied building. The existing weather-compensation curve has governed the supply water setpoint since commissioning, meaning the historical data contains almost no examples of the off-curve setpoint variations that the MPC will routinely produce. A purely data-driven model (black-box) trained on this data inherits this limitation and cannot reliably extrapolate to conditions it has never seen (Drgoňa et al. 2020). A physics-based model (white-box), on the other hand, requires detailed knowledge of every material layer and geometric property of the building, which is rarely available in practice (Drgoňa et al. 2020). The identification approach adopted in this thesis therefore takes a middle path, a grey-box model, which uses physical laws to define the structure of the equations and measured data to calibrate the parameters that first principles leave undetermined.

It is important to note that this thesis does not deploy the MPC on the physical plant. Instead, both the MPC and the legacy controller are evaluated within the same digital twin, replaying a fixed historical dataset covering the calendar year 2025. This twin-on-twin approach is a deliberate methodological choice: by subjecting both controllers to identical conditions within the same model, differences in simulated performance can be attributed to control logic alone, rather than to inconsistencies between the model and the real plant.

Chapter 3

Prior Research

This chapter reviews the research that informed the key design decisions in this thesis. It covers model predictive control in buildings, the distinction between tracking and economic formulations, grey-box thermal modelling, and state estimation for unmeasured building states. Each section connects directly to a design choice made in the methodology.

3.1 Model predictive control in buildings

While much of the literature surrounding Model Predictive Control (MPC) in buildings relies heavily on theoretical simulations, real-world implementations are important for validating the practical feasibility of predictive energy savings. An example of such a practical application is an experimental study of two-months conducted in a university building in Prague (Široký et al. 2011). The researchers implemented an MPC strategy on a ceiling radiant heating system, effectively utilizing the building's massive concrete structure as a thermal storage medium (Široký et al. 2011).

To capture the building's thermal dynamics for the controller, the authors utilized an RC-modeling approach. They extended the mathematical model with a stochastic component to estimate the unknown thermal parameters through statistical identification of measured input-output data (Široký et al. 2011). The optimization problem was formulated to minimize the total energy consumption defined linearly by the temperature difference between the supply and return water, while maintaining the reference room temperature using soft constraints to penalize comfort deviations and prevent infeasibility (Široký et al. 2011). The implementation showed that replacing a conventional weather-compensated heating curve with predictive control yielded significant energy savings ranging from 15% to over 28%, depending heavily on the building's insulation level and the prevailing outside temperatures (Široký et al. 2011). Furthermore, the MPC strategy naturally resulted in significant peak energy demand reductions by proactively preheating the concrete mass during the night (Široký et al. 2011). Despite these results, the study concluded that the most substantial barrier to the widespread practical adoption of MPC remains the extensive time, cost, and non-trivial engineering effort required to develop the initial mathematical building model (Široký et al. 2011).

This concept of predictive "load shifting" has recently been further advanced to explicitly target financial optimization. A comprehensive study conducted on Vasakronan's building "Juvelen" in Uppsala evaluated the integration of ground source heat pumps with thermal energy storage to leverage volatile day-ahead electricity spot prices (Behzadi 2025). Building on the same principles of predictive thermal storage demonstrated by (Široký et al. 2011), this study uses an MPC framework coupled with TRNSYS simulations. TRNSYS (Transient System Simulation Tool) is a simulation software used to create highly detailed models of building thermal behavior and energy systems (Thermal Energy System Specialists, LLC 2026). The mathematical objective function was designed to explicitly incorporate hourly spot market dynamics and peak power demand penalties (Behzadi 2025). To ensure the controller's resilience to real-world uncertainties, the strategy was validated through Monte Carlo simulations that introduced stochastic noise into both price and weather forecasts (Behzadi 2025). The results demonstrated that replacing traditional reactive control with MPC reduced market-responsive operational costs by over 13% and shortened the system's payback period by approximately four years, all while successfully maintaining thermal comfort (Behzadi 2025).

While comprehensive simulation environments like TRNSYS offer a very high accuracy for design evaluations, their high computational demand and complexity make them challenging to use directly into a fast, real-time Model Predictive Control loop (Drgoňa et al. 2020). For practical and continuous spot-price optimization, the controller requires a mathematical model that correctly balances physical representation with computational efficiency. Consequently, this study instead employs a data-driven Grey-box modeling approach with RC-networks (Bacher and Madsen 2011). This approach ensures that the optimization algorithm can reliably generate control signals within the strict time constraints required when using Nord Pool day-ahead prices.

3.2 Tracking or economic MPC

In general process control literature (Ellis et al. 2014; Rawlings et al. 2012) as well as in HVAC control applications, a clear distinction is made between tracking MPC and Economic MPC (EMPC). A tracking controller minimises the deviation from a fixed reference setpoint, for instance maintaining the indoor temperature at exactly 21°C at all times. This formulation leaves no room for the temperature to vary, and consequently no thermal capacity for the controller to exploit. An economic controller, by contrast, minimises a direct cost function over the planning horizon. Rather than following a fixed setpoint, it is given a comfort band, for instance between 21°C and 23°C, and is free to let the indoor temperature sit anywhere within that band without penalty (Drgoňa et al. 2020). This freedom is what enables load shifting by allowing the temperature to rise toward the upper edge of the band during cheap hours and fall down during expensive ones; the controller can concentrate compressor operation where electricity is cheapest (Halvgaard et al. 2012). The wider the comfort band, the more thermal storage capacity is available to the optimiser. In the remainder of this thesis, the term MPC refers to the economic formulation.

3.3 Grey-box thermal modelling

To predict the indoor temperature trajectory over a 24-hour horizon, the MPC requires a model that is both physically meaningful and computationally tractable. This type of model reduces the governing partial differential equations to a small set of ordinary differential equations, preserving a physical interpretation of the parameters while retaining the convexity and low dimensionality required for optimisation.

The family of RC structures relevant to building control forms a natural hierarchy of increasing complexity. The simplest member is the 1R1C model, which treats the entire building as a single thermally lumped air node. While this structure requires only a handful of parameters, (Bacher and Madsen 2011) demonstrate that it fundamentally fails to reproduce the multi-time-scale dynamics present in real buildings: a single thermal mass cannot simultaneously track the fast temperature swings driven by solar gains and internal loads and the slow discharge of the structural core. The two-state model adds an envelope node, introducing a second time constant associated with the wall mass. Three-state structures such as the 3R2C extend this further, coupling an indoor air node to both a surface mass and a second structural element. The 4R3C structure adds a fourth resistance and decouples the surface mass T_w from the deep structural core T_m, allowing the model to represent three distinct thermal time constants within a single, tractable formulation.

This decoupling is physically motivated. Thermal penetration into a massive structure is frequency-dependent as high-frequency disturbances, such as the daily solar cycle, activate only the superficial layers of the concrete, while low-frequency variations penetrate the deep structural core (Reynders et al. 2013). For an MPC that exploits the building's thermal inertia as a load-shifting buffer, the model must resolve both the fast-reacting surface mass and the slow-moving deep core; conflating them into a single node introduces severe phase errors over a 24-hour prediction horizon.

Because the RC parameters of a lumped model are fitted to measured input-output data rather than derived from first principles, they carry important limitations. The indoor air is treated as a perfectly mixed, homogeneous volume and spatial gradients are ignored. As a consequence, the identified resistance and capacitance values are effective parameters. They capture the aggregate observed dynamics, but they cannot be mapped deterministically to individual physical components (Kircher and Zhang 2015). For a $40\,000\,\mathrm{m}^2$ commercial building, this is an accepted trade-off. Attempting to resolve every material layer individually would introduce far more free parameters than the available data can constrain.

3.4 The hidden-state problem and state estimation

A critical challenge in this architecture is that the deep wall temperature (T_w) is a hidden state that cannot be measured directly by standard sensors. If the MPC simulates the future based on an incorrect initial wall temperature, the predictions will suffer from integrating errors. To ensure real-time stability, the model must be augmented with a state estimator, such as a Kalman filter. This filter continuously utilizes the measurable room temperature (T_i) to correct its estimation of the hidden structural temperature (T_w) (Drgoňa et al. 2020).

3.5 Objective function design and the role of regularisation

The design of the objective function is crucial for balancing economic performance with thermal comfort and system stability. Based on the established frameworks by (Rawlings et al. 2012) and (Ellis et al. 2014), the objective function for MPC is formulated to balance energy costs, thermal comfort, and equipment wear:

$$J = \sum_{k=0}^{N_p-1} \left(\lambda_{el} \cdot C_{el,k} \cdot P_k + \lambda_{comf} \cdot \epsilon_k^2 + \lambda_{\Delta u} \cdot \Delta u_k^2 \right). \tag{1}$$

Here, the parameters λ_{el}, λ_{comf}, and $\lambda_{\Delta u}$ represent tuning weights used to prioritize the relative importance of economic savings, thermal comfort, and equipment longevity. The formulation of Equation 1 requires consideration of its three core components: Soft Constraints and Thermal Comfort (ϵ_k), Direct Economic Cost and Actuator Wear ($C_{el,k} \cdot P_k$), and Input Regularization (Δu_k).

3.5.1 Soft Constraints and Thermal Comfort (ϵ_k)

Imposing hard boundaries on the indoor temperature introduces a risk of optimization failure. During extreme weather events or sudden system disturbances, the heating system can physically be unable to maintain the temperature within the exact boundaries. Imposing a hard limit in such a scenario would lead to infeasibility, causing the solver to fail to find a solution. To ensure the robustness of the algorithm, the comfort boundaries are implemented as soft constraints using a slack variable ($\epsilon_k \geq 0$). If the temperature deviates from the comfort band, the objective function applies a heavily weighted quadratic penalty ($\lambda_{comf} \cdot \epsilon_k^2$). This guarantees that the optimization process continuously yields a viable control strategy while strongly penalizing temperature violations (Drgoňa et al. 2020).

3.5.2 Direct Economic Cost ($C_{el,k} \cdot P_k$)

In contrast to traditional control that tracks a constant temperature setpoint, MPC directly minimizes the operational cost (Ellis et al. 2014). This term calculates the actual cost of electricity by multiplying the dynamic day-ahead spot prices ($C_{el,k}$) with the consumed power (P_k). By minimizing this linear cost over the prediction horizon, the controller automatically shifts the heating load; it pre-heats the building's thermal mass during hours with low electricity prices and reduces power consumption during price peaks (Halvgaard et al. 2012).

3.5.3 Actuator Wear and Input Regularization (Δu_k)

In MPC, the linear cost function tends to drive the control inputs to their extreme bounds to maximize financial savings, which can result in unstable, oscillatory "bang-bang" control behavior (Drgoňa et al. 2020). To mitigate this, a quadratic penalty on the input rate of change (Δu_k) is introduced. Mathematically, this term provides necessary regularization to the objective function, preventing high-frequency chattering in the control signal and adding robustness to the numerical optimization routine (Halvgaard et al. 2012; Rawlings et al. 2012). From a practical engineering perspective, this dampening is essential to prevent excessive wear on actuators and extend the operational lifetime of the mechanical equipment, such as the heat pumps and compressors.

3.6 Prediction horizon selection (N_p)

Selecting the appropriate prediction horizon (N_p) is complex, largely due to the dynamic availability of price data on the Nord Pool Day-ahead market, where tomorrow's prices are typically published between 13:00 and 14:00. The published prices cover the entire upcoming day from 00:00 to 24:00, which provides the system with a known price horizon ranging from 11 to 35 hours ahead, depending on the current time of day.

Research indicates that for heavy buildings with thermally inert systems, longer prediction horizons are required to efficiently charge the thermal mass (Oldewurtel et al. 2012). While a horizon of 24 hours is generally sufficient for most standard heating systems to achieve near-optimal performance (Oldewurtel et al. 2012), horizons of up to 48 hours are often employed in literature to fully capture the dynamics of floor heating and structural storage (Halvgaard et al. 2012).

To maintain a manageable project scope, a constant 24-hour horizon is chosen in this thesis. The lack of price data for the upcoming day during morning hours is pragmatically addressed using a persistence forecasting method, simply assuming that tomorrow's price profile will resemble today's at the corresponding hours. While perfect forecasts yield superior performance, this approach is a practical necessity when simulating continuous day-ahead market operations.

Part II

Methodology

Chapter 4

Method

This chapter describes the research design and the data used in the study. It covers the control problem and the manipulated variable, the twin-on-twin evaluation approach and why it was chosen over a comparison against measured historical consumption, and the data sources, signals, and cleaning steps applied before model identification.

