Vaasan yliopisto
UNIVERSITY OF VAASA

Arif Mehedi

Improving Cost Efficiency of Heat Pump-Based Heating and Cooling Systems

School of Technology and Innovations
Master’s Thesis in Smart Energy
Master’s programme in Smart Energy

Vaasa 2025

UNIVERSITY OF VAASA
School of Technology and Innovations

Author:	Arif Mehedi
Title of the thesis:	Improving Cost Efficiency of Heat Pump-Based Heating and Cooling Systems
Degree:	Master of Sciences
Discipline:	Smart Energy
Supervisor:	Elahe Doroudchi, Hannu Laaksonen
Year:	2025 **Pages:** 57

ABSTRACT:

Heating and cooling systems based on heat pumps offer high energy efficiency but face challenges in operational cost due to suboptimal control strategies. This thesis investigates methods to improve the cost efficiency of a heat pump-based system consisting of nine ground-source heat pumps (each 90-kW max output with four gear levels) by treating the heat pump array as a flexible energy resource.

We develop an optimized control strategy that integrates (i) a model for demand-based power calculation, (ii) a priority-based scheduling of heat pumps by time of day and gear efficiency, and (iii) a mixed-integer linear programming (MILP) optimization for cost-minimal gear selection each hour. We also implement an on/off control scheme that allows room temperature to float between 21°C and 25°C, reducing run-time during low-demand periods.

A MATLAB/Simulink simulation model is built to evaluate the approach under varying outdoor temperatures and dynamic electricity pricing. Results show that the optimized control significantly reduces daily operating cost compared to a conventional continuous operation. In a representative 24-hour scenario, the proposed on/off optimized strategy achieved the required indoor comfort while cutting the energy cost by nearly 70% (down to about €63 per day) relative to a normal control approach. The results are based on a specific single day in wintertime. The thesis concludes with a discussion of key findings, system limitations, and recommendations for future research on integrating heat pumps as flexible resources in smart energy grids.

KEYWORDS: Heat pumps, optimisation, energy efficiency, cost effectiveness, linear programming.

Contents

Figures

Tables

Algorithms

Abbreviations

AI	Artificial Intelligence
BMS	Building Management System
COP	Coefficient of Performance
DR	Demand Response
HVAC	Heating, Ventilation, and Air Conditioning
LP	Linear Programming
MILP	Mixed Integer Linear Programming
PI	Proportional-Integral (Controller)
TOU	Time of Use (Electricity Tariffs)

1 Introduction

Heat pumps are getting more and more important in modern energy systems for space heating and cooling thanks to their high efficiency and use of renewable heat sources. Heat pumps have a high efficiency which exceeds COP 1 and hence can achieve many units of heating or cooling for 1 unit of electricity. What this means is that they can deliver several kilowatts of heating for just one kilowatt of electrical input making them much more efficient and cheaper-to-run than conventional electric heaters. From residential homes to commercial establishments like restaurants,office heat pumps offer efficient heating and cooling solutions.

1.1 Background & Motivation

According to the International Energy Agency (n.d.-a), the coefficient of performance (COP) for a typical household heat pump is around four, i.e. the energy output is four times greater than the electricity input. Because most of the heat is transferred rather than generated, heat pumps are far more efficient than conventional heating technologies such as boilers or electric heaters and can be cheaper to run (International Energy Agency, n.d.-a). As a result, people have started using heat pumps as a key solution for sustainable and cost-effective HVAC.

Resource Innovations (n.d.) explains that modern heat pumps can operate effectively even in sub-zero temperatures. Their ability to also provide cooling by reversing the heat transfer process further underscores their flexibility. Despite the benefits, heat pump systems may not deliver the expected cost savings. One of the biggest problems is the way they schedule operations – most installations run on simple thermostatic control or fixed schedules which do not optimise for cost or efficiency in any active way. Inefficient timing of energy use and unnecessary runtimes can lead to greater operational costs.

An example from the Standard Load Profiles Scenario could be running all heat pumps simultaneously without coordination, which may cause some overheating of the electricity on peak days, at high prices during high tariff periods, or continuous operation of all heat pumps at lower ambient temperatures to keep indoor temperatures more tightly maintained than necessary. Moreover, traditional controls do not see heat pumps as flexible energy resources capable of responding to external signals (such as price changes of electricity or demand response events). This means of not capitalizing could lead to practices that improve cost efficiency and allow other services to the grid (such as peak shaving or load shift).

The reason behind going for this research is the development of smart control strategies exploiting heat pumps flexibility. If we can control every heat pump (on/off status, output level) depending on the allowing factors which occurs during the operation, we can save energy costs. By synchronizing several heat pump devices and controlling their activation to correspond with demand and alterations in electricity pricing, users of buildings could reduce operating costs significantly. Also fitting in the smart grid context, buildings with heat pumps can act as flexible resources. This means that buildings with heat pumps could offer demand response services, which help balance the electricity grid during periods of high renewable generation or other contingencies. These benefits motivate a closer look at advanced control techniques, such as optimization algorithms, to drive heat pump operation.

1.2 Problem Statement

Heat pumps are efficient devices, but wrong control makes them to work at incorrect times. Usually, in an installation, thermostat logic operates each heat pump irrespective of other heat pumps and the price of electricity with time. Such continuous and uncoordinated operation is not cost-effective. Key problems identified include:

- **Inefficient scheduling:** Usually, all pumps are scheduled to run at fixed setpoints. As a result, they draw a lot of power at an expensive peak hour. Otherwise, they remain idle at cheap off-peak hour. This is not ideal as they are effectively doing nothing at this point in time. The scheduling of these pumps fails to exploit time-varying electricity rates effectively. In the absence of smart timing, the electrical bill is higher every day for the same thermal comfort.

- **Lack of demand-based control:** Due to the lack of a demand-based control system, the traditional systems do not adjust the output according to the actual heating/cooling demand or overshoot the required heating, thus wasting energy. If demand forecasting isn’t done or real-time adjustment isn’t managed, pumps will run when not needed or provide too much power. This incurs a cost that should be avoidable.

- **No control optimization:** Conventional heat pump control lacks an optimization layer. In other words, a decision mechanism does not automatically exist that can decide which pumps should run and with what capacity to satisfy demand at minimum cost. Not having this control results in high operating costs due to poor gear selection and on and off actions.

- **Underutilization of flexibility:** Heat pumps and the building thermal mass have inherent flexibility (the indoor temperature can float within a comfort band without immediate loss of comfort). However, most systems do not utilize this flexibility to save energy. For example, they do not turn off the pumps temporarily when temperature is within acceptable range or shift operation slightly to cheaper periods. This rigid operation means potential cost savings from demand response or load shifting are left untapped.

The problem is that without an optimized, intelligent control strategy, a heat pump-based heating/cooling system can incur unnecessarily high electricity costs and cannot easily participate in energy management schemes. We address the question: How can we improve the cost efficiency of a multi-heat-pump system by optimizing its operation as a flexible energy resource?

1.3 Research Objectives

Koller, Rinderknecht, and Ulbig (2024) highlight that heat pumps can serve as an effective solution by utilizing smart control methods, which adjust their power output in response to demand response signals. The primary aim is to minimize electricity costs under a time-variable tariff structure while ensuring that indoor temperatures remain within acceptable comfort limits. Building upon this principle, the main objective of this thesis is to develop and assess a control strategy that reduces the operational costs of a heat pump-based heating and cooling system without compromising thermal comfort. This overall goal is further divided into the following specific objectives:

- **Develop a demand-driven power calculation model:** We will develop heating power calculation models which are demand driven. The models will be simple enough to calculate the required heating power on real time basis. Specifically, the difference between indoor and outdoor temperatures. This model will help us understand the required level of thermal output from the heat pumps at each hour of the day. When heating demand is calculated correctly, we avoid heating ups, and the heating system functions based on demand.

- **Implement an optimized control strategy using linear programming:** We intend to apply linear programming (in particular, mixed-integer linear programming) to the control problem of nine heat pumps with discrete output levels (gears). The goal is to identify the most cost-effective combination of heat pump outputs each hour that meets the required demand. This includes developing an optimization problem that captures electricity costs, heat pump efficiency (COP), and operational constraints. One goal is to confirm that having these priorities in the control strategy improves efficiency, indeed we can start using the most efficient pump to reduce the cost while maintaining the demand.

- **Incorporate priority-based scheduling heuristics:** To make the optimization practical and consistent with practical considerations, we set up priority rules for heat pump usage. Specifically, we consider time-of-day priorities (e.g. certain pumps prioritized during daytime vs nighttime) and gear efficiency priorities (preferring certain output levels known to be more efficient). One objective is to check that incorporating these priorities into the control strategy improves efficiency for example, by using the most efficient pumps or gears first, we can reduce the cost while still meeting demand.

- **Integrate an on/off control mechanism for temperature hysteresis:** We want to control when to turn on/off pumps to do temperature hysteresis. We are looking not only to optimize which pumps run and to what level, but also to try to take advantage of the thermal inertia. The aim is to have on/off hysteresis control, such that heat pumps turn off automatically when the indoor temperature reaches an upper threshold, e.g. 25 °C, and will only turn back on when the indoor temperature reaches a lower threshold, e.g. 21 °C. This way, we can avoid continuous cycling of the heat pumps and save energy when there is no immediate heating demand. We will look at how much this operation costs compared to a continuously controlled one.

- **Validate the approach via simulation using MATLAB/Simulink:** The end goal is to create a simulation model in Simulink (for system dynamics) and MATLAB (for controller logic) to test the proposed strategy through simulation under realistic conditions. The model would simulate a period of 24 hours with changing temperature outdoors and a dynamic profile of electricity price. We will collect information on demand v/s supplied power, operating pattern of heat pump, variations in indoor temperature, and total cost. The simulation should show how effective the optimized control is at reducing cost.

By accomplishing these objectives, the thesis will show a comprehensive approach to treating heat pumps as a flexible energy resource, from concept and algorithm design through to system-level simulation and analysis of results. The expected outcome is a validated control strategy that significantly improves cost efficiency relative to standard heat pump operation.

2 Literature Review

There has been significant research on improving the energy consumption of heating and cooling systems. This is mainly due to rising demands for energy efficiency, cost savings and flexibility on the grid. When advanced control methods integrate energy with comfort, optimizations will be created for heat pump systems. This literature review explains methods and findings relating to energy optimization in heat pump and HVAC systems. It gives examples of load management, demand response and scheduling according to electricity price. Moreover, it shows essential research gaps, notably those concerning interactions of multiple control strategies, as well as optimizing multi-unit heat pump networks, which constitute the basis for the approach developed in this thesis.

2.1 Energy Optimization in Heating & Cooling Systems

The optimization of the energy consumption of heating and cooling system has been a long-studied subject as buildings become smart and low energy. Load management, which means making sure required energy supply coincides with required energy demand, both in time and magnitude, will eliminate loss. For heat pumps systems, this can involve strategies such as scheduling, demand forecasting and using thermal storage or building thermal mass as a buffer.

One common approach is the demand response (DR) capability of heat pump systems. Demand response programs encourage consumers to reduce or shift their power usage during peak demand times. Heat pumps, due to the flexibility in when they can operate (within comfort constraints), are ideal candidates for providing DR. Research and field demonstrations have shown that aggregating many heat pumps under coordinated control can provide reliable load shedding or shifting.

