Developing a Method for the Operational Control of an Ecovat System

Gijs J. H. de Goeijen

UNIVERSITY OF TWENTE.

Developing a Method for the Operational Control of an Ecovat System

Dit proefschrift is goedgekeurd door:

Prof. dr. ir. G. J. M. Smit   (promotor)
Prof. dr.   J. L. Hurink   (promotor)

To decrease the emission of greenhouse gases, as well as to reduce our dependency on fossil fuels for satisfying our energy needs, we see a trend towards the use of more sustainable energy sources. While these sustainable energy sources, such as solar and wind, accomplish these mentioned goals they also present new challenges. One of these challenges lies in the fact that these energy sources are intermittent and uncontrollable. One of the consequences is that times of energy production do not necessarily coincide with times of demand. When energy is generated by fossil fuel powered energy plants, it is relatively easy to match the supply and demand of energy. However, this matching is much more difficult when relying on sustainable energy sources, such as solar and wind energy, due to their uncontrollable nature.

One of the solutions to deal with the mismatch of energy demand and production is energy storage. With storage, energy generated during times of excess production may be stored for use during times of energy shortage. In this context, energy storage may be used to cover mismatches occurring during a day, but also to cover the mismatch between different seasons. For the mismatches during a day electrical storage (i.e. batteries) can be used. However, batteries are currently too expensive for the large capacities required for seasonal storage. One of the promising options for seasonal storage is thermal energy storage. The Ecovat system is an example of such a seasonal thermal energy storage, which aims to store excess thermal energy during times of the year with high thermal and/or electrical energy production, generally during the summer, for use during times of the year with high thermal energy demand, generally during winter. The Ecovat system is designed to be able to satisfy the heat demand of a neighbourhood of houses throughout the year.

The Ecovat system consists of a large well insulated underground buffer (i.e. a large water tank), combined with a number of devices, namely photovoltaic thermal (PVT) panels, heat pumps, and a resistance heater, to charge the buffer. The buffer of the system consist of a number of segments, which although not physically separated, may be charged or discharged individually through heat exchangers integrated inside the buffer walls. The energy to charge the buffer can be obtained from locally available energy or can be bought on the energy market, preferably when the energy price is low.

In this thesis we focus on the operational control of such an Ecovat system. We develop a model to determine which of the available devices in the system

should charge which buffer segment at which point in time. Furthermore, the model also determines which buffer segment should be used to satisfy the heat demand from the neighbourhood. As the developed model should serve as the base for handling the operational control of a real Ecovat system, we are not just interested in an arbitrary model that is able to obtain charging/discharging strategies, but in a model that is able to compute these strategies in a short time (at most a few seconds).

Although we aim for a model with short computation times, we first focus on a model that does not take this restriction on the computation time into account. The goal of this first model is to get insight in the structure of a good charging/discharging strategy, i.e. a strategy that has low operational costs while satisfying the heat demand of the neighbourhood throughout the year. Furthermore, this model acts as a benchmark for other models that do satisfy the short computation time constraint. To this end, the first developed model is based on an integer linear programming (ILP) model of the Ecovat system.

Due to the long time scales involved when dealing with seasonal thermal energy storage (a year), as well as the short time interval lengths (15 minute time intervals) required to incorporate energy markets into the model, the developed integer linear programming (ILP) model can not be solved for an entire year at once. Due to this we developed an approach based on solving the ILP model in a rolling horizon fashion. Although this approach leads to a substantial reduction of computation time, we observe that solving the model in this way does not sufficiently take important seasonal effects into consideration.

To ensure that such seasonal effects are also taken into consideration by the model, we extend the model with a long-term planning step, which generates additional input for the previously developed model. In this planning step we determine daily energy targets for the buffer, based on historical data and predictions, which have to ensure that the correct seasonal behaviour is obtained. While the rolling horizon model with this extension is able to provide good charging/discharging strategies, we observe that even with these modifications the ILP model based approach is computationally still too expensive to be used in a practical situation, as in some cases it requires multiple days to determine a charging/discharging strategy for a year of operation of the Ecovat system.

Subsequently, we use the insights obtained from the ILP model based approach to develop a heuristic method to control the Ecovat system. This method is based on a number of rules of thumb, and contrary to the ILP model based approach, it does not require predictions for weather data and energy prices for future time intervals. This heuristic method requires much shorter computation times, namely it takes only a few seconds to simulate a complete year of operation of the Ecovat system. Comparing the results obtained with the heuristic method, with the results obtained with the ILP model based approach, we find that the heuristic method on average only increases the operational costs by 5.2%.

To get more insight in the practical use of the Ecovat system and the developed

approach we performed a case study, where we simulate a neighbourhood of houses including an Ecovat system in a decentralized energy management (DEM) simulation, using the developed heuristic method to control the Ecovat system. We compare the achieved results with a simulation using gas boilers to satisfy the heat demand of the neighbourhood instead. The results of this comparison show that using an Ecovat system to satisfy the heat demand leads to significant benefits in terms of energy self-consumption within the neighbourhood, as well as a decrease in CO_2 emissions compared to using gas boilers. Furthermore, the obtained results show that the developed approach is robust against prediction errors, such as e.g. a winter that is colder than predicted.

SAMENVATTING

Om de uitstoot van broeikasgassen te verlagen, en onze afhankelijkheid van fossiele brandstoffen voor onze energievoorziening te verminderen, zien we een trend richting duurzamere energie bronnen. Hoewel duurzame energiebronnen, zoals zonne- en windenergie, deze doelen behalen leiden ze ook tot nieuwe uitdagingen. Een van deze uitdagingen heeft te maken met het feit dat zulke energiebronnen onregelmatig en niet aan te sturen zijn, dit betekent dat periodes van energieaanbod en energievraag niet altijd samen vallen. Zolang energie gegenereerd wordt met behulp van fossiele brandstoffen is het relatief eenvoudig om het aanbod van energie gelijk te houden met de energievraag. Echter, wanneer we voor onze energie afhankelijk zijn van duurzame energiebronnen is dit veel moeilijker, omdat zulke energiebronnen niet aan te sturen zijn.

Een oplossing voor deze mismatch tussen energievraag en energieaanbod is energieopslag. Met behulp van opslag kan energie die geproduceerd wordt tijdens een periode met hoog energieaanbod opgeslagen worden voor gebruik tijdens een periode met een hoge energievraag. Op deze manier kunnen niet alleen mismatches gedurende een dag opgelost worden, maar ook de mismatch tussen verschillende seizoenen. Voor mismatches gedurende een dag kan elektrische opslag (batterijen) ingezet worden. Echter, de huidige kosten voor batterijen zijn te hoog voor de grote capaciteiten die benodigd zijn voor seizoensopslag. Een veelbelovende optie voor seizoensopslag is thermische energieopslag. Het Ecovat systeem is een voorbeeld van zo'n thermische seizoensopslag, met als doel om overschotten aan energie in tijden van hoge productie op te slaan, meestal gedurende de zomer, voor consumptie tijdens een periode met hoge energievraag, meestal gedurende de winter. Het systeem is ontworpen zodat het aan de warmtevraag van een woonwijk kan voldoen gedurende het hele jaar.

Het Ecovat systeem bestaat uit een grote, goed geïsoleerde, ondergrondse waterbuffer, gecombineerd met een aantal apparaten, namelijk 'photovoltaic thermal' (PVT) panelen, warmtepompen en een weerstandsverwarmer, om de buffer op te laden. De buffer bestaat uit een aantal segmenten, die ondanks dat ze niet fysiek gescheiden zijn, onafhankelijk van elkaar geladen of ontladen kunnen worden met behulp van warmtewisselaars geïntegreerd in de wanden van de buffer. De energie om de buffer te laden kan van lokale energiebronnen komen of gekocht worden op de energiemarkt, bij voorkeur wanneer de energieprijs laag is.

In dit proefschrift focussen we op de operationele aansturing van zo'n Ecovat systeem. We ontwikkelen een model dat bepaalt welke apparaten in het sys-

teem, welke segmenten in de buffer moeten laden, en op welk moment dat laden moet plaatsvinden. Daarnaast bepaalt het model welk segment van de buffer de warmtevraag van de woonwijk vervuld. Omdat het ontwikkelde model bruikbaar moet zijn voor de operationele aansturing van het Ecovat systeem in de praktijk zijn we niet simpelweg geïnteresseerd in een model dat laad/ontlaad strategieën voor het systeem geeft, maar een model dat dit kan doen in een korte tijd (maximaal een paar seconden).

Hoewel het doel is een model te ontwikkelen met lage computationele tijd, focussen we eerst op een model dat geen restrictie op computationele tijd heeft. Het doel van dit eerste model is om inzicht te krijgen in de structuur van een goede laad/ontlaad strategie, met andere woorden een strategie die leidt tot lage operationele kosten terwijl de warmtevraag van de woonwijk gedurende het hele jaar voldaan word. Daarnaast kan dit model gebruikt worden als referentiepunt voor een eenvoudiger model dat wel rekening houdt met de restrictie op de computationele tijd. Met dit als doel, hebben we eerst een 'integer linear programming' (ILP) model van het Ecovat systeem ontwikkeld.

Door de lange tijdsperiodes benodigd voor het simuleren van de werking van seizoensopslag, gecombineerd met de korte tijdsintervallen (15 minuten) die nodig zijn voor het toevoegen van een energiemarkt in het model, is het niet mogelijk om het ILP model in één keer op te lossen voor een heel jaar. Om deze reden ontwikkelen we een aanpak gebaseerd op het ILP model die een oplossing genereerd door middel van een rollende horizon aanpak. Hoewel deze aanpak leidt tot een significante afname in de computationele tijd zien we dat deze aanpak onvoldoende in staat is om seizoenseffecten mee te nemen.

Om te zorgen dat zulke seizoenseffecten voldoende meegenomen kunnen worden in het ontwikkelde model, breiden we het uit met een langetermijnplanning stap, die wordt uitgevoerd voor het eerder ontwikkelde model. In deze planning stap bepalen we dagelijkse energiedoelen voor de buffer, gebaseerd op historische data en voorspellingen. Deze energiedoelen worden vervolgens als extra input gebruikt voor de rollende horizon aanpak. Hoewel de rollende horizon aanpak met deze uitbreiding goede laad/ontlaad strategieën geeft, is de computationele tijd nog te hoog voor het gebruik in een praktische situatie, gezien in sommige gevallen een paar dagen nodig is om de laad/ontlaad strategie voor een jaar te bepalen.

Vervolgens gebruiken we het inzicht dat we verkregen hebben door de op het ILP model gebaseerde aanpak, om een heuristische methode te ontwikkelen om het Ecovat systeem aan te sturen. Deze methode is gebaseerd op een aantal vuistregels, en in tegenstelling tot de op het ILP model gebaseerde aanpak, heeft het geen voorspellingen voor het weer of de energieprijzen nodig voor toekomstige tijdsintervallen. Deze heuristische methode vereist een veel lagere computationele tijd, namelijk slechts een paar seconden voor het simuleren van de werking van het Ecovat systeem voor een jaar tijd. Bij het vergelijken van de strategieën bepaald met de heuristische methode met die bepaald door de op het ILP model

gebaseerde aanpak, zien we dat de heuristische methode gemiddeld slechts tot 5.2% hogere operationele kosten leidt.

Om meer inzicht te krijgen in de praktische toepassing van het Ecovat systeem en de ontwikkelde aanpak bekijken we een casus, waarin we een woonwijk met daarin een Ecovat systeem simuleren in een gedecentraliseerde energie management simulatie. Hierin gebruiken we de ontwikkelde heuristische methode voor het aansturen van het Ecovat systeem. We vergelijken de behaalde resultaten met een simulatie waarin gasketels worden gebruikt om aan de warmtevraag van de woonwijk te voldoen, in plaats van een Ecovat systeem. De resultaten van deze vergelijking laten zien dat het gebruik van een Ecovat systeem, in plaats van gasketels, om aan de warmtevraag van de woonwijk te voldoen leidt tot significante voordelen, namelijk een toename in de zelfconsumptie van energie in de woonwijk, en een afname in de CO_2 emissie. Bovendien laten we zien dat de ontwikkelde aanpak robuust is tegen voorspellingsfouten, bijvoorbeeld als een koudere winter dan verwacht zich voordoet.

xii

DANKWOORD

Dit is hem dan, mijn proefschrift, het resultaat van mijn onderzoek aan de Universiteit Twente de afgelopen vier jaar. In dit stukje tekst wil ik graag een aantal personen bedanken, want ondanks dat er maar één naam op dit boekje staat zijn er een heleboel mensen die direct of indirect hebben bijgedragen aan het tot stand komen van dit proefschrift.

Ik wil beginnen met het bedanken van mijn promotoren, Gerard en Johann. Gerard, tijdens mijn tijd aan de UT zat ik in een kantoor binnen de vakgroep CAES, die jij leidde tot je een paar jaar geleden met pensioen ging. Binnen deze vakgroep hangt een gezellige en open sfeer, waaraan jij naar mijn mening een grote bijdrage hebt geleverd, onder andere door er bij iedereen op aan te dringen deel te nemen aan de verscheiden formele en informele activiteiten binnen de vakgroep. Ik heb deze goede sfeer altijd als enorme meerwaarde ervaren tijdens mijn tijd aan de UT. Daarnaast stond je deur altijd open voor vragen of discussies en was je altijd bereid de tijd te nemen stukken tekst te lezen en van waardevolle feedback te voorzien. Johann, ik kan me de eerste keer dat we elkaar ontmoetten nog goed herinneren, tijdens mijn sollicitatie gesprek. Nouja, sollicitatiegesprek is misschien een groot woord. Ik geloof dat het formele gedeelte ongeveer vijf minuten duurde, waarna het al snel over Magic ging, een hobby die ik deel met jouw zoon. Net als bij Gerard, stond bij jou ook altijd de deur open voor vragen of discussies. Vooral tijdens het schrijven van mijn proefschrift hebben die discussies erg geholpen om mijn gedachten gestructureerd op papier te krijgen. Ook jij gaf altijd waardevolle feedback op stukken tekst die ik geproduceerd had, waarbij ik vaak versteld stond van hoe gedetailleerd die feedback was, zelfs als je iets al meerdere keren eerder had gelezen. Daar kwam dan wel bij kijken dat er daarna soms een hulplijn ingeschakeld moest worden om alle feedback te ontcijferen, hoewel als zo vaak ook hier oefening kunst baart. Er werd in ons kantoor wel eens gegrapt dat je pas echt kunt promoveren als je jouw feedback (soms ook wel Hurogliefen genoemd) zonder hulp kunt verwerken, gelukkig voldeed ik aan het eind van mijn promotie aan deze eis. Johann en Gerard, ik heb onze samenwerking altijd als erg prettig ervaren en daar wil ik jullie van harte voor bedanken!

De volgende groep mensen die ik wil bedanken zijn de collega's/lotgenoten met wie ik het al eerder genoemde kantoor deelde tijdens mijn tijd bij de UT. Dit kantoor, in de volksmond ook wel het energiehok genoemd, bevatte acht werkplekken die bezet werden door allerlei mensen gedurende mijn promotie, zowel mede-promovendi als studenten. Ondanks dat ik niet iedereen met wie ik dit

CONTENTS

CONTENTS

xviii

1

INTRODUCTION

Today's society is highly reliant on energy. We depend, among others, on energy for the transportation of both goods and ourselves, for heating our homes and for powering the various electronic devices present in the average modern house. Historically, most of our energy comes from the combustion of (fossil) fuels, starting with the combustion of wood thousands of years ago when fire was discovered. Later on, especially during the industrial revolution, coal became a popular source of energy. The combustion of coal allows for the evaporation of water into steam, which in turn can be used in a steam engine to perform mechanical work (or more recently, produce electricity). From that time on the amount of energy consumed by society grew at a rapid pace. Eventually other fossil fuels, in particular oil and natural gas, were added to satisfy the increasing energy demand. To give an indication of the growth in energy consumption, Figure 1.1 shows the global amount of consumed energy generated by the combustion of coal, oil and natural gas from 1800 to 2016. This figure clearly shows the immense increase in energy consumption since the industrial revolution, particularly in the past century.

In the last two centuries the combustion of fossil fuels has allowed society to make rapid progress. However, this combustion of fossil fuels has a major disadvantage, the emission of CO_2 and other greenhouse gases. It is widely believed that these greenhouse gases lead to global warming and climate change. The consensus among most climate scientists (90-100% agreement) is that the recent climate change is indeed caused by human behaviour [31]. In the last few decades this has lead to a transition towards cleaner, more sustainable ways of generating electricity in an effort to reduce greenhouse gas emissions. In 2015, during the United Nations Convention on Climate Change in Paris many countries, among which the Netherlands, signed an agreement to reduce the emissions of greenhouse gasses to keep the global temperature rise below 2 °C compared to pre-industrial levels [9].
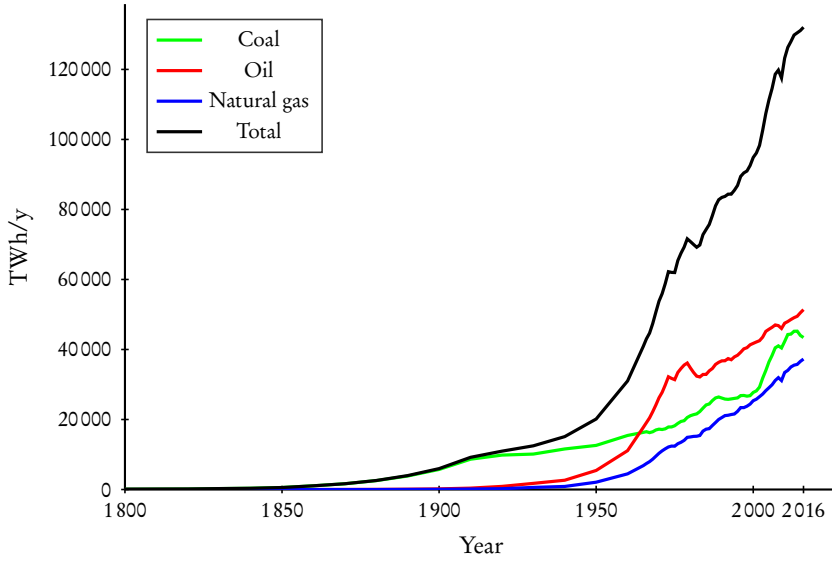
Figure 1.1: Annual global energy consumption by fossil fuel source. Data obtained from [6].

1.1 THE ENERGY TRANSITION

Aside from the aforementioned negative effects on climate change, fossil fuels are a limited resource, which are to a large extent located in politically unstable regions. For these reasons we see a shift away from such fossil fuels for satisfying our energy needs, towards sustainable options such as solar and wind energy. Figure 1.2 shows that the potential of renewables energy sources, especially solar, is in principle more than sufficient to supply the global energy needs. In fact, the yearly potential of solar energy is larger than the estimated total recoverable reserves of fossil fuels and uranium combined. Note that the area of the squares in Figure 1.2 on the left is proportional to the yearly potential and to the right on the total reserve. For reference, the yearly potential of solar energy is estimated at 23000 TWy/year, while the total reserve of coal is estimated at 830 TWy of energy [72].

The shift from generating energy by burning fossil fuels to renewable energy generation is called the energy transition. This energy transition has already lead to a dramatic increase in the installed capacity of renewable energy sources, in particular in the past decade. In Figure 1.3 the global installed capacity of several renewable energy sources for electricity generation is shown. We can see a clear increase in the installed capacities of renewables, with an especially sharp increase in installed solar capacity due to the decreasing costs of photovoltaic (PV) panels in the last years. Figure 1.4 shows the same figure but this time just

Yearly potential

Total recoverable reserves

Solar

Uranium

Natural gas

Oil

Coal

Wind

Global energy consumption 2014
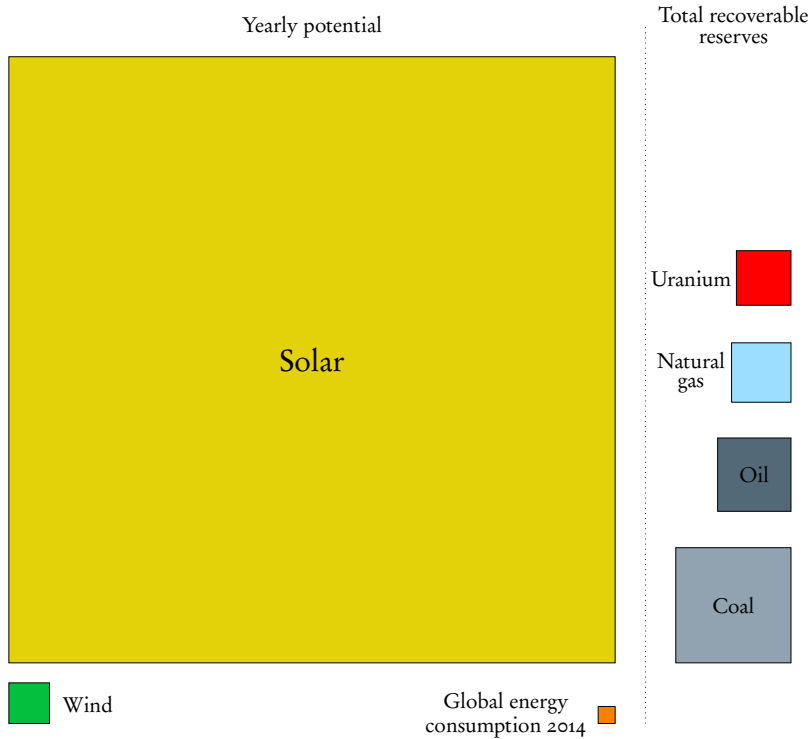
Figure 1.2: The yearly potential of solar and wind energy compared to the total estimated recoverable reserves of fossil fuels and uranium. The global energy consumption of 2014 is shown for comparison. Data obtained from [72].

for the Netherlands. We see a trend comparable to the worldwide trend, with the exception of hydro power, for which in the Netherlands the production is almost zero, due to the very small height differences in the Netherlands. While the installed capacity of renewables has increased a lot in the past decade, the Netherlands is still quite far from the targets concerning renewable energy set by the European Union [5]. The current share of renewables in the Netherlands is at 6.6% in 2017 [2] compared to the required 14% in 2020. It is expected that the installed capacities of renewable energy sources will continue to grow at a fast rate.