4.1 Research design and approach

This study employs a quantitative, model-based research design. The methodological foundation is deductive. The laws of thermodynamics and convective heat transfer impose absolute physical constraints on what the building plant can do, and these constraints are encoded directly into the prediction model and the optimisation problem. Within this deductive frame, grey-box system identification is used to calibrate the parameters that first principles leave undetermined (Bacher and Madsen 2011; Prívara et al. 2013), drawing on two years of operational data from the building management system.

The development proceeded iteratively. A focused review of building thermal modelling and MPC practice informed the choice of model structures, identification algorithms, and control formulation. Each major modelling decision was evaluated against the data before being finalised, and the architecture was revised where the initial evaluation revealed weaknesses. This iterative process is reflected in the structure of the methodology chapters that follow.

4.2 The control problem

The energy plant at the studied property operates through a hierarchy of interacting subsystems. Pump staging logic decides which compressors are active, distribution routing directs heat to the space heating, ventilation coils, and chilled-beam circuits, and zone-level controllers adjust air volumes based on carbon dioxide concentrations and occupancy. These subsystems are a functioning part of the building management infrastructure and are left entirely unchanged throughout this study.

The single signal that an external controller can write without disturbing this hierarchy is the supply water temperature setpoint for the heating distribution loop. This setpoint governs how much thermal energy is available to all downstream circuits. By adjusting it, an external controller can raise or lower the building's energy uptake without interfering with any of the local control functions below. This is the manipulated variable throughout the study, and it is the same lever used in comparable building MPC implementations (Široký et al. 2011).

4.3 A fair basis for comparison

Evaluating a predictive controller requires a comparison against a reference. The most direct reference is the historical record of measured energy consumption under the existing controller. That comparison is, however, systematically misleading for model-based controllers. The simulated MPC operates over a mathematical model of the building and the historical record reflects the physical building. Any difference between the two is not only due to the difference in control logic but also due to the gap between the model and the real system, often referred to as the plant-model mismatch (Drgoňa et al. 2020). A controller that optimises over an imperfect model will implicitly exploit that imperfection, recovering savings that have no physical counterpart. The reported savings can therefore be overly optimistic and overstate what the controller would actually achieve in deployment (Prívara et al. 2013).

This study removes the problem by replaying both the predictive controller and the legacy weather-compensation curve inside the same digital twin (Drgoňa et al. 2020). Both controllers run over the same model, receive the same exogenous inputs, and are evaluated against the same output. The only difference between the two runs is the control logic. Any difference in simulated cost or energy consumption therefore reflects a genuine algorithmic advantage, not an artefact of model imperfection.

4.4 Data collection and processing

The empirical dataset utilized in this study spans the two calendar years 2024 and 2025, and was acquired directly from Vasakronan's internal building management system (BMS) database. All operational data streams were exported at high frequency and subsequently resampled to an hourly resolution by computing the mean within each hourly window, matching the resolution of the Nord Pool day-ahead market. The year 2024 is used exclusively for model training and parameter identification, while the year 2025 is held out and used only for validation and controller evaluation. This long-term, high-quality industrial dataset is a direct outcome of the collaborative framework with Vasakronan, designed to leverage actual building performance metrics to drive scalable, data-driven decarbonization strategies across their property portfolio.

4.4.1 Sources and signals

Thermal power and temperature measurements were exported from the building management system as high-frequency trend files. Indoor temperature was computed as an aggregated average across all active zone sensors. Net thermal power delivered to the occupied zones, Q_{net}, was aggregated from space heating and mechanical cooling circuits and explicitly excludes process cooling assigned to server rooms, whose concentrated loads are not representative of the ambient thermal dynamics:

$$Q_{\text{net}} = \sum Q_{\text{heat}} - \sum Q_{\text{cool}}. \tag{2}$$

Outdoor temperature was sourced from on-site sensors and supplemented with hourly observations from the closest weather station operated by the Swedish Meteorological and Hydrological Institute (SMHI) to fill sensor outages. Global horizontal irradiance was obtained from the SMHI STRÅNG mesoscale solar radiation model (SMHI 2026). Occupancy fraction was derived from the building's motion sensor network, scaled to a continuous value between zero and one. Day-ahead electricity spot prices for the SE3 bidding zone were downloaded from the ENTSO-E transparency platform (ENTSO-E 2024). Table 1 lists the signals used in the thermal model identification.

Table 1: Signals used in thermal model identification.

Symbol	Description	Unit	Source
T_i	Indoor air temperature, zone average	°C	BMS
Q_{net}	Net thermal power to ambient zones†	kW	BMS meters
T_{out}	Outdoor air temperature‡	°C	BMS / SMHI
I_{sol}	Global horizontal irradiance	kW/m^2	SMHI Strång
φ_{occ}	Occupancy fraction	–	BMS

† Server-room process cooling excluded. ‡ SMHI backfill for flatline faults exceeding 12 consecutive hours.

4.4.2 Data cleaning

The raw data streams required three cleaning steps before use. First, physics-based range filters flagged any reading outside realistic physical bounds as missing. Thermal power meters were restricted to only positive values and solar irradiance values were capped at the clear-sky maxima for the region. The outdoor temperature sensor occasionally experienced communication failures that caused the recorded value to freeze at a fixed reading. Sequences where the temperature variance was zero for more than twelve consecutive hours were flagged as sensor outages. The removal window was extended by 72 hours to account for erratic readings preceding a flatline, and the flagged intervals were replaced with SMHI observations. Isolated instances of impossible indoor temperature drops below 19°C that lacked a corresponding compensatory heating spike in the metered data were identified as measurement artefacts and replaced by linear interpolation over gaps of up to four hours.

The resulting dataset provides a continuous, physically consistent record. The internal energy balance computed over the 2024 training year is used as the headline data-quality check before model identification begins.

Chapter 5

Modelling the building and plant

This chapter develops the prediction model used by the controller. It covers the grey-box thermal model and its parameter identification, the capacity constraints that limit heat delivery at a given supply temperature, and the thermodynamic model that translates delivered heat into compressor electricity. Together these three components form the simulation environment in which the controller operates.

5.1 Thermal model

5.1.1 RC model family

Grey-box resistance-capacitance (RC) networks are the established model class for this identification task. They reduce the transient heat transfer problem from a set of coupled partial differential equations over a complex three-dimensional geometry to a small set of ordinary differential equations, preserving a physical interpretation of the parameters while keeping the model tractable for real-time optimisation (Bacher and Madsen 2011; Drgoňa et al. 2020). A standard hierarchy of structures for buildings was established by Bacher and Madsen (2011) and has since been applied to offices and similar buildings (Berthou et al. 2014; Reynders et al. 2014).

Four candidate structures of increasing complexity are evaluated. The state vector $\mathbf{x}$, control input u, disturbance vector $\mathbf{d}$, and output y are common to all structures:

$$\mathbf{x} = [T_i, \ldots]^\top, \quad u = Q_{\text{net}}, \quad \mathbf{d} = [T_{\text{out}},\, I_{\text{sol}},\, \varphi_{\text{occ}}]^\top, \quad y = T_i. \tag{3}$$

1R1C. The simplest structure aggregates all thermal mass into a single indoor air node:

$$C_i \frac{\mathrm{d}T_i}{\mathrm{d}t} = \frac{T_{\text{out}} - T_i}{R_{ia}} + Q_{\text{net}} + A_w I_{\text{sol}} + \varphi_{\text{occ}} Q_{\text{int}}. \tag{4}$$

2R2C. A two-state model that introduces a separate envelope mass node T_e, separating the indoor climate from the slower facade dynamics:

$$C_i \frac{\mathrm{d}T_i}{\mathrm{d}t} = \frac{T_e - T_i}{R_{ie}} + Q_{\mathrm{net}} + A_w I_{\mathrm{sol}} + \varphi_{\mathrm{occ}} Q_{\mathrm{int}}, \tag{5}$$

$$C_e \frac{\mathrm{d}T_e}{\mathrm{d}t} = \frac{T_i - T_e}{R_{ie}} + \frac{T_{\mathrm{out}} - T_e}{R_{ea}}. \tag{6}$$

3R2C. A two-state model that retains a single structural node T_c but adds a direct infiltration bypass resistance R_{ia}:

$$C_i \frac{\mathrm{d}T_i}{\mathrm{d}t} = \frac{T_{\mathrm{out}} - T_i}{R_{ia}} + \frac{T_c - T_i}{R_{ic}} + Q_{\mathrm{net}} + A_w I_{\mathrm{sol}} + \varphi_{\mathrm{occ}} Q_{\mathrm{int}}, \tag{7}$$

$$C_c \frac{\mathrm{d}T_c}{\mathrm{d}t} = \frac{T_i - T_c}{R_{ic}} + \frac{T_{\mathrm{out}} - T_c}{R_{ca}}. \tag{8}$$

4R3C. The most complex candidate decouples a shallow structural surface mass T_w from a slow deep core T_m and adds a separate solar aperture A_e for the opaque facade (Reynders et al. 2014):

$$C_i \frac{\mathrm{d}T_i}{\mathrm{d}t} = \frac{T_{\mathrm{out}} - T_i}{R_{ia}} + \frac{T_w - T_i}{R_{iw}} + \frac{T_m - T_i}{R_{im}} + Q_{\mathrm{net}} + A_w I_{\mathrm{sol}} + \varphi_{\mathrm{occ}} Q_{\mathrm{int}}, \tag{9}$$

$$C_w \frac{\mathrm{d}T_w}{\mathrm{d}t} = \frac{T_i - T_w}{R_{iw}} + \frac{T_{\mathrm{out}} - T_w}{R_{wa}} + A_e I_{\mathrm{sol}}, \tag{10}$$

$$C_m \frac{\mathrm{d}T_m}{\mathrm{d}t} = \frac{T_i - T_m}{R_{im}}. \tag{11}$$

The 4R3C architecture decouples two physically distinct thermal storages. High-frequency disturbances such as the diurnal solar cycle penetrate only the surface layers of a heavy concrete structure, while low-frequency variations reach the deep core (Reynders et al. 2013). For an economic MPC that exploits multi-day structural thermal storage, capturing both time scales in the same model is important. Table 2 summarises the four candidates.

Table 2: Candidate grey-box RC structures evaluated for the thermal model.

Model	States	Parameters	Structural hypothesis
1R1C	1	6	Single lumped air node.
2R2C	2	8	Indoor air coupled to a single envelope node.
3R2C	2	9	Indoor air, structural mass, and ambient bypass.
4R3C	3	12	Indoor air, structural surface mass, and deep core.

5.1.2 Parameter identification

Each model defines a continuous-time stochastic system whose parameters are estimated by maximising the innovation likelihood evaluated through the Kalman filter after exact zero-order-hold (ZOH) discretisation at $\Delta t = 1\,\mathrm{h}$ via the matrix exponential (Madsen and Holst 1995; Bacher and Madsen 2011).

The likelihood surface is complex. Similar input-output trajectories can arise from physically distinct parameter combinations, so a single-start gradient method will routinely miss the global optimum. Identification therefore proceeds in two stages. Differential evolution (Storn and Price 1997) performs a population-based global search with a population of 15 candidates initialised without any gradient information. The best candidate returned by the global search is then passed to L-BFGS-B (Byrd et al. 1995), a quasi-Newton method that respects explicit parameter bounds, which polishes the solution to tight convergence. Four random restarts of the polishing stage are run and the best result is retained.

Physical bounds on all parameters are imposed throughout. Structural and ambient resistances are bounded from below by the inverse of the maximum permissible mean thermal transmittance given in BBR Section 9 (Boverket 2011). Node capacitances are bounded by the heavy concrete thermal mass categories in ISO 52016-1 (ISO 2017). Internal heat gains and occupancy contributions are bounded by standard operational profiles from SVEBY (SVEBY 2012). These bounds prevent non physical solutions and serve as the primary regularisation mechanism for a dataset with limited independent excitation in the setpoint direction.

5.1.3 Model Selection Metrics

To evaluate and compare the statistical performance of the candidate grey-box RC model structures during the system identification phase, three information-theoretic criteria are utilized: Negative Log-Likelihood (NLL), the Akaike Information Criterion (AIC) (Akaike 1974), and the Bayesian Information Criterion (BIC) (Schwarz 1978).

Because parameter identification proceeds by maximizing the innovation likelihood $\mathcal{L}$ evaluated through the Kalman filter algorithm (Madsen and Holst 1995; Bacher and Madsen 2011), the baseline metric for goodness-of-fit is the Negative Log-Likelihood, defined as:

$$\text{NLL} = -\ln(\mathcal{L}) \tag{12}$$

While a lower NLL indicates a superior statistical fit to the training dataset, optimizing for NLL alone introduces a structural risk of overfitting when comparing model candidates with varying degrees of freedom. To systematically account for the trade-off between statistical accuracy and parametric complexity, the AIC and BIC introduce distinct penalties based on the number of estimated parameters. The Akaike Information Criterion is formulated as:

$$\text{AIC} = 2k - 2\ln(\mathcal{L}) \tag{13}$$

where k represents the number of identified parameters in the model structure (e.g., $k = 6$ for 1R1C, and $k = 12$ for 4R3C). To apply a more stringent penalty for high-dimensional models evaluated over large datasets, the Bayesian Information Criterion is additionally computed:

$$\text{BIC} = k\ln(n) - 2\ln(\mathcal{L}) \tag{14}$$

where n is the total number of hourly observations in the estimation dataset (for 2024, $n = 8760$).