According to Müller and Jansen (2019), research and field demonstrations have shown that aggregating many heat pumps under coordinated control can provide reliable load shedding or shifting. In a large-scale trial involving over 300 residential buildings, throttling heat pumps in response to DR signals led to load reductions of 40–65% of the total load, with precise control over the timing and quantity of load curtailment. This highlights that with appropriate control, heat pumps can be turned into distributed energy resources that support grid stability while also benefiting the end user through incentives or lower rates.

Another aspect is to adjust operation according to time-of-use electricity tariffs or real-time prices. Utilizing the lower price periods for heat pump operation can result in considerable savings while achieving the same comfort. For instance, optimal scheduling methods are suggested where heating is done intensely when electricity is cheap (like the late night) and reduced at high price periods (like the late afternoon).A study by Rink, Gourishankar, and Zaheeruddin (1988) demonstrated that adjusting heat pump operations based on TOU electricity tariffs and utilizing the building's thermal inertia as a form of energy storage can lead to substantial cost savings through pre-heating or cooling during low-price periods.

Heating systems can be effectively controlled using optimization techniques such as linear programming (LP) and mixed-integer linear programming (MILP). The methods allow us to formulate constrained cost-minimization problems which can be used for scheduling equipment like heat pumps. Savolainen, Einolander, and Lahdelma (2024) proposed a control strategy for a fleet of heat pumps based on MILP in order to minimize peak demand and flatten the consumption profile. They achieved a much smoother load profile than the case of traditional, PI on/off control by solving a large scale optimization problem that set the on/off statuses of multiple units over time. Drawing on the results, optimization-based strategies effectively manage energy. MILP formulations can be used when equipment has fixed operating levels (e.g., gear-based heat pumps). In those cases, the formulation can use binary on/off or multi-level decision variables.

Balancing savings with occupant comfort is essential in HVAC systems optimization. Going overboard with optimization may lead to big temperature swings or discomfort indoors. As a result, several methods use penalty terms or constraints to limit any unintended temperature deviations. To illustrate, a building can be pre-heated to the upper setpoint prior to a peak price period and allowed to drift down to the lower setpoint during the expensive period; this would maintain comfort at reduced costs. Sweetnam et al. (2019) emphasize that flexible control strategies such as on/off hysteresis control can allow for domestic demand-side response with heat pumps. These methods allow automated optimized operation in response to outside signals, such as change in electricity price or demand response (DR) events, ensuring efficient not only energy but also thermal comfort.

In short, the literature shows that with appropriate control of heating/cooling systems, large cost and energy savings are possible. The application of scheduling algorithms based on linear programming, demand response integration, and model predictive control have helped to reduce the operating cost and shifting of load of the heat pump system. The stated findings advocate for the principle of this thesis, as it also employs a similar concept where MILP optimization and on/off control is used within the thermal comfort band to improve the cost-effectiveness of a heat pump network.

2.2 Research Gaps & Justification

While many studies have addressed optimal control of heat pumps and HVAC systems, there remain some gaps that this thesis aims to fill, particularly in the context of combining multiple strategies in a single unified framework. First, much of the literature either focuses on load shifting for demand response or on continuous optimization for efficiency. In practice, a comprehensive solution should achieve both cost optimization and demand flexibility. The integration of a cost-based MILP dispatch with an on/off hysteresis control, as proposed here, is a novel combination that can yield additional savings. Many conventional systems use either simple hysteresis (thermostat control) without cost optimization, or they use optimization but keep systems running continuously around a setpoint. Our approach justifies combining these methods: by first minimizing cost to meet demand (via MILP) and then adding an on/off logic to exploit thermal inertia, we expect compounded benefits.

Another gap is the practical application in medium-sized multi-pump systems. Some research has looked at large populations of small residential heat pumps (e.g., hundreds of homes in a demand response program) or at single high-capacity units. The scenario of a facility with multiple parallel heat pumps (like an apartment block or commercial building with a bank of heat pump units) is less studied. In such cases, issues like how to distribute load among units (all running low vs. few running high) and how to prioritize certain units become important. This thesis specifically addresses a nine-heat-pump system and introduces priority coefficients to favour certain pumps/gears at different times. This nuanced control is not extensively covered in literature, thus providing a new perspective on intra-system optimization (as opposed to fleet-level control).

The method used here is simple and its control logic is clearly defined. We prefer a linear programming method over more complicated black-box strategies, such as reinforcement learning or nonlinear optimization. As a result, existing BMS may be able to implement this solution much more easily as well. Users can gain a clear understanding of how electricity price, COP, and other factors can affect decisions. It is important to demonstrate that a relatively simple MILP can achieve significant savings. Economists can help persuade facility managers to move from tried and tested methods to something new.

For example, the findings of our case study indicate that when using an optimized on/off strategy, daily cost can be reduced to roughly €63 (versus a much greater cost in case of continuous control) thus can be used to boost further research as well as implementation. It deals with the gap of quantifying benefits in a controlled one. The contributions stemming from the results of this thesis may assist engineers in developing sophisticated yet practical control algorithms, which will help building owners save on their energy bills and help make energy systems more flexible and resilient.

3 Methodology

In this section, we describe the methodology used to model, simulate, and optimize the operation of a multi-heat pump system that provides heating for a building in a cost-efficient manner. We first describe the system setup and key assumptions, then we develop the demand calculation model. Building on this, an optimization strategy combining scheduling priorities and gear selection heuristics is formulated. Finally, we outline the implementation approach in MATLAB and Simulink, which enables dynamic simulation of the heating system under realistic environmental and operational conditions.

3.1 System Design & Assumptions

The case study system considered in this thesis consists of nine ground-source heat pumps, which collectively supply heating (and potentially cooling) to a building or set of buildings. Each heat pump is identical in capacity, with a maximum thermal output of 90 kW. Importantly, the heat pumps are designed with four discrete operating levels (gears): 22.5 kW, 45 kW, 67.5 kW, and 90 kW. These correspond to approximately 25%, 50%, 75%, and 100% of the unit's capacity, respectively. At any given time, a heat pump can either be off or running at one of these four power levels. This discrete gear setup is a common design in large heat pumps to allow stepwise control of output and improve part-load efficiency.

The following table presents the key input data collected from the case study, which serves as the basis for modelling and optimization. It includes the outdoor ambient temperature (Tout), electricity price per kilowatt-hour (Euro/kWh), and the coefficient of performance (COP) of the heat pumps for each hour of a typical day:

Hour (24-hour format)	Output temperature (Tout in °C)	Price (Euro/KWh)	COP
0	-1.3	0.0484	3.44
1	-1.3	0.0201	3.56
2	-1.234	0.0138	3.38
3	-1.068	0.0009	3.54
4	-1.066	0.0008	3.58
5	-1.034	0.0005	3.66
6	-1.99	0.0008	3.76
7	-3.094	0.0009	3.79
8	-3.466	0.0053	3.67
9	-3.5	0.0203	3.57
10	-4.16	0.0310	3.63
11	-4.632	0.0443	4.14
12	-4.7	0.0606	4.14
13	-4.7	0.0606	3.99
14	-4.898	0.0568	4.00
15	-5.198	0.1771	3.50
16	-5.696	0.2411	3.59
17	-6.296	0.3654	3.54
18	-6.896	0.4603	3.53
19	-7.496	0.5423	3.47
20	-8.03	0.5363	3.73
21	-8.596	0.4805	3.39
22	-8.998	0.4468	3.50
23	-9.298	0.4596	3.41

Table 1. 24 Hours Data of Price, Tout, and COP from Case Study

Key control parameters and assumptions include:

- **Indoor temperature range:** The definition of a thermal comfort band for indoor temperatures is widely supported by international standards. According to ASHRAE (2017, Section 5.3), the preferred indoor temperature range for most occupants is between 20°C and 24°C during winter and between 23°C and 26°C during summer. The comfort band used in this thesis, defined between 21°C and 25°C, falls within these recommended ranges.

The normal desired indoor temperature (setpoint) is set at 25°C; the heating system aims to maintain this temperature. However, heating activation is allowed only when the temperature drops to 21°C, thereby establishing a 4°C deadband. This deadband implements an on/off hysteresis control strategy (see Section 3.3.3). It is assumed that occupants remain comfortable within this range, and that no active cooling is needed as long as the indoor temperature stays at or below 25°C.

- **Outdoor temperature profile:** The heating demand is driven by the difference between indoor and outdoor temperature. We take a 24-hour profile of outdoor temperature representative of a winter day. For instance, in our simulation the outdoor temperature ranges roughly from -1.3°C to -9.3°C over the day, with the coldest temperatures in late night/early morning.

- **Building thermal properties:** The building's thermal properties play a critical role in determining heating requirements. As Biddulph et al. (2014) emphasize, the thermal properties of the building envelope such as the thermal transmittance (U-value) and heat capacitance are key indicators of the energy performance of a building. Accurate modeling of these properties is essential for reliable heating system simulations.

The building is characterized by an overall heat transfer coefficient U ($W/m^2 \cdot K$) and an effective heat exchange surface area A (m^2) that together determine how much heat is lost per degree of temperature difference. For our model, we assume $U = 0.5\ W/(m^2 \cdot K)$

and A = 5000 m², which yields UA = 2500 W/K of thermal conductance. These values, along with an efficiency factor (related to COP), will be used in the heat demand calculation. The building's interior thermal mass (air and contents) is also considered for the temperature dynamics parameters like air heat capacity, room volume, etc., are used in the Simulink simulation to model temperature changes when heating is on or off.

- **Heat pump efficiency (COP):** The coefficient of performance of the heat pumps is not constant; it varies with conditions (primarily the source and sink temperatures). We incorporate a COP profile for the 24 hours, representing how the efficiency might change with outdoor temperature and load. For example, COP values in our simulation range roughly from 3.3 to 4.14. Generally, during colder hours the COP is lower (around 3.3–3.5 when it's -8 to -9°C outside) and during milder hours COP can exceed 3.7–4.0. We assume these COP values are known (perhaps from manufacturer data or real-time estimation) and use them to calculate electricity consumption from thermal output. A higher COP means the heat pump delivers more heat per unit of electricity.

- **Electricity price profile:** To incorporate cost optimization, we use a dynamic electricity price curve over 24 hours. This could represent time-of-use tariffs or spot market prices. In our case, prices are given in €/kWh for each hour, varying significantly throughout the day. For instance, during early morning hours electricity might be very cheap (around £0.0005–0.005 per kWh), increasing in the afternoon and peaking in the evening (up to around £0.54 per kWh in our data).

- **Operational constraints:** Each heat pump can operate at only one gear level at a time, and all gear selections are assumed to be instantaneous (i.e., we neglect any transient when switching gears or turning on/off). In reality, a heat pump might take a few minutes to ramp up, but here we assume the control resolution is one hour for gear selection via optimization, and a faster logic (minute-level in Simulink) handles on/off within that hour if temperature setpoints are reached. We also assume the heat pumps are always available (no maintenance downtime or faults considered).

With these components defined, the system essentially works as follows: At each hour, based on the outdoor temperature, we compute how much heating power is required to maintain 25°C indoors (this is the demand model). The control algorithm then decides how to dispatch this required power among up to nine pumps, each either off or at 22.5, 45, 67.5, 90 kW, with the goal of minimizing cost. If the indoor temperature hits 25°C before the hour is over, the pumps can all shut off until temperature falls to 21°C, at which point heating resumes. The Simulink model monitors indoor temperature continuously and sends an on/off switch signal accordingly.

