Stochastic Optimization of Electrified District Heating Using Systems Dynamics Model

Lappeenranta–Lahti University of Technology LUT

Master's program in Business Analytics

2026

Usama Khurshid

Examiner(s): Assistant Professor Jyrki Savolainen

Professor Mikael Collan

ABSTRACT

Lappeenranta–Lahti University of Technology LUT

LUT Business School

Master's in Business Analytics (MBAN)

Usama Khurshid

Stochastic Optimization of Electrified District Heating Using Systems Dynamics Model

Master's thesis

2026

71 pages, 12 Figures, 04 Tables, 1 Appendix

Examiner(s): Assistant Professor Jyrki Savolainen and Professor Mikael Collan

Keywords: Thermal Energy Storage, District Heating, Electricity Price Uncertainty, System Dynamics Modelling.

This study examines the optimization of thermal energy storage operations in district heating when electricity prices act as the key source of uncertainty. A literature examines modelling and optimization approaches currently being employed, as well as prominent control strategies and uncertainty modelling. A particular focus is placed on understanding the need and role of dynamic modelling approaches, specifically System Dynamics modelling, and its potential in handling feedback loops that represent complex district heating operations.

An inherited System Dynamics model developed in Simulink is coupled with an external optimization layer using a Genetic Algorithm in MATLAB to determine offsets to a moving average charging/ discharging logic across three storage configurations (i) heat tank (ii) heat tank and one battery (iii) heat tank and two batteries. Electricity price uncertainty is represented by generating multiple scenarios by fitting an AR(1) GARCH(1,1) model on real electricity prices after preserving deterministic inter-temporal trends. Each scenario contains hourly electricity prices over a one-year horizon, comprising of 8784 timesteps.

The results from the literature review indicate that most district heating optimization approaches rely on closed-form mathematical representations rather than dynamic feedback loops, which are underrepresented. The empirical results indicate that external optimization added to SD models representing DH storage operations have the potential to introduce greater cost savings, however, their potential is limited by the degree of flexibility offered by the storage setup. This is supported by improved cost savings in the tank + 2 batteries storage configuration.

ACKNOWLEDGEMENTS

I would like to express my sincere gratitude to my dear family, friends, and teachers. My loving wife, Shizza Bukhari, through her unwavering support, patience, and reassurance, has helped me immensely in navigating the highs and lows of the journey towards attaining this Master's degree. My mother and siblings, though seas apart, have repeatedly kept me surefooted through reminders of my goal and offering unrelenting support. The friends I have made along the way ensured I would feel like I belong, filling in seamlessly for the friends I left back home. And last but certainly not least, I am grateful to my supervisor, Assistant Professor Jyrki Savolainen, whose help made this work possible; his valuable feedback and sincere guidance, throughout the thesis writing process, as well as the lessons taught during his course, motivate me to embark on a path to perpetual learning.

SYMBOLS AND ABBREVIATIONS

Symbols

Q_{demand}	Heat demand
$E_{tank(t)}$	Stored Energy in tank at time t
$E_{tank(t+1)}$	Stored Energy in tak at time t+1
$P_{tank(t)}^{ch}$	Tank charging power at time t
$P_{tank(t)}^{dis}$	Tank discharging power at time t
$E_{b,i(t)}$	Stored Energy in battery i at time t
$E_{b,i(t+1)}$	Stored Energy in battery i at time t+1
t	Time index
i	Battery index (1 or 2)
price(t)	Electricity price at time t
MA_Threshold	Moving-average electricity threshold
delta_tank	Threshold offset for tank storage
delta_b1	Threshold offset for battery 1
delta_b2	Threshold offset for battery 2
$J(\delta)$	Average operating cost for tested offset vector δ
δ	Offset vector being tested by the Genetic Algorithm
N	Number of training scenarios
s	Scenario Index
$C_s(\delta)$	Annual operating cost for scenario s using offset vector δ
P_g	Population set at generation g
L_p	Specified population size

G_{max}	Specified maximum number of generations
S(.)	Selection operator in Genetic Algorithm
C(.)	Crossover operator in Genetic Algorithm
M(.)	Mutation operator in Genetic Algorithm

Chemical Symbols

CO_2	Carbon dioxide

Unit Symbols

MWh	Megawatt-hour
MW	Megawatt
EUR/ MWh	Euros per megawatt-hour

Abbreviations

ACER	European Union Agency or the Cooperation of Energy
AI	Artificial Intelligence
AR	Autoregressive
AR(1)	First-order autoregressive model
ARFIMAX	Autoregressive Fractionally Integrated Moving Average With Exogenous Variables
ARMA	Autoregressive Moving Order
ARX	Autoregressive model with exogenous inputs
CHP	Combined Heat and Power
DH	District Heating

EB	Electric Boiler
FMI	Finnish Meteorological Institute
GA	Genetic Algorithm
GARCH	Generalized Autoregressive Conditional Heteroskedasticity
GARCH(1,1)	First-order GARCH model
HP	Heat Pumps
HTTES	High-Temperature Thermal Energy Storage
LP	Linear Programming
MILP	Mixed-Integer Linear Programming
MINLP	Mixed-Integer Nonlinear Programming
ML	Machine Learning
Monte Carlo	Monte Carlo Simulation
MPC	Model Predictive Control
NLP	Nonlinear Programming
P2H	Power-to-Heat
PSO	Particle Swarm Optimization
RMSE	Root Mean Squared Error
R^2	Coefficient of determination
SD	System Dynamics
TES	Thermal Energy Storage
WTES	Water-based Thermal Energy Storage

Table of contents

Abstract

Acknowledgements

Symbols and abbreviations

1 Introduction 9
1.1 Data and Methodology 10
1.2 Research Questions 11
1.3 Thesis structure 11
2 Literature Review 13
2.1 Review Methodology 13
2.1.1 Study Selection 14
2.1.2 Thematic Analysis 16
2.2 Literature Review Results 17
2.2.1 Storage-Induced Price-Responsiveness in District Heating 18
2.2.1.1 Thermal Energy Storage and District Heating Flexibility 18
2.2.1.2 Thermal storage in electrified DH settings 20
2.2.1.3 Operational Control Strategies 21
2.2.2 Modelling and Optimization Approaches in Energy Systems 22
2.2.2.1 Traditional Approaches 23
2.2.2.2 System Dynamics Based Approaches 26
2.2.3 Uncertainty Modelling Approaches 27
2.3 Synthesis 29
3 Data and Methodology 32
3.1 Data and Inputs 32
3.1.1 Electricity Price Data and Stochastic Scenario Generation 33
3.1.2 Heat Demand Function 35
3.2 Systems Dynamic Model Overview 36
3.3 Simulation-Optimization Framework 39
3.3.1 Optimization Goals 39
3.3.2 SD Model Adjustments 40

3.3.3 Optimization Approach 42
3.3.4 Choice of Optimization Algorithm 44
3.3.5 Parameter Selection 46
4 Results 50
4.1 Optimization Outcomes 50
4.1.1 Local robustness check for tank + 2 batteries configuration 54
4.2 Scenario-wise distribution of cost savings 55
4.3 Optimization Behavior 57
4.4 Results Analysis 60
5 Conclusion and Discussion 61
5.1 Answering the Research Questions 61
5.2 Limitations 63
5.3 Future Research 64

1 Introduction

District heating (DH) is an essential utility for cold climates such as in Finland where winter temperatures consistently remain below zero. In 2022, district heating contributed to 45% of the total heating market share for public buildings as well as residential and commercial structures (Statistics Finland, 2024). Currently, Finland's District Heating infrastructure consists of heating-only as well as Combined Heat and Power (CHP) plants with a significant reliance on combustion-based processes. As of 2023, carbon dioxide emissions from the district heating sector amounted to 3.1 million tons (Petersen-Dyggve, 2024). Finland's drive to become carbon neutral by 2035 (State Treasury, 2025) has driven heat providers to increasingly shift focus towards the use of renewables. As of 2023, 80.8% of the heat produced by CHP plants and dedicated heat production plants for district heating relied on fuel-based production such as natural gas, coal, oil, peat, waste, and wood/ other biofuels (Energiateollisuus ry, 2024). This number is intended to decrease swiftly over the years in order to achieve Finland's climate policy goals.

In light of this pursuit, the use of greener technologies such as electric boilers and heat pumps has increased rapidly with a 40.2% increase in adoption of these sources and heat recovery between 2022 and 2023 alone (Energiateollisuus ry, 2024). For existing CHP plants, although higher energy efficiency is principally a key benefit, the inability to independently adjust heat and electricity production leads to reduced flexibility (Christidis et al., 2012). Heat storage technologies are already widely adopted by heat production plants across Finland, offering more flexibility and enabling reserve market operations. This reserve market participation may be highly lucrative for district heating providers and help reduce emissions for those with electrified heating methods (Christidis et al.,2012). Operators have to make policy decisions that define their use such as when electricity-based heating should be used and how heat storage must be charged or discharged. Policies that optimize these decisions may significantly improve profits while also ensuring heat demand is met.

This study aims to investigate existing literature to uncover known practices regarding heat storage optimization under electricity price uncertainties, emphasizing a review of dynamic modelling approaches. The thesis builds on insights from the literature to demonstrate the combination of Monte Carlo simulation and evolutionary optimization when applied to heat

storage, given the increased relevance of the subject as Finland's district heating market shifts to such electrified, greener methods. This demonstration uses historical electricity price patterns for exploring how uncertainty effects cost operations, complementing the literature review with a demonstrated application using an existing Systems Dynamic model created by Roumpas (2025).

1.1 Data and Methodology

The study leverages hourly electricity price data for Finland to construct stochastic price scenarios for the demonstration component of the study. The simulation relies on an extracted subsystem simulating heat storage components from an existing System Dynamics model developed by Roumpas (2025). The subsystem provides sufficient technical constraints, storage dynamics, and operational relationships required to assess cost-related outcomes. Additional data used includes outdoor air temperature data that helps calculate heat demand based on the embedded heat demand function, helping construct the cost component on that basis of which the optimization is conducted.

The methodology is driven by two parts that act to complement one another. Primarily, a literature review identifies modelling approaches that are used for optimization of heat storage under price uncertainty, including system dynamics and simulation-optimization methods. Secondly, the study also demonstrates how combining Monte Carlo simulation with evolutionary optimization can be applied to the TES-enabled DH system model with a focus on optimizing an existing control decision to determine the impacts of the adopted optimization approach and analyse results to understand the impact of such optimization approaches on improving district heating operations in storage-enabled settings. Historic electricity price patterns are utilized in order to generate stochastic price trajectories, and offsets are added to the inherited charging/ discharging decision logic, which are varied across Monte Carlo runs using a Genetic Algorithm (GA). Each run produces aggregated annual costs across three storage configurations.

1.2 Research Questions

This thesis has two complementary aims: (i) seeking to synthesize existing literature on heat storage optimization under price uncertainty with a particular focus on dynamic modelling approaches and (ii) demonstrating how Monte Carlo simulation and evolutionary optimization can jointly be applied to a Systems Dynamics based heat storage subsystem to explore potential for cost efficiency under price uncertainty. This demonstration serves as an application of the specified modelling approach.

The aims specified can thus be classified into the following research questions (RQ1) and (RQ2):

RQ1: What is known about the heat storage optimization under price uncertainty in existing literature and which types of dynamic models exist for optimization?

RQ2: How can Monte Carlo Simulation with Evolutionary Optimization be utilized to enhance heat storage operations to minimize operational costs under price uncertainty?

1.3 Thesis structure

This study is structured to progressively synthesize existing knowledge, transparently document the applied methodological approach, and subsequently analyse findings to best answer the aforementioned research questions. As such, it is divided over multiple chapters and subchapters with dedicated purposes that contribute to this end.

The current chapter concerns the introduction of the study and discusses the overall scope by justifying the need to explore the topic space and outlining clear research questions that address the overall topic. It also gives a summary of the data used and methodology applied. Subsequent chapters discuss the methods applied to conduct the study in detail, actual findings across the approaches, divided into a literature-centered and empirical application sections.

The second and third chapters contribute explicitly to the literature sections wherein the second chapter defines the approach utilized to collect, evaluate, sort, and analyse literature in order to document the used approach, ensure transparency, and create room for replicating

such a search. The third chapter discusses the content of the studied literature and synthesizes findings to answer research question one while also contributing to the empirical part by informing modelling approaches and uncertainty handling.

Chapter four documents the methodology applied across the entire empirical experiment. This includes generating the electricity price scenarios to mimic electricity price uncertainty in real markets, explanations of the inherited Simulink model and made adjustments, as well as the applied optimization approach and chosen parameters. This chapter also explains the reason behind the adopted optimization approach.

The last two chapters, five and six, discuss the findings of the optimization experiment and subsequently the entire study. While chapter five dives into interpretive details regarding acquired quantitative results, chapter six synthesizes qualitative and quantitative findings to answer both research questions. Therein, limitations and directions for future research are also discussed.

2 Literature Review

2.1 Review Methodology

The literature review for this study was conducted based on a Systematic Literature Review Approach, however, some flexibility was afforded as the nature of the topic remained highly interdisciplinary and direct literature for the modelling aspect of this study was difficult to find. A structured pattern was therefore followed, which features defined search terms, screening, full-text assessment, and thematic organization of the literature selected. The approach was deemed suitable since several related areas such as DH, TES, electricity price uncertainty, optimization, and dynamic simulation are connected within the overall research topic.

As the aim extended beyond identifying studies that were in direct relation to the topic to also include how different elements of the topic have been addressed in literature, directly relevant studies on TES and DH were supplemented through methodological studies explaining uncertainty modelling or simulation-based optimization approaches.

The backbone of the search process included use of databases such as Google Scholar and LUT Primo. While Google Scholar offers access to vast amounts of literature available across a variety of databases, LUT Primo helped to acquire access to the selected papers once they had been identified for suitability and further evaluation beyond the preliminary assessment defined in subsequent parts of this chapter.

The search focused on literature from after 2010 in order to capture work that would be transferable in terms of its recent knowledge and relevance. One particular exception to this came in the form of a Systems Dynamics study from 2009 which was included for its role as an important methodological support for understanding the connection between dynamic simulation and optimization. Its role, however, was not in representing recent DH developments, only supplementing the methodological background of the thesis.

Terms related to district heating, thermal energy storage (TES), optimization, simulation, and electricity price uncertainty were used to conduct the search. Therefore, they included terms like “*heat storage*”, “t*hermal energy storage*”, “*District Heating*”, “*Optimization*”,

"*Systems Dynamic*" etc. Boolean operators AND and OR were used in order to combine these search terms to form meaningful search strings that could capture combined areas of interest that address the research question at hand. For example:

1. "District Heating" AND "Thermal Energy Storage" AND "Optimization" AND "Electricity Price"

2. "Electricity Price Uncertainty" AND ("Theramal Storage" OR "Heat Storage").

In the stated examples, distinct groups of focus are addressed. The first string captures studies drawing a direct connection between TES, DH, Optimization, and Electricity Prices; whereas the second focuses more on electricity price uncertainty and its relationship with TES.