For all three metrics, lower numerical values designate the optimal balance between physical expressiveness and parametric parsimony, guiding the final selection of the thermal prediction model for the downstream economic MPC framework.

5.2 Capacity

The thermal model tells the controller how indoor temperature responds to net heat delivery. The capacity model tells the controller how much heat the plant can actually deliver at a given supply temperature setpoint. This upper bound on heat delivery must be included as a constraint in the optimisation problem and without it the controller may request heat flows that the physical plant cannot produce. The capacity constraints in this case are derived from physical principles and observed extreme operating points rather than from a statistical fit.

5.3 Heat Delivery Limits

The thermal model tells the controller how indoor temperature responds to heat delivery. The capacity model tells it how much heat the plant can actually deliver at a given supply temperature setpoint. Without this upper bound, the controller may request heat flows the physical plant cannot produce.

The maximum heat deliverable follows from the heat exchanger relation $Q = UA \cdot \Delta T$, giving the linear constraint:

$$Q_{h,k} \leq UA_{\mathrm{em}}\big(v_k + \delta_{\mathrm{off}} - T_{i,k}\big), \tag{15}$$

where v_k is the commanded supply setpoint, $\delta_{\mathrm{off}} = 0.25\,\mathrm{K}$ corrects for the steady-state tracking error between the commanded and actual supply temperature, and $UA_{\mathrm{em}} = 25.5\,\mathrm{kW/K}$ is the hydronic emission coefficient. The latter is calibrated from the coldest observed operating hours ($T_{\mathrm{out}} < -5\,^{\circ}\mathrm{C}$), where a peak heat delivery of 490 kW at a supply setpoint of 40 °C and a room temperature near 21 °C gives $490/19 \approx 25.7\,\mathrm{kW/K}$, rounded to 25.5. This design-day calibration is preferred over a regression fit on the full dataset, where the BMS PI loop leaves too little variance in the driving temperature difference for reliable statistical identification (Prívara et al. 2013).

In addition to this setpoint-dependent limit, absolute capacity caps are applied based on the 99th percentile of observed thermal loads in the 2024 training data, reflecting the true mechanical limits of the plant.

5.4 Heat Pump Thermodynamics

The MPC minimises compressor electricity, not delivered heat. A model is therefore needed to translate between the two via the coefficient of performance:

$$E_{\mathrm{comp},k} = \frac{Q_{h,k}}{\mathrm{COP}_{h,k}}\Delta t. \tag{16}$$

A common simplification in building MPC is to treat the COP as constant or as a simple function of outdoor temperature (Halvgaard et al. 2012). This works well when the heat source temperature is stable throughout the year. At the studied plant, however, the heat source changes depending on how the plant is operating. During heating-dominant periods the heat pumps extract energy from the cold borehole field, while during simultaneous heating and cooling they recover heat from the building's own cooling circuit instead. This shifts the evaporating temperature by approximately 10°C between the two modes, which has a direct and significant effect on the COP. Treating the COP as constant would therefore misrepresent the plant's electricity consumption in each season.

The Lorentz-scaled Carnot formulation is therefore adopted (Pieper et al. 2020):

$$\mathrm{COP}_{h,k} = \eta_L^{(r)} \frac{T_{\mathrm{cond},k} + 273.15}{T_{\mathrm{cond},k} - T_{\mathrm{evap},k}}, \tag{17}$$

where $T_{\mathrm{cond},k}$ and $T_{\mathrm{evap},k}$ are the measured condenser and evaporator temperatures, and $\eta_L^{(r)}$ is the Lorentz efficiency for operating regime r. Modern inverter scroll compressors generally exhibit high overall compression efficiency compared to traditional technologies (Chua et al. 2010). Three regimes are distinguished by which thermal loads are active:

- R1, heating-dominant: cooling load below 20 kW, evaporator drawing from the borehole field.
- R2, simultaneous: both heating and cooling loads at or above 20 kW, evaporator drawing from the cooling return brine.
- R3, cooling-dominant: heating load below 20 kW.

Regime R2 accounts for approximately 84% of annual operating hours, making simultaneous heat recovery the dominant operating condition throughout the year. A separate $\eta_L^{(r)}$ is identified for each regime by least-squares regression on the 2024 compressor meter data.

5.5 Digital twin composition and validation strategy

5.5.1 Composition

The three model layers are composed into a single simulation of the plant and building system. At each hourly step, the capacity layer determines the maximum heat flows available at the current setpoint and disturbance conditions. Given those flows and the controller's setpoint, the thermal layer advances the indoor temperature state one step forward using the discretised RC model. The COP layer then evaluates the thermodynamic lift implied by the condenser and evaporator temperatures and returns the corresponding compressor electricity.

5.5.2 Validation approach

Before the digital twin can serve as a reliable simulation environment for controller evaluation, it must reproduce the measured plant behaviour within acceptable tolerances. The validation is conducted on the held-out 2025 test year and examines three aspects of model behaviour.

Closed-loop temperature prediction. The root mean square error and mean bias error of the indoor temperature are evaluated at a 1-step horizon for both the 1R1C and 4R3C structures. The 1-step metric reflects the accuracy available to the receding-horizon controller at each correction step, and is the primary acceptance criterion for the thermal model.

24-step prediction stability. The twin is run open loop for 24 steps from a reset state at the start of each day. This quantifies the prediction uncertainty the controller faces across its full planning horizon and checks whether the model drifts systematically away from the measured temperature over a full day.

Energy balance at the plant level. The simulated annual compressor energy is compared against the measured total. This step validates the composed twin as a whole rather than the thermal model in isolation, and confirms that the digital twin can serve as a meaningful basis for comparing controller costs.

5.5.3 Closed-loop evaluation mode

The primary indoor temperature metric is the closed-loop 1-step prediction. In the receding-horizon controller, the measured indoor temperature is read at the start of every hour, used to initialise the thermal state, and the optimisation is solved fresh. State error therefore does not accumulate between steps. Evaluating the twin in pure open loop over thousands of hours tests an operating mode the controller never uses and produces a thermal drift that is irrelevant to actual performance.

A 24-step rolling-restart RMSE is retained as a secondary metric. The thermal state is reset to the observed indoor temperature every 24 hours, and the twin is run open loop for the following 24 steps. This quantifies the prediction uncertainty the controller faces across the full planning horizon and informs whether the comfort constraints require tightening at longer horizons.

Chapter 6

MPC Formulation

This chapter describes the MPC formulation. It covers the manipulated variable, the objective function and its four terms, the physical constraints imposed on the optimisation, the choice of solver, and the receding-horizon execution loop including forecast assumptions and how the controller updates its plan at each hourly step.

6.1 Manipulated variables

As mentioned in earlier chapters, the building management system (BMS) accepts a single command from an external controller: a temperature setpoint for the supply water entering the heating distribution loop. The BMS inner PI loop then drives the actual supply temperature to that value by adjusting the heat pump staging. The MPC therefore operates on setpoints rather than on heat flows or compressor power. The heat delivered to the building $Q_{h,k}$ and the heat extracted from it $Q_{c,k}$ enter the optimisation as variables bounded by the hydraulic capacity envelope and are not free decision variables.

The manipulated variable at each step k of the prediction horizon is

$$u_k \;=\; T_{\text{supply},k} \;\in\; [18,\;45]\,^{\circ}\text{C}, \tag{18}$$

where the bounds correspond to the operational range of the weather-compensation curve in the BMS. With a horizon of $H = 24$ steps, the optimiser selects a vector of 24 setpoints simultaneously and executes only the first one before solving again at the next hour.

6.2 The objective function

The controller minimises a quadratic cost function over the 24-hour prediction horizon. Building on the three-term baseline established in the prior research chapter, the implemented objective contains one additional term whose necessity emerged during the design iterations:

$$J \;=\; \sum_{k=0}^{H-1} \Big[\, \underbrace{p_k \cdot E_{\text{comp},k}}_{\text{spot-price cost}} + \underbrace{\varepsilon \cdot E_{\text{comp},k}}_{\text{absolute consumption}} + \underbrace{\lambda_{\text{comf}} \left(s_{\text{hi},k}^2 + s_{\text{lo},k}^2\right)}_{\text{soft comfort}} + \underbrace{\lambda_{\Delta u}\, \Delta u_k^2}_{\text{input rate penalty}} \,\Big], \tag{19}$$

where p_k is the forecast spot price in SEK/kWh, $E_{\text{comp},k}$ is the compressor electrical energy in kWh, $s_{\text{hi},k}$ and $s_{\text{lo},k}$ are slack variables on the upper and lower edges of the comfort band, and $\Delta u_k = u_k - u_{k-1}$ is the change in setpoint between consecutive hours. The weights λ_{comf} and $\lambda_{\Delta u}$ are set to 1000 SEK/°C^2 and 0.5 SEK/°C^2 respectively.

6.2.1 Spot-price cost

The hourly spot price p_k is the day-ahead price from the Nord Pool market, augmented with a fixed grid surcharge of 0.80 SEK/kWh to mimic the difference between spot prices and what a consumer would typically pay. Minimising the product $p_k \cdot E_{\text{comp},k}$ summed over the horizon shifts compressor operation toward cheap hours and away from expensive ones, which is the load-shifting behaviour seen in the results chapter.

6.2.2 Absolute consumption penalty

The second term applies a small price-independent cost $\varepsilon = 0.03$ SEK/kWh to every kilowatt-hour of electricity used. Without it, the optimiser becomes indifferent to electricity consumption whenever p_k approaches zero, which happens several times a year on the Nord Pool market and occasionally takes negative values. The standard three-term objective then rewards unnecessary increases in compressor power at those hours, producing solutions that are economically optimal under the stated cost but operationally absurd.

6.2.3 Soft comfort band

The comfort band is implemented as a soft constraint with quadratic penalty. Indoor temperature is allowed to lie anywhere inside the comfort band [21, 23] °C without penalty, and any deviation outside activates one of two non-negative slack variables:

$$s_{\text{hi},k} \;\geq\; T_{\text{room},k} - T_{\text{hi}}, \qquad s_{\text{lo},k} \;\geq\; T_{\text{lo}} - T_{\text{room},k}, \qquad s_{\text{hi},k},\, s_{\text{lo},k} \;\geq\; 0. \tag{20}$$

Formulating comfort as a soft constraint is standard practice in building MPC (Drgoňa et al. 2020) and ensures that the solver always returns a valid control action even under extreme conditions. The comfort zone gives the optimiser maximum flexibility to exploit thermal mass for load shifting, since the room temperature can sit anywhere in the band without penalty.

6.2.4 Input Rate Penalty

The fourth term penalises the change in supply temperature setpoint between consecutive hours. The economic term in equation (19) is linear in $E_{\text{comp},k}$ and, without regularisation, drives solutions toward the extreme bounds of the feasible region producing oscillative behavior. The quadratic input rate penalty adds a cost on rapid changes and recovers a well-conditioned problem (Rawlings et al. 2012). From a practical standpoint, it also prevents the mechanical wear that would result from aggressive cycling of the heat pump compressors in response to short-term price movements. The weight $\lambda_{\Delta u} = 0.5$ allows changes of one or two degrees when needed, while changes of five or more degrees per hour become too expensive.

6.3 Constraints in the MPC

The optimisation is subject to four categories of physical constraints, each reflecting a real limitation of the building or plant. The first category governs the building's thermal response. The indoor temperature does not react instantly to a change in supply setpoint. Heat must first travel through the hydronic distribution system and into the building's structure before the air temperature changes. This thermal lag is captured by the discretised RC model, which the controller uses to predict how today's setpoint decisions will affect indoor temperature over the coming 24 hours. The controller cannot choose a setpoint sequence that would violate this physical relationship.

The second category limits how much heat the plant can actually deliver. At a given supply temperature setpoint, the hydronic distribution system can only carry a certain amount of thermal power to the building. Above this limit, the plant simply cannot produce more heat regardless of how high the setpoint is raised. The controller is therefore bounded from above on its heat delivery at every hour of the horizon.

The third category links the heat flows to compressor electricity through the COP. This is what connects the physical heat delivery to the economic cost in the objective function. A higher heat delivery requires more compressor work, and the COP determines exactly how much. Without this linkage, the controller could deliver heat at zero electricity cost, which is physically impossible.