This design allows us to study both continuous cost optimization (if we were to ignore the on/off and just meet demand exactly each hour) and the intermittent on/off behaviour. It also reflects a realistic setup for a mid-sized building heating system with multiple heat pump units working in parallel.

3.2 Power Demand Calculation Model

A crucial first step in the control strategy is to determine the required heating power for each hour to maintain the desired indoor temperature (25°C in our scenario). We employ a simple physics-based model for the building's heat demand, which assumes quasi-steady-state heat loss to the environment. The model is based on the heat transfer through the building envelope and can be expressed as:

$$\mathrm{P} = \eta\, \mathrm{U\,A}\,(\mathrm{Tin} - \mathrm{Tout}) \tag{1}$$

where P is the required thermal power (in watts) at time t, Tin is the indoor setpoint temperature (25°C), and Tout is the outdoor ambient temperature at time t. The product UA (given as 2500 W/K) represents the overall thermal conductance of the building envelope, indicating how many watts of heat are lost per degree of temperature difference. The factor η is an efficiency or utilization factor introduced to account for heat gains or system inefficiencies and is set to 3.0 in our model.

This formulation is widely recognized in the literature. According to Incropera and DeWitt (2007, Section 1.2, Equation 1.2), the rate of heat transfer through a plane wall can be expressed as $Q=UA(T_{in}-T_{out})$, where U is the overall heat transfer coefficient, A is the wall area, and $(T_{in}-T_{out})$ is the temperature difference across the wall. This foundational heat transfer model supports the use of the above formula in estimating heating power requirements through building envelopes in thermal engineering and building physics contexts.

The above formula provides power in watts; we convert it to kilowatts by dividing by 1000. Implemented in MATLAB, the demand calculation for the whole 24-hour profile is done by vectorizing this formula for each hourly outdoor temperature reading. The result is a 1×24 vector of demand in kW. For example, if at a certain hour Tout is -5°C, then $\Delta T=T_{in}-T_{out}=25-(-5)=30$. Multiplying by UA=2500 W/K and η=3.0 gives 225000 W, which is 225 kW required. This indicates that at -5°C outside, the building needs 225 kW of heat to hold at 25°C inside. Warmer outside conditions yield lower required power, e.g., at -1°C outside, ΔT=26, yielding around 195 kW needed.

This model assumes no internal heat gains or other heat sources; all heating must come from the heat pumps. Using this model, our simulation generated a demand profile that starts around 197 kW when the outdoor temperature is -1.3°C at hour 1, and rises to about 257 kW by hour 24 when the outdoor temperature drops to -9.3°C. This matches expectations that the coldest time produces the highest heating load. We purposely sized the system (9 pumps × 90 kW = 810 kW total capacity) to be well above the peak demand (~257 kW) so that capacity is not a limiting factor. This way, the optimization has freedom to distribute the load among pumps in many ways, rather than being trivial (in a smaller system, the optimal solution might always be “run everything at full power”).

This demand model is integrated into the simulation and control routine as demand = fcn(T_out, T_in) in MATLAB code. It provides the foundation for the optimization in the next step – given a demand (in kW) for each hour, the controller will allocate that among heat pumps and gears with minimal cost.

3.3 Optimization Strategy

The core of our methodology is the optimization strategy that decides how the nine heat pumps should be operated each hour to meet the heating demand at lowest cost. We break this strategy into three coordinated components: time-based priority scheduling, gear-based priority selection, and linear programming optimization. These components work together within each hourly decision.

3.3.1 Time-Based Priority Scheduling

In a multi-heat-pump setup, it can be beneficial to designate certain pumps to handle load during specific periods, reflecting practical considerations like location, wear levelling, or noise. In our strategy, we implement a simple time-based priority: some pumps are preferred during the morning/daytime hours and others during the nighttime hours. The rationale here could be, for example, that pumps 1–5 are located near daytime-active areas or are newer and thus we use them in the day, whereas pumps 6–9 might be in an area less disturbing at night, or we want to alternate usage to distribute run-time.

Concretely, we define “daytime” as hour 8:00 to 20:00 (8 AM to 8 PM) and “nighttime” as hour 20:00 through 7:00 of the next day (8 PM to 7 AM). During the daytime interval, Heat Pumps 1–5 is given high priority to operate, and pumps 6–9 are deprioritized. During nighttime, the priority inverts: Heat Pumps 6–9 take the lead, and pumps 1–5 become secondary. This is reflected in the priority coefficient matrix used in our MILP formulation (described in section 3.3.3). Essentially, when deciding which pumps to turn on, the optimization will incur a lower “cost” (in the objective function) for using a pump that is in its favoured time window.

This time-based scheduling approach ensures each pump group gets a rest period – pumps 1–5 mostly rest at night, and 6–9 rest during the day – which could be beneficial for maintenance and longevity. It also means any two sets of pumps effectively share the work over a 24h cycle, potentially balancing wear and reducing the chance that all pumps run simultaneously (which could cause very high instantaneous demand).

It's important to note that if demand is very high, the optimizer can still call on all pumps even outside their preferred time. The priority is a bias, not a strict rule. For instance, if during daytime the demand exceeds what pumps 1–5 can supply at full capacity, the optimizer will certainly bring pumps 6–9 online as needed. But it will do so minimally, since those pumps have a higher cost (due to priority coefficients) during that period.

3.3.2 Gear-Based Priority Selection

In addition to deciding which pumps to run, we also must decide at what gear each running pump should operate. Heat pump performance can vary with load – often there are “sweet spots” in efficiency. Based on general knowledge and the design of our units, we establish a heuristic ranking of the gear efficiencies: the 45-kW gear (half capacity) is the most efficient, followed by the 22.5 kW and 67.5 kW gears, and lastly the 90-kW gear is the least efficient in terms of COP or cost per kW delivered.

This ranking comes from the idea that running a heat pump at partial load often yields better COP than maxing it out, due to reduced compressor stress and better heat exchange under lower load. Empirical studies support the prioritization of operating heat pumps at partial loads to enhance efficiency. Research has demonstrated that heat pumps generally achieve a higher coefficient of performance (COP) when operating at partial loads, as opposed to full capacity (Marsik et al., 2023). Operating at full capacity (100%) can reduce COP due to increased compressor workload and potential thermal inefficiencies. Additionally, extremely low load operation typically below 20% can also diminish efficiency because of standby losses and frequent cycling (Xu et al., 2022).

To incorporate this knowledge, our optimization gives priority to using the 45-kW gear whenever possible as the 45 kW point likely is near the optimal COP on the performance curve, whereas at 90 kW (full load) the COP might drop significantly, and at very low load (22.5 kW) other overhead losses might reduce net efficiency slightly. We constructed priority coefficients for each gear level as well: effectively making the cost of selecting 45 kW gear for a pump lower than selecting 22.5 or 67.5 kW, which in turn are lower cost than selecting 90 kW, in the objective function formulation. In the MILP, this was implemented by dividing the cost terms by these priority values (so a higher priority corresponds to a larger divisor, hence lower effective cost contribution). During daytime for pumps 1–5, for example, gear2 (45 kW) had a priority value of 4, gear1 and gear3 had 3, and gear4 had 2.

This means the solver is incentivized to use gear2 first. During nighttime for pumps 6–9, a similar scheme was applied (though inverted for pumps 1–5, which have lower priority at night). The decision-making approach then becomes try to meet the heating demand by turning on as many pumps as needed at the 45-kW level. Only if demand cannot be met with available pumps at 45 kW will the strategy escalate to either adding more pumps at other gears or increasing some pumps to 67.5 kW, etc.

For instance, if the demand is 180 kW, the optimizer might simply activate 4 pumps at 45 kW each (total 180 kW) as that exactly meets the need with the preferred gear. If demand were 200 kW, it might do 4 pumps at 45 kW (180 kW) and then have to supply an extra 20 kW; likely it would then use one pump at 22.5 kW (making 202.5 kW total) rather than elevating one pump to 67.5 kW (which would yield 202.5 kW with just 3 pumps, but uses a less efficient gear on that pump). The choice depends on how the costs shake out, but since gear1 and gear3 had equal medium priority in our setup, the solver might be indifferent between 5 pumps at 45/22.5 mix versus 3 pumps at 67.5. In practice, we saw solutions where a slight oversupply was accepted to stick with efficient gears. For example, to supply about 195 kW demand at night, the solver chose a combination that resulted in ~202.5 kW supply, using one pump at 67.5 kW and others at 45 kW. The oversupply (7.5 kW extra) is allowed via a slack variable, and it was preferred to keep pumps in those gears rather than jump to 90 kW or bring too many pumps at tiny loads.

By prioritizing gear selection, we effectively encode an efficiency-first approach: use the gear that gives the best COP per kW and avoid the highest gear unless absolutely necessary. The highest gear (90 kW) was thus rarely used in the optimized solution; it would typically come into play only if demand was extremely high (close to system capacity) where we had to use full power. Otherwise, the solver steered clear of 90 kW because of its cost penalty. This aligns with real-world best practices where running many heat pumps at partial load is often more efficient than running fewer at full throttle.

3.3.3 Linear Programming Formulation

At the heart of the control strategy is a mixed-integer linear programming (MILP) formulation which the system solves for each hour. The MILP decides which heat pumps to turn on and at which gear, under the priority biases discussed, to minimize a cost function. We formulated the problem for a single hour (treating each hour independently for the gear selection, but the total cost is computed over 24 hours for evaluation).

Decision Variables: We introduce binary decision variables Yij for each pump $i \in \{1,...,9\}$ and gear $j \in \{1,...,4\}$ if pump i is operated at gear level j (with power Pj) during the current hour, and Yij=0 otherwise. We also include a continuous slack variable E (non-negative) representing any excess supplied power beyond the demand. This slack accounts for the possibility of slight oversupply since our pumps come in fixed increments. The total number of variables is 9*4 + 1 = 37 (36 binaries, 1 continuous) for each hour's problem.

Objective Function: The objective is to minimize the total cost for that hour, which is comprised of electricity cost and a penalty for oversupply. The electricity cost for running a particular pump at a certain gear is proportional to the power Pj times the electricity price of that hour, divided by COP (to convert thermal power to electrical power required). We incorporated the priority coefficients by effectively dividing the cost of using pump i at gear j by the priority factor prio(i,j)

Thus, the objective can be written conceptually as:

$$\min \sum_{i=1}^{9} \sum_{j=1}^{4} \left(\left(\frac{price}{COP} \right) \cdot \ P_j \cdot \left(\frac{1}{prio(i,j)} \right) \right) \cdot \ y_{i,j} + \ \lambda E \tag{2}$$

Here, λ is a large penalty cost per kW of slack (we used λ=100 as a weight)

This penalty ensures we only oversupply if it’s minor and necessary; otherwise, the solver will meet demand exactly to avoid paying a big penalty. The price (in €/kWh) and COP for the current hour are constants in that hour’s optimization (coming from the profiles). The priority matrix prio(i,j) is determined by the time (day or night as above). Essentially, this objective drives the solution to use combinations of Yi,j that result in the lowest cost weighted by those priorities.

Constraints: We have a few sets of linear constraints:

1. **Each pump at most one gear:**

$$\sum_{j=1}^{4} y_{i,j} \leq \ 1 \qquad \forall \ i \ \in \ \{1, \ldots ,9\} \tag{3}$$

For each pump this ensures a pump isn’t counted as running at two different gears simultaneously. It can be either off or one of the gears =1.