Sources that were available as full texts and written exclusively in English language were the only ones considered so that the complete study could be analyzed instead of drawing inferences from the abstracts which would not sufficiently cater to the scope of this study.

2.1.1 Study Selection

In order to shortlist the literature for final use in the literature review, a screening process was applied in consecutive stages. Initially, the title and abstract for each study were reviewed and used to sort out candidate literature based on its potential contribution and relevance to the study. Studies relevant to district heating, thermal energy storage, optimization, electricity price uncertainty, and simulation-optimization, particularly systems dynamics were retained. In this step, of the over 13,000 studies retrieved by the search databases across multiple string search combinations, only 200 were retained based on their titles and position in the search results. The abstracts, introductions, and conclusions for these studies were read, leading to a further shortening to 56 studies.

Once this was done, an additional screening layer was enacted wherein the papers were shortlisted through an assessment of their content by reading the full text. Specifically, studies that explained the representation of thermal energy storage, operational decisions, and electricity price uncertainty in modelling and optimizing, approaches were given particular preference.

Not every study was explicitly required to cover every part of the research topic thoroughly due to the fact that few studies that were found combined all required aspects, and the broader scope of the study covers topics that span over domains that are often diverse. Therefore, studies that made a clear contribution to one or several parts of the review were retained; some literature contributed to DH and TES discussions, whilst others were useful in understanding uncertainty modelling or optimization methods in SD.

Studies were initially added according to their relevance to at least one area of the review, with a particular preference for studies that had more than one area of interest synthesized. For example, the studies that explicitly connected TES optimization in DH with electricity price uncertainty were retained. Moreover, particularly for System Dynamics, studies that helped explain SD or simulation-based optimization were kept.

At the same time, studies that we explicitly outside the direct scope of the review or did not meaningfully contribute to understanding SD based optimization were excluded. Examples for such include those that were unrelated to DH, TES, energy system simulation, or optimization. Moreover, studies that reiterated findings beyond adding substantial value to what was already discovered through prioritized or more recent studies were discarded. As a result, the applied criteria was aimed at establishing focus whilst avoiding an overly narrow analysis space since the chosen topic spans multiple areas of research.

Moreover, backward snowballing was used, particularly with key pieces of literature cited within the work of Roumpas (2025), given its nature as the predecessor for this thesis. Thus, the reference list of key studies was used to identify further sources that could be relevant, especially in cases when certain modelling approaches, uncertainty treatments, simulation frameworks, or cross-sectional reviews could be uncovered. This process was treated as a supplement to the structured search instead of a separate method, as sources identified in this manner were also assessed and retained based on their relevance to the review themes. A total of 03 studies were added based on this additional search procedure.

A major portion of the review was dedicated to DH, TES, optimization, and uncertainty; however, literature on SD was also included to inform the direction of the employed approach in this study. Moreover, the use of SD in TES optimization and modelling was noticed as a key limitation, given scarce representation in current literature on the subject. Adjacent methodological studies in which understanding dynamic simulation models with

feedback representation and the utility potential of their optimization was supported were also considered. This consideration was particularly important for the only pre-2010 study included, as it provided methodological grounding for the SD-related part of this review.

Based on this selection process, the final number of included studies came to 25 as recoded in the Appendix 1.

2.1.2 Thematic Analysis

A theme-wise synthesis was conducted for the selected literature as the studies were grouped in accordance with their role in understanding the overall research topic. Overall, the studies reviewed had differing aims, assumptions, methods, and system boundaries, making the thematic synthesis a viable approach. The core themes for the review are the same as reflected in the literature review itself. At a higher level, these are: Storage-Induced Price Responsiveness in District Heating; District Heating Modelling and Optimization Approaches; and Uncertainty Modelling Approaches.

These themes were further divided into focused areas as defined in Figure 1 below.

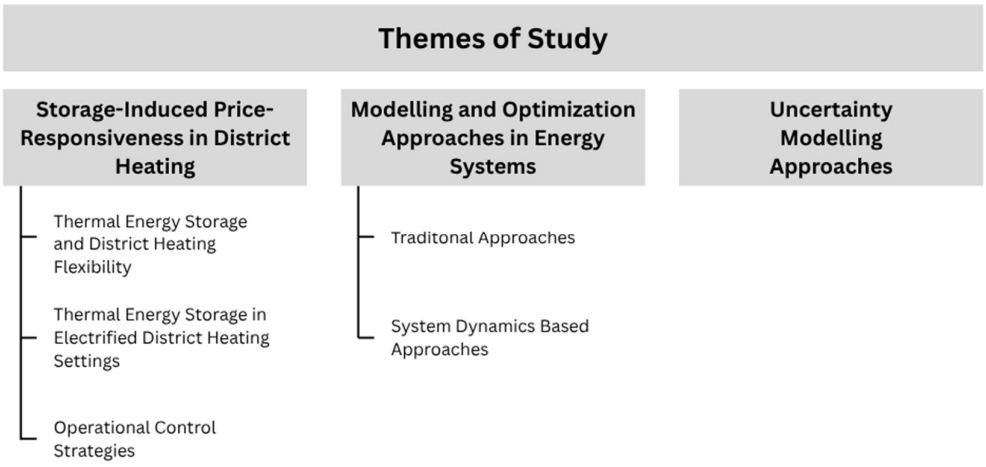

Figure 1: Thematic Areas and Subthemes of the Review

2.2 Literature Review Results

District heating is a crucial part of Finland's energy system, with building heating demand constituting 27% of energy consumption in 2022 and district heating accounting for a 45% market share across residential, commercial, and public buildings (Javanshir et al, 2024). DH has conventionally relied on combined heat and power plants (CHPs); however, greener methods are now increasingly being adopted as Finland aims to remove coal usage in energy production by 2029 and cut peat by half by 2030 (Javanshir et al, 2024). This includes shifting to the electrification of DH systems as national targets aim to reduce DH emissions by 50% before 2030 and reach carbon neutrality by 2035 (Petersen-Dyggve, 2024).

This ties crucial DH decision processes closely with electricity prices and markets. However, considering profitability concerns, selling electricity when prices are high, instead of converting to heat, is generally the better choice (Petersen-Dyggve, 2024). Considering this market-dependent operational setting, TES remains a key flexibility feature. TES involves storing heat in excess production and discharging in high-demand periods with long-term storage capabilities allowing heat storage for several months (Petersen-Dyggve, 2024). Existing literature indicates consistent operational logic for TES where periods with lower electricity prices are utilized to convert electricity into heat which is subsequently stored (Jaanto et al., 2023).

The practice of heat storage optimization perhaps becomes even more relevant when considering the relationship between climate, heat demand, and electricity prices together. Heat demand is strongly linked to outdoor temperature and highest electricity prices occur most often during coldest and darkest hours, coinciding with the highest heat demand, collectively creating a structural stress point for Power to Heat-based supply (Petersen-Dyggve, 2024).

2.2.1 Storage-Induced Price-Responsiveness in District Heating

2.2.1.1 Thermal Energy Storage and District Heating Flexibility

Thermal energy storage acts as a lever to uncouple heat load and real-time CHP production; this enables price-oriented electricity production as electricity prices tend to fluctuate over a given time period (Christidis et al., 2012). Along most literature, heat demand is most often a hard constraint that is satisfied on priority. As CHP heat production schedules vary, excess heat is produced and stored during high electricity-price intervals, since electricity generation becomes more attractive. Heat is then released in periods when electricity generation is unprofitable so that CHP can reduce output or shut down altogether (Christidis et al., 2012). This allows for thermal energy storage to serve as an operational buffer to allow CHPs to operate in a profit-maximizing fashion through liberation from demand obligations.

Literature extensively captures various layers of decision-making from operational scheduling to investment planning for thermal energy storage in variable price scenarios. One such example is shown in a Berlin-specific case discussed by Christidis et al, 2012, where unit commitment and storage tank sizing are optimized while various components of the district heating networks such as pipes, pumps, and mixers are represented in the formulation. Similarly, other works, such as the one written by Malmgren and Skolglud (2021), show both investment optimization as well as hourly dispatch over a year to represent planning models. The study by Bjørnskov et al (2021) shows operational optimization using a non-convex scheduling model for CHP and Thermal Energy Storage (TES) with 1-hour time steps, a 48-hour test period, and rolling window approach. Throughout all these papers, electricity prices are entered as an exogenous time series. For example, Bjørnskov et al (2021) uses spot market electricity prices, whereas Marmgren & Skolglud (2021) use Sweden SE3 2019 from Nord Pool in deterministic runs.

Even though studies have a spectrum of aims and scopes, optimization logic seems to be consistent in that CHP electricity earnings are tied with varying electricity prices, while thermal energy storage is constrained by balancing in the inter-temporal sense, with heat shifting enabled across time periods. Studies with cost-minimizing objectives aim to reduce annual system costs from an investment and operations lens, subtracting CHP electricity revenue subject to electricity price time series (Malmgren & Skoglud, 2021). On the other

hand, Christidis et al. (2012)'s study focuses on maximizing profit; it equates profit to the difference between electricity sales revenue and levelled-out investment costs for storage tanks, not considering heat sales, as heat demand remains stable irrespective of storage features. Across studies, thermal energy storage models vary from constant temperature, hot-water tanks whose sizes are defined within the model based on demand fluctuations to models with a high-physics focus where water temperatures are assumed to change over time and consideration given to energy leakers as well as models with additional considerations such as discharge flows.

Despite these variations, there is a consistent pattern that all literature converges on: thermal energy storage allows greater flexibility for CHP operations. Notably, the strategic approach remains the same: overproduction in times of high profitability and reduced production capacity in times of low profitability, all whilst consistently meeting heat demand. Marlmgren & Skoglud (2021) connects CHP operations strongly to price variations in a stochastic setup, emphasizing how volatility becomes more relevant as more hours become profitable when volatility creates "high price" hours.

Naturally, the shifting logic also paves way for system-level impacts. Christidis et al (2012) note that, contingent on scenarios, 29-43% of the heat supply may be derived from stored thermal power while also increasing mean returns on electrical power. Bjørnskov et al. (2021) highlight that improved initialization plays a role, yielding consistent low-cost schedules and high median savings. However, despite collective assessment in literature signalling a high positive impact on profitability in the introduction of thermal energy storage for CHP DH operations, the optimal size remains sensitive to uncertain parameters.

The literature studied highlights certain limitations in light of price uncertainty questions. Some methods rely on perfect foresight and prognosis across entire horizons from given price and heat-load timeseries, documenting the dependency of benefits on forecast quality (Christidis et al., 2012). Another alternative is represented in the form of rolling-horizons with limited look-ahead such as Bjørnskov et al (2021) that only assumes that hourly spot prices are known only 24 hours in advance. Scenario design can be stylized, even as uncertainty is introduced. A two-stage stochastic model uses generic step-function price profiles for electricity and equal probability scenarios, since likelihoods are difficult to estimate (Malmgren & Skoglud, 2021). Lastly, the sensitivity of reported results to the time resolution (e.g., the estimated benefits are higher with time resolution of an hour than with

time resolution of 4-hour steps because the storage is used more efficiently) indicate that foresight assumptions as well as time granularity can have a significant impact on the perceived value of TES. These limitations inspire further study into how storage-enabled operating policies might perform against truly uncertain, multiple electricity price paths, rather than being reliant solely on optimization runs with perfect foresight.

2.2.1.2 Thermal storage in electrified DH settings

As district heating becomes electrified through electric boilers (EB), heat pumps (HP), and power-to-heat (P2H) technologies, heat supply becomes sensitive to electricity-market conditions. Thermal energy storage, again, plays an important role in allowing greater flexibility by converting price variation into both operational as well as economic value. The reverse is observed as was with CHP+TES applications, in that electricity cost minimization becomes the predominant objective. Consumption-side arbitrage becomes the core logic, with heat production being increased with thermal stores being stocked in periods of low prices whereas thermal energy stores are discharged in times of high electricity prices (Jaanto, 2023).

Overall, the cases observed include a range of systems such as scheduling storage for tanks connected to electric boiler heat production plants, multi-asset district heating portfolios, and system-level optimizations. Sinha et al (2019)'s study assesses an EB scheduled against a hot-water tank and features an hourly decision to turn off or on over a horizon of 24 hours under storage and temperature constraints. At an entire portfolio level, Sihvonen et al (2024) explains how P2H operates with various storage types such as WTES and sand-based HTTES to optimize storage dispatch and market participation. At larger scales, electric boilers and various storage categories are treated as components of interconnected DH systems, with Hiltunen et al (2025) optimizing hourly operations over a one-year time horizon and using thermal energy storage states to time-shift heat production.

The value proposition regarding thermal energy storage seemingly remains consistent across studies: thermal energy storage allows time-shifting in heat production based on how electricity prices vary and focuses production in cheap price hours.

2.2.1.3 Operational Control Strategies

The studies examined often concern operational decision strategies in DH systems where performance is contingent to electricity market signals as well as thermal flexibility. A high-level examination could divide these strategies into rule-based, model-based, and hybrid forms of control; however, individual studies often differ in how operational decisions are made and the information that used as well as the treatment of uncertainty. A recent review by Monghasemi et al (2026) also reinforces a similar broader categorization, noting the continued prevalence of rule-based methods.

A noticeable distinction exists in heat-driven, rule-based operations and active, economic-centered control. In the reference case of Vanhoudt et al (2018), local building operation adheres to thermostat logic and heating curves whilst CHP and gas boiler actively respond to heat demand requirements instead of electricity prices. An active control strategy is employed that aims to achieve maximum CHP profit through the shifting of operations to hours where electricity prices are high. Just as discussed in this study, the thermal storage acts to decouple heat production from demand; in their specific case, building thermal mass is also utilized to this end. Thus, a more economically favorable operating pattern is achieved.

In case of the study by Quaggiotto et al. (2021), a deterministic optimization is applied wherein the operating plan is calculated from a defined set of forecast inputs. The scheduling of cogeneration plants, heat pumps, and gas boilers is formulated as a Mixed Integer Linear Programming (MILP) problem. This draws an interesting contrast where, rather than a reactive, rule-based operation, decisions are preplanned subject to expected conditions instead of being determined by instantaneously observed requirements. Through the use of an centralized thermal storage, optimization-based operations are noted to exploit additional flexibility in the presence of storage options (Quaggiotto et al., 2021).

Rolling-horizon strategies, often referred to as Model Predictive Control (MPC) are closely related. Rather than computing and sticking or a single plan, MPC recurrently updates the operating schedule in parallel to state information and forecasts becoming available. Quaggiotto et al. (2021) implements this type of look-ahead control and compares it with current rule-based operation over representative weeks. They find that their operating costs are reduced with MPC, and reduced further with the addition of centralized heat storage.

Roumpas (2025) introduces an interesting approach, which may be categorized as price responsive, employing a 24-hour moving-average threshold to which current electricity prices are controlled. When the price is lower than the threshold, the system charges the storage, whilst discharge is activated in cases where the price is higher. The study compares the cost of each set up in terms of the lowest annual operational cost with the aim of determining the most appropriate combination of storage settings out of a total of three aggregated setups. In this model, the operational decision rule is predefined, thus the performance is evaluated based on fixed control logic instead of optimized decision-making.