The fourth category enforces the comfort band. The indoor temperature is allowed to move freely within the band without penalty. If the predicted temperature exceeds the upper limit or falls below the lower limit, the corresponding slack variable becomes positive and a quadratic penalty is added to the cost. This formulation ensures the solver always finds a valid solution, even under conditions where strict comfort maintenance would be physically impossible.

6.4 Solver choice

The implementation uses CVXPY (Diamond and Boyd 2016) as the modelling layer, which separates the problem specification from the choice of underlying solver. The initial implementation used OSQP (Stellato et al. 2020) as the backend, but it failed to converge within the iteration limit on a non-trivial fraction of the simulated hours. The failures clustered on hours where the dynamics constraint became poorly scaled because the thermal-mass coefficients are of order 10^{-3}, reflecting the slow response of the building envelope, while the cost terms are of order 10^{0} or larger. The resulting ill-conditioned KKT system exceeded OSQP's adaptive scaling capacity.

The implementation therefore switched to CLARABEL (Goulart and Chen 2024), an interior-point method that handles ill-conditioned problems robustly through a homogeneous embedding. CLARABEL solved the same problems without retuning. OSQP was retained as a fallback, and in the rare cases where both solvers fail, the controller falls back to the legacy weather-curve setpoint for that hour. Over the full-year simulation, CLARABEL succeeded at the first attempt in more than 99% of cases, the remaining hours were resolved by OSQP, and the legacy fallback was never invoked. The full-year simulation completes in approximately 200 seconds on a standard workstation, confirming that the controller is suitable for real-time deployment.

6.5 Receding-horizon execution

The optimisation returns a 24-hour setpoint trajectory, but only the first value is applied. At the next hour, a new optimisation is solved using updated measurements and forecasts, and the process repeats. This receding-horizon scheme allows the controller to correct for forecast errors and unmodelled disturbances at every step (Rawlings et al. 2012; Oldewurtel et al. 2012).

The control loop at every hour t consists of five steps:

1. Read the current indoor temperature from the BMS. This value initialises the thermal state of the RC model, which resets the trajectory at every step and prevents the accumulation of

prediction errors over the horizon.

2. Build the 24-hour forecast vectors for outdoor temperature, solar radiation, occupancy, miscellaneous internal heat gains, and the condenser and evaporator temperatures used in the COP calculation.

3. Pass the current state, forecasts and price trajectory to the solver, which returns the optimal 24-hour setpoint vector $u^* = (u_0^*, u_1^*, \ldots, u_{23}^*)$ along with the corresponding heat flows, compressor energies and slacks.

4. Apply the first setpoint u_0^*. In simulation this drives the digital twin forward by one hour; in deployment it would be written to the BMS setpoint register at the start of hour t.

5. Store u_0^* as u_{prev} for the next iteration so that the input rate penalty in the next optimisation references the setpoint actually applied.

6.5.1 Forecast assumptions

The Nord Pool day-ahead market publishes hourly prices for the upcoming day between 13:00 and 14:00 the day before delivery. The controller therefore has access to true prices for 11 to 35 hours ahead, depending on the current time of day. Beyond the published window, a persistence forecast is used: tomorrow's prices are assumed to equal today's at the corresponding hours.

The controller is given perfect foresight of outdoor temperature, solar radiation, occupancy and miscellaneous heat loads over the horizon. This is the standard methodological choice in building MPC evaluations where the question of interest is economic performance under ideal information, not robustness to forecast errors (Oldewurtel et al. 2012; Drgoňa et al. 2020). The discussion chapter quantifies the expected impact of imperfect forecasts on the headline savings.

The condenser- and evaporator-side temperatures driving the COP calculation are forecast using a regime-conditional monthly median computed from historical data. These quantities are slow-varying and dominated by seasonal effects, so a monthly median is sufficient for the 24-hour horizon.

6.5.2 MPC simulation in the digital twin

The MPC is evaluated over a one year simulation inside the digital twin. The twin accepts the setpoint command at each hour, simulates the resulting thermal response and compressor energy, and provides the next indoor temperature observation. The simulation runs in closed loop meaning that the controller sees the twin's own simulated temperature at hour $t + 1$, not the historical measurement, when planning the next optimisation. This ensures that the controller's earlier decisions affect the trajectory it must manage, which ensures that the controller operates under the same conditions it would face in a real deployment.

Part III

Results

Chapter 7

Modelling

This chapter presents the validation results for the prediction model developed in Chapter 5. It evaluates the 1-step and 24-step indoor temperature predictions for both the 1R1C and 4R3C structures against the held-out 2025 test year, and validates the compressor energy model against measured consumption. The relative performance of the two model structures is assessed across all validation metrics.

7.1 Thermal System Identification

7.1.1 Model Comparison

Prior to parameter analysis, an evaluation process was conducted to analyze the primary candidate grey-box structures. The models; a 1-state lumped capacitance model (1R1C), a 2-state 2R2C model, a 3R2C structure and a 4R3C structure were evaluated against each other. The performance metrics generated during the system identification phase, including the Negative Log-Likelihood (NLL), Akaike Information Criterion (AIC), and Bayesian Information Criterion (BIC), are summarized in Table 3.

Model	States	Params	NLL	AIC	BIC
1R1C	1	6	-10696.99	-21381.97	-21339.49
4R3C	3	12	-9163.33	-18302.66	-18217.70
2R2C	2	8	-6058.00	-12100.00	-12043.36
3R2C	2	9	-6017.02	-12016.03	-11952.31

Table 3: System identification performance metrics across the evaluated grey-box RC models.

The numerical metrics show that the simplest structure, the 1-state 1R1C model, achieved the lowest NLL, AIC, and BIC scores. This demonstrates that a single thermal node provides a highly efficient statistical fit of the continuous dataset. A diagnostic parameter check showed that the internal thermal mass parameter (C_i) for the 1R1C model reached its upper physical bound of 5000.0 kWh/K. This reflects the single-state structure aggregating the total thermal inertia of the building envelope into one node.

While the 1R1C model provides a very strong statistical baseline, the 4R3C model offers a more explicit physical representation of the individual thermal layers. The statistical metrics alone do not provide a completely conclusive argument for selecting one model over the other. However, for the main purpose of this project, the exact mathematical superiority of the grey-box model selection remains secondary. The primary interest is the operational evaluation and savings potential of the downstream MPC framework. Both the 1R1C and 4R3C structures present competitive foundations for control.

7.1.2 Parameter Identification Results

The parameter optimization solver converged successfully for the evaluated structures, yielding physically plausible parameters that align with the heavy thermal characteristics of the property. The identified values for both the 1-state 1R1C model and the 3-state 4R3C model are summarized in Table 4.

Table 4: Identified physical parameters for the 1R1C and 4R3C grey-box models.

Parameter Description	**Symbol**	**1R1C Value**	**4R3C Value**	**Unit**
Internal Thermal Mass	C_i	5000.00	993.44	kWh/K
Structural Surface Mass	C_w	—	2589.17	kWh/K
Deep Core Mass	C_m	—	2483.06	kWh/K
Internal Surface Resistance	R_{iw}	—	0.0005	K/kW
Structural-Ambient Resistance	R_{wa}	—	0.3640	K/kW
Direct Ambient Resistance	R_{ia}	0.0838	0.1115	K/kW
Deep Coupling Resistance	R_{im}	—	0.1000	K/kW
Transparent Solar Area	A_w	825.42	361.38	m^2
Opaque Solar Area	A_e	—	311.13	m^2
Internal Heat Gains	Q_{int}	143.59	239.62	kW

The parameter distributions show a clear relationship between how the two architectures interpret the building envelope. For the 1R1C model, the internal thermal mass (C_i) reaches its upper predefined boundary at 5000.00 kWh/K. In the 4R3C architecture, this capacity is distributed between the internal air layer, the structural surface mass (C_w), and the deep core mass (C_m), which sum up to a total capacity of approximately 6065 kWh/K. This shows that both models recognize a very high thermal storage potential in the heavy concrete structure.

The steady state parameters also show consistent results. The direct ambient resistance (R_{ia}) is identified at 0.0838 K/kW for the 1R1C model and 0.1115 K/kW for the 4R3C model. This represents a lower overall thermal resistance for the single-state structure, which compensates for the fact that it has no alternative parallel heat loss paths through structural nodes. The transparent solar area (A_w) is larger in the 1R1C model at 825.42 square meters compared to 361.38 square meters in the 4R3C model, while the internal heat gains (Q_{int}) settle at a baseline of 143.59 kW for 1R1C and 239.62 kW for 4R3C. These structural parameters, including the internal surface resistance (R_{iw}) at 0.0005 K/kW, the structural-ambient resistance (R_{wa}) at 0.3640 K/kW, the deep coupling resistance (R_{im}) at 0.1000 K/kW, and the opaque solar area (A_e) at 311.13 square meters, are exclusive to the 4R3C architecture. They provide the multi-state model with the capacity to explicitly map the transient heat path from the outdoor environment through the external shell and into the deep concrete core. These differences demonstrate how the optimization algorithm balances solar driving forces and constant internal gains depending on the available degrees of freedom in the model equations. Ultimately, both sets of parameters yield a stable and physically reasonable representation of the overall heat transfer dynamics.

7.1.3 Model Validation

To evaluate the predictive performance of the grey-box models before closed-loop implementation, validation tests were conducted using an unseen out-of-sample dataset from the full year of 2025.

To visually demonstrate the baseline tracking capability, a representative 14-day winter period was extracted from the testing dataset. Figure 1 illustrates the rolling 24-hour ahead indoor temperature forecasts generated by all evaluated models against the actual measured temperature.

The trajectory plot reveals that both the 1-state 1R1C model and the 3-state 4R3C model track the true indoor temperature profile closely. They capture the thermal dynamics and the general behavior of the building well. In contrast, the two-state models exhibit clear amplitude deviations and struggle to maintain the correct thermal trajectory during transient changes.

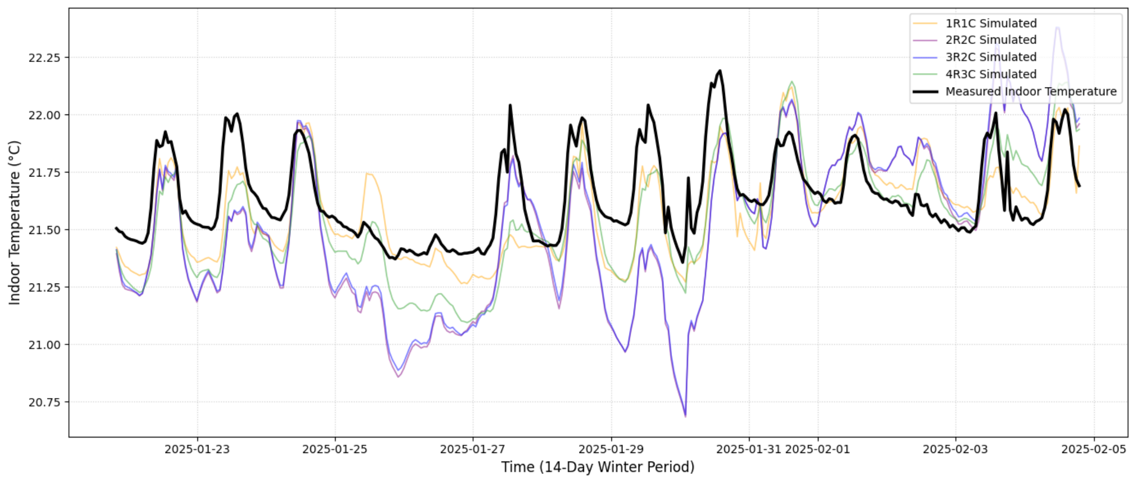

Figure 1: Simulated vs. measured indoor temperature over a 14-day winter period showing the rolling 24-hour ahead predictions of the evaluated models.

To evaluate the aggregate performance across the entire of the validation year, Figure 2 displays the distribution of the absolute prediction errors for the rolling 24-hour horizon.

The error distribution confirms that the 1R1C and 4R3C models perform significantly better than the multi-state baselines. Across the full test year, the 4R3C model maintains a highly stable performance with a total 24-hour ahead RMSE of 0.3928 degrees Celsius. The 1-state 1R1C model achieves a slightly better distribution with a lower RMSE of 0.2874 degrees Celsius and a lower variance. This low variance occurs because its internal thermal mass parameter was pinned at the upper boundary during identification, meaning the 1R1C model acts as a large thermal filter that smooths out short-term fluctuations. Both the 1R1C and 4R3C structures demonstrate acceptable predictive accuracy for control applications, whereas the 2R2C and 3R2C models exhibit much larger variances with a 24-hour RMSE exceeding 1.03 degrees Celsius.