Even though the increasing share of renewable energy sources solves a number of problems, such as reducing the amount of green house gasses, it comes with its own challenges. When energy is supplied by large fossil fuel based power plants it is relatively easy to control the energy production in order to produce just the requested demand for energy. However, with renewable energy sources such as solar and wind energy this is no longer possible, due to the intermittent and
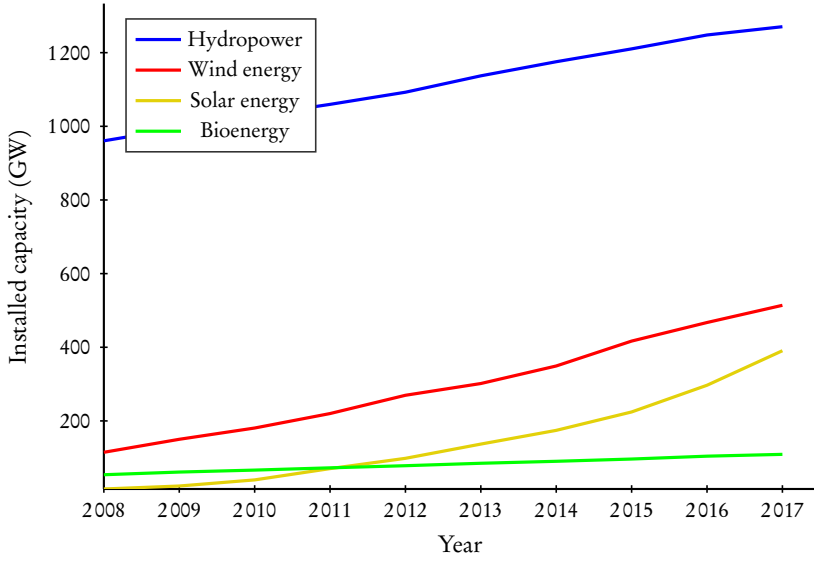
Figure 1.3: Global capacity of renewable energy sources for electricity generation. Data obtained from [8], the data is subject to copyright (©IRENA, 2018).
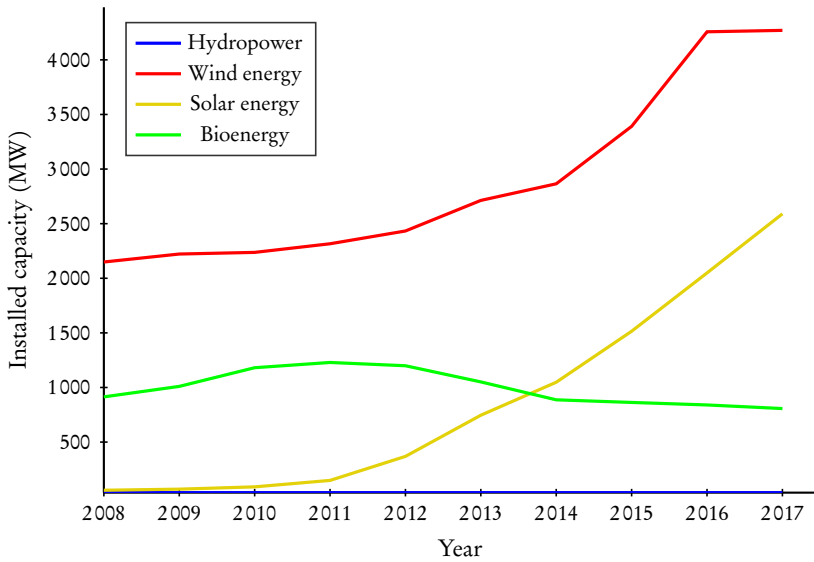
Figure 1.4: Capacity of renewable energy sources for electricity generation in the Netherlands. Data obtained from [8], the data is subject to copyright (©IRENA, 2018).
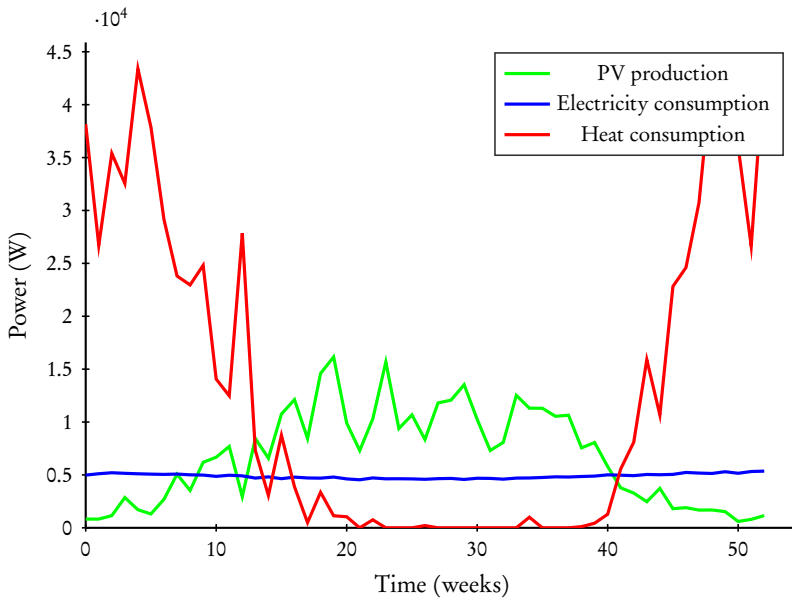
Figure 1.5: Simulated data showing the energy consumption and production from PV panels for a group of 16 houses.

uncontrollable behaviour of these energy sources. In other words, for solar and wind energy the times of production are determined by circumstances beyond our control, e.g. the time of day, the season and the weather. This often leads to a mismatch between energy supply and energy demand. For example, the energy production from PV panels in a residential neighbourhood tends to peak in the afternoon, when a lot of people are not at home and thus energy consumption in homes is low, while during the evening the energy demand peaks and the energy production from solar panels is lower/zero. Similar imbalance between supply and demand happens on a seasonal scale, where the solar energy production is highest in summer, while the time of year with the highest energy demand is the winter. Figure 1.5 shows this seasonal mismatch, with high (mostly thermal) demand in winter and high production in summer from PV panels. Note, that the differences in electricity consumption are much smaller over the year when considering weekly data such as in Figure 1.5. When the energy production from renewables in the energy system is lower than the energy consumption, this can be solved by backup power generation consisting of the traditional fossil-fuelled power plants. However, the use of such power plants is exactly what we are trying to reduce by means of the energy transition. Moreover, when the share of renewables increases this approach is no longer an economically viable option due to the high cost of the only seldom used backup power units, and other solutions for this problem need to be found.

There are a number of possible solutions to the aforementioned mismatch of energy supply and demand, from which we discuss two promising solutions here. The first is demand side management (DSM), which instead of changing the energy supply to match the energy demand tries to match the demand to the supply instead. To give an example of such an opportunity to shift the energy demand, consider a dishwasher. If someone turns on the dishwasher before leaving for work in the morning he/she tends not to care whether the dishwasher runs immediately or a few hours later, as long as it is finished by the time he/she gets home. This gives some flexibility on the demand side, which can be used to match the energy demand to the intermittent energy supply coming from renewable energy sources. The second solution is energy storage. Energy from renewable sources may be stored at times of surplus production, so that it may be used during times when these renewable sources produce insufficient energy to cover the demand. While this sounds like a simple and straightforward solution, it is quite expensive to install sufficient storage capacity into a system to simply solve the entire problem. Researchers often combine storage and DSM to obtain the better results.

While so far we have only discussed energy consumption in the form of electricity, a large portion of the energy consumption is in the form of heat. In 2016 41.2% of the total energy consumption in the Netherlands was used for heating, 20.2% for transport, 15% for electricity to power devices and the final 23.6% for the manifacturing of products using energy carriers as resources (for example using oil to produce plastic) [11]. If we look at the energy consumed within households the balance shifts even further towards heating, with 18% of the consumed energy being used to power devices and 82% being used for heating in 2018 [3]. Currently, most Dutch homes are heated with boilers using natural gas. However, the Dutch government wants to phase out heating of homes using natural gas completely by 2050 [7]. This means heating of homes needs to be done using another energy source than natural gas. One option is to use electrical means of heating, for example heat pumps, which can be combined with renewable sources of electricity generation discussed earlier. The disadvantage of this is that even more renewable energy is needed, exacerbating the problem of the supply-demand mismatch, as well as causing increased loads on the electrical grid. Another option is to rely on district heating networks combined with renewable ways of generating thermal energy, such as geothermal energy, solar collectors or photovoltaic thermal (PVT) panels (PV panels which next to electricity also produce heat), to supply the required heat. Alternatively, in some cases a district heating system may be fed with waste heat from other sectors, such as industry or agriculture. One challenge for district heating, when using renewable sources depending on the sun, is again a mismatch of energy supply and demand, but this time on a seasonal scale. During summer solar collectors and PVT panels generate more energy than is demanded, while in winter the situation is reversed (i.e. demand is higher and the production is lower), as shown in Figure 1.5. In this case DSM is not a real option, since there is very

Figure 1.6: Construction of the Ecovat prototype in Uden, the Netherlands. Image source: [4]

limited flexibility on the energy demand side. However, thermal energy storage can still provide a solution to this mismatch of energy production and demand. Advantages of thermal storage over electrical storage is that it is much cheaper to install large storage capacities, and that depending on the specific thermal storage technology, it has a much longer lifetime. The challenge, however, is to develop thermal storage solutions that can efficiently store energy on a time scale of months. One such particular seasonal thermal energy storage (STES) technology, the Ecovat system, is the focus of this work.

1.2 The Ecovat system

The Ecovat system is a novel STES solution developed by the Ecovat company [4]. It is designed to satisfy the heat demand of a neighbourhood of houses throughout the year. The most important component of the system is the Ecovat buffer, which is a large subterranean water tank. Energy is stored in the Ecovat buffer by heating the water inside the buffer. Different size Ecovat buffers are planned, the smallest of which has a diameter of 11 meters. The depth of the buffer is around 16 meters, independent of the diameter. Figure 1.6 shows a photograph of the construction of the Ecovat prototype in Uden, the Netherlands. The Ecovat buffer is divided into N_{seg} horizontal segments, which can be charged or discharged individually through heat exchangers built into the buffer walls, meaning that in principle each segment will be at a different temperature. The temperature of the buffer segments decreases from the top segment to the bottom segment to avoid mixing. It is important to note that these buffer segments are not physically separated, but merely specify different regions inside

the buffer. More precisely, the Ecovat buffer is thermally stratified, which means that segments of different temperatures do not mix but instead are seperated by a thin region with a steep temperature gradient. This phenomenon of thermal stratification is also observed in nature, for example in lakes. In that case the top region of the lake is heated by the sun while the bottom region remains cold, and the boundary between these two regions is a thin region with a steeply decreasing temperature. By using the heat exchangers in the walls of the Ecovat buffer in a smart way the buffer can remain thermally stratified, which has been shown to increase the efficiency of water tanks [62, 76].

The charging of the Ecovat buffer, i.e. storing energy in the buffer by means of heating the water in one of the buffer segments, is done through a set of devices accompanying the buffer. This set consists of PVT panels, a resistance heater, and a number of heat pumps. In general, alternative configurations of devices are possible in the future depending on the specific circumstances and requirements of the party interested in installing an Ecovat system. However, in this thesis we limit our focus on this specific configuration of devices for charging the Ecovat system. In addition to the devices accompanying the buffer other locally available heat sources (for example waste heat from industry or agriculture) may be used for charging the buffer.

The PVT panels can produce both thermal and electrical energy. The thermal energy can directly be stored in the Ecovat buffer, while the electrical energy can be used to power the other devices accompanying the buffer, thus indirectly charge the buffer. Note, that excess electrical energy produced by the PVT panels can also be sold on the energy market. Conversely, additional electricity to power the devices in the system can be purchased on the energy market as well. The resistance heater can be used to convert large amounts of electrical energy into thermal energy on a one-to-one basis in a short time. However, the resistance heater is quite inefficient at heating the buffer, since the coefficient of performance (COP) of a resistance heater is one while the COP for heat pumps can be significantly higher. Due to this the resistance heater is preferably only used during times when the energy price on the market is low or even negative. In this way the buffer can be charged while simultaneously making a profit on the energy market. The last device to charge the buffer is an air-water heat pump, which uses the ambient air as a heat source and one of the buffer segments as a heat sink.

Aside from the aforementioned devices to charge the buffer the Ecovat system also contains two water-water heat pumps for internal use. The aim of these heat pumps is to increase the energy quality of the buffer (high temperature heat is considered to have a higher energy quality than low temperature heat). This is done by using one of the buffer segments as a heat source and using another buffer segment with a higher temperature as the heat sink. The reason to aim for a high energy quality is that a certain minimum temperature is required to satisfy the heat demand of buildings connected to the Ecovat system. To satisfy

heat demand at a given demand temperature, T_{dem}, at least one of the segments of the buffer needs to be at a temperature higher than T_{dem}. In the envisioned setup of the Ecovat system there are two water-water heat pumps because the possible temperature range (5-90 °C) across the buffer segments is too large to be efficiently covered by one heat pump, since a larger temperature difference between the heat source and heat sink leads to a lower COP of the heat pump. Due to this, one water-water heat pump covers the lower part of the possible temperature range (in general only the bottom segments of the buffer) while the other covers the higher part of the temperature range (in general only the top segments of the buffer).

The Ecovat buffer is well insulated to minimize heat losses to the surrounding environment. According to a thermal analysis in an internal report [86] the heat losses are estimated at 10% or less over a period of 6 months, depending on the size of the Ecovat buffer. Note, that significant mixing between layers of different temperatures in water tanks may occur, if water is pumped into or out of the water tank. As this leads to a reduction of the temperature of the hot layer(s) in the water tank and thus reduces the ability to satisfy the heat demand as noted above, the Ecovat system uses solely the heat exchangers in the buffer walls for the charging and discharging of the buffer. By this, the amount of mixing is significantly reduced compared to other water tank designs.

1.3 Problem statement and approach

The previous section gives some insight into the operation of an Ecovat system. However, it becomes clear that there are still a lot of questions regarding the operational control of the Ecovat system. When should the Ecovat buffer be charged/discharged or which energy price should we be willing to accept to charge the buffer? Do we charge the buffer at the current time or do we wait for potentially even better circumstances, e.g. even lower energy prices, to charge in the future? If we do decide to charge the buffer, in which way should we do that or in other words, which of the devices in the system should be used to charge which segment of the buffer for the best result? Due to the very large amount of options in such a complex system, these questions do not have a simple answer. As such, research is needed to determine the best way to perform the operational control of the Ecovat system. This leads us to the main research question of this thesis:

> *How can we model and determine the operational control of the Ecovat system such that the resulting method provides good charging/discharging strategies, with a computational time short enough to be usable in a real world situation?*

The goal when controlling an Ecovat system is to supply the heat demand of a neighbourhood of houses connected to the buffer at minimum operational costs. However, due to the complexity of the system and the dependency on uncer-

tain input data, such as weather predictions and energy prices, determining the optimal charging/discharging strategy is not possible in practice. As such, our goal is to obtain good charging/discharging strategies, i.e. charging/discharging strategies that approach this optimal strategy, but that can be achieved in practice. Hereby the goal is not just to develop a model that gives charging/discharging strategies for the Ecovat system, but specifically a model which can do so on time scales short enough to be usable by an Ecovat system in the real world. For the control of such a real system a model is required which can make decisions on (sub)second time scales to be able to adapt quickly to changing circumstances, such as a change in the energy price or a change in the weather conditions. To make these decisions on such a short time scale for a system as complex as the Ecovat system it seems likely that some kind of heuristic method is required for the operational control of the system in the real world. However, related to the discussion before it is unclear what such a method should look like.

To investigate this topic in detail, a number of sub-questions are defined to guide the development of such a heuristic method. The first sub-question is:

1. *What is a possible method for determining the control of the Ecovat system, which is able to provide a good charging strategy for the Ecovat, assuming computational time is not a limitation?*

For the first step in the process of answering the main research question, we assume that computational time is not a restriction to determine a good charging strategy, and furthermore, we also assume an unlimited amount of computational resources is available. The question now is what a proper model of the Ecovat system should look like? Such a model has to be an optimization model which can give a good charging schedule for the Ecovat system, but might be too complex to be solved in a reasonable amount of time given limited computational resources. The goal in this step is to develop an unrestricted model as the basis for the development of further simpler models. This leads us to the second sub-question:

2. *How can we adapt the developed control method to be solvable given limited computational time and resources, while maintaining a good charging strategy?*

In the second step an adapted version of the model in the first step is developed. This adapted model should be solvable in a limited amount of time using a limited amount of computational resources, more specifically in time scales of hours to days on a personal computer. The goal is to achieve a model of the Ecovat system that is usable in simulations over a time frame of a complete year. This model will still be an optimization model. However, the solutions provided by this model will most likely be slightly worse due to the necessary adaptations to make it solvable given limited time and computational resources. Note, that the goal in this step is only to develop a model which gives good charging strategies

and is solvable in limited time, but which will still be too slow for use in practice. Such a model is useful to gain insight into the system, based on which we can develop a simpler heuristic model that is usable in practice. Furthermore, this model can then also be used as a benchmark to compare a simpler model against. This leads us to the third sub-question:

3. *How can we use the insights obtained from the approach developed for the second sub-question to derive a heuristic method capable of controlling an Ecovat system in real time, while maintaining a good charging strategy?*

The final step is to develop a heuristic method which can provide good charging strategies for the Ecovat system on a time scale short enough to be usable in a real world Ecovat system. This can be achieved by using the insight obtained from the model in the second step to develop a heuristic method which can generate charging strategies very quickly. Furthermore, the strategies obtained using this heuristic method can then be compared to the strategies obtained with the model from the second step. In this way we can ensure that the strategies obtained using the heuristic method are of good quality, i.e. close to the ones obtained using an optimization model. Finally, after the development of such a heuristic method for the Ecovat system, we are interested in testing the robustness of this model. This leads us to the final sub-question:

4. *How robust is the developed approach of the Ecovat system to unpredicted deviations in the input data?*

The developed model of the Ecovat system requires a number of inputs, such as data on the expected heat demand, predictions for the weather, and predicted energy prices on the energy market, to generate a charging/discharging strategy. However, some of this data is expected to be unreliable, since in general making predictions about the future is very hard, especially for volatile processes such as the energy prices. As such, it would be preferable to have a model that either does not depend on predictions at all, or if it does is robust against errors in these predictions. Therefore, the robustness of the developed heuristic model to errors in prediction is also considered in this thesis. This is done, among others, by means of a case study in which a neighbourhood of houses, including an Ecovat system, is simulated in a DSM setting.

1.4 THESIS OUTLINE

In this introductory chapter we have sketched the changes the energy system is currently undergoing, shifting from the burning of fossil fuels for generating energy to more sustainable alternatives. The challenge to match energy supply and demand arising from this transition can, among others, be addressed using energy storage. We have introduced a specific energy storage technology, the Ecovat system, which is the STES technology this work focusses on.