Collectively, the studies direct us to the idea that the space for strategy comparison does not conclude at rules and optimization; rather, distinct approaches that are split along different axes altogether. One may be the decision logic itself, which may constitute fixed thermal rules, price-responsive thresholds, deterministic optimization, and rolling-horizon MPC. A secondary axis is also the information upon which the decisions depend. For example, in the case of Vanhoudt et al (2018), the strategy depends on heating curves and local temperature. Whereas in Quaggiotto et al. (2021), it is driven by forecasts of demand, waste heat, and electricity prices. For Roumpas (2025), it depends on the relationship between the current price and moving average instead of forecasts or predictive optimization.

The strategy in question is also tied to the flexibility that makes cost-savings or profit-maximization possible: the source of thermal storage in the DH context. Vanhoudt et al. (2018) compares a central storage tank, distributed buffers, and building thermal mass whereas Quaggiotto et al. (2021) compares the thermal inertia of the network with the combined use of network flexibility as well as central storage. Roumpas (2025), on the other hand, considers scenarios around various configurations of heat storage options such as a heat tank with corresponding options of up to two batteries. Across the studies, economic payoffs is tied consistently with thermal flexibility; however, the magnitude of the benefit depends on both the flexibility itself as well as the decision framework that enables its effective use.

2.2.2 Modelling and Optimization Approaches in Energy Systems

In order to understand the impact made thus far by specific modelling choices and corresponding optimization approaches, it becomes imperative to assess existing literature.

A review of modelling and optimization choices alongside details of system behaviors serves this purpose. This section is divided into studies documenting two major observed themes: (i) Conventional DH Modelling Approaches (ii) System Dynamics Based Approaches.

2.2.2.1 Traditional Approaches

Of the reviewed literature, 16 studies demonstrated DH optimization and modelling in some capacity, of which 14 employ mathematical programming and relevant optimization-based control approaches. These studies frequently represent systems through constraints over timesteps, and the objective function focuses on maximizing or minimizing a given output, such as profit or cost. Sporleder et al (2022)'s systematic literature review discusses a notable number of linear programming or MILP approaches where studies often aim to optimize DH designs.

However, it is noted that non-linear formulations are also common, particularly when it comes to representing complex physical relationships; as such, these represent a prevalent trade-off as Non-linear Programming (NLP) models successfully model more detail whilst incurring a higher computational cost (Sporleder et al., 2022). Widely speaking, this allows drawing a high-level distinction regarding how optimization problems in the DH context are often formulated. MILP and other linear approaches are observed in cases when design choices, operational scheduling, and network decisions may be expressed in linearized balances. In contrast, Mixed-Integer Nonlinear Programming (MINLP) formulations become more suitable when stronger physical fidelity is to be preserved, and tractability becomes a less important concern. Sporleder et al (2022) continue to note that mixed-integer variables are notably introduced as DH optimization problems often need a discrete design and discrete commitment decisions, for example, if units are switched on or off; as such, MILP affords a reasonable trade between detail and solvability.

In a practical demonstration, Christidis et al. (2012) use a mixed-integer model to jointly optimize operations for nineteen power plants whilst design-sizing six heat accumulators in a Berlin-based case. They set profit maximization as the objective function, representing profit as the difference between electricity sales and the combined costs of fuel, CO_2 certificates, and levelized distribution of storage infrastructure investments. A key constraint that is applied implies that heat demand must be met in all circumstances, disregarding the

need to add heat sales to the overall revenue. Interestingly, their problem was formulated as an MINLP initially, but was linearized to an MILP and subsequently solved using CPLEX. This approach highlights an interesting, convenience-centered approach where non-linear physics is reformulated into tractable linear structures to allow the joint optimization of design and unit commitment in large systems.

Christisdis et al. (2012) further apply spatial and temporal decomposition by splitting the network into supply areas and dividing into two 6-month horizons, while timesteps were widened to 4 hours. This was done upon the realization that not all cases converged to acceptable levels of optimality. This demonstration becomes relevant when considering the appropriate use of decomposition in MILP-based optimization, such that it preserves enough structure to remain informative while improving computational convenience to allow large problems to become easily solvable.

However, in their study, Bjornskov et al (2021) retain an hourly timestep while explicitly using a non-convex MINLP modelling approach. A branch and bound method was to solve the optimization in a sliding time window and Monte Carlo Tree Search was applied to start the MINLP solver as an initialization method. This reflects a common trade-off in optimization where unlike MILP and LP where solutions arrive through a global optimum, non-convex MINLP cannot guarantee beyond a local optima. By using Monte Carlo Tree Search as a warm-start to identify relatively feasible regions prior to the brand-and-bound application, Bjornkov et al (2021) implicitly highlight the significance of not only mathematical formulation, but of search strategies too in MINLP-based optimizations.

Whilst opting for MINLP in their study, Bjornskov et al (2021) do note that linearization remains a method of transforming non-linear scheduling problems to MILP; however, finer approximations in niche components of the model improve accuracy at the cost of increasing variables and overall model complexity. Therefore, it may be inferred that the choice is not simply “a simple MILP versus a realistic MINLP”, rather between managing how complexity is distributed within the model via non-linear formulation, or externally by linear approximation, decomposition, iterative runs, or specialized initialization procedures.

Ceruti et al (2025) manage complexity using temporal aggregation, wherein period-typical features are scaled using weighted representation to represent the entire year. TES is simulated through intertemporal recursion, charging and discharging efficiencies, and

cyclical conditions as observed over periods (Ceruti et al., 2025). This temporal aggregation is an explicit modelling decision that, once again, represents a tradeoff between accuracy and runtime. This reinforces an observation by Sporleder et al (2022) regarding the use of decomposition and time aggregation approaches or reducing computation time. However, Ceruti et al (2025) highlight the effects of these choices in altering supply configuration and selected network topology. In an optimization lens, this is significant as it shows the contrasting impacts of formulation choices on both computational burden and outcomes.

Another notable approach that becomes prominent in literature is Model Predictive Control through which optimization is treated as a repeated operational decision problem. As such, the optimization extends to repeatedly updating control actions when new states and forecasts are realized. One such study demonstrates an MPC controller that is formulated as an MILP and runs on a receding horizon across hourly timesteps in order to represent a case from Veron, Italy (Quaggiotto et al., 2021).

Even within the MPC approach, uncertainty handling and system representations widely differ. In the case of Verrilli et al (2016), a two-stage stochastic optimization model is applied to define a stochastic MPC problem wherein the decisions at the first stage offer a planned schedule while the second stage has variables adjusted for realized forecast errors. This depicts MPC as an anticipatory approach where the model operates under uncertainty. In contrast, in the model developed by Zhang et al. (2023), an MPC is built on a hybrid predictive structure. This structure combines a reduced-order physical model that approximates certain physical behaviors to generate candidate control actions, with a data-driven model that is used inside the MPC framework to identify the desired solution. This particular approach demonstrates that MPC formulations can also differ based on both the mathematical program as well as the predictive model structure.

Together, the studied literature shows us that optimization technique, modelling method, simplification, and state updates collectively add to the uniqueness and practicality of observed approaches. Modelling choices largely decide the nature of optimizable decisions and the reflection of operational change as the states update. Notably, mathematical or control-oriented modelling approaches often attempt to simplify system representation for optimization to remain computationally tractable.

2.2.2.2 System Dynamics Based Approaches

Direct literature on DH systems optimization with systems dynamics models is scarce. However, from wider sources on SD optimization, especially related to energy systems, modelling approaches and comparable optimization methods can be used to help in understanding what frameworks currently exist. Optimization in SD settings has a distinct meaning as it includes both calibration as well as policy optimization; it differs from mathematical programming given the non-linear and dynamic nature of SD models (Dangerfield, 2009). System Dynamics and other simulation-based approaches are therefore extremely valuable when the intention becomes representing feedback-rich system behaviour that progresses with time rather than to directly solve closed-form optimization problems. Review studies also confirm that the domain of DH modelling and optimization features physical and blackbox models as well as deterministic and heuristic solution methods (Talebi et al., 2016).

This evidence surfaces in the work of Ziemele et al. (2016), where the stock and flow models are leveraged to model transition in DH. Rather than an embedded mathematical solver, decision support in their work arises via scenario comparison and endogenous system behaviour. Installed capacities are treated as stocks within the model, whereas flows represent depreciation and investment. In a comparable case, Roumpas (2025) evaluates system performance by repeating simulations for stochastic electricity-price scenarios; in his study, the primary operating decision revolves around a predefined threshold-based control logic that is embedded within the simulation model. In such setups, the simulation environment functions as a performance evaluator wherein candidate decision rules are tested by running the model. For Roumpas (2025), the results are observed as cost outcomes from the repeated model runs instead of being obtained from an explicit analytical optimum.

An interesting approach to couple a Systems Dynamics Model with an explicit optimization is observed in the study by Wu and Xu (2013). Although not directly related to DH, their model simulates a high-level system that represents the interconnected effects of economic growth, energy consumption, and CO_2 emissions. With the goal of maximizing GDP whilst simultaneously minimizing CO_2 emissions and energy consumption, policy variables such as growth rates of industry sectors investments poured into environmental protection and energy savings is optimized (Wu & Xu, 2013). Their work is parituclarly valuable to this

discussion by virtue of their explicit demonstration of coupling a SD model with an external optimization layer as they integrate SD with fuzzy multi-objective programming. Therein, a demonstration of applying a causal structure and dynamic evolution through SD whilst leveraging external optimization for improved decision outcomes is given.

This is consistent with Dangerfield (2009) as they highlight a wider discussion in which policy improvements are represented as payoffs that can be obtained from searching parameter spaces that improve model performance. Studies related to SD indicate that optimization in SD models has progressed beyond just calibration, now encompassing policy optimization, and mult-objective formulations by leveraging advanced approaches such as genetic algorithms to explore parameter spaces (Dangerfield, 2009).

Collectively, the reviewed literature indicates that SD based approaches are particularly useful in representing temporal behavior with embedded feedback effects. Furthermore, it can also be observed that optimization is not necessarily expressed as a single analytical problem, instead it can be applied through scenario comparison, rule-based structures, or adding separate search frameworks on top of developed SD models that represent complex system relationships and behaviours.

2.2.3 Uncertainty Modelling Approaches

As sufficiently stated, operational as well as investment decisions rely heavily on external factors that often come with problems such as uncertainty. From the studied literature, price and heat demand are major sources of uncertainty when modelling DH operations due to their direct impact on the profitability of electrified DH systems as well as meeting obligatory demand requirements (Malmgren and Skoglund, 2021). However, other uncertainties such as fuel prices, CO_2 pricing, and supplementary heat sources such as waste heat also contribute to uncertainty in observed studies; however, most such as Malmgren and Skoglund (2021), do not treat these externalities as stochastic.

Often, uncertain inputs are treated in a deterministic sense with uncertainty explored through sensitivity analysis. Malmgren and Skoglund (2021), run alternative cases including higher biofuel costs, lack of industrial waste heat, or cold temperatures as uncertainties through deterministic linear optimization. Rather than a probability distribution within optimization

problems, these approaches are straightforward to compare. However, Malmgren & Skoglund (2021), also highlight a limitation regarding missed interactions between uncertainty drivers and the resulting performance compromise in real applications, where one high-demand case saw the inability of the model to meet demand. Such consequences can lead to significant trouble in real scenarios and reiterate the costs of the tradeoff between the realism of multi-source uncertainty modelling and computational affordability of focused uncertainty representation in DH systems.

Studies such Bauer (2025), Malmgren & Skoglund (2021), and Guericke et al (2024) rely on two-stage structures to separate decisions made before and after uncertainty is realized. In Malgren and Skoglund (2021)'s study, first-stage decisions are investment capacities whereas the second-stage pertains to dispatch choices based on each scenarios for a total of 15 scenarios between 5 price profiles and 3 demand profiles. In Bauer's (2025) approach, a similar number of scenarios (15) representing uncertainties in the market (such as electricity prices), demand-side, supply-side, and regulatory environment is tested; however, Bauer (2025) runs the optimization with pre-assigned probability weights for each scenario. The scenarios are varied through structured changes via multipliers on electricity prices and heat demand (Bauer, 2025).

Uncertainty modelling seems to be even more relevant in planning operational decisions and market participation. Guericke et al. (2024) uses a stochastic MILP optimization for heat production with electricity prices, heat demand, solar heat, and waste heat as uncertain timeseries. A scenario-based framework that forces identical decisions across an initial window is used to represent the phased introduction of information and subsequent adaptive decisions; by this comparison and a subsequent calculation for the value of the stochastic solution, an improved performance of the stochastic model is confirmed (Guericke et al., 2024). Both Marlmgren & Skoglund (2021), and Guericke et al (2024) note that adding adding more stochastic variables leads to greater computational burden relative to deterministic sensitivity analysis. One approach to deal with this is observed by fitting ARX time-series models with temporal pattern inputs to simulate multivariate scenarios of equal probability, subsequently clustering them to manageable representative scenarios (Guericke et al (2024). Computational feasibility aside, the use of ARX-type time-series models to capture and simulate electricity price uncertainties is also supported in the work of Chai et al (2024) and subsequently employed by Roumpas (2025).

In Refsgaard (2021)'s DH bidding and operation model, spot prices and weather-driven variables are the main sources of uncertainty. The day-ahead electricity prices uncertainty is represented by scenarios generated at each optimization run, using on Markov chain combined with Monte Carlo simulation (Refsgaard, 2021). Historical spot prices from Nord Pool are made discrete into 100 percentile-based "states", following which 24 transition matrices, reflecting hourly price behaviour, are used to model price states evolution; prices are then recovered using uniform distribution from within the bin (Refsgaard, 2021).

While the Markov Chain Monte Carlo combination used by Refsgaard (2021) adds computational complexity and runtime, it is repeatedly observed in literature that simplified uncertainty modelling yields quicker results with reduced accuracy. A large number of DH optimization studies employ scenario testing or sensitivity analysis, with uncertainty-centered analysis being scarce (Sporleder et al., 2022). This exposes most models to the feasibility risks in real applications (Sporleder et al., 2022). Sporleder et al (2022) also suggests Monte Carlo or multiple scenarios as remedies, subject to added computational costs.

The topic of uncertainties in DH seems to follow a practical range of approaches that include deterministic as well as scenario-based uncertainty treatment. We can consistently note improved ability to reflect real operating conditions in cases where uncertainty is treated comprehensively, albeit at a higher computational cost. Implicitly, this trade is present at two levels: how many variables are considered as sources of uncertainty and how comprehensively uncertainty is represented for each variable. Moreover, electricity prices seem to be a consistently present variable when it comes to uncertainty representation, for which a timeseries modelling is a demonstrated approach for scenario generation.

2.3 Synthesis

When viewed together, the literature clearly points the flexibility that thermal energy storage offers to district heating systems. This flexibility is a key ingredient in shifting production in the temporal sense whilst remaining sensitive to necessary requirements of heat supply. Through its ability to decouple production and demand, production can be shifted to times where electricity prices are favorable. In conventional systems that rely on CHPs, the focus revolves around exploiting high electricity price instances to produce excess electricity; this

logic is reversed in cases where DH systems are electrified as electricity becomes the primary cost driver and it becomes economically more feasible to produce heat in times when electricity prices are high.