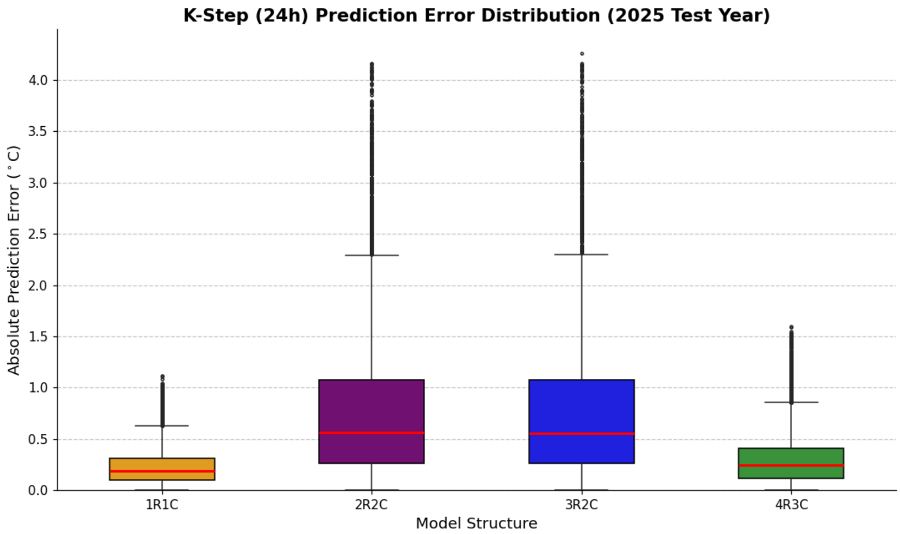

Figure 2: Absolute prediction error distribution for a 24-hour horizon across the 2025 test year.

Dynamic Synchronization and Phase-Lag Analysis

Verifying temporal synchronization is necessary for real-time model predictive control because the controller relies on accurate timing to execute load shifting before price peaks occur. A phase-lag and cross-correlation (CC) analysis was conducted on the 24-step ahead predictions across the validation dataset. The results are summarized in Table 5.

Table 5: Dynamic synchronization and phase-lag metrics for the 24-hour prediction horizon.

Model	CC Lag (h)	CC Peak Corr	FFT Lag (h)	RMSE Amp	RMSE Phase
1R1C	1	0.840	-0.532	9.898	15.827
4R3C	1	0.732	-0.967	16.041	19.111
2R2C	0	0.130	-0.048	36.492	45.773
3R2C	0	0.114	-0.070	36.016	46.104

The synchronization metrics confirm that the two-state structures, 2R2C and 3R2C, fail to capture the dynamic trajectories, resulting in very low cross-correlation coefficients around 0.12 and high RMSE.

In contrast, both the 1R1C model and the 4R3C model demonstrate strong temporal alignment with the real building dynamics. Both models share a cross-correlation peak lag of 1 hour. The single-state 1R1C model achieves the highest peak correlation coefficient of 0.840 and the lowest overall amplitude and phase errors. The 4R3C model also maintains a high correlation of 0.732.

The Fast Fourier Transform analysis (FFT) shows a phase shift of -0.532 hours for the 1R1C model and -0.967 hours for the 4R3C model. This difference is consistent with the physical structures of the models. The 4R3C model includes a decoupled deep concrete core which introduces a larger physical time lag into the mathematical structure, whereas the lumped 1R1C model responds more uniformly. These results reinforce the conclusion that both the 1R1C and 4R3C models represent high-quality options for the predictive control framework, as both successfully capture the essential daily thermal cycles of the building.

Residual Diagnostics

A residual analysis was conducted on the 1-step prediction errors across the validation dataset to evaluate the presence of unmodeled dynamics. Figure 3 illustrates the residual autocorrelation function correlogram for all four evaluated models over a 72-hour lag horizon.

The plots show that all four models exhibit a highly consistent pattern. The residuals display a slow decay and remain outside the 95 percent confidence interval across the entire evaluated period. This persistent positive autocorrelation indicates the presence of unmodeled low-frequency dynamics or autocorrelated disturbances, such as persistent weather phenomena, complex internal heat gains, or localized thermal bridging across the large building structure.

While the overall shape of the correlogram is similar for all configurations, the 1R1C model and the 4R3C model demonstrate slightly better performance than the two-state baselines, the former being the best. The autocorrelation values drop more rapidly in these two structures, with the 1-state 1R1C model reaching the lowest correlation values at longer lags, followed by the 4R3C model. This high level of similarity reinforces the conclusion that both models capture the dominant thermal trends of the property in a comparable manner. Since the absolute prediction errors remain tightly bounded, both structures remain reliable candidates for the control framework.

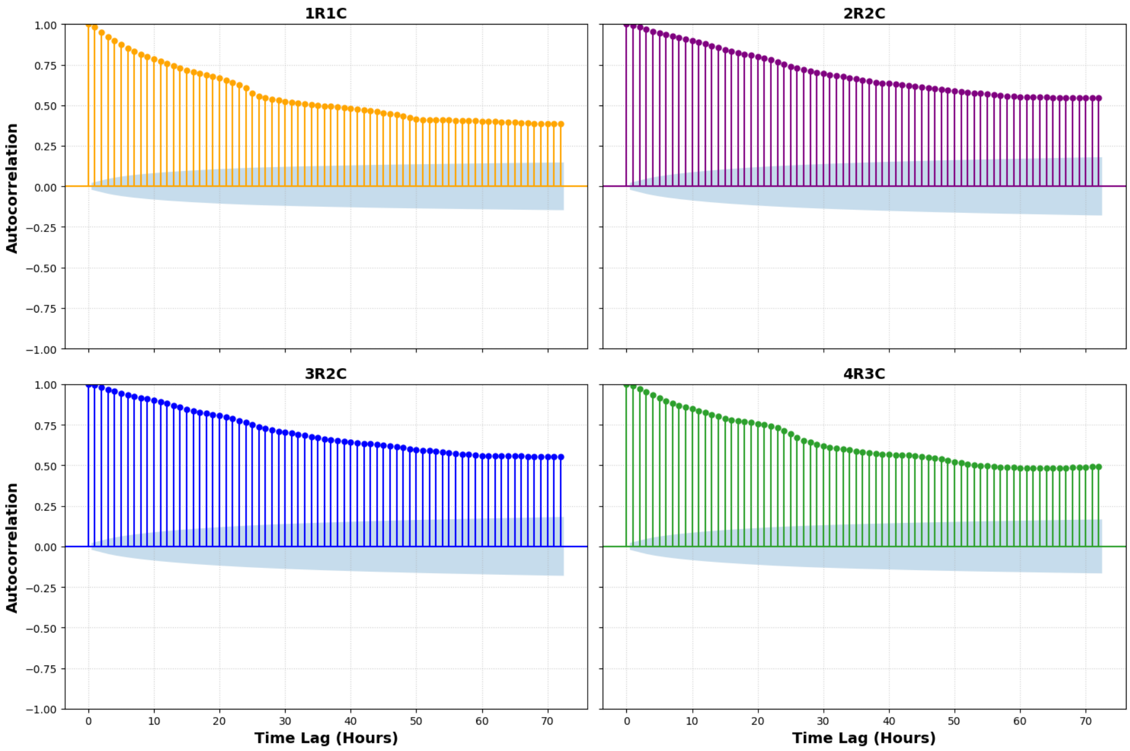

Figure 3: Residual autocorrelation function correlogram for all evaluated grey-box models over a 72-hour lag horizon.

To evaluate the directional bias and spread of the 24-hour prediction errors, the distribution of the raw residuals was analyzed. Figure 4 illustrates the error histograms and density estimations for all four evaluated models across the validation year.

The plots demonstrate a clear performance gap between the model structures. The two-state baseline models, 2R2C and 3R2C, exhibit wide distributions with large variances, where the prediction errors stretch from minus 2 degrees to over 4 degrees Celsius. This wide spread indicates frequent and substantial deviations from the true indoor temperature.

In contrast, the 1-state 1R1C model and the 3-state 4R3C model deliver highly concentrated distributions centered tightly around the zero-error line. The 1R1C model exhibits the narrowest and most symmetrical spread, while the 4R3C model displays a tight range with a minor right-skewed asymmetry. This high concentration confirms that both the 1R1C and 4R3C architectures maintain bounded and reliable prediction errors throughout the 24-hour horizon, reinforcing their shared suitability for the MPC.

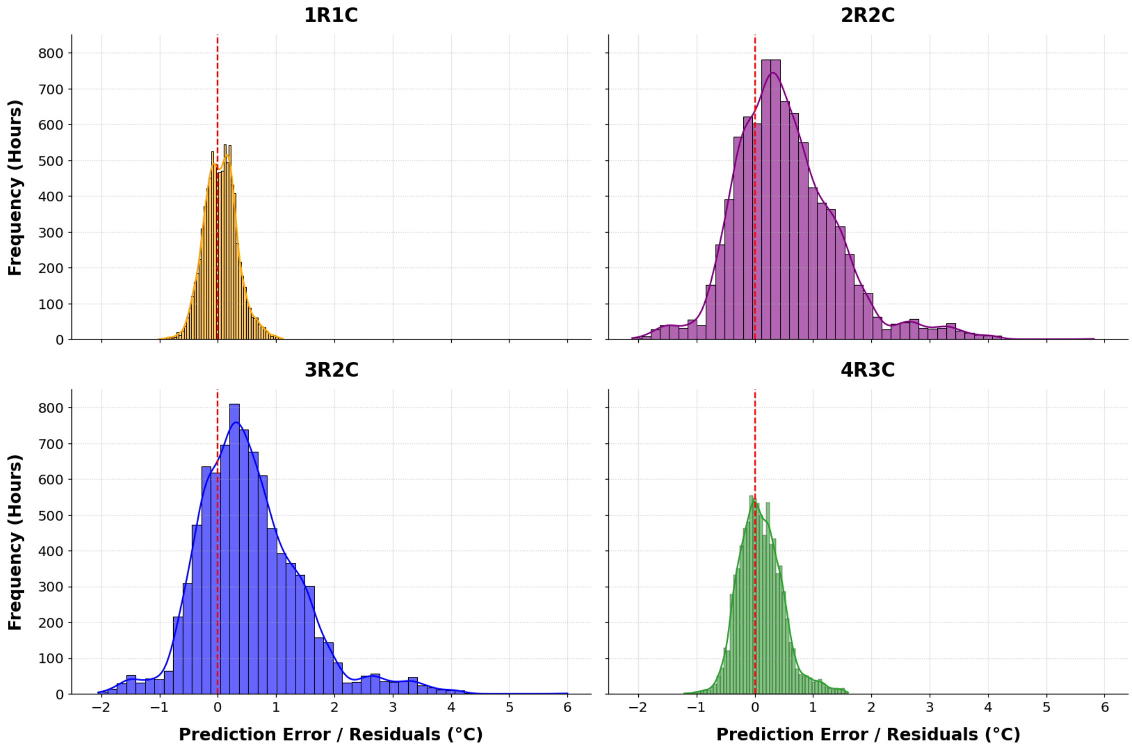

Figure 4: Distribution of raw 24-hour prediction errors for all evaluated model structures across the full validation year.

7.2 The Digital Twin

7.2.1 Indoor Temperature Prediction

The thermal sub-model's ability to track the indoor temperature was assessed using a closed-loop 1-step prediction, reflecting the operational mode of the MPC. Two RC model structures were evaluated in parallel: the three-node 4R3C formulation, which explicitly resolves wall and structural mass dynamics, and the simpler single-zone 1R1C formulation. The results for the test year 2025 are summarised in Table 6.

Table 6: Validation metrics for T_{room} (1-step prediction) for the 2025 test period.

Metric	Unit	4R3C	1R1C
RMSE (overall)	°C	0.142	0.098
MBE (overall)	°C	−0.059	−0.008
RMSE (summer)	°C	0.195	0.119
MBE (summer)	°C	−0.144	−0.005

Both models demonstrate high precision in 1-step prediction. The 4R3C formulation yields an RMSE of 0.142 °C with an MBE of −0.059 °C, indicating no systematic over- or under-estimation. The 1R1C formulation achieves a marginally lower RMSE of 0.098 °C and a near-zero MBE of −0.008 °C. While the metrics favour 1R1C, the two structures capture fundamentally different physical representations of the building: 4R3C explicitly models the thermal inertia of the envelope and structural mass through separate state nodes, whereas 1R1C lumps all dynamics into a single interior air node.