The remainder of this thesis is organized as follows. In Chapter 2 we give an overview of different thermal storage solutions. We focus in particular on water tanks, since those are the thermal storage solutions which closest resemble the Ecovat system. We conclude the chapter with some background on district heating systems and DSM, specifically discussing DSM research including a thermal storage. In Chapter 3 we describe the modelling of the Ecovat system and its control in the form of an integer linear programming (ILP) model, which gives good charging/discharging strategies for the Ecovat system but is too slow even for simulation purposes. Subsequently, in Chapter 4 we describe the modifications to this ILP that are necessary to ensure it becomes usable in simulations. Next, in Chapter 5 we use the insight gained from the ILP models developed in Chapters 3 and 4 to develop a heuristic method for controlling the Ecovat system, which is fast enough such that it is usable in practice. This heuristic method is applied to a case study described in Chapter 6 in which we simulate a neighbourhood including an Ecovat system in a DSM setting. Finally, in Chapter 7 the conclusions of this thesis are presented and potential avenues for future work are discussed.

2

BACKGROUND

ABSTRACT – *In this chapter we provide an overview of thermal energy storage technologies, with specific focus on water tanks, since that is the thermal energy storage technology which closest resembles the Ecovat buffer. In the discussion on water tanks we describe similarities and differences between these and the Ecovat buffer. Furthermore, we provide a short discussion on district heating systems and how the trend of decreasing temperatures in such systems increases the usefulness of thermal storage technologies such as the Ecovat. Finally, we conclude with a short overview of demand side management (DSM) methods, focussing in particular on previous DSM research integrating thermal energy storage.*

As described in the previous chapter, one of the solutions to resolve the mismatch between demand and supply caused by the energy transition is energy storage. Energy storage exists in many forms, such as: electrochemical storage (batteries [35]), magnetic storage (superconducting magnetic energy storage [20]), mechanical storage (pumped hydro storage [74], compressed air storage [28], flywheels [54]) and thermal storage. In this thesis we focus solely on thermal energy storage, since the Ecovat system is a thermal energy storage technology. For more information on the other mentioned energy storage technologies we refer to literature reviews such as [26] and [15].

Thermal energy storage can be divided into three categories, sensible thermal energy storage, latent thermal energy storage and chemical thermal storage. In sensible thermal energy storage, energy is stored by means of heating the storage medium, for example water or rocks. In latent thermal energy storage energy is stored by means of a phase change in the storage medium, which in the case of latent thermal energy storage are called phase change materials. Finally, in chemical thermal energy storage heat is stored by means of reversible chemical reactions, with one direction of the chemical reaction being endothermic (energy is absorbed) and the other direction being exothermic (energy is released).

In the remainder of this chapter we first describe different thermal storage technologies. After that we go into more detail on water tanks specifically, since those are the thermal storage systems which closest resemble the Ecovat system. Finally, we give some background on district heating systems and demand side management.

2.1 Thermal energy storage

In this section we present a short overview of the aforementioned three categories of heat storage, i.e. sensible thermal energy storage, latent thermal energy storage and chemical energy storage. Even though thermal energy storage can also be used for storing cold (see e.g. [63, 96]), we focus mainly on the storage of heat in this thesis, since that is the focus of the Ecovat system.

2.1.1 Sensible thermal energy storage

As mentioned above in sensible thermal energy storage energy is stored by means of increasing the temperature of the storage medium. Many different storage media can be used for thermal storage, each with their own advantages and disadvantages. The most commercially used storage medium is water [77]. The advantages of using water as storage medium are its wide availability, low cost, non-toxicity and high specific heat [19]. Other common storage media are earth materials, i.e. soil, rocks, sand, gravel etc. Like water, these earth materials are cheap, easily obtainable, non-toxic and non-flammable [19]. They can withstand higher temperatures than water, but have lower specific heats. Additionally, there are many other possibilities for storage media for sensible thermal storage, such as e.g. thermal oils, molten salts and liquid metals [19, 85]. Compared to water these materials have a larger temperature range over which they can operate. However, they have (much) lower specific heats and are more expensive. Specific heats of a number of materials used in sensible thermal storage are given in Table 2.1 [18].

In the following we discuss a number of thermal storage technologies that use water and/or earth materials as their storage media. Since the objective of the Ecovat system is to provide seasonal thermal storage for a group of houses, and since the most used storage media in seasonal thermal storage technologies are water and earth materials [73], we focus on technologies using those storage media in particular.

Water tanks: Water tanks are man-made structures and can be constructed under or above the ground. They are generally thermally stratified, meaning that there are separated regions of different temperatures within the water tank. Water tanks have to be well insulated to prevent large heat losses to their surroundings. Since a water tank is the thermal storage technology that closest resembles the Ecovat buffer we discuss water tanks in more depth in Section 2.2.

Storage material	Type	Specific heat (kJ/kg°C)
Water	Liquid	4.18
Rock	Solid	0.96
Sand	Solid	0.83
Concrete	Solid	0.85
Mineral oil	Oil	1.97
Therminol VP-1	Oil	1.55
$NaNO_3$	Molten salt	1.66
KNO_3	Molten salt	0.96
NaOH	Molten salt	0.92
KOH	Molten salt	1.34
Al	Liquid metal	0.89
Na	Liquid metal	1.3

Table 2.1: Specific heat of some materials used in sensible thermal storage, data obtained from [18]

Aquifer thermal storage: In aquifer thermal storage at least two wells (one hot and one cold well) are drilled in an aquifer (a underground layer of material permeated by water). In winter heat is extracted from the hot well and used to heat the building(s) connected to the aquifer thermal storage. The resulting cold water is then injected into the cold well. During summer this process is reversed and water from the cold well is heated by e.g. solar energy or heat from the building(s) and is injected into the hot well. Thermal energy is stored in both the groundwater as well as the material in the aquifer, this means the volumetric heat density depends on the properties of the material in the aquifer [32]. Aquifer thermal energy storage can not be employed everywhere due to its geological requirements, such as a sufficiently thick aquifer layer and low to no groundwater flow within the aquifer [93]. Even though the energy density of aquifer thermal storages is limited due to the low temperatures used (for example aquifers in the Netherlands have a maximum temperature of 15-20 °C and a maximum injection temperature of 25 °C [10]), they have a high storage capacity due to their large volumes [33]. Aquifer thermal storages are implemented worldwide, with the majority (85%) in The Netherlands [41].

Borehole thermal storage: In borehole thermal storage deep vertical shafts, called boreholes, are drilled into the soil. Inside these boreholes tubes are placed through which the heat exchanger fluid flows and exchanges heat with the surrounding soil. This is one of the differences compared with aquifer thermal storage, where usually the ground water itself is used as heat transfer fluid. Another difference is that borehole thermal storage is not dependent on the presence of an aquifer, or more generally the presence of groundwater, but can be used in most ground formations [61]. Operation of borehole thermal storages is similar to that of aquifer thermal storage. In summer the used heat transfer fluid is hot

and heats up the surrounding soil to store heat in it, conversely in winter the used heat transfer fluid is cold and extracts heat from the soil to provide energy, for example for space heating. As with aquifer thermal storage the energy density of borehole thermal storage is low. When compared to water tanks the volume of the borehole thermal storage needs to be 3-5 times larger to store the same amount of energy [93]. However, borehole thermal storage is less geographically limited than aquifer thermal storage and less expensive than water tanks [61]. Borehole thermal storage is used around the world [61], with a well know example being the Drake Solar Landing Community in Okotoks, Canada. In the Drake Solar Landing Community, a small short-term energy storage and a borehole thermal energy storage consisting of 144 boreholes supply energy for space heating of 52 energy-efficient houses throughout the year. The energy in this system is generated by PV panels. It was the first system of this kind designed to supply more than 90% of the space heating requirement of a neighbourhood throughout the year. In the fifth year of its operation a solar fraction of 97% was reported [83], i.e. 97% of the space heating requirements were satisfied using solar energy, clearly demonstrating the potential of such systems. However, that year only 36% of the heat supplied to the borehole thermal storage was retrieved for later use, showing that there are also sizeable heat losses in such a system.

Solar ponds: Solar ponds are generally 1-3m deep and designed to retain captured sunlight. Solar ponds have a reversed natural temperature gradient. In other words, the hot water is at the bottom of the pond while colder water is on top. This reversed gradient is maintained by a layer of increasing concentrations of salt towards the bottom of the pond, such that the density of the salt water at the bottom of the pond is larger than the (less salty) water on top of it, even when it is heated to a high temperature [34]. Due to the layer of increasing density convective flows upwards are suppressed and heat exchange within the pond only happens through conduction [89]. This leads to an insulating layer that allows solar irradiation to penetrate the pond but retains the heat in the bottom layer. Heat can be extracted from the pond by a heat exchanger in the bottom layer of the pond or by extracting water from the bottom layer, using an external heat exchanger, and returning the cold water back to the top of the pond [37]. Solar ponds have a number of applications such as the heating of buildings, power production, industrial process heating, desalination and salt production [37, 84]. An example of a solar pond used for heating purposes is a 2000 m^2 pond constructed in 1978 in Miamisburg, Ohio to heat a swimming pool during summer and the accompanying bathhouse throughout the entire year [34]. However, even though solar ponds show potential as thermal energy storage, a review of the literature shows only a few examples of solar ponds in use. According to [33] favourable conditions for solar ponds include much sunshine, little snow and easy availability of land, which may explain why we do not see solar ponds being used much in e.g. Western Europe. However, in [84] it is mentioned as a good option for developing countries.

Rock beds: Rock beds, or packed beds, consist of a container of rocks or other

solid materials used to store heat. This is done by circulating a heat transfer fluid, usually water or air, through the bed. To store heat the heat transfer fluid flows through the bed in one direction, while during discharge the direction is reversed. This means that contrary to water tanks, simultaneous charging and discharging is not possible [77]. When air is used as the heat transfer fluid it does not contribute to the storage, while if water is used the heat transfer fluid does contribute to the storage [73]. Advantages of using rocks instead of water as the heat storage medium include the easier containment of rocks as well as the higher temperatures they can withstand. However, as mentioned before, a significant disadvantage is the lower specific heat of rocks compared to water, which means a larger volume is required to store the same amount of thermal energy (around a factor of 3 larger [33]). While there is a lot of research being done on rock beds in an experimental setting, e.g [17, 21, 29], there are only a few examples of rock beds currently in use. One of these examples is presented in [99] were a rock bed is used to supply space heating to a dormitory and cafeteria on Qinhuang Island, China.

As evident from the discussion above sensible heat storage technologies are well developed. Different forms, mainly water tanks, aquifer thermal storage and borehole thermal storage, of sensible heat storage are currently being used in practice in seasonal thermal storage projects around the world.

2.1.2   LATENT THERMAL ENERGY STORAGE

In latent thermal energy storage thermal energy is stored by means of a phase transition in the storage medium. Storage media for latent thermal energy storage are called phase change materials (PCM). Different phase changes can be used for latent thermal energy storage. While liquid-gas PCM have a high latent heat, they suffer from a high change in volume when transitioning from one phase to the other [77]. The most used materials are solid-liquid PCM, which only suffer from limited volume variations during the phase change [77]. The usage of solid-solid PCM, where the phase change occurs between different crystalline phases, is also being researched [39].

Advantages of latent thermal energy storage over sensible thermal energy storage are the much higher energy densities involved in latent thermal energy storage [93] as well as the near constant temperature during the phase change [77]. This means that the latent thermal energy storage can be discharged at a near constant temperature. Disadvantages of latent thermal energy storage include low thermal conductivity and, depending on the specific PCM, flammability or corrosiveness [19]. In general PCM are also more expensive than sensible thermal storage materials [19].

PCM are divided into organic and inorganic materials. Organic PCM include paraffins, fatty acids, esters, alcohols and glycols [19]. Advantages of organic PCM compared to inorganic PCM are that they are not corrosive, are chemically and thermally stable and show little to no supercooling [77, 97] (supercooling is

the phenomenon of a material solidifying (freezing) at a temperature below its normal freezing point, which negatively impacts the efficiency of the PCM [25]). Disadvantages are their low thermal conductivity, low phase change enthalpy and organic PCM being flammable [97]. Inorganic PCM include salts, salt eutectics, salt hydrates and metals and their alloys [19]. The advantage of inorganic PCM is that they have greater phase change enthalpy than organic PCM. However, disadvantages include corrosiveness, supercooling and thermal instability [97]. The thermal properties of a large number of different materials, both organic and inorganic, have been studied for use as PCM. Many literature reviews have been conducted and provide listings of these materials, e.g. [19, 81, 97].

A lot of research has been done on the integration of PCM in buildings, such as in the walls and floor [51]. The benefits of this are reduced energy consumption and smoothing out the temperature fluctuations inside the building, leading to higher thermal comfort. The amount of research done on seasonal latent thermal storage is quite limited thus far. One example of seasonal latent thermal energy storage is presented in [69], where an experimental greenhouse is heated using flat plate solar air collectors and a latent thermal storage. For this system an average net energy efficiency of 40.4% was obtained. This is similar to the efficiency of 36% reported for the borehole thermal storage working in the Drake Solar Landing Community [83] (see Section 2.1.1).

Currently, latent thermal energy storage systems are mostly still in the research and development phase [77]. Latent thermal energy storage shows promise, especially in low temperature human comfort applications, such as the earlier mentioned PCM integrated in the building envelope [51]. However, the currently preferred technology for large scale systems, such as seasonal thermal storage, is sensible thermal storage [19].

2.1.3 Chemical thermal storage

In chemical thermal storage thermal energy is stored by means of chemical processes. The storage is charged by means of a endothermic process, which can be reversed to discharge the storage later. An example is a chemical reaction, where a material A is separated into two materials B and C when heat is applied to A. When B and C are mixed again at suitable conditions, energy in the form of heat is released [77, 95]. In other words, the (reversible) reaction is given by:

$$A + \text{heat} \longleftrightarrow B + C.$$
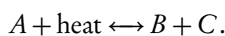
A specific example given by [95] is the reversible dissociation of ammonia:

$$2NH_3 + \text{heat} \longleftrightarrow N_2 + 3H_2.$$
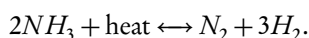
Chemical thermal energy storage is divided into two categories, one using chemical reactions (such as the example shown above) and the other using sorption

processes [95]. In sorption processes a gas is absorbed by a liquid or solid, creating a new compound (absorption), or bonds to the surface of a solid without creating a new material (adsorption) [93, 95].

One of the major advantages of chemical thermal storage is that there are no significant long-term heat losses during storage, which is a large advantage for seasonal thermal storage [73]. These low heat losses are due to the fact that as long as the reaction products are kept separated heat losses are restricted to sensible heat losses, i.e. the materials cooling down, which are usually much smaller than the amount of heat produced in the reverse reaction during discharge [77, 93]. Another large advantage of chemical thermal energy storage, compared to sensible and latent thermal energy storage, is the much larger energy densities that can be obtained in chemical thermal energy storage. To illustrate this, in [44] a comparison of the size of the storage needed to store 10 GJ of energy using different storage techniques is presented. For a storage using chemical reactions a volume of 1 m^3 is needed, for sorption chemical storage the volume required is 10 m^3, for latent thermal storage the volume grows to 20 m^3 and for sensible storage, using water with a temperature range of 70 °C, a volume of 34 m^3 is required. Disadvantages of chemical heat storage include complex reactions, low efficiency and high investment costs [95]. To be able to alleviate these disadvantages a lot of effort has been put into researching different materials, conducting numerical studies and work on reactor designs. For summaries on this we refer to [50, 95].

From the mentioned advantages of chemical thermal storage (i.e. high energy density and low long-term heat losses) it is clear that it has large potential. However, currently chemical thermal storage is in the laboratory stage and further research is required to make it commercially viable [19, 50].

2.2   WATER TANKS

In this section we provide a more in depth discussion on water tanks, since the Ecovat buffer resembles other water tank designs closely. Furthermore, we discuss the similarities and differences between the Ecovat buffer and other water tank designs.

As mentioned in Section 2.1.1 water tanks are man-made structures to store thermal energy, which is done by increasing the temperature of the storage medium, in this case water. Discharging of a water tank is done by taking hot water out of the tank, while cold water is added back to the tank. For charging the water tank different configurations exist, divided into direct and indirect heat transfer systems. In direct heat transfer water tanks hot water from a collector, which can be any device supplying hot water, is directly pumped into the water tank without any heat exchanger (top left in Figure 2.1). In indirect heat transfer systems the charging is done through a heat exchanger. The three most common configurations for these systems are [45]:
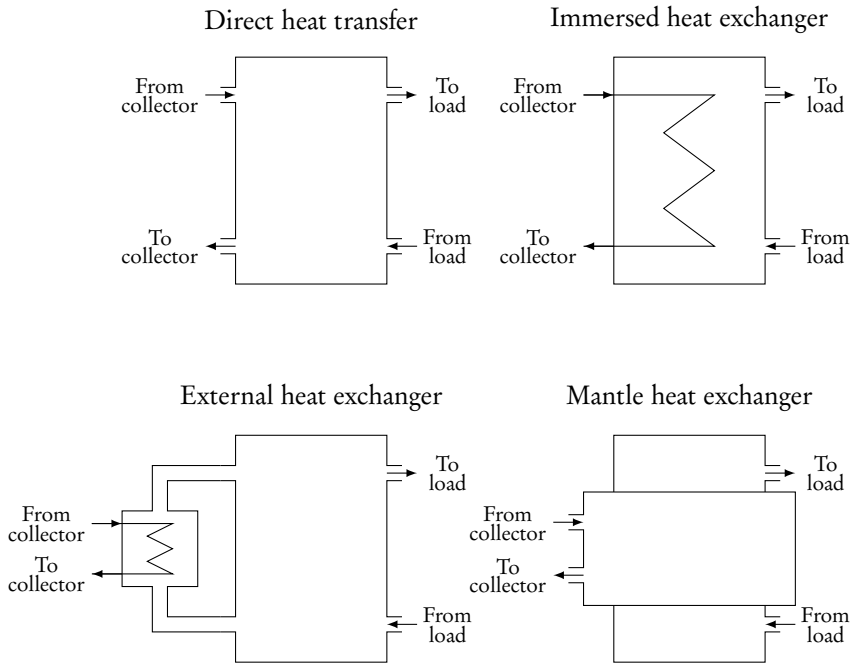
Figure 2.1: Common water tank configurations.

- » *Immersed heat exchanger*: in the immersed heat exchanger configuration the heat exchanger is placed inside the water tank itself (top right in Figure 2.1).

- » *External heat exchanger*: in the external heat exchanger configuration an external heat exchanger heats the water that is being pumped into the water tank (bottom left in Figure 2.1).

- » *Mantle heat exchanger*: in the mantle heat exchanger configuration a large part of the tank walls is covered by a so called mantle, which can be seen as a second tank surrounding a part of the first tank. This mantle functions as a heat exchanger (bottom right of Figure 2.1).

The charging of the Ecovat buffer is somewhat similar to that of a water tank using a mantle heat exchanger. However, instead of a single mantle for charging the entire water tank each segment can be charged individually through the heat exchangers in the Ecovat buffer walls. Furthermore, the heat exchangers in the Ecovat buffer work differently than a mantle heat exchanger. The heat exchange in the Ecovat buffer happens through pipes integrated into the buffer walls instead of a mantle. The water in the mantle shows recirculating flow [80], while this does not happen in the heat exchangers of the Ecovat buffer, since

those are simply pipes through which water is pumped. Another difference is in the discharging of the Ecovat buffer compared to other water tanks. In the Ecovat buffer the discharging also happens through the same heat exchangers while in most other water tanks this is done by pumping water into and out of the water tank.

2.2.1   THERMAL STRATIFICATION

Thermal stratification is the phenomenon where the water inside the tank, due to a difference in density is separated into two regions, a hot, less dense, water region at the top of the tank and a cold, more dense, water region at the bottom of the tank. The degree to which a water tank is thermally stratified ranges from no stratification, i.e. the entire tank is at more or less the same temperature, up to perfect stratification where the two regions are completely separated with no transition area between them. However, in practice a thermally stratified water tank will not show perfect stratification. Instead, there will be a generally thin transition region between the hot and cold regions with a steep temperature gradient from the temperature of the hot region to the temperature of the cold region. Opposed to most water tanks, which only have 2 different temperature regions, the Ecovat buffer has multiple segments at different temperatures. This implies that the Ecovat buffer has a transition region between each two neighbouring segments.

An increased degree of thermal stratification generally leads to a more efficient system [62, 76]. The four primary reasons for a decrease in thermal stratification according to [76] are; 1) heat losses through the tank walls/ceiling/floor to the surrounding environment, 2) heat conduction from the hot region of the tank to the cold region of the tank, 3) vertical conduction through the walls of the tank and 4) mixing inside the water tank during the charging and discharging processes. Of these four the last is in general the most important reason for a decrease in the thermal stratification. Furthermore, [76] notes that significant mixing losses may occur if the heat is stored over long periods of time.