Moreover, the economic value derived relies heavily on how the underlying production dispatch, and charging decisions are formulated. From rule-based approaches, stochastic programming, and deterministic optimization, the objective remains to exploit temporal flexibility. However, outcomes are predominantly shaped by modelling assumptions surrounding key decisions like uncertainty treatment, model foresight, and time resolution. A recurrent tradeoff is demonstrated across uncertainty treatment and model design across studied literature, where deterministic and simple formulations reduce computational costs while detailed representations lead to better representation of near-reality outcomes.

Mathematical programming approaches constitute most modelling approaches in DH optimization literature with studies focusing around MILP, MINLP, stochastic programming or MPC formulations. This remains the case despite the inherent complexities of real DH systems where feedback-oriented temporal behavior is common; as such, simulation and systems dynamics-based approaches arguably remain well-suited to represent endogenous system change. In particular, for systems dynamics-based optimization, there remains extensive room for literary contributions in the district heating context.

This leads to what is perhaps a highlight in the literature gap: the limisted evidence of direct studies that apply advanced optimization to dynamic representations of DH systems, particularly in SD models. In this sense, whilst Roumpas (2025) establishes a model that represents dynamic behavior in a TES-DH setting, the implied constraints within the model are not optimized. The studied literature shows little evidence of studies that combine system dynamic models with external optimizers in the DH context. This comes despite Dangerfield (2009)'s acknowledgements of SD-based approaches leveraging external optimizers based on AI and ML. Even in Roumpas (2025)'s case where a dynamic threshold-based rule decides charging/ discharging behavior within the DH SD model, the rule itself is not optimized beyond predefined threshold formulation.

With respect to uncertainty, electricity price remains a frequently attributed source. Despite time-series-based scenario generation as a suitable approach for representing electricity price uncertainty, most DH optimization literature relies on deterministic sensitivity analysis or

limited stochastic representation. Thus, despite the acknowledgment of its practical significance, there also remains a gap regarding how comprehensively electricity price uncertainty is demonstrated in cases of electrified DH with TES.

3 Data and Methodology

This chapter elaborates on the data, system dynamics model, optimization approach, and parameters that were used to assess the potential for Monte Carlo-based evolutionary optimization of the primary decision strategy dictating the charging and discharging behaviour of the modelled district heating system. The underlying Systems Dynamic model, created in MATLAB's Simulink environment, was originally created by Roumpas (2025) as part of his Master's Thesis. Moreover, the data collection, processing, and utilization approach was also influenced by Roumpas (2025), with specific innovations to enhance the generalizability of results.

3.1 Data and Inputs

Real-world data collected for the optimization exercise included data on hourly wholesale (day-ahead) electricity prices acquired from Ember (2025) organization. Ember is an independent think tank that specializes in evidence-led power sector analysis. The specific dataset utilized was the European Wholesale Electricity Price Data, which captures day-ahead electricity prices with monthly, daily, and hourly frequencies. The horizon of the original dataset covers multiple years until the present day.

The dataset contains two essential ingredients required for our study: (i) Electricity Prices (represented in Euros per Megawatt hours) and (ii) corresponding timestamps represented in local time for the concerned regions.

The primary function of electricity prices within the SD model is to influence the decision for charging/ discharging the heat storage for each of the configurations of the overall DH system. These configurations are emphasized further in the Base Model Overview section of this thesis. Furthermore, the prices are also multiplied by the power consumed at a given timestamp to calculate the instantaneous cost for the overall system operation. This instantaneous cost is then, in turn, aggregated across the entire year for respective configurations to compute annual operational costs for each. The objective of the optimization is to minimize annual costs.

The electricity prices used in this study were based on stochastic simulations created by fitting an AR(1) GARCH(1,1) model to a unified electricity price regime based on actual prices from 2023 to 2025. Subsequently, student-t innovations applied on residuals after preserving deterministic hour-of-day and day of week preservations to generate numerous scenarios, each consisting of 8784 timesteps representing hourly prices across an entire year. A more detailed explanation of the price scenario generation is contained in the subsequent subsection.

These scenarios were subsequently fed into the SD model as detailed in the Stochastic Price Scenarios section of this thesis. The hourly day-ahead wholesale electricity prices for Finland for full years 2023 to 2025 were considered. The purpose of selecting these years was to specifically ensure the simulation model captured contemporary price dynamics without compromising regime structures and going beyond 2023 to previous years would bring about drastic structural change as present before the 2022 energy crisis.

Outdoor air temperature data is also an important input in the DH model and feeds directly into the computation of heat demand, which consequently represents the primary constraint for DH systems: heat demand must always be met. The data for hourly outdoor temperature with local timestamps was acquired from the Finnish Meteorological Institute's (FMI) open data. The FMI data is recorded in real-time, every ten minutes from dedicated observation stations across Finland, and resolutions could be narrowed based on need. FMI Datasets are available over wide horizons, and hourly temperature is computed by averaging the overall temperature across 10-minute interval recordings for a given hour.

3.1.1 Electricity Price Data and Stochastic Scenario Generation

While various methods could be used to create stochastic electricity prices, it was understood that the approach applied by Roumpas (2025) to fit an AR(1)-GARCH(1,1) model to electricity prices would be the best fit. The approach enables realistic replication of electricity price behaviour which is often highly autocorrelated and has time-varying volatility. ARMA and GARCH type models are extensively observed in studies related to electricity price forecasting (Chai et al., 2024). The AR(1) model helps capture the conditional mean and the GARCH (1,1) model helps in recording the volatility clustering element. When used in conjunction, the models compensate for the limitations arising from

their individual applications (Chai et al., 2024). Chai et al. (2024) also discusses the wide application of this combination across different price generation studies, such as application of the ARFIMAX-GARCH-type model for forecasting electricity prices in Italy by Bille et al. (2023, as cited in Chai et al., 2024). Their results support the use of such simpler autoregressive structures that often outperform complex specification by also adapting better to market volatility.

Roumpas (2025) also references the study by Chai et al. (2024) as the motivation for their choice to employ the augmented AR(1) GARCH(1,1) model to simulate electricity prices. Roumpas (2025) further explicitly states their choice for Student-t distribution for the stochastic term to allow for extreme spikes as witnessed in modern electricity price markets.

However, this approach has its limitations. Particularly, as highlighted by Roumpas (2025), these models assume stationarity in the data and treat positive and negative fluctuations as symmetrical. Although the simplicity of the model still advocates for its use, this study applied additional cautionary measures for electricity price simulations. In its current form, the AR(1) GARCH (1,1) model can simulate extreme negative values.

While negative prices have been observed in recent times, it is important to highlight their relative infrequency and containment within real world limits. For Finland's case, its membership of and abidance to policies laid out by the European Union Agency for the Cooperation of Energy Regulators offers insight into observed response to extreme price events. Negative prices are effectively floored at –500 EUR/ MWh and capped at 2800 EUR/ MWh (ACER, 2023), with conditional extension offered only for the price ceiling that has, thus far, not reached its limit. Reflecting these real-life limitations in simulated price scenarios is essential to effectively demonstrate potential for key decision optimization subject to stochastic electricity prices.

Roumpas (2025) does not not explicitly hint at any attempts to preserve key deterministic features that particularly define electricity price behaviour over different time periods. The preservation of these features improves proximity of simulated scenarios to realistic patterns and offers better application of experiment results.

The price scenarios generated in this thesis were intentionally generated by innovating beyond the AR (1) GARCH (1,1) approach adopted by Roumpas (2025) and extensively justified by literature as implied by Chai et al. (2024). The aforementioned precautions such

as preserving important deterministic elements and establishing realistic limits for extreme values were therefore applied.

Historical price series was decomposed into deterministic and stochastic components with deterministic components capturing significant patterns such as hour-of-day and day-of-week. The residual component was retained as the stochastic variation in prices and subsequently modelled using an AR(1)-GARCH(1,1) process with Student-t distributed innovations.

Multiple stochastic realizations of the residual process were then generated based on the estimated model, each of which was then supplemented by adding back the deterministic component in order to reconstruct full price trajectories. This additional process ensured that generated scenarios preserve systematic temporal structures as well as embody stochastic characteristics of historical electricity prices.

The price paths were then subjected to bounds post-simulation. With –500 EUR/ MWh as the lowest observed value and corresponding to the policy-induced floor mentioned above, the prices were tuncated below the minimum observed historical average. Positive spikes were also truncated above 2800 EUR.

3.1.2 Heat Demand Function

The same values for the demand function were used as employed by Roumpas (2025) in his thesis. The consequent function, despite having inherently data-dependent coefficients, remains appropriate for reuse in the optimization context, as heat demand is treated as an exogenous input that shapes operating conditions rather than the primary subject of investigation; the focus of the present study is on demonstrating the potential for improving costs by optimizing the offsets to an existing decision logic based on a chosen optimization approach rather than producing location-specific heat demand forecasts. The entire SD model developed by Roumpas (2025) is not explicitly tied to a specific district heating network, instead representing a generalized system that is informed by empirical data. The functional form also complies with the expected inverse relationship between outside temperature and heat demand.

The empirical basis for the function, as stated by Roumpas (2025), lies in estimating heat demand by setting it as the target of a simple linear regression, with outside temperature as the explanatory variable. The resulting R2 value of 0.949 and RMSE of 4.82 coupled with a high statisitcal significane ($p<0.001$) made it an ideal candidate over a relative second-degree polynomial, owing its applicability partly also to simpler interpretability (Roumpas, 2025).

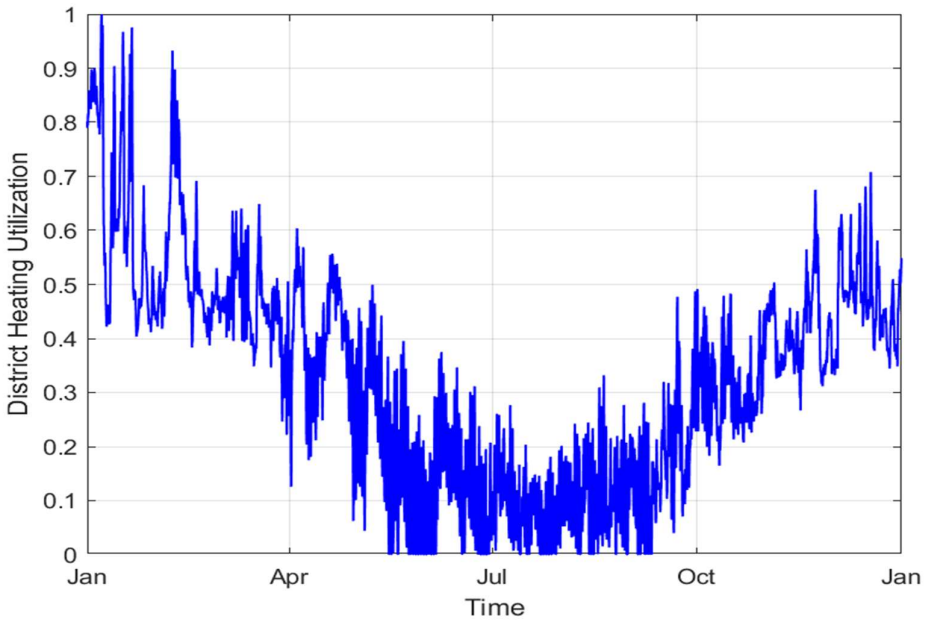

Figure 2: District Heating Utilization

3.2 Systems Dynamic Model Overview

At the core of this optimization study is a Systems Dynamic model that is created in MATLAB's Simulink environment by Roumpas (2025). This section explains the operational concepts of the Systems Dynamics Model. The model, as adopted in its original form, serves the simulation environment. The simulation environment is where the different performances are evaluated for the operational strategies. No structural modifications are made to the original model, allowing it to retain its logic and representation of dynamic DH system behavior under varying electricity prices.

Fundamentally, the interactions between energy storage, heat demand, and stochastic electricity prices are captured through a stock-and flow structure as is conventional of SD modelling. The three primary components of the system include:

(i) A thermal energy storage unit (heat tank)

(ii) Electricity grid

(iii) Head demand module

(iv) Additional storage units in the form of two batteries as add-on configurations of the system

In dynamic behavior of storage components is dictated by energy balance relationships. Stored energy in the thermal storage tank develops over time accodring to the following relationship:

$$E_{\text{tank}\,(t+1)} = E_{\text{tank}(t)} + P^{ch}_{\text{tank}(t)}\Delta t - P^{dis}_{\text{tank(t)}}\Delta t \quad (1)$$

Where:

$E_{\text{tank}(t)}$represents stored thermal energy at time t

$P^{ch}_{\text{tank}(t)}$ represents charging power

$P^{dis}_{\text{tank}(t)}$ represents discharging power

For batteries, a similar function is applied where:

$$E_{b,i(t+1)} = E_{b,i(t)} + P^{ch}_{b,i(t)}\Delta t - P^{dis}_{b,i(t)}\Delta t \quad (2)$$

Where:

$E_{b,i(t)}$ represents the stored thermal energy at time t with *i* indicating the battery represented (1 or 2)

$P^{ch}_{b,i(t)}$ represents the charging power of the respective battery

$P^{dis}_{b,i(t)}$ represents the discharging power of the respective battery

The energy balance relationships ensure the evolution of storaged energy is consistent with operational constraints as well as physical limits along the simulation. In the original model,

physical restrictions of the system are defined as fixed parameter values, particularly representing the limited storage capacities of each storage medium set at 400MWh. The charging saturation for the tank is set at 15 MW, while both batteries have the same charging/ discharging saturations set at 40 MW. These represent system limits as caused by assumed physical restrictions of the represented technologies.

Each simulation step is restricted by the system's design to ensure that demand, as an exogenous input derived from outdoor temperature, is satisfied. A hierarchical energy dispatch logic is followed, wherein heat demand is first met by discharging the heat storage tank, with the remaining power drawn directly from the batteries, if present, and then directly from the grid. The same charging preference is established with the tank charged first, prior to the batteries if applicable. Roumpas (2025) defines this relationship as:

$$Q_{demand(t)} = P^{dis}_{\text{tank}(t)} + \Sigma_i P^{dis}_{b,i(t)} + P_{grid(t)} \tag{3}$$

This ensures domain is satisfied in all cases while the system also exploits available energy prior to directly supplying heat by using grid electricity. A rule-based decision strategy governs how the storage components operate, based on current electricity price signals that are compared against a dynamic threshold logic that compares current price to the average price of electricity over the past 24 hours. When the price is below the threshold, storage across all applicable mediums is charged and during incidents where price is higher than the threshold, energy across storage mediums is discharged in line with the established charging/ discharge hierarchies.

The three configurations of the model are evaluated: (i) system with only heat storage tank (ii) system with tank + 1 battery (iii) tank + 2 batteries. Primarily, this allows studying the cost-based performance of added storage capacity and resulting flexibility. The model in its original form, and in the context of this present study, adheres to certain assumptions. Primarily, it is assumed that the system has identical external conditions for each configuration. It is also assumed that no heat losses occur during storage or conversion and that these processes are 100 percent efficient. Degradation effects or other operational costs such as overheads are also not assumed.