The performance is visually presented in Figures 5 and 6, contrasting the 1-step and 24-step predictions across representative months for each season. In both models, the 1-step prediction displays an exceptional fit to the measured data, while the 24-step prediction demonstrates model stability over a full planning horizon. Although the residuals naturally widen as the horizon increases, neither model exhibits significant systematic drift.

Looking at the seasonal breakdown, both models display a degradation in prediction accuracy during the summer months, as reflected in Table 6 and the Summer subplots in Figures 5 and 6. This increased variance is likely attributable to the higher complexity of modelling cooling loads, where unmeasured internal heat gains from solar radiation and occupant behaviour are more pronounced than during the heating season. The effect is present in both structures, though it is slightly reduced in the 1R1C model, as reflected by a summer RMSE of 0.119 °C compared to 0.195 °C for 4R3C. For the MPC, this implies a higher degree of uncertainty when planning cooling sequences.

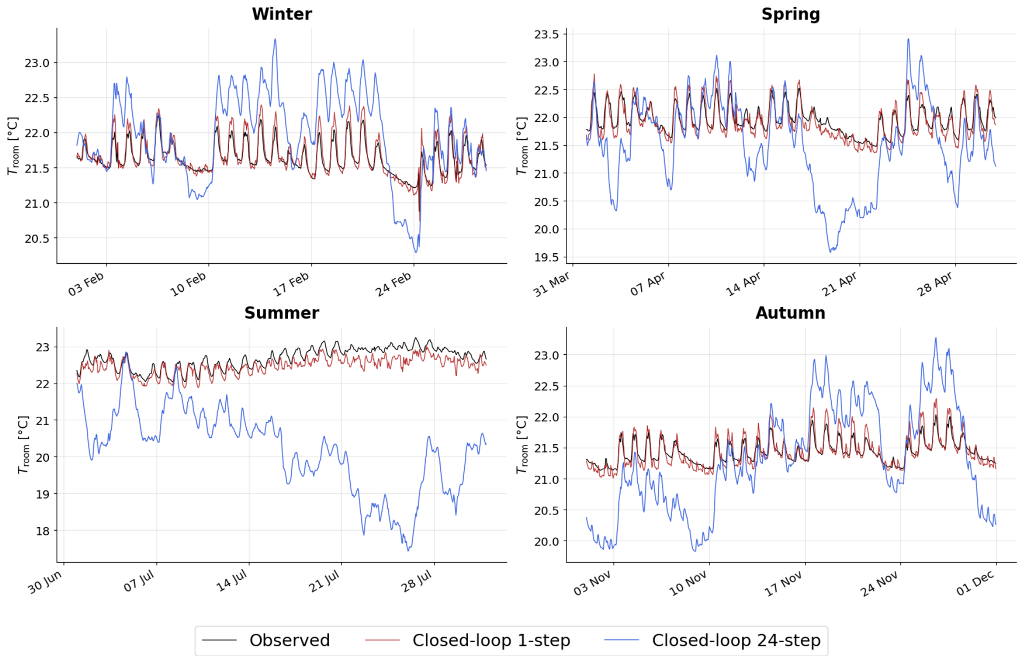

Figure 5: Validation of indoor temperature (T_{room}) for the 2025 test period, 4R3C model: 1-step and 24-step closed-loop predictions.

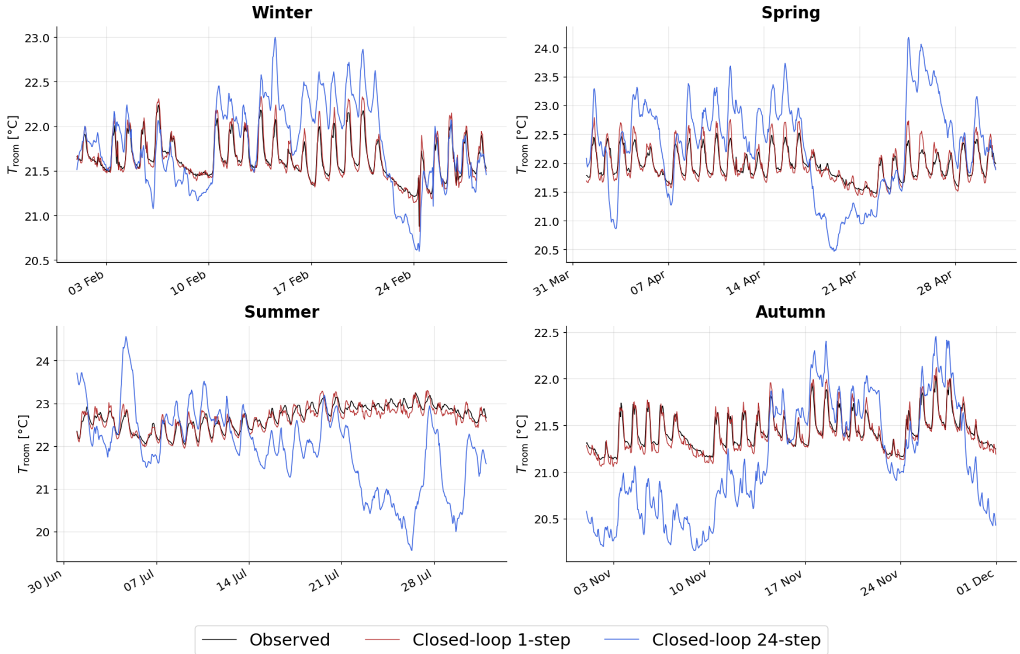

Figure 6: Validation of indoor temperature (T_{room}) for the 2025 test period, 1R1C model: 1-step and 24-step closed-loop predictions.

7.2.2 Compressor Energy Consumption

Unlike the indoor temperature prediction, the compressor energy model does not depend on the choice of RC structure. It operates directly on measured heat pump loads and temperatures, meaning the validation results are identical for both the 4R3C and 1R1C formulations and are therefore reported only once below.

The energy sub-model was validated by comparing the simulated compressor power (E_{comp}) with the measured electrical power. The aggregate energy results for 2025 are presented in Table 7. The annual energy ratio of 0.973 demonstrates that the digital twin captures the overall magnitude of consumption with high accuracy, while the RMSE of 18.6 kW reflects an acceptable tracking of the plant's highly dynamic hourly power loads.

Table 7: Compressor energy validation metrics for the 2025 test period.

Metric	Value	Unit
Annual Energy Ratio ($E_{\text{sim}}/E_{\text{meas}}$)	0.973	-
RMSE (Power)	18.6	kW

The accuracy of the electrical load prediction is illustrated in Figure 7. The model effectively tracks the operational cycles and peak loads of the heat pump plant. Notably, even at a 24-step horizon, the digital twin maintains a high correlation with the measured small power spikes, which is key for the MPC to accurately evaluate price-shifting opportunities and power tariff constraints. However, the sudden and unpredicted power spikes pose a significant challenge, as these anomalies could affect the MPC cost evaluation. An unexpected power surge during an assumed "cheap" hour may increase the actual operating cost, thereby lowering the financial efficiency of the load-shifting strategy.

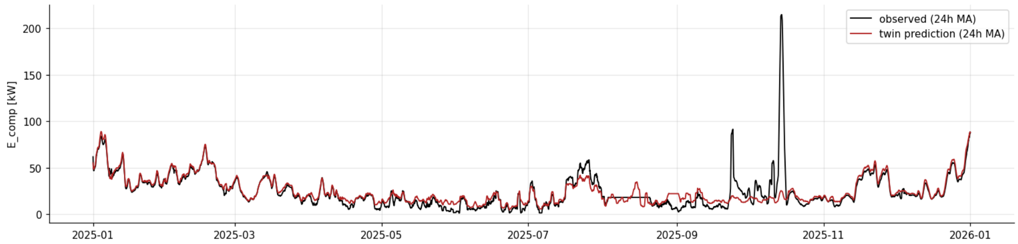

Figure 7: Validation of compressor energy consumption (E_{comp}) for the 2025 test period.

Chapter 8

MPC Performance

This chapter presents the controller evaluation results over the full 2025 simulation year. Both MPC formulations, based on the 1R1C and 4R3C thermal models, are evaluated against the legacy weather-compensation curve across four dimensions. These cover overall energy consumption and operating cost, adherence to the thermal comfort band, the ability to shift compressor load toward cheaper hours, and the seasonal variation in controller behaviour. Where the results require immediate interpretation, this is provided alongside the figures rather than deferred to the discussion chapter.

8.1 Overall Economic and Energy Performance

The results show that the economic model predictive control significantly reduces both energy use and operating costs compared to the legacy system. Table 8 presents the total energy consumption and cost across all four seasons.

Over the full year, both controllers achieve near identical economic returns. The 4R3C controller reduced energy consumption by 57.7 % and total cost by 72.6 %, while the 1R1C controller achieved a 56.0 % energy reduction and a 72.4 % cost reduction. These figures represent an upper bound under idealized simulation conditions and both controllers are evaluated within the same digital twin, not against measured historical consumption.. They demonstrate a theoretical maximum potential rather than a guaranteed real-world outcome.

The performance varies depending on the season. Winter presents the highest energy demand and the most volatile electricity prices. During winter, the 4R3C controller saved 58.8 % in costs alongside a 36.0 % reduction in energy use. The 1R1C controller achieved a 56.9 % cost reduction with a more moderate 25.7 % energy reduction. This difference reflects the 4R3C model's explicit representation of the deep structural mass, which allows the controller to store more thermal energy during cheap night hours.

Spring and autumn show large reductions in both energy and cost for both controllers. During summer, the savings approach 100 %. This occurs because both MPC formulations actively utilise the passive free cooling from the borehole system instead of running the mechanical compressors.

Table 8: Seasonal performance summary comparing the legacy system to the 1R1C and 4R3C MPC controllers.

	Legacy		**1R1C**		**4R3C**	
Season	Energy (MWh)	Cost (kSEK)	Energy (MWh)	Cost (kSEK)	Energy (MWh)	Cost (kSEK)
Winter	84.9	59.0	63.1 (−25.7%)	25.5 (−56.9%)	54.3 (−36.0%)	24.3 (−58.8%)
Spring	38.5	21.2	9.6 (−75.0%)	3.0 (−85.8%)	17.0 (−55.7%)	4.7 (−77.8%)
Summer	37.6	14.9	0.0 (−100%)	0.0 (−100%)	0.3 (−99.1%)	0.1 (−99.7%)
Autumn	55.6	47.8	22.5 (−59.5%)	10.9 (−77.1%)	20.0 (−64.1%)	10.1 (−78.8%)
Full Year	216.6	142.9	95.2 (−56.0%)	39.4 (−72.4%)	91.7 (−57.7%)	39.2 (−72.6%)

However, the summer figures require a closer examination. Both controllers completely deactivate the mechanical compressors during the summer months, relying entirely on passive free cooling from the borehole system. While this is physically plausible for limited periods, it is unlikely to hold across an entire season in practice. To provide a more conservative estimate, the historical legacy summer consumption (37.6 MWh and 14.9 kSEK) is reinstated into the MPC totals. Under this worst-case summer scenario, the annual adjusted savings remain substantial. The 1R1C controller yields a 38.7 % reduction in energy (132.8 MWh) and a 62.0 % reduction in total operating costs (54.3 kSEK), while the 4R3C controller yields a 40.4 % energy reduction (129.3 MWh) and a 62.1 % cost reduction (54.1 kSEK).

8.2 Thermal Comfort and Boundary Adherence

Another aspect of the controller evaluation concerns thermal comfort and compliance with the temperature boundaries. Figure 8 shows the indoor temperature duration curve for the full year, where all operating hours are sorted from the highest to the lowest temperature.

All controllers manage to keep the indoor temperature tightly within the required boundaries for almost the entire year. The sorted curves for the models flatten out within the comfort zone. The minor deviations at the extreme ends of the curves represent rare hours where the soft constraints allowed very small temperature shifts to prevent massive power peaks or solver issues. Notably, the 4R3C controller operates consistently towards the lower end of the comfort band compared to 1R1C, which partially explains its higher energy reduction during winter. These results prove that the extensive economic savings described in Section 8.1 were achieved while maintaining or even improving the indoor thermal climate.

8.3 Proof of Economic Load Shifting

The core mechanism behind the cost savings is the ability to move energy consumption to cheaper hours. Figure 9 illustrates the mean compressor power as a function of the spot price percentile over the entire year. The horizontal axis ranks all hours from the cheapest (0) to the most expensive (100).