The effect of design parameters of water tanks on the thermal stratification has been studied extensively. For the position of the inlet, i.e. the point where cold water enters the water tank, the best position was found to be as close as possible to the bottom of the tank [62]. Furthermore, it was found that a low flow rate is beneficial to the degree of stratification in the water tank [52, 62, 94, 101]. Systems with multiple or flexible inlets were found to lead to a better stratification as well. For example [62], shows that a copper tube with holes connected to the inlet inside the water improves thermal stratification by depositing the incoming water at the vertical position inside the tank with a similar temperature, while [40] indicates that the use of a flexible tube connected to the inlet of the water tank shows the same improvement. Similarly, [42] shows that having two outlets to draw water from during discharge improves thermal stratification, leading to better performance. The amount of mixing is suppressed

if the difference in temperature between the hot water in the water tank and the cold water injected in the inlet is large [52, 62, 94, 101]. Finally, a higher aspect ratio (heigth/diameter) of the water tank is found to increase the thermal stratification [52, 62]. The above discussed design parameters are relevant for all the water tank configurations shown in Figure 2.1, where discharging is done by drawing hot water from the top of the water tank while inserting cold water is done at the bottom of the tank. Since mixing in the water tank is the most important reason for a decrease in thermal stratification [76] it makes sense that design choices which limit the amount of mixing in the tank lead to better thermal stratification.

For the charging of a system using a mantle heat exchanger, which closest resembles the way the Ecovat buffer is charged, similar research on thermal stratification inside the water tank and the mantle itself has been done. One of the conclusions from [58] is that a mantle heat exchanger promotes thermal stratification inside the water tank even if the temperature at the inlet of the mantle is lower than the temperature of the inner tank at the same height. In [57] different positions of the inlet of a mantle heat exchanger are researched, showing that for a high temperature at the inlet of the mantle the inlet should be placed at the top of the mantle, while for a lower temperature a low position of the inlet improves performance. Similar to the discussion above, it is observed that minimizing the amount of mixing in the water tank in systems using a mantle heat exchanger leads to better thermal stratification.

In the Ecovat buffer the amount of mixing is minimized by never pumping water into or out of the buffer. The charging and discharging of the buffer is done solely through the heat exchangers integrated into the buffer walls. During charging and discharging there will be some flow in the buffer due to temperature differences at the wall compared to the center of the buffer. The expectation is that this flow will be restricted to the segment that is being charged/discharged at that point in time as long as the temperature difference between the water in the heat exchanger and the water in the buffer segment is not too large. As such the thermal stratification between segments would not be disturbed. This expectation is confirmed in a preliminary two-dimensional numerical study [38]. However, a verification of this behaviour using experimental data obtained from an Ecovat in practice is not available yet, since such data was not available at the time of writing.

2.2.2 Modelling of water tanks

Many different models for water tanks exist, ranging from very simple one-dimensional models to complex three-dimensional computational fluid dynamics (CFD) models. In general, complex multi-dimensional models describe the thermal processes inside the water tank better. However, this comes at the cost of a higher computational effort. In this section we present an overview of some of these models.
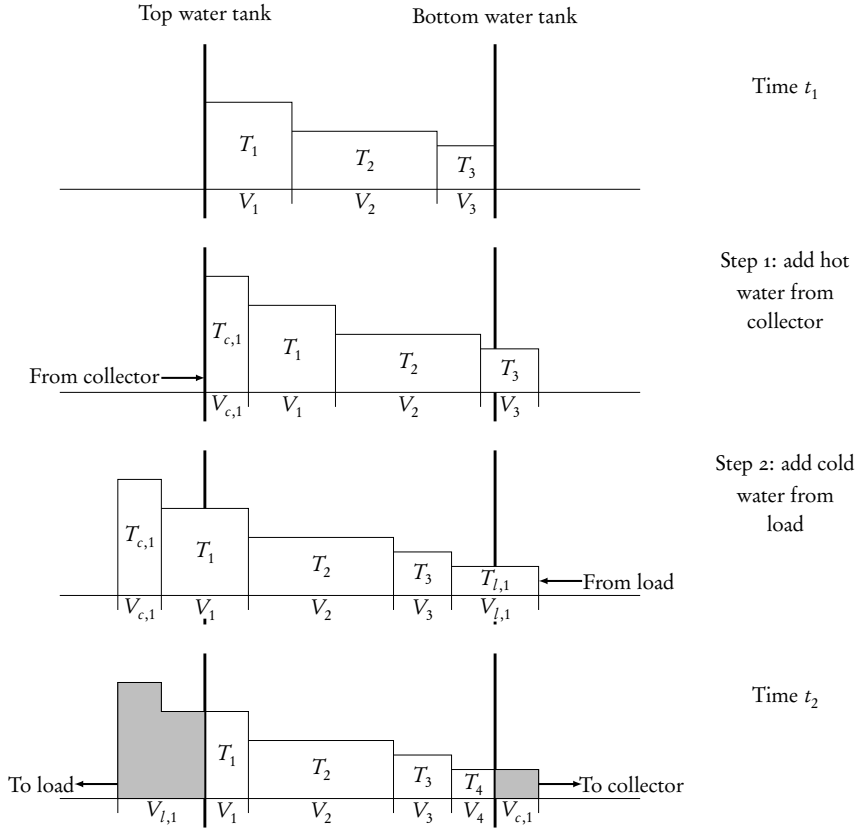
Figure 2.2: Operation of a plug flow model.

Two one-dimensional modelling approaches are presented and compared in [56]. The first is a multinode approach, in which the buffer tank is divided into a number of nodes, with an energy balance equation describing the energy, or equivalently the temperature, of each node. The second approach is a plug flow model, which uses a varying number of segments of varying sizes, with different temperatures. The operation of such a model is depicted in Figure 2.2. In this model for every time step a volume $V_{c,t}$ of hot water, with temperature $T_{c,t}$, coming from the collector is added to the top of the tank (step 1 in Figure 2.2). If $T_{c,t}$ is larger than the highest temperature segment T_1 in the water tank, a new segment is created. If $T_{c,t} \leqslant T_1$, $V_{c,t}$ and V_1 become one segment with an averaged temperature. Simultaneously a volume of cold water coming from load is added to the bottom of the tank in a similar way (step 2 in Figure 2.2). These incoming volumes shift the temperature profile in the tank up or down, depending on the relative sizes of the volumes. The parts of the profile that 'fall outside' the boundaries of the water tank are returned to the collector ($V_{c,1}$ of

cold water) and load ($V_{l,1}$ of hot water). When comparing the approaches the multinode model performs better in general than the plug flow model [56].

In [68] the authors compare a one-dimensional and a two-dimensional model for a water tank. The one-dimensional model is again a multinode model, in this case combined with a computational artifact to prevent temperature inversions (i.e. a colder node on top of a hot node) at the top of the tank. This can be done by either swapping the two nodes involved in such a temperature inversion, or by assigning them both the mean temperature of the two nodes. The two-dimensional model uses the finite volume method [91] in cylindrical coordinates. The two-dimensional model is validated by experiments and used as a benchmark for comparison with the one-dimensional multinode model. Even though there is no physical base for using one of the computational artifacts mentioned earlier, [68] finds that the one-dimensional multinode model incorporating such an arte-fact show results close to those obtained with the two-dimensional model. The authors conclude that for long-term simulation of solar water heating systems a one-dimensional model is sufficient, while a two-dimensional model is useful to get an understanding of the thermal phenomena taking place in water tanks.

In [94] a three-dimensional CFD model of a water tank is presented. The authors validate their model based on experimental data from literature. The developed CFD model is then used to study the thermal processes in a water tank for differ-ent tank parameters. The authors conclude that the developed CFD model is a useful tool to support the design of efficient water tanks. The trade off between complexer models requiring higher computational times is nicely illustrated, as the authors report simulation times of 9 to 24 hours for cases simulating 1000 to 3700 seconds of water tank operation, making it infeasible for simulations over very long time horizons. Another example of a CFD model, this time for a water tank with a mantle heat exchanger, is presented in [79]. The predictions made by this CFD model are compared with experimental data, showing good agreement. A simpler model for a water tank with mantle heat exchanger is presented in [53]. The goal of [53] is to develop a model with shorter simulation times than CFD models, such that it is feasible to simulate long time periods. The authors use a zonal model for a water tank with a mantle heat exchanger, which divides the simulated volume into a number of zones of constant tempera-ture. The results show a difference between experiments and the zonal model of less than 7%. The authors conclude that the model is useful for long simulations, where CFD models may be computationally too expensive.

To conclude, we see that to specify the thermal processes in water tanks com-plex models, such as CFD models, are very useful. However, in cases where the exact thermal behaviour inside the water tank is not as important simpler models are often sufficient. This is especially true when dealing with simulations over long time periods (e.g. simulation of an entire year for seasonal thermal storage), where for example CFD models may be infeasible due to their high computational times.

2.3   District heating systems

District energy systems consist of a system of pipes connecting groups of houses with the goal of supplying those houses with thermal energy. In general these systems are used for heating, cooling or both. Here, we focus exclusively on such systems used for heating, i.e. district heating systems. There are many advantages of a district heating system compared to heating through for example individual gas boilers (the currently most common way of space heating in the Netherlands), which include [75]:

» A reduced consumption of fuel, which leads to lower costs and emissions.

» Ease of using local energy sources (e.g. waste heat from industry).

» Safer operation.

» Increased reliability.

Disadvantages include [75]:

» Large up front investment costs to realize a district heating system.

» Requiring a heat source close by.

» Requiring sufficient consumers nearby.

Based on this, district heating may not be economically feasible everywhere [60, 75]. Nevertheless, district heating networks are expected to be an important part of the future energy system in Europe [70].

A Dutch policy study on district heating systems [46] arrives at the same conclusion, namely that district heating has large potential to be an important piece of the Dutch energy system of the future. A report from 2009 [78] shows only a small amount of district heating systems in the Netherlands at that time. However, recently the interest in district heating systems has increased. For example the municipality of Enschede (home of the University of Twente) is extending its district heating system with the goal of supplying at least 50% of the city's heat load [14]. Furthermore, two older district heating systems which previously used a combined heat and power (CHP) unit and gas boilers as their heat sources will use a more sustainable energy source in the future. This energy source is the Twence waste incineration plant, which will be extended to include a biomass plant (biomass from green waste material) to supply the extended district heating network in Enschede.

The history of district heating systems is generally divided in several generations of such systems. The current district heating systems fall into the third generation of district heating systems. In [66] an overview of the first three generations of district heating systems is provided as well as an outlook to the future fourth generation of district heating systems, both of which we summarize below. The first generation of district heating systems, established between 1880 and 1930,

used steam as their heat carrier. Two major disadvantages from using steam as heat carrier are the high thermal losses as well as the danger of steam explosions, which have lead to a fatal incident as recent as 2007 in New York [23], whose district heating system still uses steam as its heat carrier. The second generation (1930-1970) switched to pressurized water, with temperatures generally above 100 °C, as the heat carrier. Even though this eliminates the danger associated with using steam as a heat carrier, the heat losses are still relatively high due to the high temperature of the water. The third generation, which includes all recently developed district heating systems, still uses pressurized water as the heat carrier, but generally at temperatures below 100 °C. This leads to a decrease in the heat losses. In the future fourth generation of district heating systems the trend of lowering the temperature of the water in the systems in an effort to reduce heat losses is expected to continue. The water temperatures in the fourth generation of district heating systems is expected to be between 30 and 70°C [66]. One issue with such very low temperatures in a district heating system is that an additional device is needed within individual households to increase the water temperature for use as domestic hot water to avoid the risk of legionella contamination, which may offset the economic benefit gained from the reduced heat losses in the system. However, in [24] an economic analysis is presented showing that (ultra-)low temperature (40 °C) district heating systems are economically feasible at least in Germany.

The trend towards lower temperatures in district heating systems allows them to be combined more effectively with low temperature thermal storage technologies, such as the Ecovat system. Assuming the output temperature of the thermal energy storage (TES) is not increased before being delivered to the district heating system, the effective capacity of the storage is given by:

$$C_{eff} = m_{st} c_p (T_{max,st} - T_{dhs}), \tag{2.1}$$

where, C_{eff} is the effective capacity of the energy storage, m_{st} is the mass of the storage medium, c_p is the specific heat coefficient of the storage medium, $T_{max,st}$ is the maximum allowed temperature in the storage and T_{dhs} is the temperature of the district heating system. The intuition behind the definition of the effective capacity as given by Equation 2.1 is that without a device to increase the output temperature of the storage technology any thermal energy in the storage at a temperature lower than the demanded temperature can not be used to satisfy this energy demand. When a TES, such as the Ecovat system, is connected to a district heating system, a lower temperature in the district heating system leads to a larger effective capacity of the thermal storage, which makes it more effective in shifting energy from times of surplus production to times of surplus demand. To achieve a better performance of the TES a smart control strategy is required, as demonstrated for example in [90], where a model predictive control approach based on a mixed-integer linear programming model is used to show that both

in simulations and experiments a decrease in the costs of supplying energy with a district heating system is achieved.

2.4 DEMAND SIDE MANAGEMENT

As mentioned in Chapter 1 the massive introduction of renewable energy sources in the energy system leads to an increase in the difficulty of matching the supply and demand of energy. Energy storage is one of the technologies that supports the continuous matching of energy supply and demand, by storing energy during times of excess supply and supplying energy during times of excess demand. Another means of maintaining the balance between supply and demand is demand side management (DSM). DSM is defined as any measure taken on the demand side of the energy system to increase the performance of the energy system [71]. Some examples of DSM include load shedding or shifting the energy demand of certain devices, like electric vehicles, to times with high energy supply or less total energy consumption. Since DSM is not the main focus of this thesis we only refer to other work containing extensive discussions on DSM. For an overview of DSM methods we refer to e.g. [71]. DSM in a broader context, including for example monitoring and communication systems for DSM, is surveyed in [82].

Traditionally, DSM is often used in the development of smart grids. The definition of a smart grid by the International Energy Agency [1] is:

"A smart grid is an electricity network that uses digital and other advanced technologies to monitor and manage the transport of electricity from all generation sources to meet the varying electricity demands of end-users. Smart grids co-ordinate the needs and capabilities of all generators, grid operators, end-users and electricity market stakeholders to operate all parts of the system as efficiently as possible, minimising costs and environmental impacts while maximising system reliability, resilience and stability."

As we can see, this definition strictly concerns the electricity grid. However, recently an increased interest in so called smart energy systems is observed [67]. The definition of a smart energy system is given by [64] as:

"A smart energy system is defined as an approach in which smart electricity, thermal and gas grids are combined with storage technologies and coordinated to identify synergies between them in order to achieve an optimal solution for each individual sector as well as for the overall energy system."

The reason to prefer a smart energy system over an approach just focussing on one energy carrier is that a strong reduction in costs may be achieved by using a smart energy system approach. An example of this is shown in [65], where the author shows that for transforming the heating sector of Denmark

to a 100% renewable energy solution a smart energy system approach leads to significantly lower investment costs than a smart grid approach just focussed on electricity. One of the reasons for this is that, at least currently, thermal energy storage is much cheaper than electricity storage. As such it makes sense to not only consider electricity, but also include other energy carriers, such as thermal energy, on a smaller scale, such as a microgrid of a neighbourhood, as well.

Since we are focussing on thermal energy storage in this thesis, we limit the rest of our discussion to DSM that includes thermal energy storage. Many of the thermal storage technologies discussed in Section 2.1 are potential candidates for use in DSM systems [22]. We conclude this section by discussing some examples of DSM research incorporating thermal energy storage.

In [59] a microgrid consisting of a single house containing a number of household devices and both electrical and thermal storage is simulated. The microgrid is modelled as a mixed integer linear programming model. For the thermal storage a very simple model is used that just considers the energy content of the storage at every time step. The conclusion is that by using DSM and storage a reduction of costs between 4.7 and 7.6% can be achieved for the simulated case. In [98] a stochastic non-convex optimization model is presented for the optimization of a microgrid. This model is then simplified to a linear programming model to be able to solve it. In this model, the thermal energy storage is again modelled by just considering the energy content of the storage at every time step. The model is applied to historical data from a college in San Francisco and a 27.5% reduction in costs is reported when comparing this model with a greedy approach. Three buildings, each with their own microgrid are considered in [49]. A comparison is made between a case where each microgrid operates on its own and a case where an energy management system is used to combine the operation of the three microgrids. In the second case a larger network in which the three microgrids are embedded is considered. The larger network includes thermal storage, which is modelled by just considering the energy content of the storage. Using the energy management system, including the thermal storage, leads to a reduction in costs of 45.5% during summer and 38.5% during winter. In [16] the benefit of adding thermal storage to a microgrid consisting of a building of 12 apartments and a restaurant was investigated. The authors found that replacing part of the battery storage with thermal storage, while keeping the total capacity equal, lead to a reduction in costs of 8%. Similar to the previously discussed cases, the thermal storage is simply modelled by just the energy content of the storage.

From the above mentioned examples we see that using an extremely simplified model of the thermal storage is often used in demand side management research that incorporates thermal storage. In this model the thermal storage is simply described by a single parameter, which gives the energy content of the storage. Research on demand side management that includes a thermal storage technology is of course not limited to the examples discussed above. However, from a survey of the literature it seems that little to none of this research uses a more complex

model for the water tank thermal storage.

We end this section with a few remarks on decentralized energy management (DEM). As the name implies, DEM focusses on decentralized energy management, as opposed to the traditional centralized approach when large scale power plants were used for the energy generation. Such a decentralized approach is becoming more important, due to the large scale introduction of renewable energy sources. DSM and DEM are in some ways similar, however, where DSM focusses solely on the demand side, DEM considers both the demand and supply side of the energy system. As such, for the case study discussed in Chapter 6 we feel it is more appropriate to use the term DEM instead of DSM, since in that case study a neighbourhood is considered that also contains energy generation.

2.5 CONCLUSION

In this chapter we provided an overview of thermal storage technologies, as well as some background on district heating systems and demand side management. From our discussion on water tanks it can be concluded that to determine the exact thermal processes in a water tank complex models, such as CFD models are beneficial or even necessary. However, for cases where the exact thermal behaviour of the water tank is less important, simpler models can often be sufficient. This was confirmed when discussing research combining demand side management methods and thermal storage, in which generally the simplest model possible, namely just a single parameter describing the energy content of the water tank, is used to model the thermal storage. In this type of research the exact thermal behaviour inside the water tank is not as important, while using a simpler model for the thermal storage reduces the complexity of modelling the entire system (of which the thermal storage is just a part), which in turn reduces the simulation times of such systems.

3

MODELLING OF THE ECOVAT SYSTEM AND ITS CONTROL

ABSTRACT – *In this chapter we present an integer linear programming (ILP) model for the operation and control of the Ecovat system. The decision variables in this model are the possible control actions to charge/discharge the Ecovat system. The restrictions and components of the Ecovat system are discussed and translated into a set of constraints on the possible decisions of the control of the Ecovat system. The objective of controlling the Ecovat system, and therefore the ILP model, is to satisfy the heat demand of the neighbourhood connected to the Ecovat system against minimal operational costs of the system. The output of the ILP model consists of the charging/discharging strategy for the Ecovat system as well as the temperature evolution within the Ecovat buffer over the time horizon.*

In Chapter 1 we introduced the Ecovat system, which is a seasonal thermal storage technology designed to supply the heat demand of a neighbourhood of houses. To operate such a storage system, a control strategy is needed which decides at which time how much energy should be stored or supplied by the system. This control strategy has to ensure that all the heat demand of the neighbourhood is supplied throughout the year at minimum operational cost. To determine such a control strategy asks for a proper model of the Ecovat system.

To develop a control strategy for an Ecovat system that is usable in a practical setting, the control strategy is required to make very fast (within seconds) decisions. Due to this requirement such a control strategy will most likely not be

---

This chapter is based on [G:1].

able to determine the optimal charging/discharging strategy for a given input. However, to be able to judge the quality of such a practical strategy, it is desirable to have a benchmark model, which gives a (near) optimal strategy, to compare the results of the practical control strategy against. This benchmark model does not have requirements on the computational time since it will not be used for the operational control of an Ecovat system.

In this chapter we describe such a benchmark model. It is based on an ILP model of the Ecovat system, which models the Ecovat system as well as its control. This model is used as a benchmark to determine the performance of a simpler model, which is usable to control an Ecovat system in a practical setting, i.e. in real time.

The remainder of this chapter is structured as follows: in Section 3.1 the physical model used for the Ecovat buffer is discussed. In Section 3.2 the ILP model of the Ecovat system is described. Finally, in Section 3.3 a summary of the model is given and conclusions are presented.