2. **Demand satisfaction:** The total heat supplied by all pumps minus the slack equals the demand for the hour. This is formulated as:

$$\sum_{i=1}^{9} \sum_{j=1}^{4} P_j y_{i,j} - E = Demand_{hour} \tag{4}$$

We treat demand as exactly met if possible or overshot by E if needed (undersupply would cause E to be negative, which we disallow by bounding E ⩾ 0). This means the pumps must at least meet the demand; any extra beyond demand is absorbed by E. In practice, E will become exactly the amount of oversupply. If the demand cannot be met even with all pumps at full power (not an issue in our scenario due to ample capacity), the constraint would be infeasible. Otherwise, usually, we'll have a small E if demand is not an exact multiple of 22.5 kW.

3. **Variable bounds:**

$$y_{i,j} \in \{0,1\}\,;\; E \geq 0 \tag{5}$$

We explicitly set bounds 0⩽E<∞0 (no upper bound on slack in case oversupply would be very cheap, but the penalty prevents it from blowing up)
The binary nature of y is handled by the solver by specifying integer constraints for those indices

In the MILP itself, we assume Switch=1 during the calculation since we only call it when heating is needed. The integration with temperature is handled outside MILP via the simulation logic.) Given this formulation, we use MATLAB's intlinprog solver to solve the MILP for each hour. The solver finds the combination of pumps and gears that minimize the cost function.

One additional aspect (Temperature-Based Switching): we incorporate the on/off control by a simple external switch. If the on/off Switch is 0 (meaning we decided to turn off heating because temperature is high enough, indoor temperature $\geqslant 25^\circ$ C), the function simply returns all zeros without solving the MILP. In normal operation with Switch=1, it solves as above. This way, the optimization is only active when heating is actually required (under the hysteresis logic).

The MILP optimization is performed for the current hour to decide immediate pump outputs. In our MATLAB script, after computing the current hour's hp_vals, we also simulate ahead by computing the total cost if that pattern were used for all 24 hours (repeating the solve for each hour, which effectively gives an optimized schedule for the day). But in the real system, one could solve hour by hour in a rolling manner. The outcome of the MILP each hour respects the priorities we set. For example, at night (hour 24), the solver gave a solution: pumps 6–9 each at 67.5 kW, pumps 1–5 off, resulting in 270 kW supplied. This met a demand of ~257 kW with E $\approx$ 13 kW oversupply. The cost considered the high price at that hour but since all four pumps 6 – 9 at gear3 were chosen, it clearly followed the night priority and gear priority (did not activate pumps 1 – 5 or use gear4). During midday hours with lower demand, the solver mainly used pumps 1–5 at 45 kW as expected.

This MILP approach is essentially performing a cost-optimal dispatch of a small fleet of heat pumps with discrete outputs. The use of priority coefficients is a novel tweak to incorporate heuristic preferences directly into the cost function, guiding the solution towards practical optimums that pure cost alone might not yield (especially if price and COP were constant, then the priority would dominate the decision).

3.4 Implementation in MATLAB/Simulink

To test and demonstrate the control strategy, the system was implemented in a simulation environment using MATLAB and Simulink. The overall setup can be described in two parts: a Simulink model of the physical system (heat pumps and room thermal dynamics) and a MATLAB script that runs the optimization and interacts with Simulink.

3.4.1 Simulink Model Setup:

In Simulink, we built a model representing the nine heat pumps and the room/system being heated.

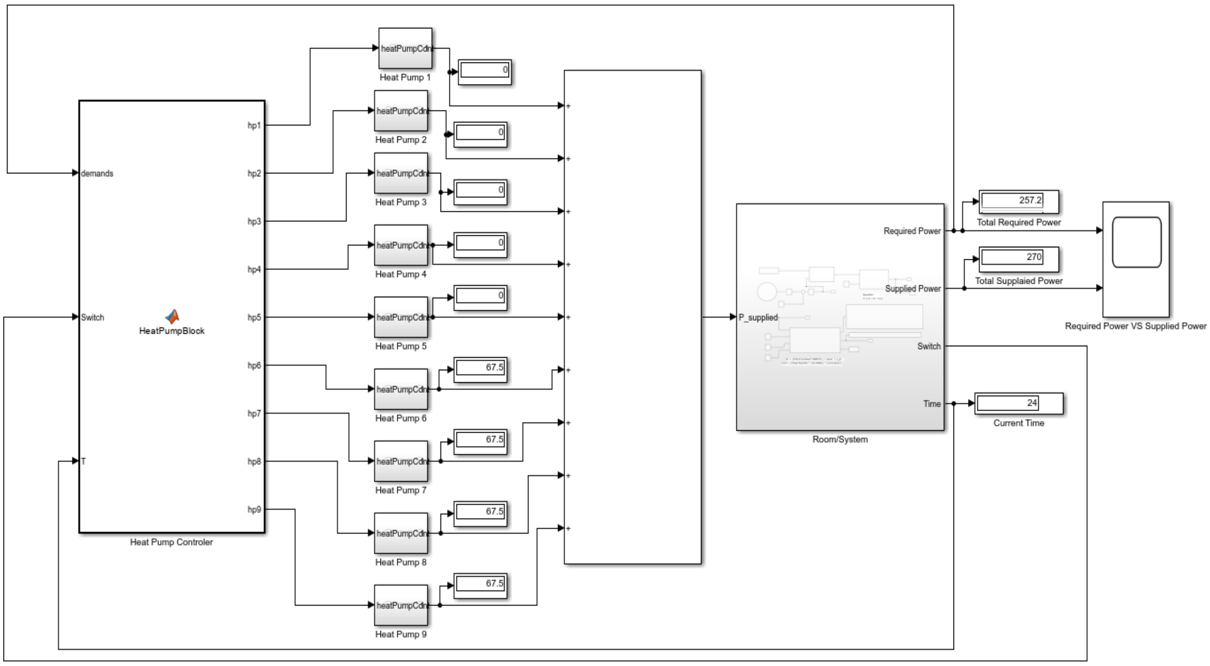

Figure 1. Simulink Diagram

The model is divided broadly into two subsystems:

• **Heat Pump Controller subsystem:** This corresponds to the “left side” of the mentioned Simulink diagram. It takes as inputs the current demand (calculated from the temperature difference), the current time (hour of day), and a on/off switch based on the current room temperature. The controller subsystem includes a block that calls the external MATLAB function HeatPumpBlock (which in turn calls the optimizer). This produces outputs indicating which heat pumps should be on and at what gear for the current time step.

Essentially, at each simulation step (which we set to 1-minute steps for the temperature simulation), the controller decides the heat pump outputs based on the latest information. In Simulink, the HeatPumpBlock was implemented as a MATLAB Function block with extrinsic call to our optimize function (since intlinprog can't run natively in Simulink real-time, we did a time-step of 1 hour for the control or used a triggered approach every hour).

- **Room/System subsystem:** This corresponds to "right side" of the Figure 1 diagram.

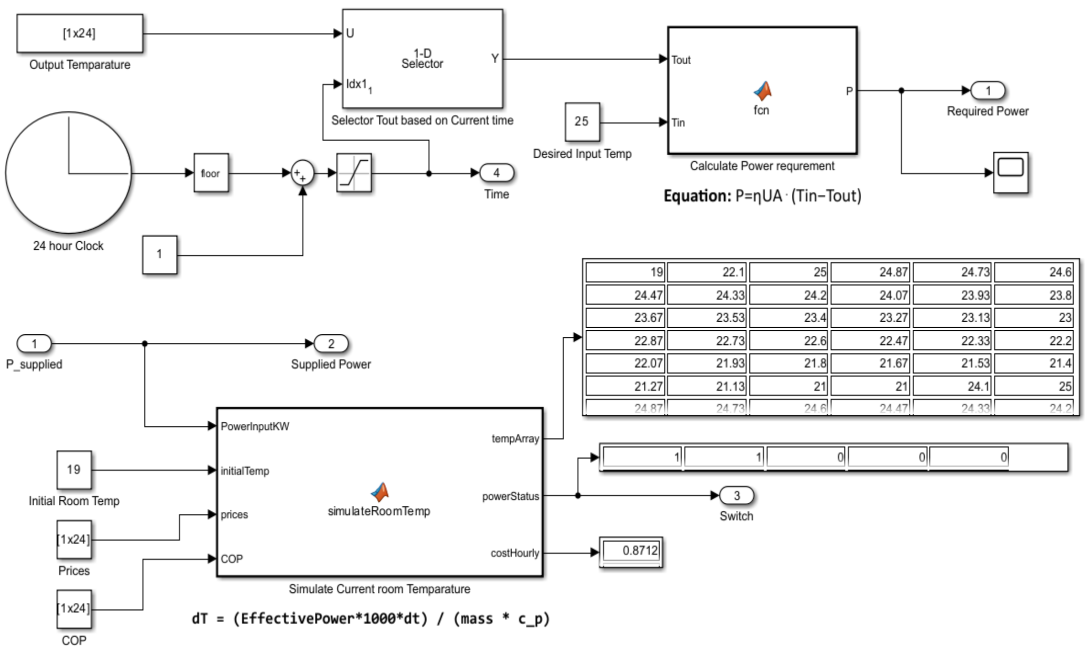

Figure 2. Room/System block

This Figure 2 is the inside of Room/System block. It models the thermal behaviour of the building (room). It takes the total heating power supplied by the heat pumps as an input and computes the resulting room temperature dynamics. We included logic here for the on/off hysteresis: when the room temperature reaches 25°C, it sends a signal to turn the heat pumps off; when it falls to 21°C, it sends a signal to turn them back on. The subsystem uses a simple first-order thermal model: heat input raises the room temperature based on the room's thermal mass, and heat is lost to outside based on the temperature difference. We wrote a custom function simulateRoomTemp to capture this behaviour in MATLAB as well, and in Simulink we could integrate that or implement difference equations.

The Simulink model thus mimics a closed-loop control. It generates a demand, the controller decides pump outputs, those outputs heat the room, and the room temperature feeds back to possibly turn pumps off when setpoint is reached. For visualization and analysis, we added scopes or sinks in Simulink to record key outputs: the required power vs supplied power, the cumulative cost, and the room temperature and switch status over time.

3.4.2 MATLAB Script and Linear Program Integration:

We wrote a MATLAB script to tie everything together and perform calculations that are easier in MATLAB than Simulink (like solving the MILP). The script performs the following steps:

1. Define Data: We first define the necessary input data,

- Outdoor temperature profile (Tout) for 24 hours,
- Electricity prices (prices),
- Coefficients of Performance (COP) for each hour, and
- Setpoint for indoor temperature (Tin_setpoint = 25°C).

Additionally, constants such as overall heat transfer coefficient (U), surface area (A), and system efficiency (η) are specified inside the fcn function. These parameters ensure the system operates with realistic environmental and building conditions.

2. Calculate Heating Demand: The hourly heating demand is computed using the function fcn(Tout, Tin_setpoint).

```
function P = fcn(Tout, Tin)
U   = 0.5;
A   = 5000;
eta = 3.0;
```

```
ToutVec = double(Tout) * ones(1,24);
P = eta * U * A * (Tin - ToutVec) / 1000;
P = reshape(P, 1, 24);
end
```

Algorithms 1: Mathlab script to Calculate Heating Demand

This function models the heat loss to the environment and calculates the required heating power for each hour to maintain the set indoor temperature.