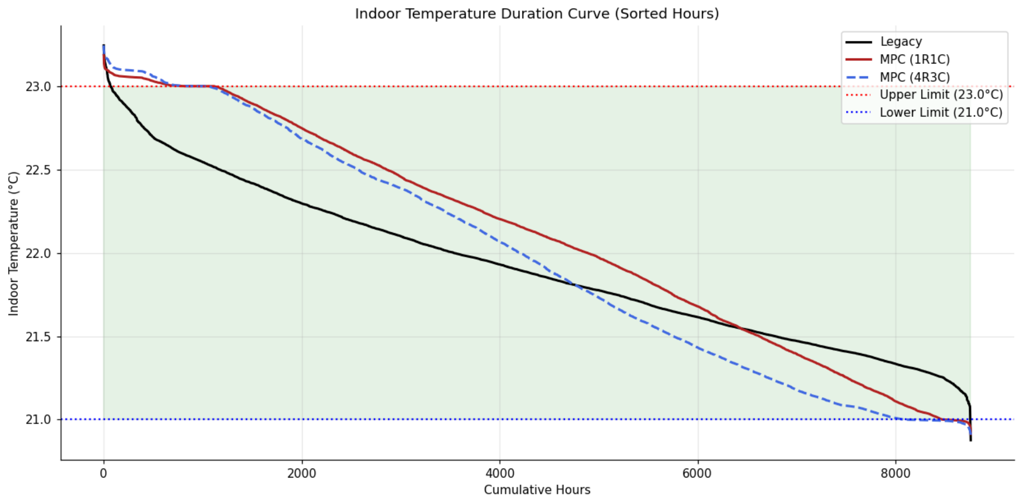

Figure 8: Indoor temperature duration curve sorting all operating hours from highest to lowest for the legacy system, MPC with 1R1C, and MPC with 4R3C.

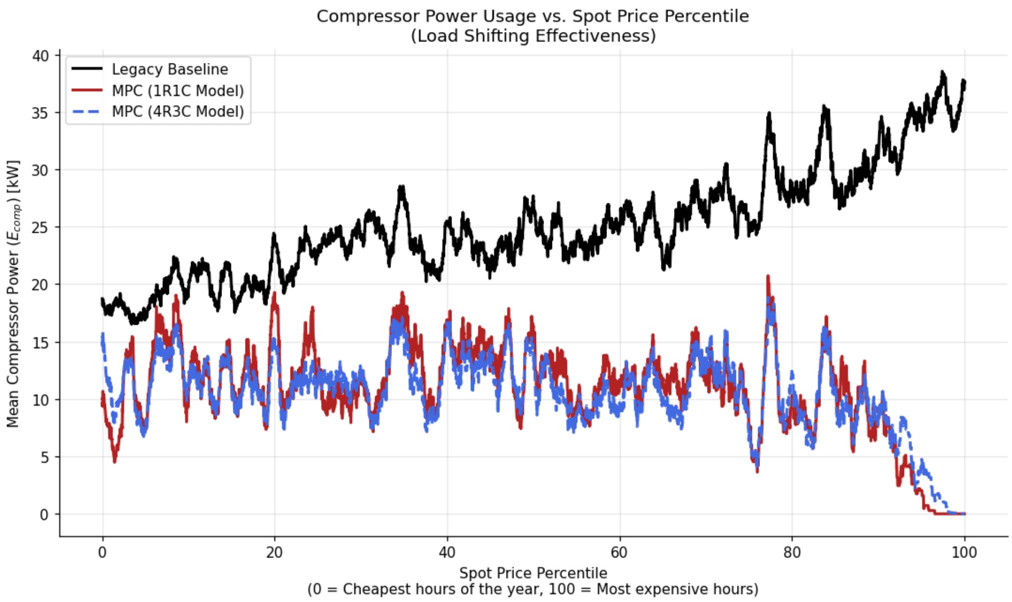

Figure 9: Mean compressor power plotted against the electricity spot price percentile. A percentile of zero represents the cheapest hours of the year, while 100 represents the most expensive.

The legacy system distributes its power usage relatively evenly across all price percentiles and even slightly more usage of more expensive hours. It reacts only to the current weather and indoor conditions and completely ignores the cost of electricity.

Both predictive models display a very similar behavior. As the price percentile increases, the compressor power drops steadily for both controllers. During the most expensive 10 % of the hours, both controllers operate the compressors very close to zero. Notably, the 1R1C and 4R3C formulations produce near identical load shifting profiles throughout the entire price range, confirming that the economic behavior of the MPC is driven by the objective function rather than the choice of thermal model. This graph confirms that both algorithms successfully transform the building into a thermal battery, charging it when electricity is cheap and avoiding compressor operation during peak price hours.

8.4 Dynamic System Behavior and Seasonal Strategies

To understand how the predictive controllers operate on a shorter time horizon, Figure 10 shows a detailed three-day simulation window for each season. This visualization contrasts the reactive nature of the legacy system with the proactive behavior of the MPC.

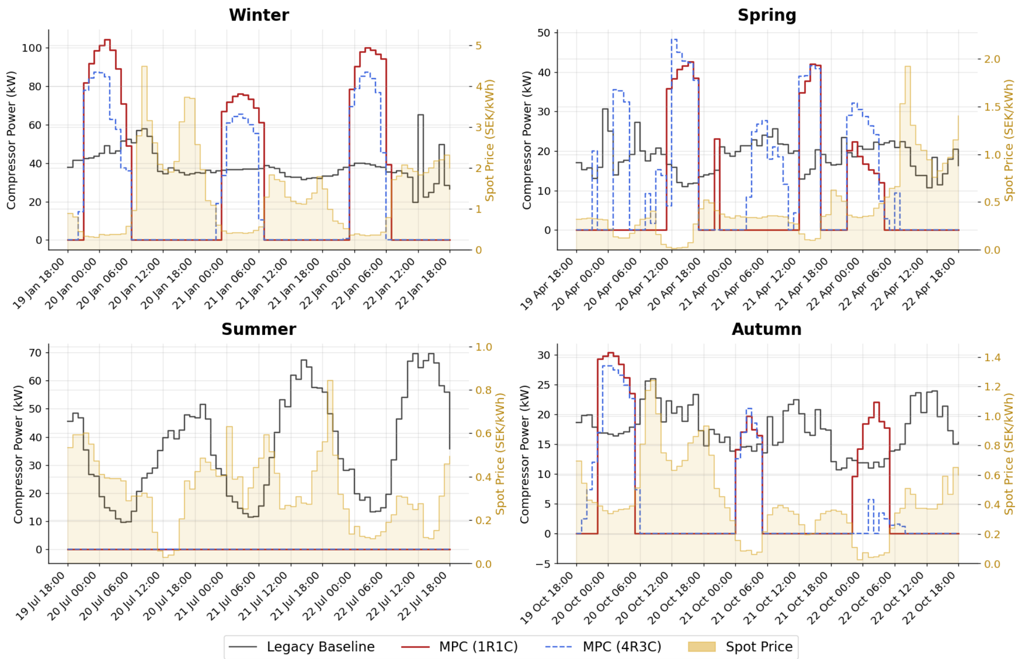

Figure 10: A three-day zoomed analysis of compressor power consumption and electricity spot prices across all four seasons.

During winter, spring, and autumn, both the 1R1C and 4R3C controllers demonstrate clear load shifting capabilities. When electricity prices are low, typically during the night, the predictive models increase the compressor power to preheat the building structure. When the morning price peaks occur, the controllers reduce the compressor power to zero and allow the building to utilize the stored thermal energy. The legacy system follows a completely reactive pattern, tracking the outdoor temperature variations and often peaking directly during the most expensive hours. Notably, both MPC formulations produce near identical load shifting behaviour across all heating and cooling seasons.

The summer profile reveals a unique operating strategy. During these 72 hours, the compressor power for both MPC models remains at zero. The predictive controllers completely deactivate the mechanical cooling compressors and instead fulfil the indoor comfort requirements using passive free cooling from the borehole system. This behaviour explains why the summer energy and cost savings described in the previous sections approach 100 %.

The underlying supply water temperature setpoint behaviour is illustrated in Figure 11 for a representative winter week. The legacy baseline remains nearly flat throughout, reflecting the weather compensation curve. Both MPC formulations instead show a clear diurnal pattern, driving the setpoint aggressively upward during cheap night hours to precharge the structural mass, then allowing it to drop well below the legacy level during the day. The 1R1C controller produces sharper and higher overnight peaks, while the 4R3C controller exhibits a more damped profile, consistent with its explicit representation of the deep structural mass retaining heat over longer time scales.

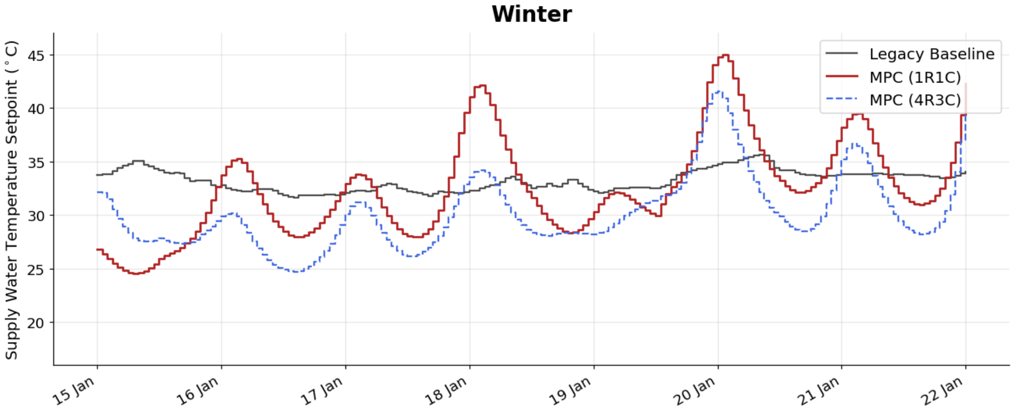

Figure 11: Supply water temperature setpoint (v_k) over a representative winter week for the legacy baseline, 1R1C, and 4R3C controllers.

Part IV

Interpretation

Chapter 9

Discussion

This chapter interprets the results in three parts. It positions the headline figures against comparable studies in the building MPC literature, examines where the simulation assumptions favour the controller and how large those effects are, and discusses what the results mean for a practical deployment at this and similar properties.

9.1 The numbers in context

The winter comparison between the two MPC formulations points to something interesting about the savings. The 1R1C controller consumes more compressor energy than the 4R3C controller over the heating season (63.1 MWh versus 54.3 MWh), yet both achieve nearly identical cost reductions. The dominant driver of financial performance is therefore not how much electricity is consumed, but when it is consumed. A controller that concentrates compressor operation in the cheapest hours captures savings that are largely independent of total energy use, and this distinction matters for how the figures should be read.

Those figures, a 39 % reduction in compressor energy and a 62 % reduction in operating cost, sit at the high end of what the building MPC literature reports. Široký et al. (2011) achieved 15 to 28 % energy savings experimentally using a radiant ceiling system, and Behzadi (2025) reports a 13 % cost reduction at a comparable property in Uppsala. Two primary factors explain why the present results exceed both. The studied plant is a reversible system, which provides more degrees of freedom than the heating-only configurations common in the literature. Simultaneous heating and cooling recovery operates during approximately 84 % of annual hours, meaning savings compound across both sides of the refrigeration cycle. The SE3 spot prices during 2025 also showed higher volatility than the years covered by the older literature, which directly amplifies the value of any load-shifting strategy. These factors together provide a plausible explanation, but they do not fully resolve the gap. A poorly tuned legacy baseline may inflate apparent savings regardless of how well the MPC performs.

The architecture of the legacy controller also contributes to the gap. A weather-compensation curve must be set conservatively enough to guarantee comfort under the worst conditions it will ever face, which means it routinely delivers more heat than the current hour actually requires. The MPC faces no such constraint. Rather than a single setpoint for each outdoor temperature, it selects a full 24-hour trajectory across multiple dimensions simultaneously, and its advantage comes precisely from dropping below the conservative level whenever the building has stored enough heat and prices are high. Some part of the apparent saving therefore reflects this structural difference rather than the quality of the optimisation, though the two contributions cannot be separated cleanly from the available data.

During summer, the supply water setpoint loses effective authority over the plant state, and the near zero compressor consumption the MPC achieves in those months reflects an unmodelled operating mode rather than genuine optimisation. The same gradual loss of authority extends into the neighboring seasons, and its precise reach into the spring and autumn figures is difficult to establish from the available data. The controller also operated with perfect weather foresight throughout the simulation, whereas real forecasts carry meaningful errors beyond six to twelve hours. Both effects act in the same direction: the headline figures represent an upper bound under idealised conditions.