3.1 Physical model

The model to determine the charging/discharging strategy of the Ecovat system should be based on a physical model that describes the temperature changes inside the Ecovat buffer. If the goals is to accurately describe the temperature at every point in the buffer a proper approach would be based on solving the Navier-Stokes equations [30]. The solution of these equations describes the motion of the fluid, which can in turn be used to obtain the temperature distribution inside the buffer. Such a computational fluid dynamics (CFD) model [30] would in theory give a very accurate description of the temperature evolution inside the buffer. However, for the following reasons it is impractical to use this method for a model of the Ecovat system. The combination of buffer size and time scales considered (we would like to determine a charging/discharging strategy for an entire year) makes such a model computationally too expensive, especially considering the fact that we want to combine this physical model of the Ecovat buffer with an optimization model to determine a good charging/discharging strategy for the Ecovat system. Such a combination of a CFD model with an optimization model would in theory be preferable, but is infeasible in practice. To illustrate this: in Chapter 2 we discussed the long simulation times associated with CFD modelling, e.g. the CFD model for a water tank introduced in [94] requires 9 to 24 hours of simulation time to simulate 1000 to 3700 seconds of water tank operation. From this it becomes clear that using a CFD model to simulate a year of operation of the Ecovat system is impossible, especially considering that for our purpose the CFD model would also need to be combined with an optimization model leading to even longer simulation times.

Aside from the fact that a CFD model is computationally too expensive, it is likely that for the required accuracy in the temperature distribution inside the

buffer needed for our optimization, a less accurate model is also sufficient. Since we are interested in the operational control of the system, the decrease in accuracy when compared to a more complex model is not a problem as long as this decrease does not significantly influence the decisions made in the optimization. To clarify this: as long as the decisions made by the optimization model do not significantly differ for two different, but similar, temperature distributions (coming from two different models), the less accurate model still fits our purpose. Furthermore, given that some of the other inputs of the optimization model are also not known exactly, but instead are based on predictions, a simple approximation of the temperature distribution inside the buffer should be sufficient.

Based on the above considerations, we decided to use a very simple physical model for the temperature distribution inside the buffer in which every segment is assumed to have a uniform temperature. Obviously this is a simplification of reality, however, the resulting model is computationally very inexpensive compared to more complex models. Furthermore, as we have discussed in Chapter 2, simple models are often used, especially in DSM related research, because they are considered sufficient if the exact thermal processes inside the water tank are not too important. When comparing a simple model with a more complex model of the Ecovat buffer, there may be some specific cases, for example when two buffer segments have temperatures very close to each other, in which a difference in temperature distributions between these models may lead to different decisions being made in the optimization model. However, we estimate this effect to be small. In other words, in the majority of cases we expect the decisions resulting from the optimization to be indifferent to small differences in temperature distributions between the different models. Furthermore, we expect that whether a single average temperature is taken for a segment, or whether a temperature distribution with a temperature gradient throughout the buffer segment is taken (with for example slightly higher temperature at the walls than in the center of the segment after charging), has a low impact on the decisions made by the optimization model.

Additionally, we expect that even though within a segment there might be slight temperature variations due to the charging/discharging of this segment, that there is little to no water flow, and thus little to no temperature mixing, between different segments due to low amounts of turbulence inside the buffer. The reason for this is that the Ecovat buffer is thermally stratified and by using a smart way of charging/discharging the buffer we can ensure that the temperature remains decreasing from the top to bottom segment of the buffer preventing segments from mixing. Ideally, this assumption, as well as further assumptions made later on in this thesis, should be confirmed by measurement data from an actual Ecovat buffer. However, at the time of writing such data is not yet available.

As the buffer segments are heated from the buffer walls, this will induce some circulation within a segment. However, assuming the temperature difference

between the water in the heat exchangers and the water in the buffer is not too large, the water circulation is expected to not disturb other segments in the buffer. As mentioned in Chapter 2, this behaviour was confirmed in a preliminary two-dimensional numerical study [38]. Even though within the optimization model we do not specifically restrict the temperature difference between the water in the heat exchangers and the water inside the buffer, the water flow through the heat exchangers can be varied in practice, meaning that a low temperature difference can be realised.

For the reasons given above we expect to only have some minor water flow within a given segment, but not between different segments. As mentioned before, we furthermore expect that the small temperature differences within a given segment, induced by such flows, only have a minor impact on the decisions made by the optimization model. This expectation is based on the fact that the decisions made by the optimization model concern the energy content of the different buffer segments, which is important for satisfying the heat demand of the neighbourhood throughout the year. As such, we expect that slight temperature differences within a buffer segment do not lead to significantly different decisions, compared to taking an average temperature over the entire buffer segment, as both cases result in the same energy content.

One concern with using models, simple or complex, over longer time horizons in a practical situation is that over time errors in the temperatures (when comparing models with reality) may add up and as such the impact on the decisions of the optimization model may start to become larger over time. However, during practical operation of an Ecovat system this can be solved easily by a few temperature sensors situated inside the buffer or the buffer walls. These sensors can then be used to calibrate the temperatures in the model to the actual temperatures in the buffer, for example at the end of every day. In this way the temperatures of the model and the actual buffer should always be similar. This recalibration of the model temperatures would also significantly reduce the effect of using a simplified physical model instead of a complex model.

Due to the reasons outlined in this section, we expect that a simple physical model is sufficient for our purpose. To be able to directly quantify the effect that using a simple physical model instead of a complex model has on the decisions being made, one would have to use different physical models in combination with the optimization model described in the remainder of this chapter. The results of this could then be used to compare how often the optimization model makes the same decisions versus how often the decisions differ when using a different physical model. However, this research falls outside the scope of this thesis.
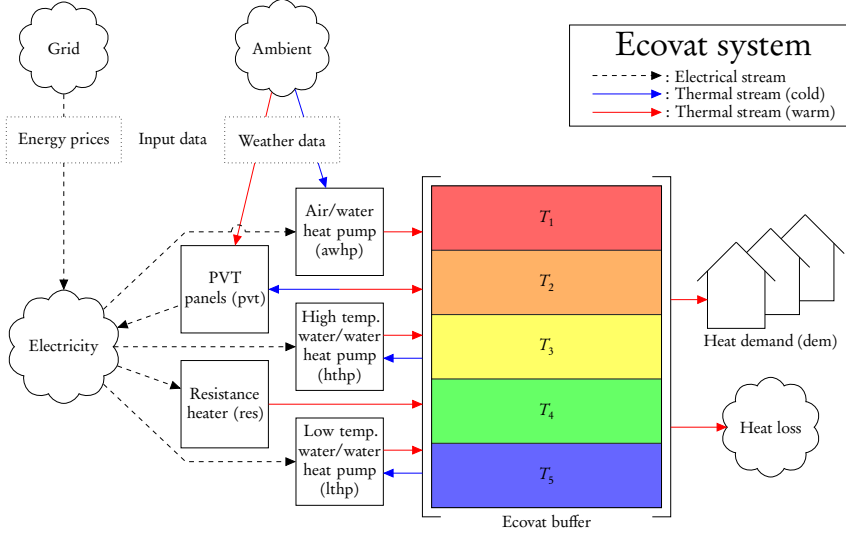
Figure 3.1: Schematic overview of the Ecovat system.

3.2 INTEGER LINEAR PROGRAMMING MODEL

In Section 1.2 we gave a description of the Ecovat system consisting of the Ecovat buffer and a set of devices charging this buffer. In this section we introduce an ILP model to optimize the control of such an Ecovat system. A schematic overview of the different components of this ILP model is depicted in Figure 3.1, which includes the Ecovat buffer and the set of devices that can be used to charge the buffer as described in Section 1.2.

The presented model is based on a discretization of the overall time horizon, which in our case is a complete year, into N_{int} time intervals of equal length. We define the set of time intervals in the time horizon as $\mathcal{I} = \{1, ..., N_{int} + 1\}$. The last time interval $N_{int} + 1$ is used to characterize the state of the system at the end of the time horizon [1].

The input of the developed ILP model consists of a set of vectors containing the energy prices $\mathbf{p} = (p_1, p_2, ..., p_{N_{int}})$, the ambient temperature $\mathbf{T}^{amb} = (T_1^{amb}, T_2^{amb}, ..., T_{N_{int}}^{amb})$, the global irradiation $\mathbf{G} = (G_1, G_2, ..., G_{N_{int}})$ and heat demand of the neighbourhood $\mathbf{D} = (D_1, D_2, ..., D_{N_{int}})$ for every time interval in the time horizon. For the developed ILP model we assume that predictions for this input data are available for the entire time horizon on beforehand. Later

---

[1]Although the presented model can be used with any time interval length, all simulations in this thesis are performed using 15 minute time intervals

on, when developing a simpler heuristic model we will consider a more realistic case in which such predictions are not available on beforehand.

The output of the ILP model consists firstly of a charging/discharging strategy of the Ecovat system over the time horizon, which is given by the decision variables of the model. These variables indicate whether the devices in the system are used during a given interval, and if so in which way they are used. Note, that we assume the state of a device (e.g. whether it is turned on or not) does not change within a time interval. Secondly, the output of the model includes the temperature evolution inside the Ecovat buffer over the time horizon. The temperatures of the buffer segments are calculated using the values for the input parameters and the values of the decision variables.

The Ecovat buffer itself is modelled as a set of segments $S = \{1, ..., N_{seg}\}$, where N_{seg} is the number of segments in the Ecovat buffer. The segments of the buffer are numbered from top to bottom as shown in Figure 3.1. In general any number of segments can be accommodated by the model, however, in this thesis we only consider the case $N_{seg} = 5$ since this is the number of segments the built proto-type of the Ecovat buffer has, as well as the number of segments the currently proposed versions of the Ecovat buffer will have. The temperature of segment s at the start of time interval i is denoted by $T_{i,s}$, where we assume that a given segment has a uniform temperature throughout the entire segment, as described in Section 3.1. Finally, we define the set of devices used for charging/discharging the buffer as $D = \{pvt, awhp, lthp, hthp, res, dem\}$, where the labels designate the different devices in the Ecovat system as shown in Figure 3.1. Note, that we also include the heat demand in the device set D.

3.2.1 Decision variables

We model all the devices in the Ecovat system to be on/off devices. This implies that when a device is turned on it will always run at its full capacity. This means that all the decision variables in the ILP model, which designate the state of a given device, can be modelled as binary variables. All decision variables are written in the following form: $x_{i,s}^{dev}$, which designates whether device dev is connected to buffer segment s during time interval i. If $x_{i,s}^{dev} = 1$ then device dev is connected to buffer segment s during time interval i, while $x_{i,s}^{dev} = 0$ means device dev is not connected to buffer segment s during time interval i. If for a time interval i^* all the decision variables for a given device dev^* are equal to zero, i.e. $\sum_{s \in S} x_{i^*,s}^{dev^*} = 0$, device dev^* is turned off during that time interval.

Additionally, the design of the Ecovat system puts some constraints on the use of the devices in the system. First, for any given time interval any device may only be connected to at most one segment, i.e:

$$\sum_{s \in \mathcal{S}} x_{i,s}^{dev} \leqslant 1 \quad \forall i \in \mathcal{I},\, dev \in \mathcal{D}. \tag{3.1}$$

Second, at most one of the devices may be connected to a given segment during any time interval:

$$\sum_{dev \in \mathcal{D}} x_{i,s}^{dev} \leqslant 1 \quad \forall i \in \mathcal{I},\, s \in \mathcal{S}. \tag{3.2}$$

All further, device specific, constraints are discussed in the corresponding subsections, dealing with the specific devices.

3.2.2 PVT PANELS

Photovoltaic thermal (PVT) panels generate both electrical and thermal energy from incident solar irradiation. The efficiency of such panels depends, among others, on the temperature of the panels, with a lower temperature leading to higher efficiency [100]. Therefore, the PVT panels connected to an Ecovat system are cooled to increase their efficiency. This is done by connecting their thermal output to the bottom segment of the Ecovat buffer. The choice for this segment is natural since it has the lowest temperature. In this way the PVT panels charge the bottom segment of the buffer, while simultaneously being cooled by that segment. The maximum allowed temperature for the bottom segment of the buffer has to be kept low, to be able to increase the efficiency of the PVT panels in this way.

Whether the PVT panels are connected during a given time interval to the bottom segment of the Ecovat buffer depends on their output temperature. The output temperature of the PVT panels needs to be higher than the temperature of the bottom buffer segment, $T_{i,N_{seg}}$, to be able to charge the buffer. If the output temperature of the PVT panels is lower than $T_{i,N_{seg}}$ the PVT panels are not allowed to be connected to the buffer. However, even if the output temperature is high enough the model should have the option to not connect the PVT panels if that would lead to a better overall result. The reason for this overall better performance could be that connecting another device to the bottom buffer segment at that time interval leads to a better result, or that the connecting of the PVT panels to the buffer in that time interval would restrict the options of the system in future time intervals. To specify whether the PVT panels are connected a binary decision variable w_i is introduced that designates whether the output temperature of the PVT panels is sufficiently high to charge the buffer (if $w_i = 1$ the output temperature is sufficiently high, if $w_i = 0$ it is not). The correct value of this variable is set by the following two constraints:
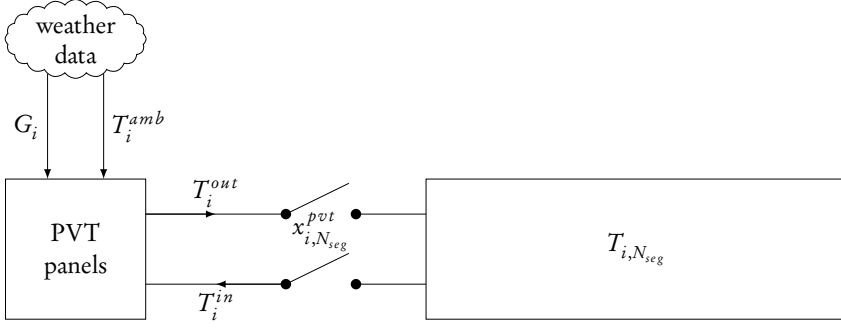
Figure 3.2: Schematic representation of the PVT panels connected to the bottom Ecovat buffer segment.

$$T_i^{out} \geqslant T_{i,N_{seg}} \; w_i \qquad\qquad \forall i \in \mathfrak{I}, \qquad (3.3)$$

$$T_i^{out} \leqslant T_{i,N_{seg}} + M \; w_i \qquad\qquad \forall i \in \mathfrak{I}, \qquad (3.4)$$

where M is a sufficiently large constant. To allow for the option of not connecting the PVT panels even if $w_i = 1$, the following constraint is added:

$$x_{i,N_{seg}}^{pvt} \leqslant w_i \quad \forall i \in \mathfrak{I}. \qquad (3.5)$$

In this constraint $x_{i,N_{seg}}^{pvt}$ is the decision variable indicating whether the PVT panels are connected to the bottom segment of the buffer in time interval i. Note, that the PVT panels can only be connected to the bottom segment of the buffer. However, for ease of formulation we also introduce such variables for the other segments and set them to 0; i.e.:

$$x_{i,s}^{pvt} = 0 \quad \forall i \in \mathfrak{I}, \; s \in \mathcal{S} \setminus \{N_{seg}\}. \qquad (3.6)$$

Figure 3.2 gives a schematic representation of how the PVT panels can be connected to the bottom segment of the Ecovat buffer. The input temperature of the panels during time interval i is denoted by T_i^{in} and the output temperature of the panels during time interval i is denoted by T_i^{out}. Note, that if $x_{i,N_{seg}}^{pvt} = 1$ the panels are connected to the bottom segment, and in this case the output temperature of the PVT panels depends on the global irradiation G_i and the ambient temperature T_i^{amb}, as well as on the input temperature T_i^{in}.

To determine the thermal and electrical energy produced by the PVT panels we need to consider the thermal and electrical efficiencies of the panels. The steady

state thermal efficiency of a PVT panel, η^{th} is defined as the yield of the panel divided by the global irradiation incident on the panel [55]:

$$\eta^{th} = \frac{\dot{m}\, c_p\, (T^{out} - T^{in})}{G\, A}, \qquad (3.7)$$

where $\dot{m}$ is the mass flow rate of the heat transfer fluid running through the PVT panel, which for simplicity we assume to be water (in reality this might be mixed with another liquid to prevent freezing), c_p is the specific heat of water and A is the surface area of the PVT panel. The input temperature T^{in} of the PVT panels is equal to the temperature of the bottom layer of the Ecovat buffer. The heat transfer fluid of temperature T^{out} is fed back to the bottom segment of the buffer, which charges the buffer due to it now having a higher temperature than the water in the bottom segment of the buffer. Simultaneously, this lowers the temperature of the heat transfer fluid back to $T^{in} = T_{i,N_{seg}}$ and thus it can further cool the PVT panels, increasing their electrical efficiency. Here we make the simplifying assumption that by transferring energy from the heat transfer fluid to the buffer segment, the temperature of the heat transfer fluid will always be cooled down to the temperature of the bottom buffer segment. Note, that by this charging of the bottom buffer segment over time the temperature in the bottom segment will slowly increase. To counteract this temperature increase, one of the water-water heat pumps may be used at some point in time to transfer energy from the bottom segment to a higher segment of the buffer. This is done during time intervals in which the PVT panels are not connected to the bottom segment. We refer to Section 3.2.4 for further information on the modelling of the water-water heat pumps.

If we look at Equation 3.7 we note that it does not give us the thermal efficiency yet, since T^{out} is also unknown. We solve this problem by deriving another expression for the thermal efficiency of the PVT panels. This equation can be combined with Equation 3.7 to obtain the thermal efficiency η^{th} and the output temperature of the PVT panels T^{out}.

To obtain a second expression for the thermal efficiency of the PVT panels we use the reduced temperature, which is a quantity defined as the difference between the mean temperature in the PVT panels and the ambient temperature divided by the global irradiation [36], i.e.:

$$T^{red} = \frac{\frac{T^{in} - T^{out}}{2} - T^{amb}}{G}, \qquad (3.8)$$

where T^{red} is the reduced temperature and T^{amb} is the ambient temperature. The thermal efficiency of the PVT panels can be approximated as a linear function of the reduced temperature, T^{red} (see for example [36, 100]):
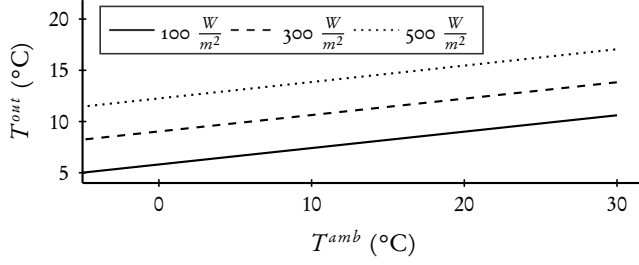
Figure 3.3: The output temperature of the PVT panels as a function of the ambient temperature, T^{amb}.

$$\eta^{th} = \eta_0^{th} - a_{th}\, T^{red}, \qquad\qquad (3.9)$$

where η_0^{th} is the thermal efficiency of the PVT panels at a reduced temperature, T^{red}, of zero and a_{th} is the thermal loss coefficient of the PVT panels. This gives us a second expression for the thermal efficiency, leading to two equations (Equations 3.7 and 3.9) and two unknowns, η^{th} and T^{out}. By combining Equations (3.7), (3.8) and (3.9) we obtain an expression for the output temperature of the PVT panels, which is used in the ILP model:

$$T_i^{out} = \frac{2\,\dot{m}\,c_p\,T_i^{in} - a_{th}\,A\,T_i^{in} + 2\,A\,\eta_0^{th}\,G_i + 2\,a_{th}\,A\,T_i^{amb}}{a_{th}\,A + 2\,\dot{m}\,c_p} \qquad \forall i \in \mathfrak{I}, s \in \mathcal{S}.$$

$$(3.10)$$

In Figure 3.3 the output temperature of the PVT panels as a function of the ambient temperature is shown for 3 different global irradiations. In the figure the input temperature of the PVT panels, T_i^{in}, is set at 5°C. The solid line corresponds to a global irradiation, G_i, of 100 $\frac{W}{m^2}$, the dashed line to a global irradiation of 300 $\frac{W}{m^2}$ and the dotted line to a global irradiation of 500 $\frac{W}{m^2}$. The mass flow rate of the heat transfer fluid in the PVT panels considered in this figure is equal to its value in the Ecovat system, which is $\dot{m} = 0.018$ kg/s.

As mentioned before, the PVT panels will not be connected to the buffer in case the output temperature of the panels is lower than the temperature in the bottom segment of the buffer. As one would expect, this occurs at times when the ambient temperature is low and/or the level of global irradiation is low, as confirmed by Figure 3.3.