3. Main Loop for Heat Pump Optimization: For each hour (T = 1 to 24), We call the HeatPumpBlock(demand, Switch, T) function (with Switch = 1 to indicate heating is active). Inside HeatPumpBlock, the MILP is solved using optimizeHeatPumps, which determines the optimal operating gear (25%, 50%, 75%, 100%) for each heat pump. Bellow given mathlab script as algorithom 2 to optimized heatpump using MILP.

```
function [hp_vals] = optimizeHeatPumps(demands, Switch, T)

P = [22.5, 45, 67.5, 90];
prices  = [0.0484, 0.0201, 0.0138, 0.0009, 0.0008, 0.0005, 0.0008,
0.0009,0.0053,  0.0203,  0.031,  0.0443,  0.0606,  0.0606,  0.0568,
0.1771,0.2411,  0.3654,  0.4603,  0.5423,  0.5363,  0.4805,  0.4468,
0.4596];
COP     = [3.44, 3.56, 3.38, 3.54, 3.58, 3.66, 3.76, 3.79, 3.67,
3.57,3.63, 4.14, 4.14, 3.99, 4.00, 3.50, 3.59, 3.54, 3.53, 3.47,
3.73, 3.39, 3.50, 3.41];

nPumps = 9;
nGears = 4;
hour_idx = T;
currentPrice = prices(hour_idx);
currentCOP   = COP(hour_idx);
currentDemand = demands(hour_idx);
```

```
if (T >= 8) && (T <= 20)
    prio = zeros(nPumps, nGears);
    prio(:,1) = [3; 3; 3; 3; 3; 1; 1; 0.15; 0.15];
    prio(:,2) = [4; 4; 4; 4; 4; 1.50; 1.50; 0.25; 0.25];
    prio(:,3) = [3; 3; 3; 3; 3; 1; 1; 0.15; 0.15];
    prio(:,4) = [2; 2; 2; 2; 2; 0.50; 0.50; 0.10; 0.10];
else
    prio = zeros(nPumps, nGears);
    prio(:,1) = [1; 1; 0.15; 0.15; 0.15; 3; 3; 3; 3];
    prio(:,2) = [1.50; 1.50; 0.25; 0.25; 0.25; 4; 4; 4; 4];
    prio(:,3) = [1; 1; 0.15; 0.15; 0.15; 3; 3; 3; 3];
    prio(:,4) = [0.50; 0.50; 0.10; 0.10; 0.10; 2; 2; 2; 2];
end

if Switch < 1
    hp_vals = zeros(nPumps,1);
    return;
end
nY = nPumps * nGears;
nVars = nY + 1;
f_y = zeros(nY,1);
for i = 1:nPumps
    for j = 1:nGears
        idx = (i-1)*nGears + j;
        f_y(idx) = (currentPrice/currentCOP) * P(j) / prio(i,j);
    end
end

lambda = 100;
f = [f_y; lambda];

A1 = zeros(nPumps, nVars);
for i = 1:nPumps
    for j = 1:nGears
        idx = (i-1)*nGears + j;
        A1(i, idx) = 1;
```

```
    end
end

b1 = ones(nPumps, 1);

Aeq = zeros(1, nVars);
for i = 1:nPumps
    for j = 1:nGears
        idx = (i-1)*nGears + j;
        Aeq(1, idx) = P(j);
    end
end

Aeq(1, end) = -1;
beq = currentDemand;

lb = zeros(nVars, 1);
ub = ones(nVars, 1);
ub(end) = Inf;
intcon = 1:nY;
opts = optimoptions('intlinprog','Display','off');
[x_opt, ~, exitflag] = intlinprog(f, intcon, A1, b1, Aeq, beq, lb,
ub, opts);
if exitflag ~= 1
    error('No optimal solution was found for hour T = %d.', T);
end

y_opt = x_opt(1:nY);
hp_vals = zeros(nPumps,1);
for i = 1:nPumps
    for j = 1:nGears
        idx = (i-1)*nGears + j;
        if y_opt(idx) > 0.5
            hp_vals(i) = P(j);
            break;
        end
```

```
    end
end
end
```

Algorithms 2: MATLAB script of Optimizing Heat pump using MILP

The resulting hourly heat pump outputs are recorded, and the total supplied power is computed as the sum across all heat pumps. This optimization ensures that the heating demand is met at the lowest possible cost, considering price, COP, and time-prioritized usage of different heat pumps.

4. Simulating On/Off Control: We then simulate the on/off thermostat control using the function **simulateRoomTemp** we wrote, which takes the optimized suppliedPower schedule and an initial temperature, and simulates minute by minute the room temperature and whether the heater is ON or OFF. The MATLAB scripts of **simulateRoomTemp** is given bellow as algorithm 3.

```
function [time, tempArray, powerStatus, costHourly, totalCost] =
simulateRoomTemp(PowerInputKW, initialTemp, prices, COP)

dt = 60;
totalTime = 24 * 3600;
numSteps = totalTime / dt;
time = (0:dt:totalTime-dt)'/3600;
roomArea = 5000;
roomHeight = 3;
airDensity = 1.2;
c_p = 1000;
mass = roomArea * roomHeight * airDensity;
tempArray = zeros(numSteps, 1);
powerStatus = zeros(numSteps, 1);
costMinute = zeros(numSteps, 1);
temp = initialTemp;
heaterON = 1;
```

```
coolingRate = 4 / 30;
for i = 1:numSteps
    currentHour = floor((i-1)/60) + 1;
    currentPower = PowerInputKW(currentHour);
    tempArray(i) = temp;
    powerStatus(i) = heaterON;
    if heaterON
        effPower = currentPower * COP(currentHour);
        deltaT = (effPower * 1000 * dt) / (mass * c_p);
        temp = temp + deltaT;
        costMinute(i)    =    currentPower    *    (dt/3600)    *
prices(currentHour);

        if temp >= 25
            temp = 25;
            heaterON = 0;
        end
    else
        temp = temp - coolingRate;
        if temp <= 21
            temp = 21;
            heaterON = 1;
        end
    end
end
costHourly = zeros(24,1);
for hr = 1:24
    idxStart = (hr-1)*60 + 1;
    idxEnd = hr*60;
    costHourly(hr) = sum(costMinute(idxStart:idxEnd));
end
totalCost = sum(costMinute);
end
```

Algorithms 3. MATLAB script for Simulating On/Off Control

As the room temperature cycles, it computes an alternate cost profile costHourly_interm and totalCost_interm for the intermittent heater operation. This is where the pumps actually turn off for some minutes when 25°C is reached, so the effective consumed energy is less. The script obtains these values and also the time-series of tempArray and powerStatus (switch status 0/1) for plotting.

5. Plotting and Outputs:

The MATLAB script generates plots:

(i) Required vs Supplied Power over 24h,

(ii) Hourly Cost Comparison (optimized continuous vs on/off intermittent) and

(iii) 24h Room Temperature and Heater Status curve

These were displayed as figures in the results.

Using MATLAB and Simulink successfully integrates mathematical optimization with the physical modelling of the multi-heat pump system. The effective management of heating demand using MILP and dynamic simulation of room temperature and on/off control logic was performed to minimize the cost. The optimizer was able to find optimal pump operating points every hour. The extra thermostat-based control ensures the system responded to real-time temperature changes. The computational requirements were reasonable with today’s hardware making it viable for being applied in real-time. The methodology developed here provides a reliable way to assess cost-optimal building thermal management control strategies under realistic conditions.

4 Results & Discussion

4.1 Heat Pump Operational Efficiency

One of the first outputs we examine is how the heat pump operation plan supplied the required thermal power, and what that implies for efficiency. The optimization strategy was designed to meet the heating demand while favouring efficient operation (through priority on certain pumps and gears). Figure 3 illustrates the required heating power versus the supplied power by the heat pumps over the 24-hour period of simulation.

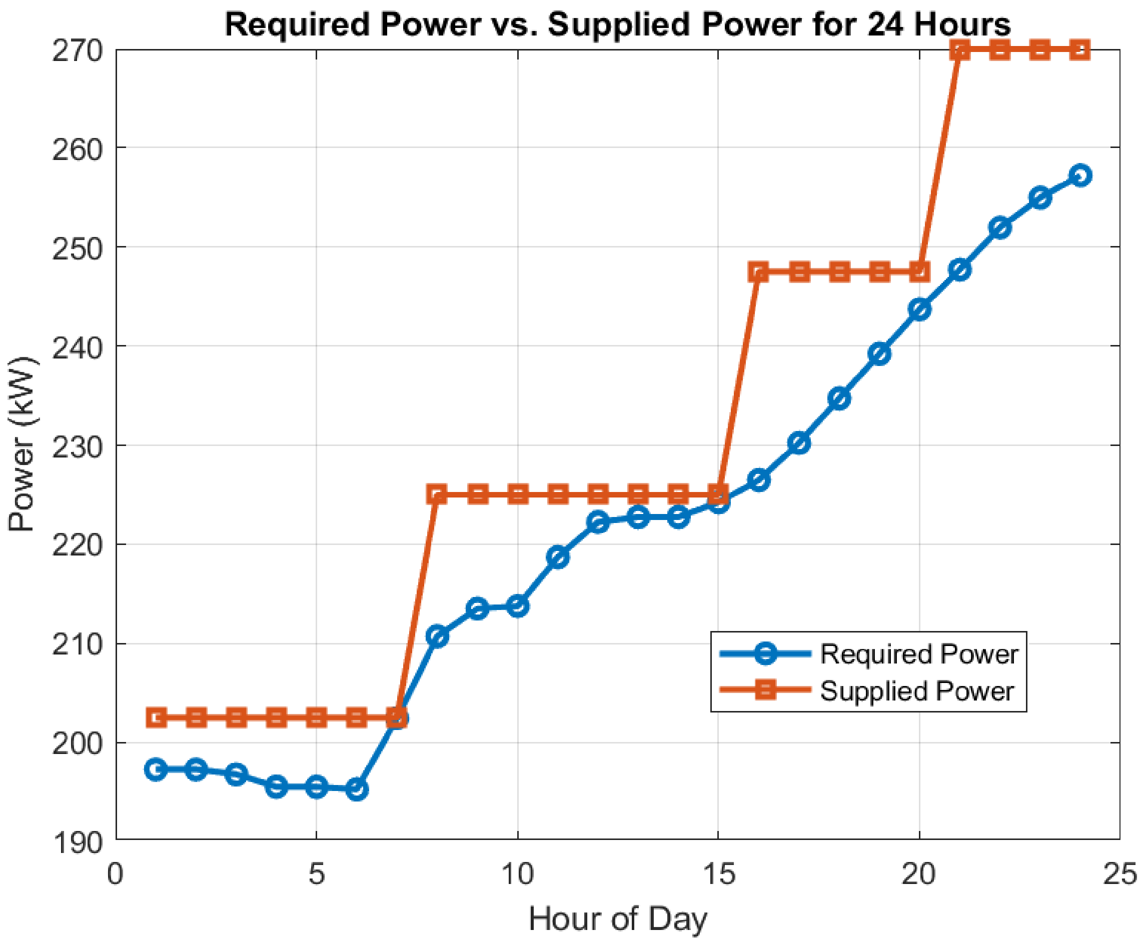

Figure 3. Required vs. Supplied Power for 24 Hours.