The SD model offers a robust and dynamic representation of DH operations. However, its decision rule is predefined and is not adaptive to various stochastic conditions beyond the moving threshold. The current study introduces a simulation-optimization framework that

targets parameters dictating the decision strategy, optimizing them under simulated, stochastic electricity prices. The SD model developed by Roumpas (2025), therefore, serves as the underlying simulation engine, complemented by an additional optimization layer to enhance its ability to identify operating policies that deliver better economic performance.

3.3 Simulation-Optimization Framework

In order to address the highlighted gaps in literature regarding the scarce application of advanced optimization algorithms as additional optimization layers for key parameters in SD-based models in the thermal energy storage-based DH context, this experiment aims to demonstrate the suitability and potential for applying such advanced methods to an existing SD-based model created by Roumpas (2025). Subsequent sections define the optimization goal, adjustments to the core SD model, the layered application method used and the suitability of using a genetic algorithm for the problem.

3.3.1 Optimization Goals

As stated earlier, there seems to be scarce evidence of applying advanced metaheuristc or derivative- free optimization algorithms to optimize SD-based DH representations, particularly in terms of addressing the suitability or potential for improvement in predefined decision logics employed in most DH studies. Often, these rules are defined based on assumed optimality or to reflect decision logic that is considered suitable for specific cases without iterative attempts to refine the logic or consider the potential for improved outcomes on top of the defined rule.

In the case of the underlying decision logic considered for the chosen base model, as developed by Roumpas (2025), the dynamic threshold is established based on a moving average consideration of the past 24-hour values of observed electricity prices. The decision to charge or discharge storage is made such that if the current price of electricity is lower than the average, storage is charged; however, if the price is higher than the avearge, the storage is discharged. This attempts to take advantage of the flexibility offered by the different storage configurations to leverage favourable electricity prices and maximize

profitability through cost reduction, assuming heat demand must be met and prices for heat are stable as well as no efficiency losses.

For the purpose of this study, this decision logic is slightly modified by adding offsets to the threshold and demonstrating the potential for advanced optimization techniques in improving decision rules for dynamic DH models. Thus, the charging and discharging decisions are modified across all three modelled storage setups: (i) thermal storage tank only (ii) storage tank + 1 battery (iii) storage tank + 2 batteries. The new logics saved within the charging/ discharging decision nodes for each storage block can thus be defined as:

If MA_Threshold + delta_tank > price(t): charge storage

If MA_Threshold + delta_b1 > price(t): charge storage

If MA_Threshold + delta_b2 > price(t): charge storage

If MA_Threshold + delta_tank < price(t): discharge storage

If MA_Threshold + delta_b1 < price(t): discharge storage

If MA_Threshold + delta_b2 < price(t): discharge storage

Where

MA_Threshold = moving average threshold established by Roumpas (2025) based on the average value of electricity prices over the past 24 hours

delta_tank = threshold offset for the storage tank

The optimization put in place aims to determine the appropriate values for delta_tank, delta_b1, and delta_b2 across all three operating cases, i.e. with only the storage tank; with the tank and one battery, as well as with the tank and both batteries. The objective is set to minimize the average annual operational costs of the DH-TES setup over the entire set of simulated electricity price scenarios.

3.3.2 SD Model Adjustments

To enable the core SD model to accommodate the threshold offsets, new blocks were added specifically to represent variables from the workspace using from workspace blocks.

Subsequently, the outputs of these from workspace blocks and the existing discrete FIR filter that calculated the moving average threshold were connected to a newly inserted sum block, the output of which represented the new offset-induced threshold. For a simple demonstration of the structural changes induced, Figures 1a and 1b show the different configurations for the original and adjusted structures within the storage tank subsections of the SD model respectively.

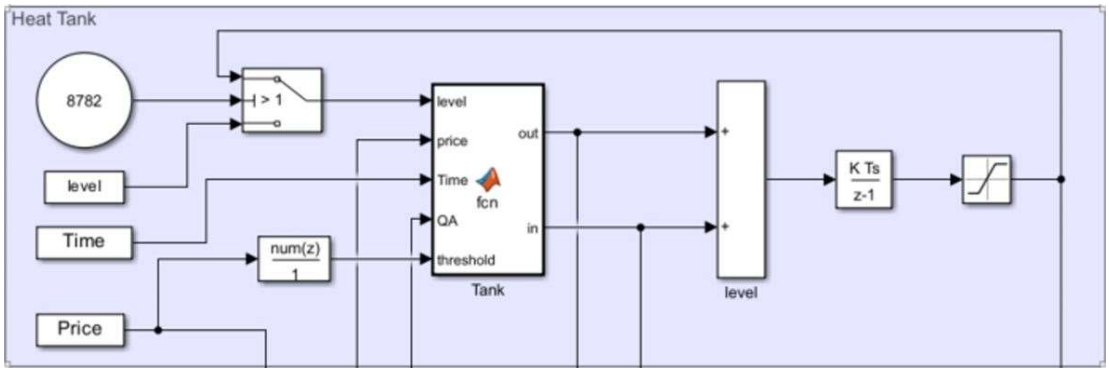

Figure 3a (ref: Roumpas (2025)): Heat tank Simulink Model

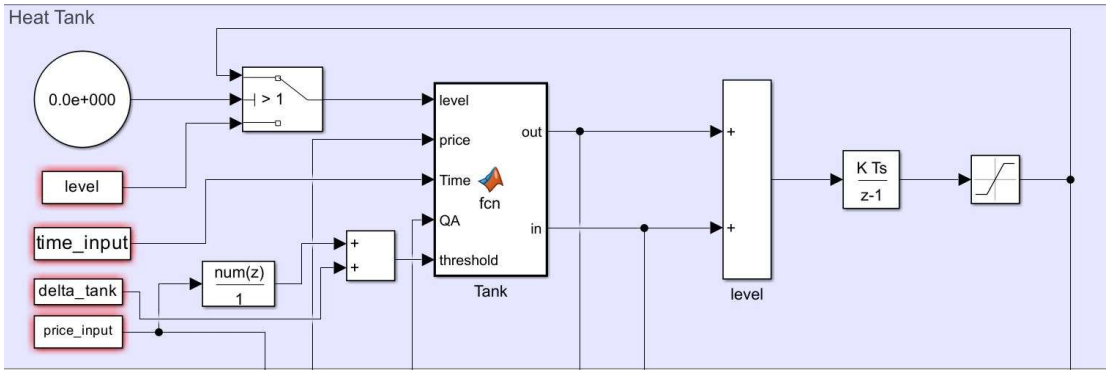

Figure 3b: Adjusted Heat Tank Simulink Model (adapted from Roumpas (2025))

3.3.3 Optimization Approach

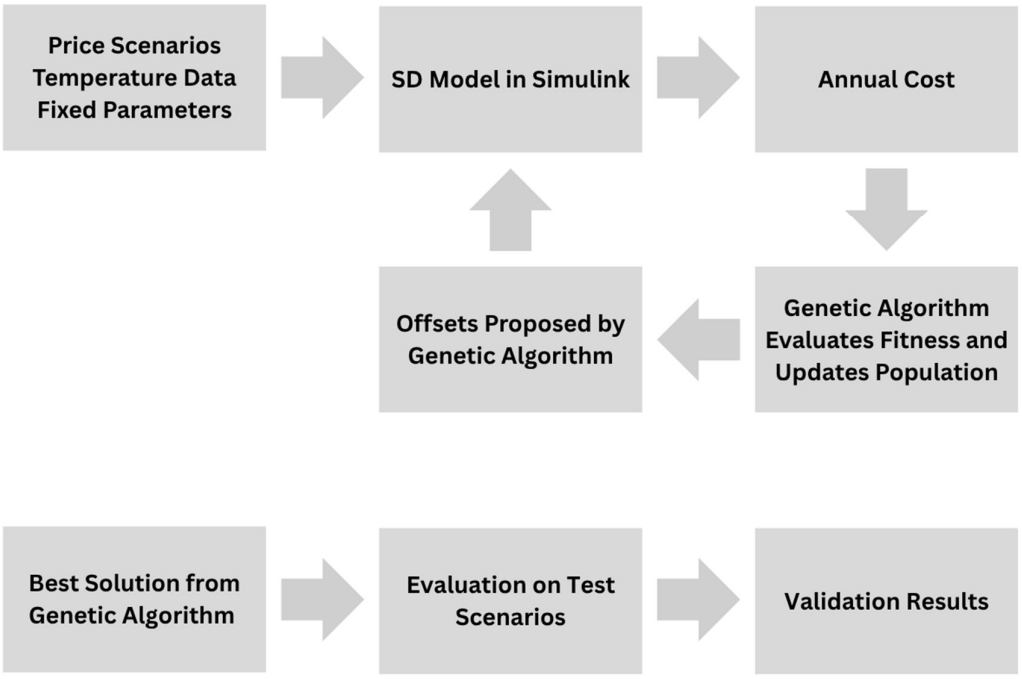

Figure 4: Representation of the simulation-optimization framework

This study uses an optimization framework that follows a layered structure, as shown in Figure 4. Within this structure, the SD model and optimization algorithm interact externally instead of being formulated into a single analytical program. The utilized SD model functions as the simulation engine and preserves the relationships between storage operations, heat demand, and electricity prices that influence the operating costs. Placed outside the SD model, the optimization component searches for improved values of the offsets that were introduced to the original moving average decision logic.

Therefore, the framework does not alter the existing stock-and-flow structure of the Simulink-based SD model. Rather, through repeated simulation runs, the performance of alternative parameterizations of the operating rule is evaluated under stochastic electricity price scenarios. The inputs used for each simulation run include hourly outdoor temperature data, hourly electricity-price scenarios over an 8784-hour horizon, fixed storage capacities, and fixed charging and discharging limits. Optimization is conducted independently for each of the three operating configurations, i.e., tank only, tank with one battery, and tank with

two batteries. Across each configuration, the offset parameters represent decision variables, and the remainder of the model structure, as well as exogenous inputs, are kept constant, allowing the optimization to focus explicitly on the impact of modifying the threshold-based control rule.

Inside the simulation layer, the Simulink model is given the candidate offset values as proposed by the optimizer alongside the exogenous inputs. The baseline threshold, based on the moving-average rule set by Roumpas (2025), is calculated, and the offset parameters are subsequently added to the baseline threshold, and the sum is introduced to the relevant storage-control blocks as the adjusted threshold. Resultantly, the hourly charging and discharging decisions are governed by this adjusted threshold.

The model then decides at each timestep whether the storage medium is to be charged or discharged based on the aforementioned defined rule while all imposed limits, including capacity, charging/ discharging limits, and the priority discharging rule between the storages remains strictly respected. These decisions are then propagated within the simulation to cover the full annual horizon, creating cumulative operating cost outcomes for the selected storage configuration.

Candidate solutions are evaluated at the optimization layers through repeatedly passing the solutions into the simulation and seeing the resulting annual cost performance. Within each candidate parameter vector, the model is run across a number of stochastic price scenarios, and the average annual cost across all scenarios is the set objective value. This is explained analytically in equation (4) in the following subsection.

The optimizer itself, however, is not based on a closed-form objective or direct mathematical information, interacting with the SD model as a black-box evaluator. The performance of the proposed offsets is assessed through simulation, with the resultant cost values being returned to the optimizer for each next search step. This remains an iterative process until the stopping criteria of the optimization algorithm trigger a halt.

The scenarios themselves are divided into training and testing sets using a 70/30 split. This split is enacted at the beginning of the optimization run and is consistent across all three optimization exercises in terms of the constituents of the train and test sets, across the thermal energy storage configurations, to ensure fair comparison. This adds an additional layer of structure and serves to validate the results of the optimization through out-of-sample

performance assessment, albeit within the same scenario-generation framework. Through this, it is also intended to discover whether optimized operating rules are generalizable beyond the calibration scenarios.

This overall framework constitutes a simulation-based policy search procedure featuring an inherited SD model that continues to represent the temporal and feedback-driven behavior of the DH system.

3.3.4 Choice of Optimization Algorithm

Given the layered structure that preserves the internal logic and stock-and-flow relationships of the underlying SD model created by Roumpas (2025), the optimization problem itself may be defined as a simulation-based policy search problem instead of a typical analytical optimization problem. As such, the decision variables are not derived from a closed-form mathematical formulation. Rather, offset parameters that alter the threshold-based logic of the inherited SD model are optimized. The overall behaviour of the objective within the optimizer may be summarized through the following:

$$\min J(\delta) \;=\; \frac{1}{N}\sum_{s=1}^{N} C_s(\delta) \tag{4}$$

Where:

$J(\delta)$ is the average operating cost for a tested offset vector

δ is the offset vector being tested by the GA

N is number of training scenarios

s is the scenario currently being tested

$C_s(\delta)$ is the annual operating cost using offset vector δ, calculated by a complete Simulink run over the entire time horizon of 8784 hours/ timesteps for a given scenario s

For each offset vector, the model must be executed across various stochastic electricity-price scenarios, and once all scenarios are run, the annual costs are averaged. Therefore, the objective function is not available to the optimizer in an explicit algebraic form and is instead revealed through only the simulation. This is highly similar to the definition of simulation optimization as provided in the review by Amaran et al. (2017), wherein objective values

and constraints may both be evaluated through simulation, sometimes under stochastic outputs and no direct access to derivatives. The review also notes that these settings are often black-box in nature, thus derivative information is either unavailable or unreliable when evaluations are noisy and have high computational costs.

Optimization literature in systems dynamics also frames the problem structure in a similar way. Dangerfield (2009) mentions how SD-based optimization is distinct in meaning from mathematical programming as the optimized model is both dynamic and often non-linear. Instead of solving static problems with explicit equations and single time-point solution, optimization in SD often includes parameter space searching to improve performance of the system in alignment with a chosen payoff or objective function.

Given this structure with the problem at hand, derivative-free, metaheuristic methods were considered suitable candidates. Genetic Algorithms (GA) and Particle Swarm Optimization (PSO) are both members of the broader family of metaheuristic techniques which are frequently employed in cases when mathematical methods have difficulties due to nonlinear, non-convex, and highly constrained search spaces (Papazoglou and Biskas, 2023). Papazoglou and Biskas (2023) describe GA as an evolutionary method wherein an initial population of candidate solutions is generated and evaluated against a fitness function, subsequently being evolved through selection, mutation, and crossover until a defined, maximum number of generations is reached; in contrast, PSO sees solutions represented as particles described by their respective velocities and positions and are moved across the search space based on knowledge of their own best position and the global best position as identified by the swarm.

Papazoglou and Biskas (2023) also highlight a small trade-off that PSO may involve a relatively lower computational burden whereas GA may have a slight advantage in accuracy in certain cases. Simultaneously, they also highlight that metaheuristics can better handle complex optimization problems, but require application-specific hyperparameter tuning which can be computationally demanding. Between the two identified options, choice of either depends heavily on the context.

The primary reason for the adoption of GA in this study is its alignment with the nature of the search itself. As population-based search methods, GAs do not require inherent assumptions such as smoothness, convexity, linearity or continuity of the objective function,

making it suitable for the developed framework, as the candidate offset vectors may be evaluated on the basis of their simulated performance. The optimizer's understanding of the internal SD model equations is not required, as it only requires a fitness value i.e. the average annual operating cost across the entire training set of scenarios, in our case.