To obtain the thermal and electrical outputs of the PVT panels in the ILP model we need to add constraints on the reduced temperature, thermal efficiency and electrical efficiency to specify their dependencies. The constraint for the re-

duced temperature is as based on Equation 3.8, with an exception added to avoid dividing by zero when the global irradiation is zero:

$$T_i^{red} = \begin{cases} \dfrac{\dfrac{T_{i,N_{seg}} + T_i^{out}}{2} - T_i^{amb}}{G_i}, & \text{if } G_i > 0 \quad \forall\, i \in \mathcal{I}, \\ 0, & \text{otherwise.} \end{cases} \qquad (3.11)$$

To make sure that the thermal efficiency takes the correct value in the ILP model multiple constraints are required. This is because the thermal efficiency is restricted to the range of 0 to η_{max}^{th}, where η_{max}^{th} is the maximum thermal efficiency the PVT panels can reach. We define the following constraints to specify the thermal efficiency of the PVT panels:

$$\eta_i^{th} \leqslant \eta_0^{th} - a_{th}\, T_i^{red} + M\, y_i^{th} \qquad\qquad \forall i \in \mathcal{I}, \qquad (3.12)$$
$$\eta_i^{th} \leqslant \eta_{max}^{th}\, (1 - y_i^{th}) \qquad\qquad \forall i \in \mathcal{I}, \qquad (3.13)$$
$$\eta_i^{th} \geqslant 0 \qquad\qquad \forall i \in \mathcal{I}, \qquad (3.14)$$

where y_i^{th} is a binary indicator variable that indicates which of the Constraints (3.12) or (3.13) is active during interval i and M is again a sufficiently large constant. There are three different cases depending on the values for y_i^{th} and T_i^{red}:

» $y_i^{th} = 0$ and $T_i^{red} > 0$: in this case Constraint (3.12) gives a tighter upper bound than Constraint (3.13), i.e. $\eta_i^{th} \leqslant \eta_0^{th} - a_{th}\, T_i^{red}$.

» $y_i^{th} = 0$ and $T_i^{red} \leqslant 0$: in this case either Constraint (3.12) or Constraint (3.13) gives a tighter upper bound to η_i^{th}.

» $y_i^{th} = 1$: in this case $\eta_i^{th} = 0$ due to Constraint (3.13) and Constraint (3.14).

Due to the way the objective function is constructed, which is discussed in Section 3.2.9, the model will always set $y_i^{th} = 0$ if possible and will always select the upper bound for the thermal efficiency, in other words we have $\eta_i^{th} = \eta_0^{th} - a_{th}\, T_i^{red}$ or $\eta_i^{th} = \eta_{max}^{th}$ depending on the value of T_i^{red}.

Finally, we approximate the electrical efficiency of the PVT panels in a similar way as the thermal efficiency. This leads to a linear function of the reduced temperature, similar to Equation (3.9). To model this, similar constraints as for the thermal efficiency are defined for the electrical efficiency:

$$\eta_i^{el} \leqslant \eta_0^{el} - a_{el} \, T_i^{red} + M \, y_i^{el} \qquad\qquad \forall i \in \mathcal{I}, \qquad (3.15)$$

$$\eta_i^{el} \leqslant \eta_{max}^{el} \, (1 - y_i^{el}) \qquad\qquad \forall i \in \mathcal{I}, \qquad (3.16)$$

$$\eta_i^{el} \geqslant 0 \qquad\qquad \forall i \in \mathcal{I}, \qquad (3.17)$$

where η_i^{el} is the electrical efficiency of the PVT panels, η_0^{el} is the electrical efficiency at a reduced temperature of zero, η_{max}^{el} is the maximum electrical efficiency, a_{el} is the electrical loss coefficient of the PVT panels and y_i^{el} is a binary indicator variable that indicates whether Constraint (3.15) or Constraint (3.16) is active during interval i. The thermal and electrical output of the PVT panels, which can be obtained using the thermal and electrical efficiencies, is discussed in Section 3.2.8.

3.2.3 AIR/WATER HEAT PUMP

A heat pump is a device that transfers thermal energy from a heat source to a heat sink, where the heat source has a lower temperature than the heat sink. To accomplish this transfer of thermal energy, which opposes the natural flow of thermal energy, electrical energy is needed to operate the heat pump. An air/water heat pump uses the ambient air as its heat source, and in the case of the Ecovat system it uses one of the buffer segments as the heat sink.

The air/water heat pump can only charge a buffer segment s that has a temperature, $T_{i,s}$, within a limited temperature range: $T_{min}^{awhp} \leqslant T_{i,s} \leqslant T_{max}^{awhp}$, where T_{min}^{awhp} and T_{max}^{awhp} are the minimum and maximum temperatures the air/water heat pump can supply. To make sure that the air/water heat pump is only connected to a segment within this temperature range, the following constraints are added to the ILP model of the Ecovat system:

$$(1 - x_{i,s}^{awhp}) \, M + T_{i,s} \geqslant T_{min}^{awhp} \qquad\qquad \forall i \in \mathcal{I}, \, s \in \mathcal{S}, \qquad (3.18)$$

$$x_{i,s}^{awhp} \, T_{i,s} \leqslant T_{max}^{awhp} \qquad\qquad \forall i \in \mathcal{I}, \, s \in \mathcal{S}, \qquad (3.19)$$

where $x_{i,s}^{awhp}$ is the decision variable that indicates whether the air/water heat pump is connected to segment s during interval i. These constraints ensure that if the air/water heat pump is selected to charge a buffer segment, that buffer segment has a temperature within the temperature range of the air/water heat pump.

It should be noted that Constraint (3.19) contains a multiplication of the continous variable $T_{i,s}$ and the binary variable $x_{i,s}^{awhp}$, which means the constraint is

not linear and in principle can not be used in this form in an ILP model. However, using the well known method (described for example in [92]) of modelling the multiplication of the two variables as a new variable itself and placing a number of constraints on this new variable, such a quadratic constraint can be linearised. Note, that also all further quadratic constraints introduced in this chapter can be linearised using the same method.

To give some insight in this linearization of quadratic constraints, consider as example the multiplication $x_{i,s}^{awhp}\, T_{i,s}$ from Constraint (3.19). First, we define a new variable $\tilde{x}_{i,s}^{awhp} = x_{i,s}^{awhp}\, T_{i,s}$. The following constraints ensure that this variable $\tilde{x}_{i,s}^{awhp}$ always takes the correct value:

$$\tilde{x}_{i,s}^{awhp} \geqslant 0 \qquad\qquad \forall i \in \mathcal{I},\ s \in \mathcal{S}, \qquad (3.20)$$

$$\tilde{x}_{i,s}^{awhp} \leqslant M\, x_{i,s}^{awhp} \qquad\qquad \forall i \in \mathcal{I},\ s \in \mathcal{S}, \qquad (3.21)$$

$$\tilde{x}_{i,s}^{awhp} \leqslant T_{i,s} \qquad\qquad \forall i \in \mathcal{I},\ s \in \mathcal{S}, \qquad (3.22)$$

$$\tilde{x}_{i,s}^{awhp} \geqslant T_{i,s} - M\, (1 - x_{i,s}^{awhp}) \qquad\qquad \forall i \in \mathcal{I},\ s \in \mathcal{S}, \qquad (3.23)$$

where M is again a large constant. Constraint (3.20) and Constraint (3.21) make sure $\tilde{x}_{i,s}^{awhp} = 0$ if $x_{i,s}^{awhp} = 0$, while Constraint (3.22) and Constraint (3.23) make sure $\tilde{x}_{i,s}^{awhp} = T_{i,s}$ if $x_{i,s}^{awhp} = 1$.

3.2.4 WATER/WATER HEAT PUMPS

The water/water heat pumps in the Ecovat system work similarly to the air/water heat pump, except that instead of using the ambient air as the heat source they use one of the other buffer segments as the heat source. The buffer segment used as the heat source has to have a lower temperature than the buffer segment used as the heat sink. In this way the water/water heat pumps increase the energy quality of the buffer which is important for satisfying the heat demand as discussed in Section 1.2. The Ecovat system is equipped with two water/water heat pumps, which have different temperature ranges, to be able to cover the entire temperature range possible in the buffer. One heat pump covers the lower part of the temperature range, while the second heat pump covers the higher part. We label these two heat pumps as low temperature heat pump (lthp) and high temperature heat pump (hthp) respectively. In this section we only specify the constraints required to model the lthp, since the constraints for the hthp are exactly the same (they can be obtained by replacing all lthp labels with hthp labels in the constraints discussed in this subsection).

As mentioned, the lthp has a limited output temperature range: $T_{min}^{lthp} \leqslant T_{i,s} \leqslant T_{max}^{lthp}$, where T_{min}^{lthp} and T_{max}^{lthp} are the minimum and maximum temperature

that the lthp can supply. For the lthp heat pump the temperatures of both the buffer segment used as heat source and the buffer segment used as heat sink are required to be in this range. To accommodate for the fact that the lthp needs to be connected to two buffer segments we define the labels, lthp_ and lthp$_+$, for the heat source and heat sink respectively. This means we extend the set of devices to $\mathcal{D} = \{pvt, awhp, lthp\_, lthp\_+, hthp\_, hthp\_+, res, dem\}$. The decision variables that indicate whether segment s is used as a heat source or heat sink for the lthp during interval i are then $x_{i,s}^{lthp\_}$ and $x_{i,s}^{lthp\_+}$ respectively.

The constraints to ensure that the temperatures of the segments selected as heat source and heat sink are within the temperature range of the lthp are similar to those given for the air/water heat pump:

$$(1 - x_{i,s}^{lthp\_}) M + T_{i,s} \geqslant T_{min}^{lthp} \qquad \forall i \in \mathcal{I}, \ s \in \mathcal{S}, \qquad (3.24)$$

$$x_{i,s}^{lthp\_+} T_{i,s} \leqslant T_{max}^{lthp} \qquad \forall i \in \mathcal{I}, \ s \in \mathcal{S}. \qquad (3.25)$$

The buffer segment selected as heat sink must have a higher temperature than the buffer segment selected as heat source, which is ensured by adding the following constraint:

$$\sum_{s \in \mathcal{S}} x_{i,s}^{lthp\_+} T_{i,s} > \sum_{s \in \mathcal{S}} x_{i,s}^{lthp\_} T_{i,s} \quad \forall i \in \mathcal{I}. \qquad (3.26)$$

Finally, we need to ensure that a buffer segment is selected as heat sink if and only if another buffer segment is selected as heat source. To make sure of this the following constraint is added:

$$\sum_{s \in \mathcal{S}} x_{i,s}^{lthp\_+} = \sum_{s \in \mathcal{S}} x_{i,s}^{lthp\_} \quad \forall i \in \mathcal{I}. \qquad (3.27)$$

Constraint (3.1) combined with the extension of the device set $\mathcal{D}$ ensures that at most one buffer segment can be selected as heat source and at most one buffer segment can be selected as heat sink.

3.2.5 RESISTANCE HEATER

The resistance heater is a simple device that converts electrical energy to thermal energy on a one-to-one basis. The resistance heater can cover the entire range of temperatures possible in the Ecovat buffer, so no constraints are necessary to restrict its use. Furthermore, the resistance heater only connects to a single buffer segment, which is already covered by Constraint (3.1). Due to this no further constraints are necessary for the modelling of the use of the resistance heater.

3.2.6 HEAT DEMAND

As mentioned before, one of the inputs of the ILP model is the predicted heat demand for every time interval in the time horizon. This demand is specified by the amount D_i of thermal energy demanded and a temperature T_{dem} this demand should be supplied at, which we call the demand temperature. If during a given time interval there is a non-zero heat demand, the model has to ensure that this demand gets supplied. In other words, if there is a demand during a given interval, a buffer segment has to be selected to fulfil that demand. A consequence of this restriction is that the model must always ensure that the buffer contains enough energy to satisfy the heat demand.

To be able to supply the heat demand the buffer segment selected to fulfil the demand must have at least a temperature equal to the demand temperature. The following constraint ensures the buffer segment that is selected to supply the heat demand has a sufficiently high temperature:

$$T_{i,s} \geqslant T_{dem} \, x_{i,s}^{dem} \quad \forall i \in \mathcal{I}, \, s \in \mathcal{S}, \tag{3.28}$$

where $x_{i,s}^{dem}$ is the decision variable indicating whether segment s is supplying the heat demand during time interval i. Furthermore, for the heat demand constraint (3.1) is modified as follows:

$$\sum_{s \in \mathcal{S}} x_{i,s}^{dem} = \begin{cases} 1, & \text{if } D_i > 0 \\ 0, & \text{if } D_i = 0 \end{cases} \quad \forall \, i \in \mathcal{I}, \tag{3.29}$$

which ensures that precisely one of the buffer segments is selected to supply the heat demand if there is any demand during time interval i, and no buffer segment is selected if there is no heat demand.

3.2.7 HEAT LOSSES

Over time the Ecovat buffer will lose some of its energy content to the surrounding environment. If the heat loss coefficient U of the Ecovat buffer is known one can use the following constraint to determine the heat loss of the Ecovat buffer to its surroundings:

$$Q_{i,s}^{loss} = U \, A_s \, (T_{i,s} - T_{gw}) \quad \forall i \in \mathcal{I}, \, s \in \mathcal{S}, \tag{3.30}$$

where $Q_{i,s}^{loss}$ is the heat lost to the environment from segment s during time interval i, A_s is the surface area of the walls (and lid for the top segment) of segment s and T_{gw} is the temperature of the environment, which in the case of the Ecovat buffer is the ground water temperature.

However, at the time of writing no measurements concerning the heat loss of the Ecovat buffer were available. For this reason estimates for the heat losses presented in an internal report of the company were used to determine the heat losses. However, the given data only specified an estimated heat loss over a period of 6 months [86]. Because of this, the following constraint was used to determine the heat loss of the Ecovat buffer instead:

$$Q_{i,s}^{loss} = \left(1 - (1 - \beta)^{\frac{2}{N_{int}}}\right)(T_{i,s} - T_{gw})\, m_s\, c_p \quad \forall i \in \mathcal{I},\ s \in \mathcal{S}, \qquad (3.31)$$

where β is the given heat loss coefficient of the Ecovat buffer over 6 months, m_s is the mass of segment s and the factor $\frac{2}{N_{int}}$ is equal to one over the number of time intervals in half a year ($\frac{1}{N_{int}/2} = \frac{2}{N_{int}}$).

3.2.8 Temperature evolution

In the previous subsections we have described the modelling of the devices in the Ecovat system. The decision variables $x_{i,s}^{dev}$ that describe whether these devices run during a give time interval i and if so, to which buffer segment s they are connected, determine when and how the buffer is charged/discharged. This leads to changes in the temperature of the buffer segments. In this subsection we define constraints that specify these temperature changes in the Ecovat buffer. The temperature changes over the entire time horizon, i.e. the temperature evolution in the Ecovat buffer, are one of the outputs of the ILP model.

To describe the temperature changes in the Ecovat buffer we take a differential equation similar to the one used in [68] for a one-dimensional water tank model:

$$m_s\, c_p\, \frac{dT}{dt} = Q_{heat} - Q_{cool} - Q_{loss}, \qquad (3.32)$$

where Q_{heat} is the amount of energy added to the buffer (by devices in the Ecovat system), Q_{cool} is the amount of energy extracted from the buffer (when used as heat source for one of the water/water heat pumps or to supply the heat demand of the neighbourhood) and Q_{loss} is the amount of energy lost to the surroundings of the buffer. Note, that for simplicity we have omitted the inclusion of a term to describe the heat exchange/mixing between different segments of the buffer, since we expect those effects to be minor compared to the heating and cooling terms from the devices in the system. However, if from experiments it turns out that this heat exchange between segments is significant it can easily be added to the model, at the cost of longer computational times. After discretizing Equation 3.32 and some rewriting, we obtain a constraint specifying the temperature changes in the buffer:

$$T_{i+1,s} = T_{i,s} + \frac{\Delta t}{m_s \, c_p} \Big(\sum_{dev \in \mathcal{D}} Q_{i,s}^{dev} - Q_{i,s}^{loss} \Big) \quad \forall i \in \mathcal{I} \setminus \{N_{int}+1\}, \ s \in \mathcal{S}, \ (3.33)$$

where Δt is the length of a time interval, $Q_{i,s}^{dev}$ is the energy change in segment s during interval i induced by device dev and $T_{1,s}$ is the initial temperature distribution in the buffer. The $Q_{i,s}^{dev}$ terms differ for the different devices, these terms are given below:

$$Q_{i,s}^{pvt} = \begin{cases} 0, & \forall i \in \mathcal{I}, \ s \in \mathcal{S} \setminus \{N_{seg}\}, \\ \eta_i^{th} \, G_i \, A \, x_{i,s}^{pvt}, & \forall i \in \mathcal{I}, \ s = N_{seg}, \end{cases} \tag{3.34}$$

$$Q_{i,s}^{awhp} = C^{awhp} \, COP^{awhp} \, x_{i,s}^{awhp} \qquad \forall i \in \mathcal{I}, \ s \in \mathcal{S}, \tag{3.35}$$

$$Q_{i,s}^{lthp-} = -C^{lthp} \, (COP^{lthp}-1) \, x_{i,s}^{lthp-} \qquad \forall i \in \mathcal{I}, \ s \in \mathcal{S}, \tag{3.36}$$

$$Q_{i,s}^{lthp+} = C^{lthp} \, COP^{lthp} \, x_{i,s}^{lthp+} \qquad \forall i \in \mathcal{I}, \ s \in \mathcal{S}, \tag{3.37}$$

$$Q_{i,s}^{hthp-} = -C^{hthp} \, (COP^{hthp}-1) \, x_{i,s}^{hthp-} \qquad \forall i \in \mathcal{I}, \ s \in \mathcal{S}, \tag{3.38}$$

$$Q_{i,s}^{hthp+} = C^{hthp} \, COP^{hthp} \, x_{i,s}^{hthp+} \qquad \forall i \in \mathcal{I}, \ s \in \mathcal{S}, \tag{3.39}$$

$$Q_{i,s}^{res} = C^{res} \, x_{i,s}^{res} \qquad \forall i \in \mathcal{I}, \ s \in \mathcal{S}, \tag{3.40}$$

$$Q_{i,s}^{dem} = -D_i \, x_{i,s}^{dem} \qquad \forall i \in \mathcal{I}, \ s \in \mathcal{S}, \tag{3.41}$$

where COP^{dev} is the coefficient of performance of device dev and C^{dev} is the capacity of device dev, which equals the consumption of device dev when it is turned on, since all devices in the Ecovat system are modelled as on/off devices in the ILP model. The factor $(COP^{dev}-1)$ in Constraints (3.37) and (3.39) is due to the fact that even though $C^{dev} \, COP^{dev}$ units of thermal energy are transferred to the heat sink, only $C^{dev} \, (COP^{dev}-1)$ units of energy are coming from the heat source. The remaining C^{dev} is the electrical energy consumed by the heat pump to transfer the thermal energy. It should be noted that we assume, for simplicity, that there are no energy losses in the operation of the heat pump. Furthermore, we assume a constant COP for the heat pumps. In reality the COP of a heat pump will depend on the input and output temperatures of the heat pump. However, including this non-linear dependence in an ILP model is hard and would significantly increase the complexity. For this reason we assume a constant COP for each heat pump, the value of which will be based on an estimated average COP of that heat pump.

Finally, a few additional constraints need to be added to model the temperatures of the buffer segments. The first ensures that the temperature inside a buffer segment can not exceed the maximum allowed temperature in that segment:

$$T_{i,s} \leqslant \bar{T}_s \quad \forall i \in \mathfrak{I}, \, s \in \mathcal{S}, \tag{3.42}$$

where $\bar{T}_s$ is the maximum allowed temperature in segment s. The second ensures that the temperature gradient from the top segment to the bottom segment of the buffer is always decreasing:

$$T_{i,s} \geqslant T_{i,s+1} \quad \forall i \in \mathfrak{I}, \, s \in \mathcal{S} \setminus \{N_{seg}\}. \tag{3.43}$$

3.2.9 OBJECTIVE FUNCTION

The objective of the Ecovat system is to minimize the operational costs of the system under the restriction that the heat demand must always be satisfied. Note, that we only consider the operational costs of the system. We do not include e.g. maintenance costs or the profit made by supplying the heat demand of the neighbourhood, since we assume these do not depend on the specific charging/discharging strategy used.