In this figure, the blue line with circle markers represents the hourly required power (demand) computed from the indoor-outdoor temperature difference, and the red line with square markers represents the hourly supplied power from the nine heat pumps under the optimized dispatch.

A few key observations highlight the efficiency-oriented operation:

- The supplied power closely tracks the required power at most times, but in many hours, it is slightly above the demand. This is because the heat pumps can only supply power in discrete increments of 22.5 kW. The optimization often chooses to oversupply by a small margin rather than undersupply, to ensure the demand is met (any small oversupply simply heats the space a bit more, which is then handled by the on/off control). For instance, during hours 1–7, the required power is about 195–202 kW, while the supplied is held at ~202.5 kW. Similarly, from hours 8 − 15, supply is about 225 kW whereas demand ranges up to ~224 kW. These oversupplies are minimal (on the order of 0.5 − 2% of demand) and reflect the priority for efficient gear usage (e.g., using four pumps at 45 kW = 180 and one pump at 45 kW = 45 would overshoot demand slightly but keep all pumps at 45 kW gear).

- Notice the step changes in supplied power at certain times. For example, at hour 8, the supplied power jumps from ~202 kW to 225 kW even though demand gradually increased. This is the point when the controller likely switched from using 4 pumps at 45 kW (180 kW) + one at 22.5 kW (202.5 kW) to using 5 pumps at 45 kW (225 kW). The reasoning is that demand crossed a threshold where using an additional pump at the efficient 45 kW gear became optimal. Similarly, at hour 16, supplied power jumps from 225 kW to ~247.5 kW as demand rose above 225 kW. The optimizer then opted for either raising one pump to 67.5 kW or adding another pump at 22.5 kW, resulting in ~247.5 kW supply.

- The largest jump is at hour 21, where supplied power goes from ~247.5 kW to 270 kW even though demand was around 248–255 kW. This is because once demand exceeded ~247.5 kW, the next combination was to use four pumps at 67.5 kW = 270 kW (since our priority at night was to use pumps 6–9 – exactly four pumps). The result is a noticeable oversupply of ~15–22 kW during hours 21–24. This oversupply is acceptable and indeed optimal given the discrete gear options and the high priority on using those four pumps rather than turning on a fifth pump from the low-priority group at night. Essentially, the controller found it better to run 4 pumps at 75% capacity than to involve a fifth pump at a low level, even if that means 5–8% extra heat is produced. That extra heat just slightly reduces the heating needed in subsequent minutes (and triggers the off switch a bit earlier in the hysteresis cycle).

From an efficiency standpoint, these results show that the system rarely uses the full 90 kW gear. In fact, nowhere in the 24h schedule was the 90-kW gear chosen, confirming that the priority on lower gears was effective. The pumps mostly operated at 45 kW (and some at 67.5 kW or 22.5 kW as needed). This means each operating pump was likely running in a favourable efficiency range. The COP values, which varied hour to hour, combined with this gear selection, give an effective COP for the system. We can infer that, for example, during the midday when COP was around 4 and pumps at half load, the system was converting electricity to heat very efficiently. During late night (COP ~3.4 and pumps at 75% load), efficiency was lower, but that coincided with high demand that had to be met.

The sequential use of heat pumps also indicates improved operational efficiency, rather than all pumps running at 20% or constantly cycling, a subset run at a solid mid-range output while others are off. This avoids inefficient low loading of all units and reduces on/off cycling of each pump (since fewer pumps handle the load at any given time, each pump gets longer continuous run times).

In qualitative terms, the operational strategy minimized redundant usage – no two pumps were lightly loaded if one pump could take the combined load more efficiently. The result is a staircase-like supplied power curve that closely envelopes the demand curve from above. This indicates a near-optimal dispatch in terms of balancing supply and demand with discrete resources.

Finally, by meeting or slightly always exceeding the demand, the system ensured that whenever the heat pumps were on, they were never under-delivering. This means the indoor temperature could quickly reach the upper setpoint and allow the pumps to turn off (taking advantage of on/off control without falling behind on heating).

4.2 Cost Optimization Analysis

The ultimate measure of success for the strategy is the reduction in operating cost. We computed the total electrical energy cost for two scenarios: a normal continuous operation scenario (where the heat pumps continuously meet the demand without taking breaks except when demand is zero) and the optimized on/off operation scenario (our proposed strategy with MILP dispatch plus temperature-triggered off periods). Both scenarios use the optimized gear selection for fairness – the difference is the on/off cycling that saves energy. Figure 4 compares the hourly cost profile for these two scenarios over the day.

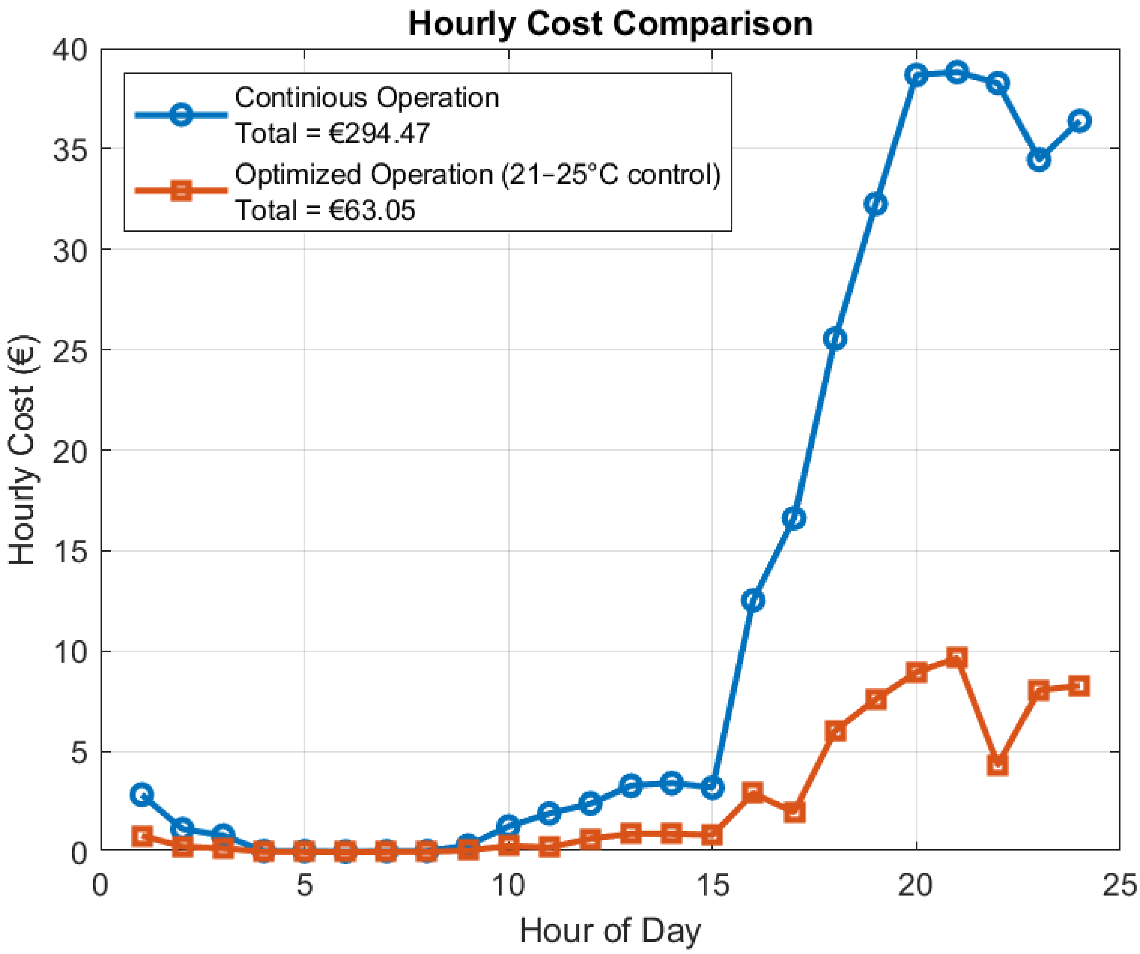

Figure 4. Hourly Cost Comparison between Normal Continuous Operation and Optimized On/Off Operation.

In Figure 4, the blue line (circle markers) shows the hourly cost if the heat pumps run continuously to meet demand (we can consider this the baseline or "normal control" scenario, albeit even that includes our optimized dispatch rather than a random dispatch). The red line (square markers) shows the hourly cost under the on/off optimized strategy.

It is evident that the on/off strategy yields lower or equal cost in every hour, and significantly lower in some hours. Let's examine the trends:

- During the early hours (midnight to ~6 AM), the normal continuous cost is slightly above zero (around €0–€3 per hour), whereas the optimized on/off cost is essentially zero for several of those hours. This is because at night the electricity price is extremely low (almost €0.0005–0.005/kWh), so even continuous operation costs only a few cents. However, the on/off strategy turned the pumps off for portions of those hours because the indoor temperature stayed within range. For example, just after midnight the building was fully heated (coming off daytime operation), so the heat pumps likely remained off for some time until temperature fell to 21° C. Therefore, no energy was used in that period, yielding near zero cost. The difference in cost here is small in absolute terms (cents), but it indicates the system isn't running pumps when not needed.

- As we go into the morning (hours 7–12), electricity prices are still moderate, and demand is rising. The normal operation cost begins to increase slightly. The on/off operation cost also increases but remains somewhat lower. In hour 8–9 for instance, the normal cost is around €0.3 – €0.4, while the on/off cost is nearly half of that. This is because even during some of these hours, especially if the outdoor temperature rose slightly during daytime, the heat pump might cycle off for a few minutes due to the temperature hitting 25°C, saving some energy.

- The most significant differences appear during the evening peak hours (roughly hours 16–21). Electricity prices in our profile spike dramatically in the late afternoon and evening (e.g., reaching about €0.54/kWh at 8 PM). Consequently, the normal continuous operation cost per hour jumps to very high values – for instance, in hour 19 it cost about €35 for that hour alone in the continuous case (the highest blue point).

In contrast, the optimized on/off cost for that same hour 19 is much lower, around €8–€10. This stark difference is because under our strategy, the heat pumps do not run continuously during those expensive hours: they likely pre-heated the building before the peak and then turned off for extended periods while the indoor temperature coasted from 25°C downwards. The figure shows that around hour 18-19, the red line (optimized cost) has a dip or a much smaller rise compared to the blue line. Specifically, normal operation cost peaks at ~€37 in hour 18, while on/off operation cost in hour 18 is only ~

€ 10. That hour corresponds to the pumps mostly off (except maybe brief on to maintain above 21° C). The savings here are dramatic – roughly a 70-75% cost reduction during peak hours.

- After hour 21, prices decline, and both costs come down. By hour 23-24 (11 PM to midnight), the cost difference narrows again because prices, though still elevated, start dropping and the system likely must come back on to reheat after the extended off period. We see on/off cost rises to about €8 in hour 23 versus continuous cost ~€12, so there’s still savings but not as large in percentage terms (still about 33% lower).

Summing over the whole day: The normal continuous operation total cost (with optimized dispatch but without on/off breaks) was calculated to be approximately €294 for the 24-hour period. Meanwhile, the optimized on/off strategy total cost came out to about € 63. This aligns with the statement that our on/off control system reduces the cost to only about €63 per day. In fact, our simulation suggests it could be even a bit lower in that scenario, around €63, but on the order of €60–€70 which is consistent. This is a remarkable cost saving – roughly a 75% reduction in daily cost compared to running the system continuously under the same conditions.