Moreover, covered literature also offers applied examples of the use of GA in energy systems literature. In their study of heat source systems with thermal energy storage, Yu and Nam (2016) highlight GA as a suitable choice in problems with nonlinear optimization due to its ability to handle complex design factors. Their work also stresses the role of mutation that helps the algorithm avoid convergence to a partial optimum and supports a global search process.

Broader energy optimization literature further adds weight to the use of GA as a search method; however, limits and cautions are also addressed. According to Ganjehkaviri et al. (2017), GA is widely accepted in energy-systems optimization, where complex systems increase the appeal of evolutionary approaches. However, a key limitation is also to be considered that repeated runs may not reproduce similar solutions unless parameterization and convergence are carefully handled (Ganjehkaviri et al., 2017).

For the purpose of this study, PSO and other derivative-free methods also represent promising candidates. GA was, however, selected as a suitable and defensible optimization method for the present study, given its compatibility with black-box simulation, derivative-free search, and established use in energy-systems optimization literature.

3.3.5 Parameter Selection

The selected method, i.e. GA, may be represented in the applied case as a population of offset vectors at generation g as:

$$P_g = \left\{x_{g,1},\ x_{g,2},\ \dots,\ x_{g,L_p}\right\},\quad g = 1,\ 2,\ \dots.,\ G_{max} \tag{5}$$

Where

$$x_{g,i} \in [-200,\ 200]$$

P_g is the population at generation g

$x_{g,i}$ is the candidate solution vector

L_p is the specified population size

G_{max} is the specified maximum number of generations

Therein, when moving from a population at generation g to generation g+1, mutation, crossover, and selection are applied to previous populations to evolve towards a better solution. This is represented by:

$$P_{g+1} = M\left(C\left(S(P_g)\right)\right) \tag{6}$$

Where

S(.) represents the selection

C(.) represents the crossover

M(.) represents the mutation

Selecting the parameters to be applied within the Genetic Algorithm-based optimization framework represents a practical trade-off between optimality and computational expense. In the scope of the current study, the parameter selection was not chosen for theoretical optimality but instead treated as practical implementation choices that suit the structure of the problem at hand.

The primary influencing factor was the high computational costs associated with running the experiment. Given the simulation-based nature of the optimization, each candidate solution is evaluated through running the Simulink-based SD model repeatedly across stochastic electricity price scenarios. As a result, the parameters for the GA had to balance computational tractability with search capability.

Table 1: Chosen GA Parameters

Hyperparameter	**Setting**
Offset bounds	-200 to 200
Population Size ()	15
Maxgenerations	10
Selection function (S())	Uniform stochastic default

Mutation function (M())	Adaptive feasible default
Crossover function (C())	Scattered crossover of fraction 0.8 (default)

Since the optimization was conducted independently for the three system configurations: (i) tank only (ii) tank and one battery (iii) tank and two batteries, the dimensionality of the search space remained low since each case had one, two, and three decision variables, respectively. Across all three cases, the offsets introduced to the storage logic were the optimized variables and remained bounded within an interval from -200 to 200. This selection enabled a wide enough range to keep the search space meaningful whilst also restricting the search space to be computationally manageable and ensuring symmetry around zero. Furthermore, given that the offsets represent an added value to average electricity prices, which in observed data seldom exceeds these bounds, the choice was aimed to reflect changes applicable to realistic electricity price averages.

A population size (L_p in formula 5) of 15 with a maximum of 10 generations (G_{max} in formula 5) was used for all three optimization exercises. This choice was made to reach a meaningful compromise between diversity of search and runtime burden. While a greater number of generations would expand the exploratory space, the repeated call for simulations within each fitness evaluation would call multiple simulated scenarios across the full year hourly horizon.

Given the low-dimensionality of the search space and the expensive nature of each fitness function, a modest GA configuration was deemed suitable for this study as it aims to demonstrate the potential for improvement in DH-based SD models through evolutionary optimization with stochastic electricity prices. This is in line with what is reported in broader literature that highlights the influence of GA settings on computation time, convergence, and diversity whilst acknowledging that no parameterization is optimal universally as hyperparameter selection remains contingent to the specific application case (Papazoglou and Biskas, 2023).

The remaining, lower-level GA settings, particularly selection, mutation, and crossover used MATLAB's default settings with stochastic uniform selection, adaptive feasible mutation, and Scattered crossover of fraction 0.8 representing S(), M(), and C() in equation (6)

respectively. These default settings were kept due to their conventional acceptability and computational manageability, as their relative contribution to runtime is significantly less than the population size and number of generations.

Additional nonlinear constraints were not passed separately to the GA, rather feasibility was handled through the endogenous restrictions of the SD model wherein operational constraints such as charging/ discharging limits, priority discharge and charging logics, and storage capacities remained intact.

4 Results

This chapter aims to show the results obtained from the threshold-offset optimization exercise using a genetic algorithm optimization layer on top of the DH model developed by Roumpas (2025) across all three system configurations: tank only, tank with one battery, and tank with two batteries. This was performed on a training scenario set and subsequently evaluated over the test scenario set. The discussion involves the results of the optimization through cost-saving potentials and proposed offset values for the individual threshold rules for each storage media across all three configurations; scenario-wise savings; optimization behaviour and convergence, and out-of-sample validation.

4.1 Optimization Outcomes

Table 2: Configuration-wise outcomes for optimization and Scenarios

	Tank Only	**Tank + 1 Battery**	**Tank + 2 Batteries**
Baseline Cost on Training Set (Million EUR/ year)	11.537	10.314	9.807
Optimized Cost on Training Set (Million EUR/ year)	11.536	10.314	9.471
Cost Savings on training set (Million EUR/ year)	Negligible	Negligible	~ 0.34
Relative Improvement on Training Set	~ 0%	~ 0%	~ 3.4%
Percentage of Scenarios Improved in Training Set	50%	51.4%	92.8%
Baseline cost on Test Set (Million EUR/ year)	11.096	9.824	9.378

Optimized Cost on Test Set (EUR/ year)	11.086	9.832	8.935
Cost saving on Test Set (Million EUR/ year)	Negligible	Negligible	0.44
Relative Improvement on Test set	~ 0%	~ 0%	4.7%
Percentage of Scenarios Improved in Test Set	70%	46.7%	100%

Table 2 above shows the parameter values identified as best by the genetic algorithm, mean annual operating costs between baseline and optimized rules, and the share of scenarios where the optimized rule outperformed the baseline. As noted, the tank-only case shows a reduction in mean operating cost of 0.0067% on the training set; on the test set, this increases slightly to 0.0848%. In either case, the improvement is barely noticeable.

Across scenarios, 50% of training scenarios witnessed an improvement when shifting from the baseline rule to the optimized rule, whilst this improvement was 70% in the test set. As a whole, these values indicate a very small and negligible effect of the optimization in improving cost-saving potential for the tank-only configuration.

A similar negligible effect was noticed in the tank with one battery configuration, where the mean annual operating costs between the baseline and optimized rule only improved by 0.0042% on the training set and, in fact, worsened by 0.0803% (-0.0803%) in the test set. Moreover, while 51.4% of scenarios were improved in the training set, only 46.7% saw improvement over the test set. This too shows almost no effect of the optimization exercise.

The third configuration, i.e. tank with two batteries, optimization yielded much more promising results as an improvement of 3.432% over the baseline costs for the training; this improvement further increased to 4.7194% for the test set. Collectively, the test and training results showed improvement potentials of 336,585.91 and 442,500.08 EUR respectively.

The magnitude of differences in performance across the training and test sets across each configuration may be understood as a reflection of their overall potential for improvement over the baseline moving average threshold, as in the inherited SD model. In cases of the

tank only or the tank with one battery, where negligible differences were noticed, and offsets identified by the GA remained close to zero, the available operational flexibility was likely not sufficient to further capitalize on the stochastic variations between the training and test scenario sets. However, in the case of the tank with two batteries, a significant potential was witnessed across both training and test sets, with further improvement witnessed in performance over the test set scenarios. This substantiates the idea that greater storage flexibility offers greater room to exploit price variations more efficiently. The stated improvements are further referenced in relative terms in Figures 5a and 5b below.

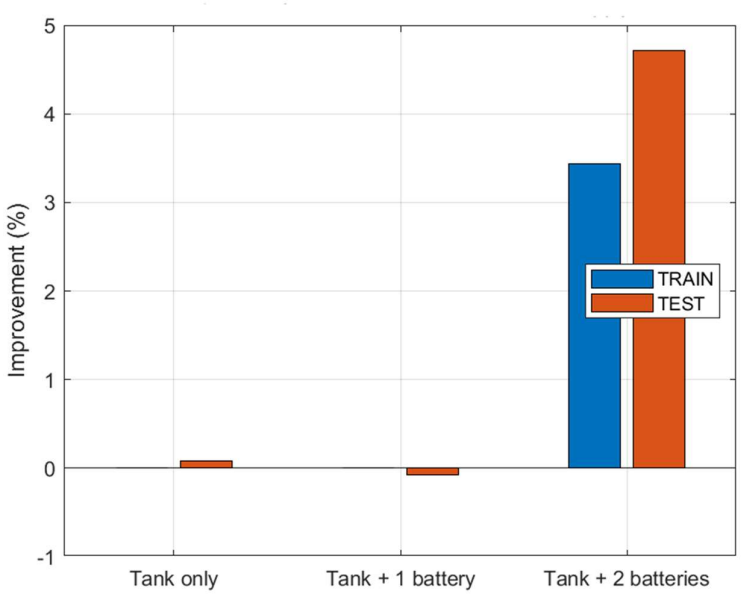

Figure 5a: Configuration-wise percentage improvement in cost across training and test scenarios

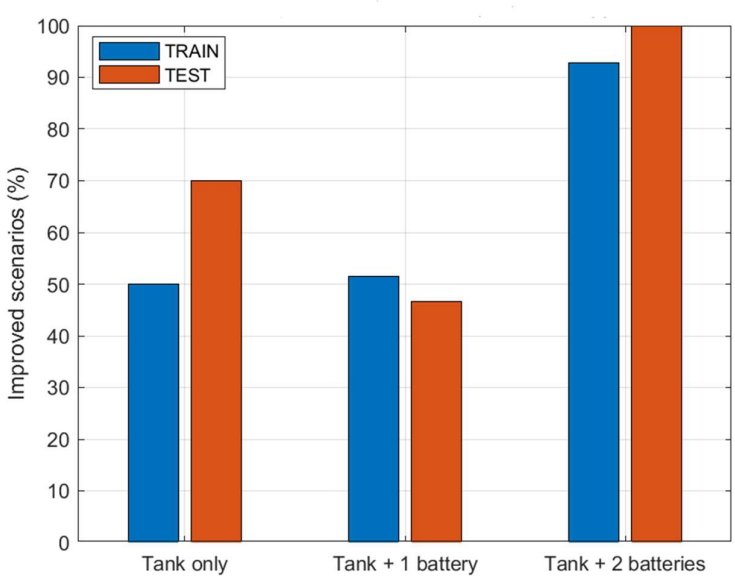

Figure 5b: Configuration-wise percentage overall improvements in scenarios

Apart from the financial improvements noted in the preceding section, the actual decision rule is influenced by the selected offset values as identified by the GA optimization. The offsets are added to the moving average threshold rule, which was originally part of the SD model's internal decision framework, and are optimized using an external optimization layer with the GA in MATLAB's Global Optimization Toolbox acting as the optimizer. Here, delta_tank represents the offset added to the threshold rule in the tank storage simulation model, where as delta_b1 and delta_b2 represent the same for batteries 1 and 2 respectively. Logically, delta_b1 is only present in configurations including the tank and one battery and the tank + 2 batteries while delta_b2 is only present in the tank + 2 battery configuration; all the while, delta_tank is present across all three storage configurations.

The best threshold offsets as determined by the optimizer are noted in Table 3 below.

Table 3: Best threshold offsets identified by optimization

Configuration	**delta_tank**	**delta_b1**	**delta_b2**
Tank Only	-1.0800	-	-
Tank + 1 Battery	-0.9623	1.0839	-
Tank + 2 Batteries	-1.4623	0.8339	-23.5023

In the tank-only configuration, the best offset is estimated to be -1.08, whereas for the tank + 1 battery configuration, the best offset vector is (-0.9623, 1.0839). These are notably small values in relation to the search interval of [-200, 200] and correspond to negligible differences from the baseline moving average rule.

Again, a considerable difference was only noticed in the tank + 2 batteries configuration where, whilst the best offset values delta_tank and delta_b1 remained close to zero, the value for delta_b2 was far from zero and the overall best offset vector was identified as (-1.4623, 0.8339, -23.5023). This shows that the largest deviation from zero is associated with the second battery and the threshold rule applied to battery 2 is shifted downwards by 23.5 cost units (EUR/ MWh). Therefore, the optimized parametrization involves a substantially larger threshold adjustment for battery 2 than the tank or battery 1.

4.1.1 Local robustness check for tank + 2 batteries configuration

In order to assess the level of optimality as identified by the GA optimization, a local robustness check was conducted for the tank + 2 batteries. This was exclusively performed for this configuration due to its position as the only setup that displayed meaningful results. The test was applied by adding positive or negative Perturbations (EUR) to the GA-optimized offsets in order to see The results of this robustness check are displayed in Table 4 below.

Table 4: Local Robustness Check Results for Tank + 2 batteries configuration

Offset Targeted	Perturbations added to the GA-optimized threshold offset (EUR)	Difference from GA-Selected offset costs (Percentage)	Improvement from Baseline (Percentage)
delta_tank	-10	0.153	4.574
delta_tank	-5	-0.343	5.046
delta_tank	5	1.376	3.408
delta_tank	10	3.181	1.688

delta_b1	-10	-0.519	5.214
delta_b1	-5	-0.524	5.219
delta_b1	5	1.053	3.716
delta_b1	10	2.370	2.461
delta_b2	-10	1.247	3.532
delta_b2	-5	0.429	4.311
delta_b2	5	-0.078	4.794
delta_b2	10	0.891	3.871

From the attained results, limited room for further improvement was noted in shifting values of delta_tank, delta_b1, and delta_b2. However, noticeably, the results remained more favourable closer to the GA-optimized offset values as witnessed via better cost reduction in –5 and 5 perturbations compared to –10 and 10. Furthermore, the changes compared to the baseline remain significant, whereas the additional improvements over the GA-optimized offset value remain less than 1%. This highlights that the optimized values identified by the GA lie in a generally favourable region; however, as the optimization setup limited the number of generations in light of computational costs, its value in improving cost performance remains justified.

4.2 Scenario-wise distribution of cost savings

To examine the gains or losses from the GA optimization and their distribution across individual scenarios, scenario-wise cost savings were calculated based on the difference between the baseline annual cost and the optimized annual costs. Therein, positive values indicate a reduction in annual cost because of the optimization, whilst negatives show the baseline performed better than the proposed optimized solution.