The operational costs consist of the cost to run the devices that charge the Ecovat buffer, i.e. the air/water heat pump, the two water/water heat pumps and the resistance heater. The cost to run any electricity consuming device during time interval i depends on the electricity price during that interval, p_i, which may even be negative, the length of a time interval Δt and the capacity of the device C^{dev}. The cost to run a device dev during time interval i is then given by $\sum_{s \in \mathcal{S}} p_i \, \Delta t \, C^{dev} x_{i,s}^{dev}$. The operational costs are reduced by the amount of electrical energy E_i^{pvt} produced by the PVT panels in the Ecovat system, which is given by:

$$E_i^{pvt} = \eta_i^{el} \, G_i \, A \sum_{s \in \mathcal{S}} x_{i,s}^{pvt} \quad \forall i \in \mathfrak{I}. \tag{3.44}$$

The operational costs K_i incurred during time interval i, are then given by:

$$K_i = p_i \, \Delta t \left(\sum_{s \in \mathcal{S}} (C^{awhp} \, x_{i,s}^{awhp} + C^{lthp} \, x_{i,s}^{lthp-} + C^{hthp} \, x_{i,s}^{hthp-} + C^{res} \, x_{i,s}^{res}) \right.$$
$$\left. - E_i^{pvt} \right) \quad \forall i \in \mathfrak{I}, \tag{3.45}$$

There are a few things to note; first, in Constraint (3.45) we have used $x_{i,s}^{lthp-}$ and $x_{i,s}^{hthp-}$ to determine whether the water/water heat pumps are running during time interval i, and are thus contributing to the costs during that interval.

Whether we use those decision variables or $x_{i,s}^{lthp_+}$ and $x_{i,s}^{hthp_+}$ does not matter, since it gives the same results due to Constraint (3.27). Second, in Constraint (3.45) we make the assumption that any electrical energy that is generated by the PVT panels and is not consumed by the Ecovat system can be sold for the electricity price at that time. This assumption is based on the current net-metering (salderen in Dutch) laws in the Netherlands [12]. However, the current net-metering law is planned to be replaced by a different regulation, which decreases the feed-back tariff for consumers based on the prices of PV panels at that point in time [13] (the goal of this regulation is to keep the payback time of PV panels around 7 years to still incentive consumers to invest in PV panels). When this changes in the future it is straightforward to include this different feed-back tariff in the model.

As stated before, the objective of the Ecovat system is to minimize the total operational costs over the time horizon, i.e $\min \sum_{i \in \mathcal{I}} K_i$. However, a few more terms need to be added to the objective function to ensure the model shows the desired behaviour. First, we want the model to favour the charging of buffer segments situated higher in the buffer, in other words those of higher temperature. The reason this is important is that, as discussed in Section 3.2.6, at least one segment in the buffer is required to have a temperature at least as high as the demand temperature. For time intervals inside the time horizon the model will allow only solutions where there is always at least one buffer segment with sufficiently high temperature to satisfy the heat demand. However, we would like to ensure that this is also true after the time horizon has passed. The reason for this is that in a real world situation the system would be operating continuously. As such, it is desirable to favour the charging of high temperature segments, to ensure the system always has the means to satisfy the heat demand of the neighbourhood, even after the time horizon. To ensure this behaviour the term $-c_1 \sum_{s \in \mathcal{S}} w_s \, T_{i,s}$ is added to the objective function, where the weights w_s are chosen such that $w_s > w_{s+1}$. This choice for the weights ensures that a segment with higher temperature is favoured over a segment with lower temperature. In our case we choose the weights as $w_s = 6 - s$. Second, we need to ensure that the upper bound on the thermal efficiency constraints discussed in Section 3.2.2 is always selected. To do so we add the term $-c_2 \, Q_{i,s}^{pvt}$ to the objective function. A higher thermal efficiency leads to a higher value of $Q_{i,s}^{pvt}$, and as such, the term $-c_2 \, Q_{i,s}^{pvt}$ ensures that the highest thermal efficiency allowed by the constraints discussed in Section 3.2.2 is always selected. In most cases the model will select the upper bound on the thermal efficiency even without the inclusion of this term. However, cases can occur in which it would be beneficial for the model to select a thermal efficiency of zero while simultaneously having a high electrical efficiency. This would lead to a situation in which the PVT panels have a high electrical efficiency due to being cooled, but without heating up the heat transfer fluid that is doing this cooling, which is physically not possible. The addition

of this second term to the objective function prevents such cases from occurring in the model. The full objective function is then given by:

$$\min \sum_{i \in \mathcal{J}} \left(K_i - c_1 \sum_{s \in \mathcal{S}} (6-s) \; T_{i,s} - c_2 \; Q_{i,s}^{pvt} \right), \tag{3.46}$$

where c_1 and c_2 are small, positive constants.

3.3 CONCLUSION

In this chapter we have modelled the Ecovat system and its control as an ILP model. The objective of the model is to find a minimum cost charging/discharging strategy for the Ecovat system over a given time horizon, under the restriction that the requested heat demand from the connected neighbourhood must always be satisfied. This charging/discharging strategy is determined by specifying for each time interval in the time horizon which devices are turned on, and to which buffer segment they are connected in case they are turned on. In the different subsections of this chapter we have described these devices and presented a set of variables and constraints to properly model the restrictions of each device. Next to the device specific restrictions, a number of general restrictions apply to the system among which the most important are:

- » A device may only be connected to at most one buffer segment during each time interval (Constraint 3.1).
- » A buffer segment may have only one device connected to it during each time interval (Constraint 3.2).
- » Each buffer segment must have a temperature lower than a specified maximum temperature at all times (Constraint 3.42).
- » Each buffer segment must have a lower temperature than the buffer segment above it at all times (Constraint 3.43).

As output the presented model gives a charging/discharging strategy for the Ecovat system over the time horizon. This strategy is specified by the decision variables of the model, which describe whether a device is connected to a specific segment during a given time interval. Additionally, the model provides the temperature evolution within the segments of the buffer over the entire time horizon as output. As input the model requires predictions for the energy prices, the heat demand of the neighbourhood, the ambient temperature and the global irradiation over the entire time horizon.

While this model in theory finds a good charging/discharging strategy for the Ecovat system given a certain input, it is not yet usable in practice. Since the Ecovat systems is a seasonal thermal storage it is important to consider a sufficiently long time horizon to include seasonal effects, which will be important for such a

system. For this reason, a time horizon of a year is considered in this thesis. Furthermore, we wish to incorporate an energy market into the model, in our case this is the Dutch imbalance market. Since the Dutch imbalance market operates on time intervals of 15 minutes, this is also the time interval length we consider in this thesis. However, due to the combination of a short time interval length and a long time horizon, the number of variables and constraints in the ILP model presented in this chapter is too large to be solved in a feasible amount of time. As such, some modifications are required to make the ILP model solvable in a shorter time (at most around a week of simulation time). The modifications required to achieve this are discussed in the next chapter.

54

4

Method for operational control in simulations

Abstract – *In this chapter we adapt the ILP model presented in the previous chapter to be able to solve it in a reasonable amount of time. To achieve this, we first introduce a rolling horizon approach. However, it turns out that this approach insufficiently incorporates seasonal effects, and as a consequence often fails to generate a good solution. To properly incorporate seasonal effects we extend the model with a long-term planning step and present the results of this extended model, which significantly improves the results obtained without this long-term planning step. Furthermore, we analyse the effect of prediction errors in the energy prices on the results achieved with the presented model.*

In the previous chapter we have presented an ILP model to describe the Ecovat system. The output of this model includes a charging strategy for the Ecovat system. However, as noted the model from Chapter 3 can not be solved for reasonable instances. Due to the complexity of the model, the long time scales involved in seasonal thermal storage, and the short time intervals required to incorporate the energy market into the model, the computational time required to solve the ILP model would be unrealistically long (see Chapter 3). Our goal in this chapter is to limit the solving time of the model to at most a few days when optimizing a year of operation for an Ecovat system. Therefore, in this chapter we start by introducing a rolling horizon approach to ensure the presented ILP model becomes solvable in the previously specified amount of time, i.e. at most around a week of simulation time on a personal computer.

This chapter is based on [G:1] and [G:2].

4.1   Rolling horizon

As mentioned, due to the very large number of variables in the ILP model and the short time intervals required it is not possible to optimize over a whole year at once. For this reason a rolling horizon approach is introduced in which the year-long optimization is divided into smaller overlapping periods of k days, whereby from an optimization of each of these periods only the first part is realised. This process is depicted in Figure 4.1 for the situation where we optimize for $k = 2$ days and only realise the first day. In each of the optimizations, the initial conditions of the Ecovat buffer are taken based on the already realised days and combined with predictions for the next days (in the case of Figure 4.1 the next two days) to optimize the charging/discharging strategy for the Ecovat system. From this solution, only the realization for the first part of the period is used. The temperature distribution in the buffer at the end of this first part is then used as the initial conditions for the next period which starts at the end of that first part. This process is repeated until the whole time planning period of a complete year has been considered. The effect of different rolling horizon lengths, i.e. different values of k, on the resulting charging strategies has been investigated and the corresponding results are presented in Section 4.1.2. It is clear that the total objective value for the entire optimization using this method will be worse compared to an optimization over the whole year at once. However, the iterative approach employed in the rolling horizon approach better simulates reality, as for example good weather predictions are only available for a few days in advance.

As even for the shorter periods it may happen that solving the corresponding model takes quite long, we have introduced the following stopping criteria for the optimizations over a period using the ILP model discussed in Chapter 3:

» The gap between the best lower bound and the best solution found by the solver is less than 0.2 %.

» The absolute value of the gap between the best lower bound and the best solution found by the solver is 1 € or less.

» The solver takes longer than one hour to reach one of the previous two conditions.

4.1.1   Optimization setup

The ILP model of the Ecovat system described in Chapter 3 requires a number of input parameters and profiles. The first profile that is required is the heat demand of the buildings connected to the Ecovat system. For this heat demand we consider a case described in a study on the sizing of the Ecovat system [27]. In this case 78 apartments distributed over five apartment buildings are connected to the Ecovat system. To obtain a heat demand profile for the entire year a typical heat demand profile for a winter, a spring, an autumn and a summer day
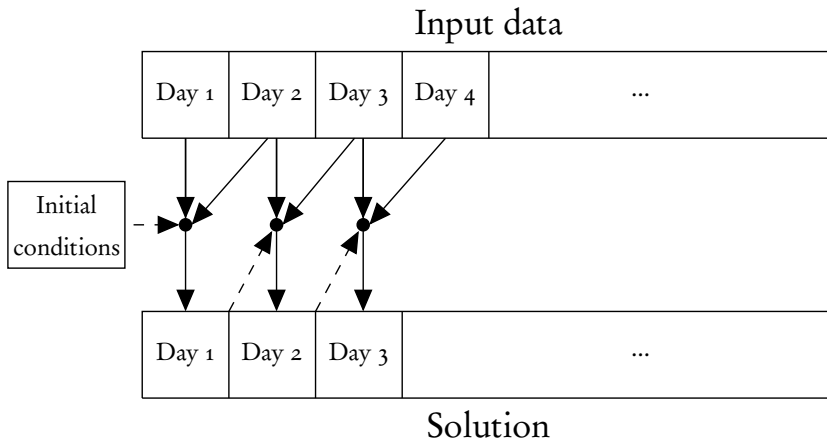
Figure 4.1: Scheme showing the build up of the solution using a rolling horizon approach. In this case a rolling horizon length of 2 days is used.
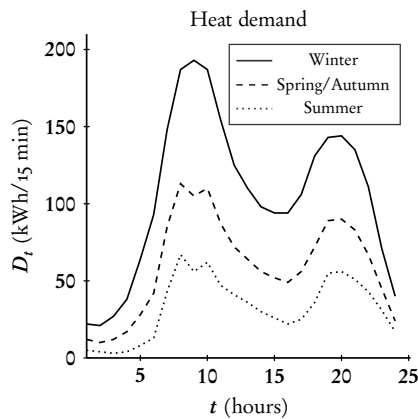
Figure 4.2: The daily heat demand profiles used as input for the model of the Ecovat system.

is constructed based on historical data from the housing corporation owning the apartment buildings [27]. As the spring and autumn days are quite similar we use the same profile for these two seasons. Figure 4.2 shows the resulting daily heat demand profiles for the different seasons. As the profiles include both space heating and tap water demand there is still a significant heat demand during the summer. From these daily profiles a yearly heat demand profile is then constructed in the following way.

The optimization starts on January 1, the first 60 days are winter days, followed
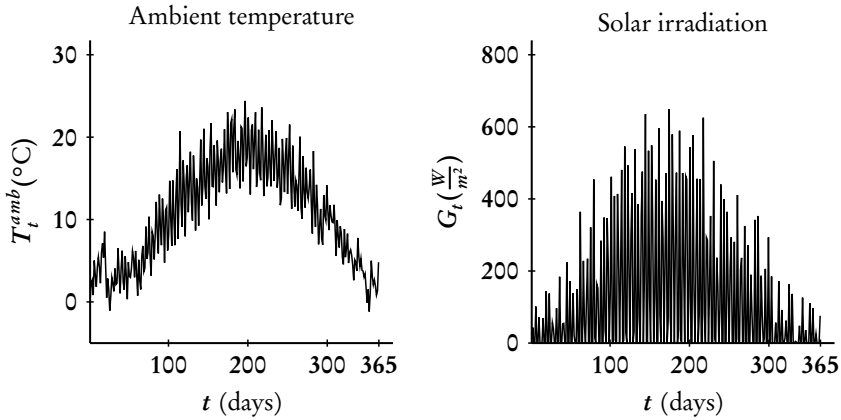
Figure 4.3: The weather data used as input for the model of the Ecovat system.

by 91 spring days, 91 summer days, 91 autumn days and finally 32 winter days for a total of 365 days. Every day uses the heat demand profile from Figure 4.2 corresponding to the season of that day. The heat demand temperature T_{dem} was taken constant throughout the year. In the cases described in this chapter we consider the same heat demand profile for every day in a given season, since we are primarily interested in determining whether the developed model is capable of generating good strategies, and not yet in simulating a realistic use case. As such, simple heat demand profiles are sufficient for the cases considered in this chapter. In Chapter 6 we consider a more realistic case in which every day has a different heat demand profile. Furthermore, in that chapter we investigate the effect of increased heat demand, due to for example a cold winter, on the performance of the developed model.

For the evaluation in this chapter we consider two different heat demand temperatures, namely $T_{dem} = 40$ °C and $T_{dem} = 60$ °C. These values were chosen since the expectation is that the heat demand temperature in future district heating systems will be lower than is currently common. As discussed in Chapter 2, fourth generation district heating systems are expected to operate at temperatures between 30 and 70 °C.

Two other required input profiles are related to the weather data, namely a profile for the ambient temperature and a profile for the solar irradiation. For these profiles averages based on historical data for the city of Eindhoven, where the apartment buildings from the considered case are located, are used. The used profiles are shown in Figure 4.3. The considered weather data profiles have a resolution of one hour.

The final input profile that is required is a profile for the energy prices. For this we use the energy prices from the Dutch imbalance market for the year
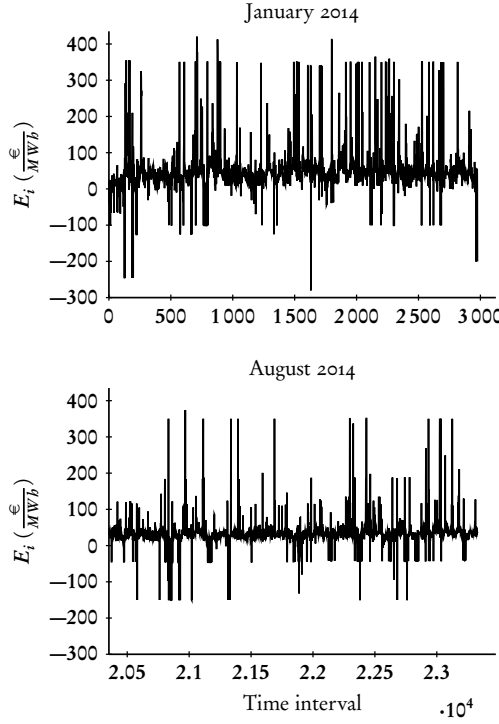
Figure 4.4: The energy prices of January and August 2014, which are used as input for the model of the Ecovat system.

2014. The imbalance market is chosen for the high volatility in its energy prices, including negative energy prices, since this market presents the most interesting business case for an Ecovat system. In this case the system can charge during time intervals with very low (preferably negative) energy prices, leading to lower operational costs. In Figure 4.4 the energy prices for January and August 2014 are shown to indicate the difference in energy prices between a winter and a summer month. We observe that extremely high energy prices are much more frequent in January, and that the average energy cost is higher in January (an average price of 50.94 $\frac{€}{MWh}$ in January versus an average energy price of 34.19 $\frac{€}{MWh}$ in August). Furthermore, from this figure we can see that negative energy prices are more abundant in August, but even in January there are time intervals in which the energy price is negative. This indicates that time intervals with negative energy prices occur around the year. As mentioned, these intervals are the most interesting times to re(charge) the Ecovat system from an economical point of view.

The three previously discussed input profiles are expected to be similar from year to year, with only limited variations. While some years may for example have

a colder winter than other years, the shape of the heat demand profile or the ambient temperature profile will be very similar from year to year. However, for the energy prices this is not the case. Here, the exact moment at which extreme prices occur (in either positive or negative direction) will differ from year to year. Furthermore, we expect that the profile for the energy prices has a strong influence on the outcome of the optimizations. As such, it does not make sense to use an average profile for the energy prices, but to consider instead specific years.

For the cases considered in this section we only consider the energy prices from 2014, while in later sections we also consider the energy prices from different years. Furthermore, the influence of a higher than expected heat demand on the performance of the developed model is investigated in Chapter 6. However, the effect of variations in the weather data are not explicitly studied in this work. The weather data, in particular the solar irradiation, mostly determines the performance of the PVT panels in the Ecovat system. However, those contribute only a small portion to the total energy of the system over an entire year. As such, small variations in the weather data are expected to have only a minor effect on the charging strategies obtained.

Aside from the input profiles a number of input parameters have to be specified, see Table 4.1. The values for the thermal efficiency of the PVT panels at a reduced temperature of zero degrees η_0^{th}, the thermal loss coefficient of the PVT panels a_{th}, the electrical efficiency of the PVT panels at a reduced temperature of zero degrees η_0^{el} and the electrical loss coefficient of the PVT panels a_{el} are obtained from measurements given in [36], where the reduced temperature is defined as the difference between the mean temperature in the PVT panels and the ambient temperature divided by the global irradiation. Values for the maximum efficiencies, η_{max}^{th} and η_{max}^{el} have been selected slightly higher than those at zero reduced temperature, because of the positive effect cooling has on the PVT panels. The capacities and temperature ranges of the devices are based on one of the setups described in [27]. The heat loss factor β of the buffer is based on an estimation made in [86], since no measurements of a real Ecovat system were available at the time of writing.

As noted in Chapter 3, the COP values in the ILP model have been taken as constants in this research. To determine these values a heuristic model for the Ecovat system presented in the aforementioned study on the sizing of the Ecovat system [27] is used. While this heuristic is too simple to be used for the operational control of an Ecovat system, it does include a non-linear calculation of the COPs of the heat pumps. A year long simulation was run using this heuristic and an average COP for each heat pump was determined over this year. These average COPs were used as input values for the ILP model. From a quick comparison between two simulations by this heuristic, one using non-linear COPs, and one using the averaged COPs, we find that the difference in performance is very small, namely less than one percent. This indicates that, at least for a

Input parameters for modelling the Ecovat system			
Parameter	Value	Parameter	Value
N_{seg}	5	A	149.4 m^2
N_{int}	35040	η_0^{th}	0.73
Δt	900 s	η_{max}^{th}	0.75
m_s	{1.04e6, 1.04e6, 1.04e6, 9.11e5, 9.11e5}* kg	a_{th}	7.25
		η_0^{el}	0.1
$\dot{m}$	0.018 kg/s	η_{max}^{el}	0.15
c_p	4186 J/(kg K)	a_{el}	0.44
β	0.08	C^{awhp}	9 kW
$T_{1,s}$	{90, 75, 50, 30, 5} °C	C^{lthp}	15 kW
T_{dem}	40 or 60 °C	C^{hthp}	15 kW
T_{min}^{awhp}	0 °C	C^{res}	1000 kW
T_{max}^{awhp}	59 °C	COP^{awhp}	2.686
T_{min}^{lthp}	0 °C	COP^{lthp}	2.851
T_{max}^{lthp}	49 °C	COP^{hthp}	3.681
T_{min}^{hthp}	48 °C	c_1	1e-5 $\frac{\text{€}}{\text{°C}}$
T_{max}^{hthp}	79 °C	c_2	1e-5 $\frac{\text{€}}{\text{kWh}}$
$\bar{T}_s$	{90, 90, 78, 48, 5} °C	M	150
T_{gw}	15 °C		

Table 4.1: Listing of the input parameters for the modelling of the Ecovat system. *The top 3 segments have a length of 3.3 m and the bottom 2 segments a length of 2.9 m, the diameter of the modelled buffer is 20 m and the density of water is taken as $1000 \frac{kg}{m^3}$.

simpler model, taking an average COP over the entire year does not perform significantly worse than using a non-linear model to determine the COP.