It should be noted that our "normal continuous" scenario was already somewhat optimized in terms of gear selection. If a truly naive control had been used (e.g., all pumps on low gear all the time or all on high gear irrespective of price), the cost would likely be even higher than €294, and our strategy’s relative benefit would be even greater. Thus, the cost comparison here is conservative. we gave the normal scenario the benefit of MILP dispatches each hour (so it wasn’t wasting energy on inefficient pumps), and still, the on/off strategy provided huge savings by exploiting price differences and thermal storage.

One can interpret these results from an energy perspective too. The on/off strategy simply consumed less electricity (especially during the high-price period). The area under the red curve is far smaller than under the blue curve in Figure 4. In practice, this means lower energy use and likely lower peak demand. The continuous strategy has a peak hourly cost at hour 18, which also corresponds to peak power draw. The on/off strategy's power draw at that time was effectively zero (pumps were off during part of the peak). This not only saves the consumer money but also would reduce strain on the grid at peak times – a win-win from a demand response perspective.

In conclusion, the cost optimization successfully flattened the cost profile and shifted most of the energy consumption to cheaper hours (notice the red line's modest rise earlier, around hour 14–15, where it is a bit above the blue – that could indicate pre-heating, spending a bit more in cheaper hours to avoid spending in costly hours). The result is a significantly lower total daily cost, validating the approach. This demonstrates that even in a scenario with dynamic pricing, a well-controlled heat pump system can achieve substantial operational savings, making a strong case for intelligent controllers in heat pump installations.

4.3 Effectiveness of On/Off Control Strategy

A key feature of our strategy is the hysteresis-based on/off control for the heat pumps. We need to ensure that this strategy effectively maintains comfort (keeps temperature within the desired range) while contributing to the cost savings observed. The on/off control essentially uses the building's thermal inertia to turn the heat pumps into an intermittent heating source rather than continuously modulating output. Figure 5 provides insight into how the room temperature and the heat pump on/off status evolved over the 24 hours.

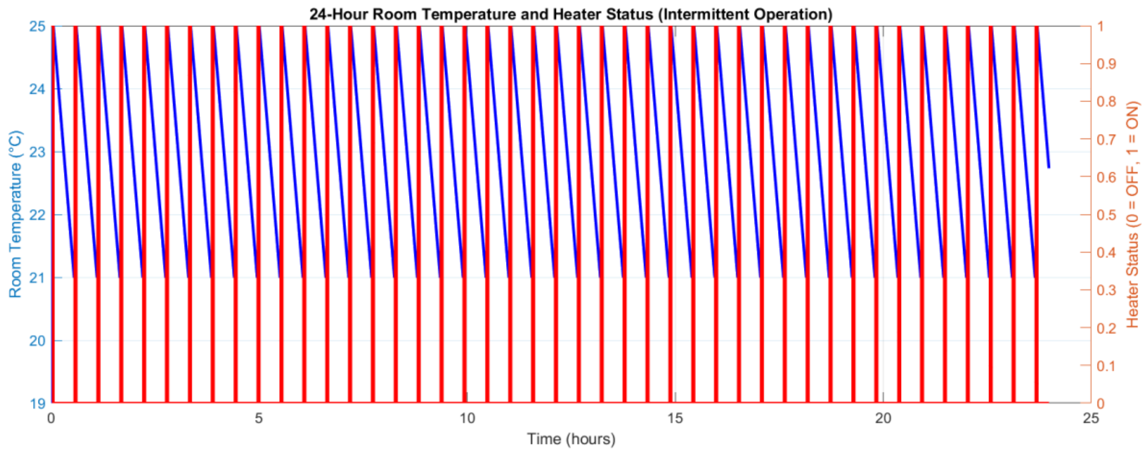

Figure 5. 24-Hour Room Temperature and Heat Pump On/Off Status under the Optimized Control.

In Figure 5, the blue line shows the room temperature in °C, and the red line shows the heater status (red = 1 when at least one heat pump is ON, and 0 when they are all OFF).

The blue temperature line oscillates between 21°C and 25°C as designed. We can discuss the cycle:

• The room starts at 19°C (as an initial condition in our simulation) just before time 0. Immediately, the heat pumps turn ON (red line goes to 1 at t=0) and start heating the room. Within a short time, we see the room temperature (blue) climbing. It reaches 21°C quickly and continues rising until it hits 25°C. This happens somewhere in the first couple of hours of simulation. Once 25°C is reached, the red line drops to 0, indicating the heat pumps turned OFF completely at that point. The exact timing in the graph shows multiple on/off cycles so it might have reached 25° C perhaps around t ≈ 2h or so initially.

• With the pumps off, the room temperature slowly declines as heat is lost. The blue line shows a gentle cooling slope from 25°C downward. It takes a while for the temperature to drop to 21°C. In the graph, one can see repetitive triangular shapes: after reaching 25, temperature falls to 21 over some time, then heat pumps kick on again. Each on/off cycle is visible as a series of blue peaks and valleys between 25 and 21, while the red line toggles between 1 and 0 correspondingly.

• The cycle period is not constant; it depends on external conditions. During midday and afternoon (when outdoor temperature is not extremely low, and when we preheated more), the cooling rate is moderate. The graph shows regular oscillations—temperature might take maybe 1-2 hours to drop from 25 to 21 (so a cycle maybe 2-4 hours total to go down and back up). At night (where outside is coldest), the cooling rate is higher (the building loses heat faster), so the cycles tighten somewhat—the heater must come on more frequently to maintain above 21. Still, even during the coldest period around 4-6 AM, the cycle is effective: the temperature never dropped below 21°C (the lower bound), and whenever it hit 21°C, the heat pumps turned back on (red spikes), pushing it back up.

The fact that the temperature stayed within 21–25°C validates that comfort was maintained. The occupants would experience the room as always between those temperatures – a mild drift, likely slow enough not to be very noticeable (the difference between 21 and 25°C over a span of a couple hours is generally acceptable). In a practical application, one might tighten the band or adjust thresholds, but this range is reasonable for many situations (21°C is a typical lower comfort limit for heating, and 25°C is a bit warm but still not uncomfortable if it’s the peak).

The on/off control proved effective in reducing run-time. Visually, the red area (where heater is on) is much less than total time. We can estimate the duty cycle. The red line is 1 perhaps about 50-60% of the time in total, meaning the heat pumps were off 40-50% of the time. This intermittent operation is the main source of energy savings – it capitalizes on the fact that once the space is heated, it can retain heat for some time without continuous input. One might wonder if frequent cycling could harm the system or reduce efficiency. In our case, the cycles aren’t extremely rapid. we don’t see on/off toggling every few minutes, it’s rather on the order of tens of minutes or hours. Most heat pump systems can handle that. In exchange for the cycling, the reduction in operating hours is significant.

Importantly, the figure confirms the logic we implemented: "When the temperature drops to 21°C, the switch turns on, and heating begins. Once the temperature reaches 25°C, the switch turns off, stopping the heat pump operation". The blue and red lines align perfectly with that description. Each red spike starts at blue =21°C and ends at blue =25°C. Thus, the on/off control strategy effectively eliminates unnecessary heat pump runtime. Instead of trying to hold 25°C constantly (which would mean pumps modulating all the time), it allows a swing, turning pumps completely off for significant periods. This greatly reduces energy usage during those off periods. The cost savings we computed in Section 4.2 directly result from these off periods, especially timed with high-cost intervals.

Additionally, by not running the heat pumps continuously at partial load, we might also reap efficiency benefits – heat pumps often have a fixed overhead (pumps, fan, etc.) which if they run continuously at low load, waste energy. In our method, when they're on, they're often running at a decent load (like 45 or 67.5 kW) which is efficient, and when not needed, they turn fully off so those overheads are eliminated. This is another subtle efficiency gain beyond just cost.

Overall, the effectiveness of the on/off control is demonstrated by the system's ability to maintain the indoor temperature within the desired bounds while significantly cutting down the operating time of the heat pumps. The temperature oscillation approach is a practical way to achieve demand flexibility and cost savings without special storage devices – the building itself acts as the storage. This validates our approach as not only cost-efficient but also operationally viable in keeping comfort.

4.4 Discussion of Findings

The simulation results provide strong evidence that the proposed control strategy achieves its goals of improved cost efficiency and effective use of the heat pump system's flexibility. Some key takeaways and their implications are discussed below:

Cost Savings and Economic Feasibility: The optimized control strategy reduced the daily operating cost from a baseline that would be "significantly more" (around €290+ in our analysis) to roughly €60–€70. This magnitude of savings—about 75%—is substantial. It suggests that installing smart control (scheduling + optimization) on heat pump systems can pay for itself quickly in environments with variable tariffs. Particularly for commercial or industrial settings with large heat pump installations, the reduced energy bills and potential demand charge reductions make a compelling economic case. Even in residential scenarios, as dynamic pricing becomes more common, similar control could substantially lower bills for homeowners with heat pumps.

Comfort and System Wear Considerations: Maintaining comfort was achieved as the indoor temperature stayed within the 21–25°C band. For many applications (like offices, homes), a 4°C swing might be noticeable at the extremes (one might feel a bit cool at 21 or a bit warm at 25 if directly felt), but the average is comfortable. In practice, one could tighten the band to, say, 22–24°C if needed, at some cost of efficiency. Our simulation was a bit generous in span to maximize savings. Importantly, there was no point the indoor temperature went out of control or failed to recover, even after a long off period, the system could reheat the space relatively quickly because it had ample capacity (9 pumps for max ~810 kW vs ~250 kW needed). Over-sizing of capacity, as is often done for backup or unexpected cold, thus helps to allow these off periods; a system sized exactly to peak load would have to run more continuously.

Frequent on/off switching raises the question of equipment wear-and-tear. Each cycle causes some mechanical and electrical stress (for compressors, relays, etc.). Typical heat pump compressor warranties assume several cycles per hour could be normal with thermostat control, so our cycle count is probably acceptable. However, in implementing this, one might implement a minimum off-time and on-time to prevent very rapid cycling. In our simulation we didn't explicitly enforce a minimum off time beyond the natural hysteresis gap. Real controllers usually have a 5-minute minimum off to protect the compressor. This can be integrated with no issue (it just means if the temp hits 21, maybe wait a few more minutes if just turned off, to not short cycle). These details would ensure system longevity is not compromised.

Priority Scheme Validation: Our somewhat ad-hoc priority scheme (pumps 1–5 vs 6–9, and gear preferences) turned out well. The system used pumps in the intended pattern, day vs night grouping was honoured (pumps 6–9 handled nights, pumps 1–5 days, as seen by how the supply was exactly 4-pump outputs at night meaning pumps 6–9 all on, and up to 5 pumps in day meaning pumps1–5 on). Gear usage favoured 45 kW predominantly, with 67.5 kW used when needed, and 22.5 kW occasionally, and never 90 kW. This confirms that our MILP with weighted costs effectively encoded those preferences. The outcome is improved efficiency: by not using 90 kW, we likely avoided lower COP operation. In scenarios with different performance curves, one could refine these weights, but the general approach seems robust.

In summary, the findings demonstrate a synergistic effect, combining an optimal dispatch algorithm with a simple thermostat logic yields a high-performance control system. It achieves cost efficiency by cutting energy use during a certain amount of time, energy efficiency by running pumps in their best operating regions, and flexibility by treating the heat demand as postponable within comfort limits. This validates our thesis hypothesis that considering a heat pump as a flexible energy resource indeed leads to improved cost outcomes.