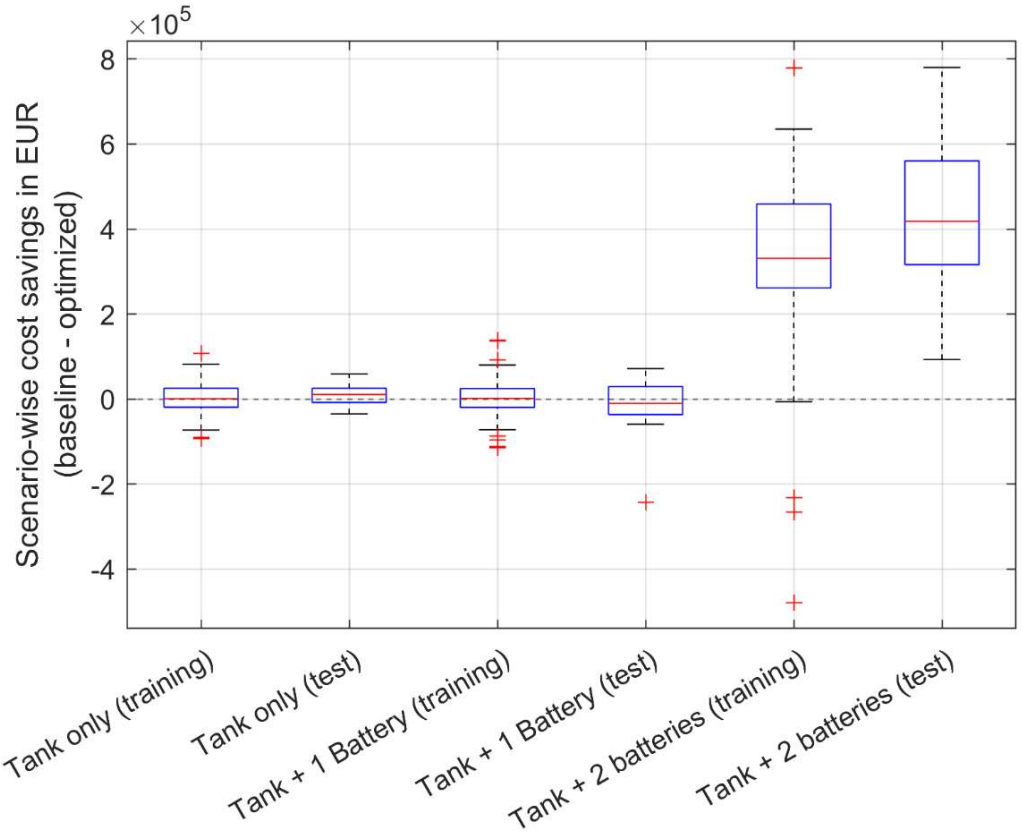

Figure 6: Spread of cost savings by configuration and scenario set

Figure 6 highlights the scenario-wise cost saving distributions across configurations as well as training and test scenarios. There is a concentration of values around zero in the tank-only case, where the test scenarios show a slight positive shift in the distribution; however, it shows that overall operating cost is changed only marginally across scenarios. For the case of the tank and one battery, a mild negative shift is witnessed in the test case, while both distributions remain centred close to zero.

Again, a major difference is noted in the case with the tank and two batteries, where an upward shift is visible in both training and test scenarios, more pronounced than the other two cases. There seems to be a small share of adverse cases in the training set Figure, yet the test scenario plot shows the median and interquartile range remaining exclusively in the positive region; this shows the optimized rule reduces annual cost across the central part of the held-out distribution.

The optimization objective currently employed aims to reduce average annual costs across storage configurations by testing various offset vectors added to the inherited moving average threshold logic without explicitly redesigning the entire dispatch logic. As such, the potential for cost savings in individual scenarios as well as across entire sets depends on the flexibility offered within the storage configuration. In the tank only as well as tank + 1 battery case, this flexibility is limited and offsets remain close to zero, allowing only marginal changes in annual costs across scenarios and an overall concentration of scenario-wise savings around zero in both cases.

This is not the case with the tank + 2 batteries configuration, where the second battery offers the optimizer an additional storage component to control. This also means that the threshold rule can more efficiently differentiate between use of each storage component while ensuring heat demand constraints within the SD model are also met. As such, the substantially more negative offset for battery 2 suggests its distinct role in managing cost-responsive dispatch. Thus, this contributes to greater positive saving values for the tank and two batteries configuration and serves as a key takeaway regarding configuration-wise performance differences in the optimization.

4.3 Optimization Behavior

Although the optimization was performed with modest hyperparameters to account for limited computational resources, studying the convergence behavior allows us to understand how the search progressed. Figures 7a to 7c show the evolution of the best and mean population objective values for 10 generations across the three storage configurations. Notably, a decline over successive generations can be witnessed for the mean population values; this tells us that the quality of the candidate solutions did improve during the search. However, the behavior of the objective values did differ across all three storage configurations.

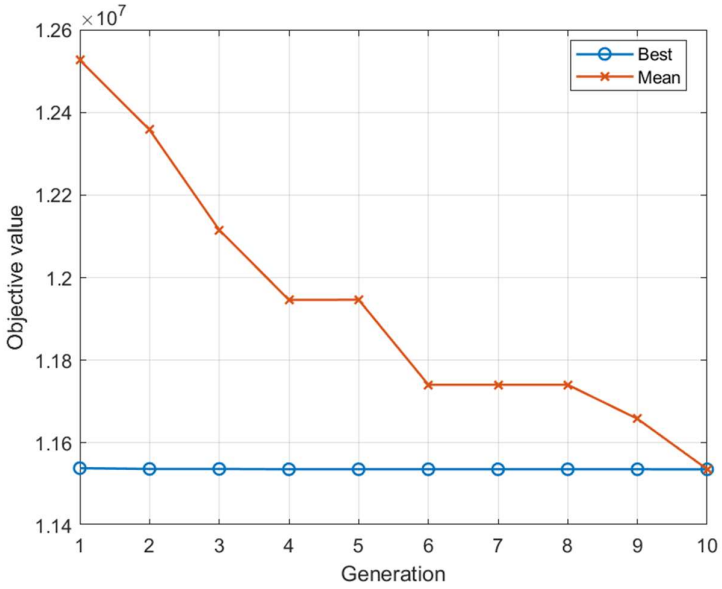

Fig 7a. GA Optimization Convergence plot for Tank Only

The tank only setup shows the best value objective changes barely once the early generations were conducted, but the mean population objective decreased reasonably. This suggests that early in the run, a near-best solution was identified and later generations led to improvement in the average population quality and not a meaningfully better solution.

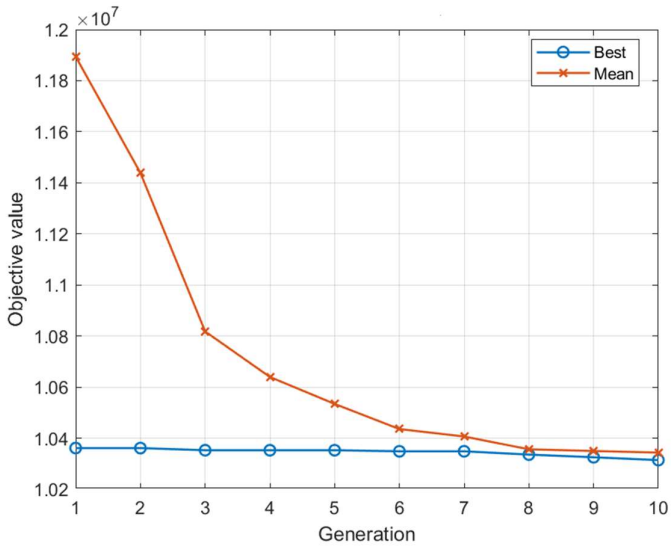

Fig 7b. GA Optimization Convergence plot for Tank and 1 battery

A similar behavior for the best objective was noted, albeit with a minor downward slope, especially towards later runs. However, the overall intensity of improvement in the best

objective was mild. The reduction in the mean population objective, however, was more pronounced than in the tank-only case, which shows that the candidate population saw a continued improvement.

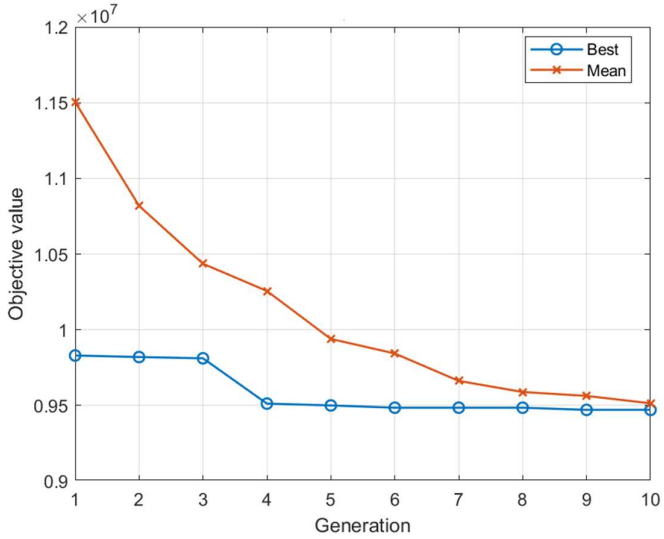

Fig 7c. GA Optimization Convergence plot for Tank and 2 batteries

The clearest improvement in the best objective value was observed in the tank and 2 batteries case, as shown above. The mean population objective and the best objective both showed declines over the run, where, early in the run, the best objective value showed a large improvement before eventually becoming flatter.

The decreasing distance between the mean and best objective values suggests that the population was progressively concentrating around better solutions, and there was noticeably less dispersion at the later stages of the run. This and the flattening of the best objective curves hint towards a practical stabilization being achieved within the computational budget, despite a global optimum not being conclusively reached. These findings can only be interpreted within the context of the chosen GA hyperparameters while clearly showing that the search progressed towards lower-cost regions within the selected parameter space.

4.4 Results Analysis

The optimization exercise focused on demonstrating the potential of applying an external optimization layer on top of the SD model and proposing offsets to the moving average threshold logic. As such, the demonstrated results must be interpreted within this context. The absence of meaningful improvements in the tank-only and tank + 1 battery case contrasts with the visible improvement potential as witnessed in the tank + 2 batteries case. These outcomes are consistent in both the training and test cases for each configuration.

This is accompanied by the near-zero offset values for the tank and first battery across all three configurations; however, the optimized offset for the second battery is highly negative, showing that benefit is identified by the optimizer in using this additional layer of controllable flexibility. This indicates that the degree of flexibility may greatly impact how well optimization benefits DH operations. The reliability of the observed improvement is backed by the local robustness check applied to the optimized offsets for the tank and two-battery configuration. While a global optimum cannot be claimed, the test shows how the attained offsets lie in a generally favourable region. This comes despite a modest budget being allocated to the GA optimizer, especially in terms of the number of maximum generations.

The results attained remain subject to the design of the optimization exercise, its computation budget, and other limitations such as simulation-level constraints and sources of uncertainty. However, the potential for improvement using an external optimization layer on top of an SD-based simulation model is demonstrated in the tank + 2 battery case. The results are in line with the general logic uncovered in literature, particularly the works of Vanhoudt et al. (2018), which notes the role of storage configurations in determining profit, as well as the work of Quaggiotto et al. (2021), which emphasizes the value of storage being tied to the system's ability to leverage it.

5 Conclusion and Discussion

This thesis studied thermal energy storage optimization in district heating when electricity prices are uncertain. A layered optimization approach was created and applied to an inherited system dynamics model that served as the simulation engine representing storage management operations in district heating. The results showed that the simulation optimization framework helped improve the economic performance of TES-enabled DH operations; however, this was true only for the case of the tank + 2 batteries, showing that the benefit depends on the degree of flexibility offered by the storage configuration. This idea, although subject to the particular optimization setup, is largely consistent with the notion of TES-enabled flexibility and resulting economic benefits as identified in literature.

5.1 Answering the Research Questions

Answer to Research Question 1

Research Question 1 was regarding the exploration of existing knowledge about heat storage optimization under electricity price uncertainty and the types of dynamic models used for optimization. According to the findings, the optimization of thermal energy storage heavily includes using the flexibility offered by storage to shift production of heat to times when electricity prices offer economic benefits. Whereas the approach for conventional CHP systems includes increased production during times of high electricity prices in order to capitalize on high revenue potential, this approach is reversed in the case of electrified DH systems as cost becomes the main component. Thus, the value of storage is operational just as much as it is technical, as it depends on the control strategy or decision as much as it does on the technical specifications of the system. Furthermore, electricity price uncertainty has widely accepted importance but is not represented comprehensively; instead losing favor in place of uncertainty simplification through deterministic cases or limited scenario structures.

At present, optimization studies in this context are most formulated through mathematical optimization or control-oriented approaches that offer various methods to represent storage. However, the value remains subject to modelling choices, which seldom depict time-evolving, feedback-oriented interactions, especially in the DH-TES context. SD modelling

approaches offer potential to create models that offer this additional lens, despite their limited demonstrated application in studies in this domain. Even less so are SD models representing DH optimized using external optimization, even though established SD literature explains the potential of jointly using SD models with external optimizers. This gap in methodological representation is also what motivates the empirical part of this thesis.

Answer to Research Question 2

The second research question focused on exploring how a combination of Monte Carlo simulation and evolutionary optimization may be used to improve the economic performance of TES operations under price uncertainty. To achieve this, the thesis generated electricity price scenarios using AR(1) GARCH(1,1) fitted onto the residuals of electricity prices after inter-temporal deterministic patterns were conserved. The scenarios were further bounded based on real-life policies that limit extreme ends of electricity price values, even as the market becomes more volatile. This approach built on top of existing work as observed across literature as well as the motivating research conducted by Roumpas (2025), whose original SD model was retained as the simulation engine for this optimization exercise.

A genetic algorithm was used as an external optimization layer on top of the underlying SD simulation engine and stochastic price scenarios were injected. The objective of the optimization remained to minimize mean annual electricity costs across a set of training scenarios by adjusting offsets added to the underlying, moving average charging/ discharging rule. This setup helps understand the impact of optimizing the charging/ discharging rule as well as the component-wise behavior across the three storage configurations that would minimize electricity costs.

The results show that the adopted approach can improve economic performance; however, it may not be uniformly achieved in all cases of storage setups. Particularly, the results obtained from the specific experimental setup show that optimization benefits may best appear when a higher degree of storage flexibility is present in the system, as only the configuration with the highest TES flexibility, i.e. with tank and two batteries, showed meaningful improvement. Furthermore, the value of the particular optimization setup also lies in uncovering the role of component-wise control behaviour under stochastic price conditions – this direction may be very useful in informing operating strategy design in multi-configuration setups similar to the ones in this study. The offsets optimized were

separate, not being optimized or set uniformly across all three components; this revealed that the best performing setup only had meaningful behavioural change implemented in the second battery to achieve best results within the current optimization scope and setup.

5.2 Limitations

The findings of this study must be interpreted with caution and due consideration to its limitations. In terms of the assessed literature, whilst underrepresentation in SD-based DH optimization studies highlights a noteworthy gap that this study aimed to contribute to, that gap also restricts the comparability of the applied approach. Moreover, the argumentative justification for the chosen approach relies on drawing connections from adjacent literature.