One further observation from the results deserves to be mentioned. Despite their fundamentally different physical representations, the 1R1C and 4R3C controllers produce nearly identical load shifting results and economic outcomes across all seasons. This suggests that the financial performance of the MPC is driven primarily by the price signal and the total thermal inertia of the building, rather than by how that inertia is distributed across model states. For anyone considering MPC for a similar building, this is a useful finding as a simpler model structure may be sufficient to capture most of the economic benefit. Given that the 1R1C structure is also easier to identify and maintain in practice, it might represent a more pragmatic choice for deployment at this as well as other facilities with similar infrastructure. This conclusion should however be read with care. The comparison is made against a reactive legacy controller that carries no price information whatsoever. Against a stronger baseline, or under realistic forecast errors, the more accurate representation of multi-timescale thermal dynamics in the 4R3C structure may prove advantageous.

9.2 The simulation-reality gap

A recurring methodological issue in the building MPC literature is that simulated controllers are benchmarked against measured historical consumption from the real plant. A model based controller optimises over its own representation of the plant, and wherever that representation diverges from reality, the controller concentrates or avoids operation in ways that appear as savings when scored against a measured baseline. None of that necessarily reflects what the controller would achieve in deployment.

This study illustrates the effect with its own summer results. The unadjusted twin-on-twin comparison yields a cost reduction of approximately 72 % and an energy reduction of approximately 57 %. Once the summer months are treated more conservatively by reinstating the legacy consumption for a period where the supply water setpoint demonstrably loses authority over the plant, the figures fall to approximately 62 % in cost and 39 % in energy. The gap between the two scenarios arises from a single unmodelled operating mode. The twin-on-twin protocol at least ensures that both controllers are subject to the same modelling limitations, which is why the comparison is meaningful at all. The headline figures in this study should therefore be read as upper bounds under the specific conditions of the simulation.

9.3 Implications for practice

Even under conservative adjustment, the financial case for predictive control at a property of this scale is clear. A 62 % reduction in compressor operating cost represents a material saving for a building operator exposed to volatile day-ahead prices. The practical path to deployment is also shorter than the modelling effort in this thesis might suggest. The supply water setpoint is already a writable register in the building management system, and no re-commissioning of downstream equipment is required. The principal prerequisite is connecting a Nord Pool day-ahead price feed and a weather forecast to the controller input pipeline. What this thesis provides is the validated prediction model and the evidence that the economic optimisation performs as intended within the simulation environment.

However, if the thermal model underestimates the building's heat losses, the real indoor temperature may run lower than the controller predicts, even when the optimisation satisfies the comfort constraints. One way to address this in a field deployment would be to use a heat delivery scaling parameter as a tuning lever. Increasing it compensates for eventual heat losses that the model does not capture, while preserving the economic logic of the controller.

The methodology is not specific to this building. The combination of grey-box thermal identification and economic MPC can in principle be applied to any commercial building with a hydronic heating system and a writable supply temperature setpoint. The identification procedure requires only standard BMS trend data, and the model structures identified in this report suggests that the modelling effort is modest even for large properties. The savings potential will however vary. Buildings with lower thermal mass offer less storage capacity for load shifting, and plants without simultaneous heating and cooling recovery will see smaller gains from the same price signal. The results presented here should therefore be read as indicative of what is achievable under favourable conditions rather than as a general benchmark.

These requirements translate into a compact procedure for transferring the approach to other buildings. The first step identifies a low-order grey-box thermal model from historical records of heat and cooling delivery together with the corresponding weather data, where the single-state 1R1C structure examined in this report is sufficient to capture the dominant dynamics. The second step locates the writable setpoint that governs the heating plant and adopts it as the single decision variable. The third formulates the economic optimisation around that variable and connects it to a day-ahead price feed and a weather forecast. The final step calibrates the heat delivery scaling parameter against measured indoor temperatures, raising it if the building runs cold and lowering it if it runs warm, after which the realised cost reduction can be evaluated against the historical baseline.

Chapter 10

Conclusions and Future Work

This chapter answers the three research questions and outlines the most important directions for future work. These concern evaluating the controller under realistic forecast conditions, conducting a field trial during the heating season, and examining whether model complexity affects performance against a stronger baseline or under imperfect information.

10.1 Conclusion

This thesis set out to investigate the techno-economic potential of economic model predictive control for a reversible heat pump plant with borehole thermal energy storage in a large commercial office building in Stockholm. The three research questions are addressed in turn.

The first question asked how the thermal dynamics of a heavy-mass office building can be identified and represented for control purposes. Four grey-box RC structures were evaluated against two years of operational data. Both the simplest structure, 1R1C, and the most complex, 4R3C, produced accurate 24-hour temperature predictions with an RMSE below 0.4 ℃ on the held-out test year and were 0.1 ℃ within each other. The two intermediate structures failed to capture the building's thermal dynamics adequately. The identification procedure, combining differential evolution with a quasi-Newton polish, converged reliably and produced physically plausible parameters for both selected structures. Grey-box RC modelling is therefore a viable and practical approach for buildings of this type.

The second question asked to what extent MPC can reduce operational energy costs and how total energy consumption is affected. Under the twin-on-twin evaluation, the MPC reduced compressor energy by 39 % and operating cost by 62 % compared to the legacy weather-compensation curve. The dominant mechanism is temporal. The controller concentrates compressor operation during the cheapest hours and suppresses it during price peaks, capturing cost savings that are largely independent of total energy consumption.This distinction matters practically, as a building operator exposed to volatile day-ahead prices benefits more from shifting consumption than from reducing it.

The third question asked how the MPC affects indoor thermal comfort. Both controllers maintained the indoor temperature within the defined comfort band of 21 to 23 ℃ for almost the entire evaluation year. The minor deviations at the seasonal extremes were small and short-lived. The economic savings reported above were therefore achieved without any deterioration in thermal comfort.

Taken together, the results demonstrate that the studied building possesses substantial unused flexibility that a predictive controller can exploit. The methodological contribution of this work is the twin-on-twin evaluation protocol, which ensures that the comparison between controllers reflects a genuine algorithmic advantage rather than an artefact of model imperfection. That the headline saving survives the most conservative adjustment available, reinstating the full legacy summer consumption, supports the conclusion that the economic benefit is real. Quantifying it precisely under real operating conditions requires a field trial.

10.2 Future Work

Three directions follow naturally from the limitations of this study. The most immediate is to evaluate the controller under realistic forecast conditions. Throughout this simulation, the controller was given perfect knowledge of weather and electricity prices, which is not available in practice. Re-evaluating the controller under realistic forecast errors would establish how much of the simulated saving survives imperfect information and give a more grounded estimate of real-world performance.

The second direction is a field trial limited to the heating season, roughly October through April, where the supply water setpoint retains full authority over the plant. This would provide the first measured evidence of whether the simulated savings translate to the physical building, while avoiding the summer operating mode where the current implementation is known to be unreliable.

The third direction concerns model complexity. This study found no meaningful economic difference between the 1R1C and 4R3C controllers when evaluated against a reactive legacy baseline. Whether that equivalence holds against a stronger baseline or under realistic forecast errors remains an open question, and a controlled comparison under those conditions would clarify whether the additional complexity of the 4R3C structure is warranted in practice.

Bibliography

Akaike, Hirotugu (1974). "A new look at the statistical model identification". In: *IEEE Transactions on Automatic Control* 19.6, pp. 716–723.

Bacher, Peder and Henrik Madsen (2011). "Identifying suitable models for the heat dynamics of buildings". In: *Energy and Buildings* 43.7, pp. 1511–1522.

Behzadi, Amirmohammad (2025). "Smart Design, Control, and Optimization of Thermal Energy Storage in Low-Temperature Heating and High-Temperature Cooling Systems". PhD thesis. Stockholm, Sweden: KTH Royal Institute of Technology.

Berthou, Thomas et al. (2014). "Development and validation of a gray box model to predict thermal behavior of occupied office buildings". In: *Energy and Buildings* 74, pp. 91–100.

Boverket (2011). *Boverkets byggregler (BBR) – BFS 2011:6 med ändringar.* Tech. rep. Accessed incorporating current amendments for commercial building regulations. Karlskrona, Sweden: Boverket (Swedish National Board of Housing, Building and Planning).

Byrd, Richard H. et al. (1995). "A limited memory algorithm for bound constrained optimization". In: *SIAM Journal on Scientific Computing* 16.5, pp. 1190–1208.

Chua, K. J., S. K. Chou, and W. M. Yang (2010). "Advances in heat pump systems: A review". In: *Applied Energy* 87.12, pp. 3611–3624.

Diamond, Steven and Stephen Boyd (2016). "CVXPY: A Python-Embedded Modeling Language for Convex Optimization". In: *Journal of Machine Learning Research* 17.83, pp. 1–5.

Drgoňa, Ján et al. (2020). "All you need to know about model predictive control for buildings". In: *Annual Reviews in Control* 50, pp. 190–232.

Ellis, Matthew, Helen Durand, and Panagiotis D Christofides (2014). "A tutorial review of economic model predictive control methods". In: *Journal of Process Control* 24.8, pp. 1156–1178.

ENTSO-E (2024). *ENTSO-E Transparency Platform: Day-ahead prices, SE3 bidding zone.* `https://transparency.entsoe.eu/`. Historical hourly data downloaded for calendar years 2024–2025.

Goulart, Paul J. and Yuwen Chen (2024). "Clarabel: An interior-point solver for conic programs with quadratic objectives". In: *arXiv preprint.* arXiv:2405.12762.

Halvgaard, Rasmus et al. (2012). "Economic model predictive control for building climate control in a smart grid". In: *2012 IEEE PES Innovative Smart Grid Technologies (ISGT).* IEEE, pp. 1–6. DOI: `10.1109/ISGT.2012.6175631`.

ISO (2017). *SS-EN ISO 52016-1:2017 Energy performance of buildings - Energy needs for heating and cooling, internal temperatures and sensible and latent heat loads - Part 1: Calculation procedures.* Tech. rep. Stockholm, Sweden: Swedish Standards Institute.

Kircher, Kevin J and Max Zhang (2015). "On the lumped capacitance approximation accuracy in RC network building models". In: *Energy and Buildings* 108, pp. 454–462.

Madsen, Henrik and Jens Holst (1995). "Estimation of continuous-time models for the heat dynamics of a building". In: *Energy and Buildings* 22.1, pp. 67–79.

Oldewurtel, Frauke et al. (2012). "Use of model predictive control and weather forecasts for energy efficient building climate control". In: *Energy and Buildings* 45, pp. 15–27.

Pieper, Henrik et al. (2020). "Comparison of COP estimation methods for large-scale heat pumps used in energy planning". In: *Energy* 205.117994.

Prívara, Samuel et al. (2013). "Building modeling as a crucial part for building predictive control". In: *Energy and Buildings* 56, pp. 8–22. DOI: 10.1016/j.enbuild.2012.09.036.

Rawlings, James B, David Angeli, and Colin N Bates (2012). "Fundamentals of economic model predictive control". In: *2012 IEEE 51st IEEE Conference on Decision and Control (CDC)*. IEEE, pp. 3851–3861.

Reynders, Glenn, Jan Diriken, and Dirk Saelens (2014). "Quality of grey-box models and identified parameters as function of the accuracy of input and observation signals". In: *Energy and Buildings* 82, pp. 263–274.

Reynders, Glenn, Tomas Nuytten, and Dirk Saelens (2013). "Potential of structural thermal mass for demand-side management in dwellings". In: *Building and Environment* 64, pp. 187–199. DOI: 10.1016/j.buildenv.2013.03.010.

Schwarz, Gideon (1978). "Estimating the dimension of a model". In: *The Annals of Statistics* 6.2, pp. 461–464.

Široký, Jan et al. (2011). "Experimental analysis of model predictive control for an energy efficient building heating system". In: *Applied Energy* 88.9, pp. 3079–3087. DOI: 10.1016/j.apenergy.2011.03.009.

SMHI (2026). *STRÅNG Mesoscale Model for Solar Radiation*. https://strang.smhi.se/. Accessed: 2026-05-29.

Stellato, Bartolomeo et al. (2020). "OSQP: an operator splitting solver for quadratic programs". In: *Mathematical Programming Computation* 12.4, pp. 637–672.

Storn, Rainer and Kenneth Price (1997). "Differential Evolution – A Simple and Efficient Heuristic for Global Optimization over Continuous Spaces". In: *Journal of Global Optimization* 11.4, pp. 341–359. DOI: 10.1023/A:1008202821328.

SVEBY (2012). *Brukarindata kontor – Användardata för energiberäkningar i kontorsbyggnader*. Industry Standard. Stockholm, Sweden: Sveby (Standardisera och verifiera energiprestanda i byggnader). URL: https://www.sveby.org/wp-content/uploads/2012/01/Brukarindata-Kontor.pdf.

Thermal Energy System Specialists, LLC (2026). *TRNSYS: Transient System Simulation Tool*. URL: https://www.trnsys.com/ (visited on 04/22/2026).