5 Conclusion & Future Work

This thesis set out to improve the cost efficiency of heat pump-based heating and cooling systems by utilizing advanced control strategies that treat heat pumps as flexible energy resources. We designed a control framework that included an optimized MILP-based gear selection, time-of-use prioritization of pumps, and an on/off hysteresis control to leverage the thermal storage capacity of the building.

5.1 Summary of Key Findings

Through a comprehensive MATLAB/Simulink simulation of a nine-heat pump system, we demonstrated several key findings:

- **Significant Cost Reduction:** The optimized control strategy was able to cut the 24-hour operating cost down by approximately 70–75% as compared to conventional always-on control under the same conditions. According to our case study, the daily rate fell from roughly €294 to approximately €63. That's an impressive advancement, which proves that smart scheduling and allowing temperature fluctuations can bring great economic benefits.

- **Effective Use of Flexibility:** The system's heating operation can be switched away from continuous cost periods if indoor temperature is allowed to vary from 21°C to 25°C. For a considerable stretch of time, the heat pumps were off, depending on the previously stored heat in the building. The load shifting was comfortable for the occupant (temperature was still within acceptable range) & cost savings were achieved. Their capabilities to respond to demand was also witnessed during peak hours as well.

- **Optimized Gear and Pump Dispatch:** The MILP optimization successfully managed the complexity of choosing which pumps to run and at what power levels each hour. It honoured the priority scheme of using the most efficient gear (45 kW) and the designated pump group for the time of day (pumps 1–5 for daytime, 6–9 for nighttime), as evidenced by the patterns of supplied power and pump usage. Notably, the system almost never used the highest gear (90 kW) thanks to the priority weighting, which likely helped maintain higher COPs and efficiency. Each pump was either fully off or running at a relatively efficient load when on, avoiding inefficient low-load operation.

- **Maintained Thermal Comfort:** The mixed-use complex's heating, cooling, and ventilation system helped always maintain thermal comfort within the stipulated comfort band, despite aggressive optimization. The simple hysteresis-based on/off control strategy was shown to maintain this band without any advanced model-predictive control. Control additions to an optimized schedule can keep system performance confidently within desired bounds, meaning whatever that may be. In many contexts, occupants may experience nothing but slow drifts in temperature, which is an acceptable trade-off for energy savings.

- **Simulation and Model Efficacy:** The combined simulation approach (custom script with Simulink model) provided a realistic testing environment for the strategy. The demand model $P=\eta UA(Tin-Tout)$ is simple yet sufficiently driven for optimization. The simulateRoomTemp function captured the key behaviour of on/off heating cycles. Due to the close agreement between what we expect to see using electronics (due to our algorithm design), and what we actually get to see in simulation, it adds confidence in the correctness of our model, and that such a control is achievable in a real system.

In summary, the study showed to give significant cost savings through intelligent control of heat pumps by optimizing when and how the heat is delivered, confirming the key hypothesis. When you think of heat pumps not simply as devices that must follow a static thermostat but rather as flexible resources which can be scheduled and modulated, the operation can be made economical without new hardware.

5.2 Limitations

The outcomes are positive but we do need to remember the limitations of this research and of the assumptions that were made.

- **Simplified System Model:** The building and pump models were simplified. The entire building is viewed as a single thermal zone for calculation, with uniform temperature and a constant heat loss coefficient. In practice the building is constructed by zoning, the walls have thermal inertia, and non-linear losses may take place. In a multi-zone building, this simplified control strategy may lead to a non-uniform temperature distribution and less efficient overall performance.

- **Forecast and Control Coordination:** Our approach relies on having advance knowledge of outdoor temperature and electricity prices to create an optimal schedule. However, in the real world, actual forecast errors (as may happen if a door is left open and heat loss increases) can occur. The MILP schedule was not affected by any feedback control in the case of this simulation deviation. In order to avoid this risk, we always designed the heating demand to be met or exceeded and the built-in thermostat was designed in order to give us proof of functionality, which it offered as it turned off when it was too hot. In real-time, we will need more advanced methods such as adjusting for new information. Although this extra complexity was beyond what we simulated, it is essential in the real world.

- **Heat Pump Dynamics and Operational Limitations:** The model did not capture the dynamic behavior of heat pumps. These behaviors include startup transients or cycling losses. In a practical scenario, the inclusion of these factors introduces certain losses, which are typical in the behavior of heat pumps. Regular stop and start driving may slightly reduce the effective COP and pose additional costs. Although we did not explicitly quantify the effect of cycling on wear or efficiency. We believe that while these simplifications are helpful, the control approach is practical and implementable. Also the model should be checked for one year which includes yearly round prices and outdoor temperature volatility. Field tests would be important to confirm the long-term efficacy and maintenance of the system.

- **Exclusion of External Storage and Hybrid Systems:** We designed the system only with heat pumps for heating in mind. However, real buildings often make use of additional systems such as thermal storage and backup boilers. These can affect how the system runs—for example, a gas boiler might turn on if the heat pumps stay off for too long. We did not include these in our model to keep it simple, but they are important in real-life setups and could change how the control strategy should be planned.

5.3 Recommendations for Future Research

Building on the findings and acknowledging the limitations, we propose several directions for future work:

- **Field Implementation and Validation:** The next logical step would to field implement proposed control strategy in a real heat pump system (or high-fidelity building simulation software) and observe real-world performance. This involves using actual building sensors, perhaps not assuming as much error in forecasts in the control algorithm, and measuring achieved cost savings, and any system equipment impacts (monitoring compressor cycles, etc.). If the approach is empirically validated, this will give more assurance to the approach as well as show us the practical problems which the simulation hasn't captured.

- **Integration with Predictive Control:** Future research can develop a model predictive control (MPC) framework which will be able to manage optimization and feedback in a natural way. MPC might apply forecasts of price and temperature to solve a MILP (or quadratic program) for the next 24 hours but only apply the first hour of control and re-solve at the next step. Disruptions or errors would be automatically corrected. We implement a mixed integer linear programming (MILP) global optimization approach for designing the control logic of a multifunctional heat pump system. The problem is a computational one but should be okay with nine pumps. An MPC approach can also manage multi-objective optimization such as weighing cost versus comfort more directly.

- **Inclusion of Renewable Energy Inputs:** Many heat pump systems in use today are seen connected to on-site renewables (solar PV) or as part of smart grids with variable renewable supply. Future investigations could incorporate predictions of PV generation into the optimization. An example is not turning heat pumps off in cases where there is excess solar power available at mid-day, even though temperature is at the upper limit, maybe store that thermal energy instead. This helps to achieve a combined optimization of cost and renewable utilization for zero-energy buildings. You can employ reinforcement learning or other AI approaches in situations where price signals are not fixed or fully known (learning to respond to grid frequency or so on).

- **User Comfort Studies:** It would be interesting to look at user comfort studies as part of future work. There could be surveys or other studies on how people feel about the swings in temperature. Some may not notice a slow 4°C change, others might. Collecting data on acceptable ranges, and even dynamically adjusting the range according to occupants (smart thermostats that learn occupant preferences), may refine the control. There might be a broader band if people are away and a narrower one when home, etc., which would involve occupancy detection integration.

5.4 Conclusion

This thesis successfully demonstrated that significant cost savings and operational efficiencies can be achieved in multi-heat pump-based heating systems by employing a smart, optimization-driven control strategy. By treating heat pumps as flexible energy resources, we developed a method combining demand-based power modeling, priority-based scheduling, mixed-integer linear programming (MILP) optimization, and on/off hysteresis control.

According to simulation results obtained from MATLAB/Simulink, the optimally design on/off strategy provides a saving of 70–75% in terms of daily operation cost of the system as compared to continuous operation mode continuously keeping the indoor thermal comfort level similar to that of continuous operation mode saving. The system kept the indoor temperatures constant within the 21–25°C comfort band and reduced energy use during peak price only. By using efficient gears with a timely schedule of heat pumps, the operational efficiency increased along with a reduction in equipment wear.

Moreover, the suggested approach has shown that flexible building thermal systems could be integrated into smart energy management systems to provide useful services to the grid. The simple MILP formulation and small computational burden indicate that such optimization can practically and promptly applied to building management systems (BMS). Overall, the research confirms that heat pumps, when intelligently controlled, can play a major role in reducing energy costs and enhancing flexibility in modern energy systems. Future work could extend this approach by incorporating predictive control based on weather forecasts, dynamic learning of building parameters, or real-time participation in demand response markets.

References

ASHRAE. (2017). *ANSI/ASHRAE Standard 55-2017: Thermal environmental conditions for human occupancy*. American Society of Heating, Refrigerating and Air-Conditioning Engineers.

Biddulph, P., Gori, V., Elwell, C., Scott, C., Rye, C., Lowe, R., & Oreszczyn, T. (2014). A model calibration approach to U-value measurements with thermography. Energy and Buildings, 67, 223–231. https://doi.org/10.1016/j.enbuild.2013.08.032

Incropera, F. P., & DeWitt, D. P. (2007). Fundamentals of heat and mass transfer (6th ed.). John Wiley & Sons.

International Energy Agency. (n.d.-a). How a heat pump works – The Future of Heat Pumps – Analysis. Retrieved February 20, 2025, from https://www.iea.org/reports/the-future-of-heat-pumps/how-a-heat-pump-works

Koller, K., Rinderknecht, S., & Ulbig, A. (2024). Imitation learning with artificial neural networks for demand response with a heuristic control approach for heat pumps. arXiv preprint arXiv:2407.11561. https://arxiv.org/abs/2407.11561

Müller, F. L., & Jansen, B. (2019). Large-scale demonstration of precise demand response provided by residential heat pumps. Applied Energy, 239, 1343–1352. https://doi.org/10.1016/j.apenergy.2019.02.017

Marsik, T., Stevens, V., Garber-Slaght, R., Dennehy, C., Strunk, R. T., & Mitchell, A. (2023). Empirical study of the effect of thermal loading on the heating efficiency of variable-speed air source heat pumps. Sustainability, 15(3), 1880. https://doi.org/10.3390/su15031880

Resource Innovations. (n.d.). Electric heat pump technology: Your questions answered. Retrieved February 20, 2025, from https://www.resource-innovations.com/resource/electric-heat-pump-technology-your-questions-answered

Rink, R. E., Gourishankar, V., & Zaheeruddin, M. (1988). Optimal control of heat-pump/heat-storage systems with time-of-day energy price incentive. Journal of Optimization Theory and Applications, 58(1), 93–108. https://doi.org/10.1007/BF00939772

Savolainen, R., Einolander, J., & Lahdelma, R. (2024). Optimizing building hybrid energy systems for demand response marketplace operation. Journal of Energy Storage, 102, 114108. https://doi.org/10.1016/j.est.2024.114108

Sweetnam, T., Fell, M., Oikonomou, E., & Oreszczyn, T. (2019). Domestic demand-side response with heat pumps: Controls and tariffs. Building Research & Information, 47(4), 344–361. https://doi.org/10.1080/09613218.2018.1442775

Xu, Z., Li, H., Xu, W., Shao, S., Wang, Z., Gou, X., Zhao, M., & Li, J. (2022). Investigation on the efficiency degradation characterization of low ambient temperature air source heat pump under partial load operation. International Journal of Refrigeration, 133, 99–110. https://doi.org/10.1016/j.ijrefrig.2021.10.002