At its core, the optimization exercise relied on the feedback-oriented SD model as developed by Roumpas (2025). No extensive changes were applied to the model and the inherited assumptions were carried over to the results of the optimization experiment. These include assumptions of no efficiency losses, technology failure, or wear and tear as associated cost outcomes are not modelled. Moreover, the uncertainty was exclusively modelled for electricity prices, whereas other factors like heat demand and external cost components also account for considerable uncertainty and may significantly impact economic performance in actual settings.

The optimized parameters concerned only offsets to the existing moving average threshold as an adjusted control rule. Therefore, a new control strategy was not fully optimized. The optimized rules, therefore, should not be interpreted as directly applicable in real life strategies unless validated further. Instead, the optimization exercise should primarily be viewed through a comparative lens that demonstrates the economic potential, limitations, and component-specific behavioral implications for the optimization setup and, by extension, similar setups. Furthermore, the optimization concerned only operational-level decisions; economic impact may potentially be significant even prior at the system-design phase and is not optimized in this study.

In terms of the optimization itself, key parameters chosen for GA, such as population size and maximum generations, were extremely modest due to the limited computational budget at hand. This restricts the depth of the search and does not allow the full search potential o

the GA to be assesed. By virtue of this, as well as inherent heuristics-based nature of GA, the solution is not a guaranteed global optimum, even if, in cases such as the tank + 2 batteries, a meaningful improvement is witnessed.

Whilst validation was performed by testing the performance of optimized offsets in an out-of-sample scenario set, these test scenarios themselves were part of the same scenario pool as the training set. This does not constitute an exhaustive validation of the obtained results, even as it serves the purpose of this study. A fully independent, real-world future dataset was not used to test actual performance.

Conclusively, these limitations signal that the results and conclusions of this study must be interpreted in the context of the applied framework and are not universally generalizable. Their value lies in a comparative assessment of optimization potential.

5.3 Future Research

While this study possessed inherited and budget-induced limitations, future research could aim to build on the proposed simulation-optimization framework using a more detailed heating model that captures greater realism and incorporates multiple sources of uncertainty. Such an extension could bring derived solutions closer to real-life applicability and serve to uncover nuances and limitations not yet highlighted. Moreover, the setup currently adopted may be extended meaningfully by extending computational budgets and using larger GA configurations, repeated optimization trials, or alternative optimization methods such as PSO.

It is also worthwhile to investigate the impact of such an optimization framework on the design aspect of DH storage configurations, instead of only the operational rule. Even within the operational rule, alternative control strategies or constant threshold values could be optimized to see and compare economic performance improvements. Their findings may also be compared across the lens of equipment reliability and costs associated with adherence to proposed strategies; this could help identify strategies that are not only cost-reducing in simulation, but also technically and economically feasible.

The current setup could also be expanded to capture the optimization-induced behavior across similarly sized additional storage media or varying storage capacities of current

media. This, in particular, may help understand the impacts of priority rule setups and associated flexibilities versus simply increasing storage capacities. This can also help investigate why, in the current study, the second battery showed the most meaningful change. In all, the space for exploring DH-TES optimization using SD and external optimization layers is relatively untapped. Therefore, a multitude of directions exist where future research may contribute. Greater exploration of this domain may become increasingly relevant as electrified DH systems coupled with TES become more common.

References

Amaran, S., Sahinidis, N.V., Sharda, B. and Bury, S.J. 2016. Simulation optimization: a review of algorithms and applications. Annals of Operations Research, 240, pp. 351–380. Available at: https://doi.org/10.1007/s10479-015-2019-x

Bauer, F. 2025. Two-stage stochastic optimization of district energy systems: a case study from Norway. Master's thesis. TU Wien, Institute of Energy Systems and Electrical Drives.

Billé, A.G., Gianfreda, A., Del Grosso, F. and Ravazzolo, F. 2023. Forecasting electricity prices with expert, linear, and nonlinear models. International Journal of Forecasting, 39(2), pp. 570–586. Available at: https://doi.org/10.1016/j.ijforecast.2022.01.003

Bjørnskov, J., Mortensen, L.K., Filonenko, K., Shaker, H.R., Jradi, M. and Veje, C. 2021. Optimization of district heating production with thermal storage using mixed-integer nonlinear programming with a new initialization approach. Energy Informatics, 4(Suppl. 2), 34. Available at: https://doi.org/10.1186/s42162-021-00150-y

Ceruti, A., Lambert, J. and Spliethoff, H. 2025. Integrating renewable energy and thermal storage in district heating networks: A design optimization approach. Energy Conversion and Management, 345, 120323. Available at: https://doi.org/10.1016/j.enconman.2025.120323

Chai, S., Li, Q., Abedin, M.Z. and Lucey, B.M. 2024. Forecasting electricity prices from the state-of-the-art modeling technology and the price determinant perspectives. Research in International Business and Finance, 67, 102132. Available at: https://doi.org/10.1016/j.ribaf.2023.102132

Christidis, A., Koch, C., Pottel, L. and Tsatsaronis, G. (2012) 'The contribution of heat storage to the profitable operation of combined heat and power plants in liberalized electricity markets', Energy, 41, pp. 75–82. doi: 10.1016/j.energy.2011.06.048.

Dangerfield, B. 2009. System dynamics models, optimisation of. In: Meyers, R.A. ed. Encyclopedia of Complexity and Systems Science. New York: Springer, pp. 9034–9043. Available at: https://doi.org/10.1007/978-0-387-30440-3_542

Ember. 2026. European wholesale electricity price data. Ember. Available at: https://ember-energy.org/data/european-wholesale-electricity-price-data/ (Accessed: 24.2.2026).

Energiateollisuus ry (2024) District heating in Finland 2023. Helsinki: Finnish Energy. Available at: https://energia.fi/en/statistics/statistics-on-district-heating/district-heating-and-cooling/ (Accessed: 4.11.2025).

European Union Agency for the Cooperation of Energy Regulators. (2023, January 10). Decision No 01/2023 on the Nominated Electricity Market Operators proposal for the harmonised maximum and minimum clearing price methodology for the single day-ahead coupling. Decision No 01/2023

Finnish Meteorological Institute. 2026. Hourly outdoor temperature observations, 2024. FMI Open Data. Available at: https://en.ilmatieteenlaitos.fi/download-observations (Accessed 24.2.2026).

Ganjehkaviri, A., Mohd Jaafar, M.N., Hosseini, S.E. and Barzegaravval, H. 2017. Genetic algorithm for optimization of energy systems: solution uniqueness, accuracy, Pareto convergence and dimension reduction. Energy, 119, pp. 167–177. Available at: https://doi.org/10.1016/j.energy.2016.12.034

Guericke, D., Schledorn, A. and Madsen, H. 2024. A generic stochastic network-based formulation for production optimization of district heating systems. Energy Systems. Available at: https://doi.org/10.1007/s12667-024-00698-

Hiltunen, P., Lindroos, T.J. and Rämä, M. 2025. The impact of electric boilers and heat storages in the Nordic power markets and district heating systems. Cleaner Engineering and Technology, 27, 101028. Available at: https://doi.org/10.1016/j.clet.2025.101028

Jaanto, J., Sihvonen, V., Honkapuro, S., Riikonen, J., Price, A., Ylönen, M. and Kivioja, V. 2023. Thermal storage deployment in the framework of current electricity market design. In: 2023 19th International Conference on the European Energy Market (EEM). Lappeenranta, Finland, pp. 1–7. Available at: https://doi.org/10.1109/EEM58374.2023.10161744

Javanshir, N. and Syri, S. (2023) 'Techno-economic analysis of a highly renewable and electrified district heating network operating in the balancing markets', Energies, 16(24), 8117. doi:10.3390/en16248117

Malmgren, F. and Skoglund, H. 2021. Optimization of district heating systems for future uncertain conditions: Deterministic and stochastic linear modelling. Master's thesis. Chalmers University of Technology, Department of Space, Earth and Environment, Division of Energy Technology.

Papazoglou, G. and Biskas, P. 2023. Review and comparison of genetic algorithm and particle swarm optimization in the optimal power flow problem. Energies, 16(3), 1152. Available at: https://doi.org/10.3390/en16031152

Petersen-Dyggve, N. (2024) A review of non-combustible technologies in district heating production in Finland. Master's thesis. Lappeenranta–Lahti University of Technology LUT, LUT School of Energy Systems, Energy Technology.

Quaggiotto, D., Vivian, J. and Zarrella, A. 2021. Management of a district heating network using model predictive control with and without thermal storage. Optimization and Engineering, 22, pp. 1897–1919.Available at:https://doi.org/10.1007/s11081-021-09644-w

Refsgaard, J. 2021. Dealing with uncertainties in the operational planning of district heating plants. Master's thesis. Aalborg University, Sustainable Energy Planning and Management.

Roumpas, A. 2025. Economic feasibility of heat storage systems under price uncertainty – A system dynamic approach. Master's thesis. Lappeenranta–Lahti University of Technology LUT, LUT Business School, Business Analytics.

Sihvonen, V., Ollila, I., Jaanto, J., Grönman, A., Honkapuro, S., Riikonen, J. and Price, A. 2024. Role of power-to-heat and thermal energy storage in decarbonization of district heating. Energy, 305, 132372. Available at: https://doi.org/10.1016/j.energy.2024.132372

Sinha, R., Bak-Jensen, B., Radhakrishna Pillai, J. and Zareipour, H. 2020. Flexibility from electric boiler and thermal storage for multi energy system interaction. Energies, 13(1), 98. Available at: https://doi.org/10.3390/en13010098

Sporleder, M., Rath, M. and Ragwitz, M. 2022. Design optimization of district heating systems: A review. Frontiers in Energy Research, 10, 971912. Available at: https://doi.org/10.3389/fenrg.2022.971912

State Treasury (2025) Carbon neutral Finland 2035. Helsinki: State Treasury. Available at: https://www.treasuryfinland.fi/investor-relations/sustainability-and-finnish-government-bonds/carbon-neutral-finland-2035/ (Accessed: 3.11.2025).

Statistics Finland (2024) 12vl – District heating consumption by sector, 1970–2023. StatFin database: Energy supply and consumption. Available at: https://pxdata.stat.fi/PxWeb/pxweb/en/StatFin/StatFin__ehk/statfin_ehk_pxt_12vl.px/ (Accessed 3.11.2025).

Talebi, B., Mirzaei, P.A., Bastani, A. and Haghighat, F. 2016. A review of district heating systems: modeling and optimization. Frontiers in Built Environment, 2, 22. Available at: https://doi.org/10.3389/fbuil.2016.00022

Vanhoudt, D., Claessens, B.J., Salenbien, R. and Desmedt, J. 2018. An active control strategy for district heating networks and the effect of different thermal energy storage configurations. Energy and Buildings, 158, pp. 1317–1327. Available at: https://doi.org/10.1016/j.enbuild.2017.11.018

Verrilli, F., Parisio, A. and Glielmo, L. 2016. Stochastic model predictive control for optimal energy management of district heating power plants. In: 2016 IEEE 55th Conference on Decision and Control (CDC). Las Vegas, NV, USA, pp. 807–812. Available at: https://doi.org/10.1109/CDC.2016.7798376

Wu, Z. and Xu, J. 2013. Predicting and optimization of energy consumption using system dynamics-fuzzy multiple objective programming in world heritage areas. Energy, 49, pp. 19–31. Available at: https://doi.org/10.1016/j.energy.2012.10.030

Yu, M.G. and Nam, Y. 2016. Study on the optimum design method of heat source systems with heat storage using a genetic algorithm. Energies, 9(10), 849. Available at: https://doi.org/10.3390/en9100849

Zhang, Z., Zhou, X., Du, H. and Cui, P. 2023. A new model predictive control approach integrating physical and data-driven modelling for improved energy performance of district heating substations. Energy and Buildings, 301, 113688. Available at: https://doi.org/10.1016/j.enbuild.2023.113688

Ziemele, J., Gravelsins, A., Blumberga, A., Vigants, G. and Blumberga, D. 2016. System dynamics model analysis of pathway to 4th generation district heating in Latvia. Energy, 110, pp. 85–94. Available at: https://doi.org/10.1016/j.energy.2015.11.073

Appendix 1: Literature Covered in Literature Review

Sr No.	Title	Author, Year
1	The risks of electrified district heating in Finland's cold climate	Javanshir et al., 2024
2	A REVIEW OF NON-COMBUSTIBLE TECHNOLOGIES IN DISTRICT HEATING PRODUCTION IN FINLAND	Petersen-Dyggve, 2024
3	Thermal storage deployment in the framework of current electricity market design	Jaanto et al., 2023
4	The contribution of heat storage to the profitable operation of combined heat and power plants in liberalized electricity markets	Christidis et al., 2012
5	Optimization of District Heating Systems for Future Uncertain Conditions	Malmgren and Skoglund, 2021
6	Optimization of district heating production with thermal storage using mixed-integer nonlinear programming with a new initialization approach	Bjornskov et al., 2021

7	Flexibility from Electric Boiler and Thermal Storage for Multi Energy System Interactio	Sinha et al., 2019
8	Role of power-to-heat and thermal energy storage in decarbonization of district heating	Sihvonen et al., 2024
9	The impact of electric boilers and heat storages in the Nordic power markets and district heating systems	Hiltunen et al., 2025
1o	Control strategies for district heating and cooling systems: a comprehensive review across production, distribution, and end-user levels	Monghasemi et al., 2026
11	An active control strategy for district heating networks and the effect of different thermal energy storage configurations	Vanhoudt et al., 2018
12	Management of a district heating network using model predictive control with and without thermal storage	Quaggiotto et al., 2021
13	ECONOMIC FEASIBILITY OF HEAT STORAGE SYSTEMS UNDER PRICE UNCERTAINTY - A SYSTEM DYNAMIC APPROACH	Roumpas, 2025

14	Design optimization of district heating systems: A review	Sporleder et al, 2022
15	Integrating renewable energy and thermal storage in district heating networks: A design optimization approach	Ceruti et al., 2025
16	Stochastic Model Predictive Control for Optimal Energy Management of District Heating Power Plants	Verelli et al., 2016
17	A new model predictive control approach integrating physical and data-driven modelling for improved energy performance of district heating substations	Zhang et al., 2023
18	System dynamics models, optimisation of	Dangerfield, 2009
19	A Review of District Heating Systems: Modeling and Optimization	Talebi et al., 2016
20	System dynamics model analysis of pathway to 4th generation district heating in Latvia	Ziemele et al., 2016
21	Predicting and optimization of energy consumption using system dynamics-fuzzy multiple objective programming in world heritage areas	Wu and Xu, 2013

22	Two-stage stochastic optimization of district energy systems: a case study from Norway	Bauer, 2025
23	A generic stochastic network-based formulation for production optimization of district heating systems	Guericke et al., 2024
24	Forecasting electricity prices from the state-of-the-art modelling technology and the price determinant perspective	Chai et al., 2024
25	Dealing with uncertainties in the operational planning of district heating plants	Refsgaard., 2021