4.1.2 Results

In this section we present the results of the rolling horizon approach applied to the ILP model described in Chapter 3 combined with the input profiles and parameters presented in this chapter. We look at different scenarios, based on the demand temperature T_{dem}, and the rolling horizon length. As mentioned, we consider two different values for the demand temperature, $T_{dem} = 40°C$ and $T_{dem} = 60°C$. For the horizon length we consider lengths of two, three and five days. We do not consider horizon lengths longer than five days, because for those the computational time becomes too long. For all considered horizon lengths only the first day is realised for each step of the optimization as discussed in Section 4.1. The performance of the model in the presented cases is evaluated based on two metrics. First and most importantly, the objective value of the

ILP model is used, which gives the total operation costs over the time horizon. The second metric is the amount of useful energy that is contained in the Ecovat buffer at the end of the time horizon. We define useful energy to be the energy that can be used to satisfy the heat demand, i.e. energy at temperatures higher than the given demand temperature T_{dem}. This means the useful energy, U_i at the start of time interval i is given by:

$$U_i = \sum_{s \in S} m_s \; c_p \; \max\{T_{i,s} - T_{dem}, 0\}, \tag{4.1}$$

where m_s is the mass of segment s, c_p is the specific heat of water and $T_{i,s}$ is the temperature of segment s at the start of time interval i. Using equation (4.1) we can also obtain the maximum useful capacity $\bar{U}$ of the buffer by substituting the maximum temperature of the buffer segments $\bar{T}_s$ instead $T_{i,s}$. As we can see from (4.1) the useful energy content, and thus the useful energy capacity depends on the demand temperature. This indicates that it is not useful to consider higher demand temperatures, since in that case the useful energy capacity of the buffer is very low.

The useful energy content of the buffer at the end of the time horizon is an important quantity when the Ecovat system will continue to operate also after the considered time horizon, which will be the case for a real world system. In such a case starting the following year with a very low useful energy content would most likely lead to much higher operation costs during that year. In such a case more energy needs to be bought on the energy market to compensate for the low useful energy content, and thus higher energy prices might need to be accepted to charge the buffer enough to be able to satisfy the heat demand of that next year.

We first consider the results for the cases where $T_{dem} = 40$ °C. In Figure 4.5 the temperature evolution inside the Ecovat buffer throughout the year is shown for rolling horizon lengths of two, three and five days. The different coloured lines in the figure give the temperatures of the different segments in the buffer. As long as at least the top segment, i.e. the red line, is above the demand temperature, in this case 40 °C, the buffer has energy to supply the heat demand. The further a line is above the demand temperature the more useful energy that segment contains. Additionally, the lines in Figure 4.5 are not allowed to cross one another due to the constraint that a buffer segment must have a temperature at least as large as the buffer segment below it, to avoid situations which can not occur in practice.

We can see from Figure 4.5 that the temperature evolution, and thus the useful energy content of the buffer throughout the year is very similar for all three considered rolling horizon lengths. Furthermore, we can see that there is always more than enough useful energy content to be able to supply the heat demand. The objective value, the useful energy content of the buffer at the end of the time horizon and the optimization time for the different rolling horizon lengths

Hor. length (days)	Obj. value (€)	Useful energy (kWh)	Opt. time (s)
2	-28131	74420	6462
3	-28161	73578	12093
5	-28175	72711	36355

Table 4.2: The objective value, the useful energy content at the end of the time horizon and the optimization time for different rolling horizon lengths and $T_{dem} = 40$ °C.

are shown in Table 4.2. From this table we can see that the objective value gets slightly lower with increasing rolling horizon length, at a cost of slightly lower useful energy content at the end of the time horizon. This indicates that for a demand temperature of 40 °C the different rolling horizon lengths show similar performance. However, the optimization time increases significantly with increasing rolling horizon length. Due to this, we conclude that for a demand temperature of 40 °C using longer rolling horizon lengths is not really beneficial.

The second set of cases we consider are those where $T_{dem} = 60$ °C. The temperature evolution inside the Ecovat buffer throughout the year for rolling horizon lengths of two, three and five days is shown in Figure 4.6. We can see that for these cases there are extended periods of time during which the useful energy content of the buffer is almost zero. Since the heat demand must always be satisfied, a situation where the buffer is almost empty for extended periods of time is undesirable. The reason for this is that such a situation may lead to increased costs if no time intervals with low energy prices are available during that period of time. In this case the model is forced to buy energy at a higher energy price to be able to satisfy the heat demand.

In Table 4.3 the objective function, the useful energy content at the end of the time horizon and the optimization time for the different rolling horizon lengths is shown. We can see that the objective values are higher than in the scenarios where the demand temperature is 40 °C, which is expected because the useful capacity of the buffer is lower. This means that less useful energy can be stored during summer and more charging has to be done during the winter months, which in general have less time intervals with low energy prices. We see the same trend as for the cases where a demand temperature of 40 °C was considered, namely that longer horizon lengths lead to lower objective values, but the differences being very small. Additionally, we can see that the useful energy content of the buffer at the end of the time horizon is close to zero. This means that during the next year of operation the costs are expected to be much higher, as discussed previously. Furthermore, we see that the optimization times for these cases are much longer compared to the cases with a demand temperature of 40 °C, requiring around 10 to 12.5 days of optimization time.
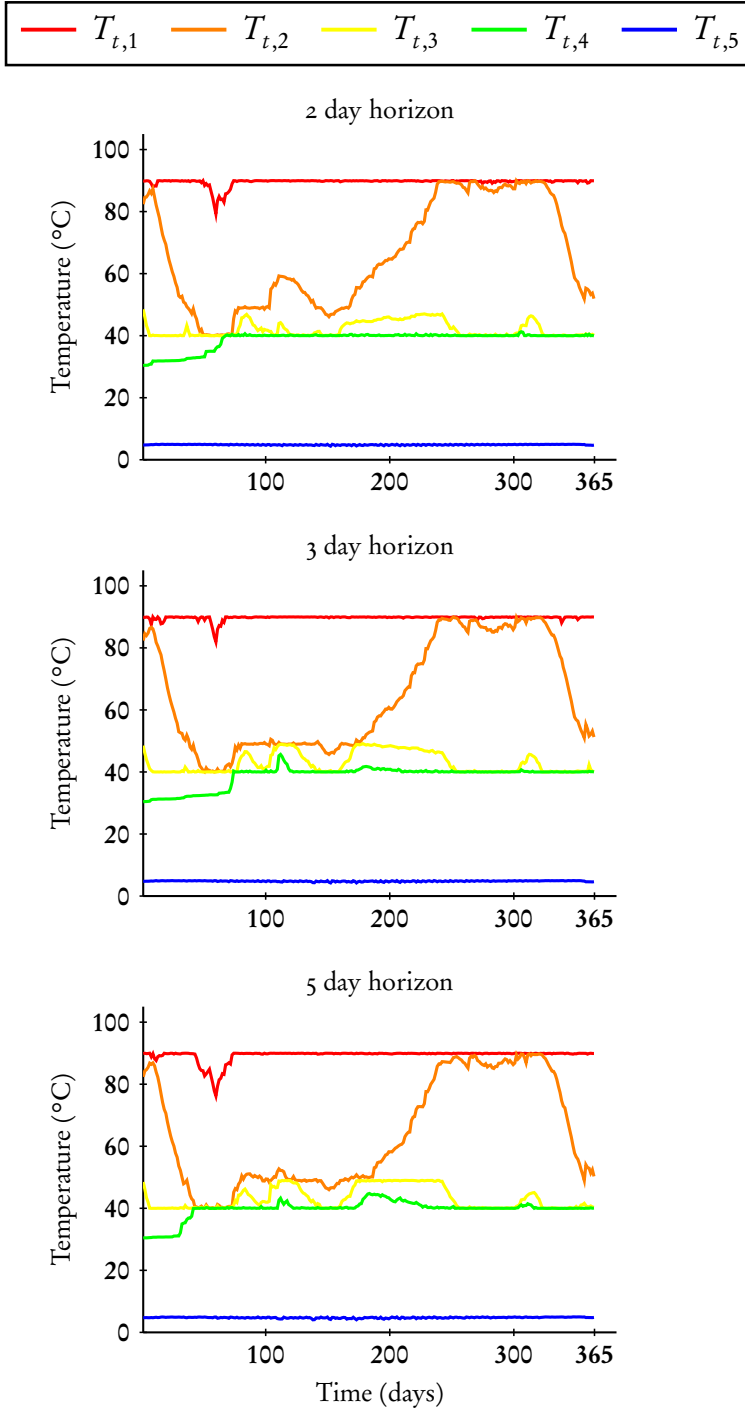
Figure 4.5: Temperature evolution inside the Ecovat buffer for different rolling horizon lengths in the ILP model for $T_{dem} = 40$ °C.
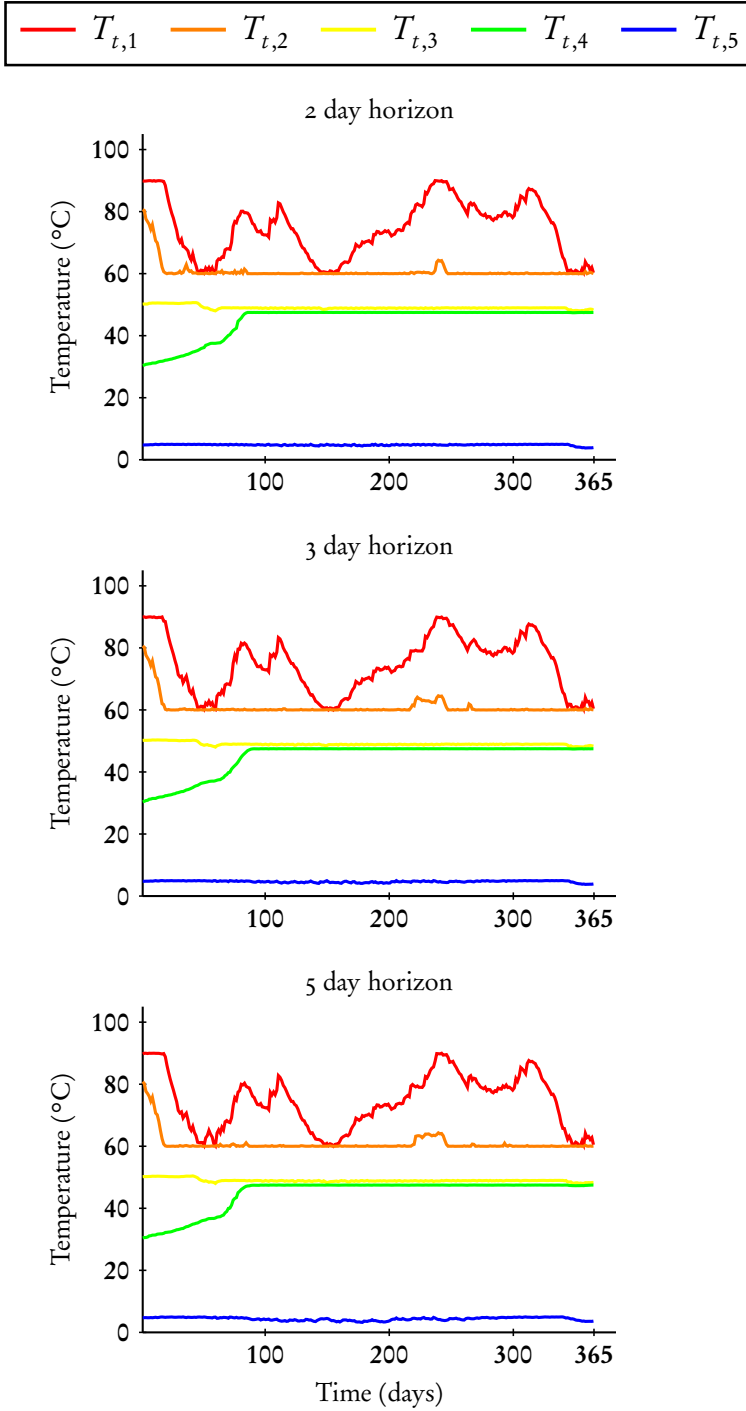
Figure 4.6: Temperature evolution inside the Ecovat buffer for different rolling horizon lengths in the ILP model for $T_{dem} = 60$ °C.

Hor. length (days)	Obj. value (€)	Useful energy (kWh)	Opt. time (s)
2	-27694	291	853646
3	-27743	272	989700
5	-27804	312	1072847

Table 4.3: The objective value, the useful energy content at the end of the time horizon and the optimization time for different rolling horizon lengths $T_{dem} = 60°C$.

From the results presented in this section we observed that the ILP model presented in Chapter 3 in combination with the rolling horizon approach introduced in this section performs well when a demand temperature of 40 °C is requested. In that case the buffer always has more than enough energy to supply the heat demand throughout the year. However, when the demand temperature is increased to 60 °C the performance of the model decreases significantly. In this case long time periods are observed where the useful energy content of the buffer is low. This may lead to severely increased costs if a period of high energy prices would come up at such a point in time. In this case the model would not have the option to delay charging of the buffer due to the fact that the heat demand must be satisfied at all times. The reason that such situations of low useful energy content occur is due to the fact that the ILP model solved in a rolling horizon fashion can only look a few days ahead, i.e. the rolling horizon length, for future problems and opportunities. In other words, seasonal effects are not properly taken into account in the ILP model yet, since the model can not look far enough ahead to do so. This leads to a situation where the buffer is almost empty long before the winter months are over. To include seasonal effects using the approach presented here, we need to consider a horizon length of a few months, which is not possible due to the already long computational time observed for the short horizon lengths of two to five days. Therefore, in the next section we extend the ILP model with a long-term planning, which allows it to properly take seasonal effects into account without requiring very long computational times.

4.2 Long-term planning

In the previous section it got clear that the ILP model described in Chapter 3 when solved using the introduced rolling horizon approach is not able to incorporate seasonal effects well. This shortcoming was especially visible for the case in which a higher temperature (60 °C) was demanded. In that case, the buffer was not sufficiently charged to deal with time periods of high heat demand. This meant that the energy content of the buffer stayed very low for a large part of the year, which may lead to higher costs for satisfying the heat demand. This

issue may be prevented if we extend the model with a method that is able to plan further ahead. As such, we present an extension to the ILP model based approach in this section that adds a long-term planning before solving the ILP model in rolling horizon fashion. In this long-term planning daily targets for the useful energy content of the Ecovat buffer are generated based on predictions for the energy prices and heat demand. The ILP model described in Chapter 3 is then slightly modified, to make sure the useful energy content in the buffer stays close to these generated targets. The remainder of this section is built up as follows. In Section 4.2.1 we state the problem that needs to be solved to generate the mentioned daily useful energy targets, while we discuss the algorithm designed to solve this problem in Section 4.2.2. The input data used to solve the problem of generating these targets will be discussed in Section 4.2.4. Finally, the results of adding this long-term planning to the ILP model based approach are presented in Section 4.2.5.

4.2.1   PROBLEM DEFINITION

To overcome the problem of the ILP model being unable to incorporate seasonal effects, we generate, in advance, a target for the useful energy content of the Ecovat buffer at the end of every day over the complete time horizon (in the cases considered in this thesis, a year). These targets have to ensure that enough useful energy will be available throughout the year and that this energy is produced during beneficial times (in other words, preferably at times the energy price is negative). However, allowing targets that are close to 0 or the maximum capacity of the buffer may lead to problematic situations where the ILP model ends a day with a useful energy content close to one of these extremes (i.e. almost empty or almost full). In such a situation the ILP model only has flexibility in one direction, which is especially problematic if the useful energy content of the buffer is very close to 0 at the end of a day. If the useful energy content of the buffer at the end of a day is close to 0, and during the next day only time intervals with high energy prices are available, unnecessarily high operational costs may be incurred. As such, we introduce predetermined minimum and maximum values for the targets of the useful energy content of the buffer, C_{min} and C_{max}, that are close to 0 and the maximum useful energy capacity of the buffer respectively, but still allow for some flexibility when solving the ILP model. Furthermore, we allow flexibility in the model by not requiring that the ILP model exactly achieves a given target value for each day, but only that it stays close to the targets by way of an extra incentive term in the objective function. Additionally, we add a constraint that the target for the useful energy content of the buffer at the end of the time horizon is at least equal to the useful energy content at the start. This ensures that the useful energy content of the buffer at the end of the year will be, equal to or higher than the useful energy content of the buffer at the start of the year. In this way we make sure that the buffer is prepared for the next year of operation, which will be very important when controlling an Ecovat in a real world situation.

As we specify the targets for the useful energy content of the buffer only for one time interval of a day, we introduce a new set, which represents the days in the time horizon and is defined as: $\mathcal{J} = \{1, ..., 365\}$. Since we consider time intervals of length 15 minutes, a day contains 96 time intervals. This means day j contains the time intervals $\{(j-1) \cdot 96 + 1, ..., j \cdot 96\}$. We set a target for the useful energy of the buffer at the end of the last time interval of each day, or equivalently at the start of the first time interval of the next day. As such, we define the time interval at which we set the target for day j as: $\mathrm{tar}(j) = j \cdot 96 + 1$. As the input of the problem to generate such targets, we consider predictions for the energy prices and the heat demand during the optimization period. We use the same vectors $\mathbf{p}$ and $\mathbf{D}$ as in Chapter 3, which contain predictions for the energy prices and heat demand for each interval in the time horizon.

The output of the target generation problem are the targets for the useful energy content of the Ecovat buffer at the end of every day in the time horizon. These targets are obtained by determining during which time intervals the buffer is charged. The decision on whether to charge the buffer during a give time interval depends on the energy prices, the heat demand of the neighbourhood and an estimate for the amount of useful energy that can be stored during a time interval.

The decision variables of the problem are given by the vector $\mathbf{z} = (z_1, z_2, ..., z_{N_{int}})$, where z_i is the binary decision variable that determines if energy is stored in the Ecovat buffer during time interval i or not. Hereby, a value of one for z_i means energy is stored in the Ecovat buffer during time interval i, while a value of zero means no energy is stored during time interval i. If we decide to store energy during a time interval i (i.e. $z_i = 1$), the amount of useful energy that gets stored is a given estimated amount e_i. We consider two values for e_i, one value for time intervals in which the energy price is at most $0 \frac{\text{€}}{MWh}$, which we call e_-, and one value for time intervals in which the energy price is larger than $0 \frac{\text{€}}{MWh}$, which we call e_+. The difference between e_- and e_+ lies in the fact that for simplicity, we assume that the resistance heater accompanying the buffer is only used when the energy price is non-positive, while the heat pumps are also used during intervals with positive energy prices if required (concrete choices for these parameters are given in Section 4.2.4). Note that we do not create operational plans for the different devices when solving the target generation problem, but that we only decide if during a time interval the buffer is charged. The amount of useful energy that can be stored during time interval i is defined as:

$$e_i = \begin{cases} e_-, & \text{if } p_i \leqslant 0 \\ e_+, & \text{if } p_i > 0, \end{cases} \tag{4.2}$$

where e_- and e_+ are the mentioned given values.

The resulting problem to be solved is to minimize the total monetary energy cost under the given constraints, which we call the Target Generation (*TG*) problem:

Problem *TG*

$$\min_{z_i} \sum_i p_i \, z_i \, e_i, \tag{4.3a}$$

$$\text{s.t.} \quad V_i = V_{i-1} + e_i \, z_i - D_i \qquad\qquad \forall \, i \in \mathcal{I}, \tag{4.3b}$$

$$V_{N_{int}} \geqslant U_0, \tag{4.3c}$$

$$V_{\text{tar}(j)} \geqslant C_{min} \qquad\qquad \forall \, j \in \mathcal{J}, \tag{4.3d}$$

$$V_{\text{tar}(j)} \leqslant C_{max} \qquad\qquad \forall \, j \in \mathcal{J}, \tag{4.3e}$$

$$V_0 = U_0 \tag{4.3f}$$

$$z_i \in \{0,1\}, \tag{4.3g}$$

where U_0 is the initial useful energy content of the Ecovat buffer and V_i is the target for the useful energy content of the buffer at the start of interval i. Constraint (4.3b) describes the change in the amount of useful energy due to energy being stored in the buffer, and heat demand being satisfied. Constraint (4.3c) makes sure that the target for the useful energy content of the buffer at the end of the time horizon is at least as high as at the start of the optimization. Finally, Constraints (4.3d) and (4.3e) ensure that the targets for the amount of useful energy stay within the specified minimum and maximum bounds, respectively.

It should be noted that we neglected heat losses to the surroundings of the buffer in Problem *TG*. Incorporating these losses in Problem *TG* would change the results, as due to the losses, some extra intervals have to be selected to store energy to compensate for these losses. However, we estimate the difference in results to be small, assuming that the time intervals with low energy prices are somewhat evenly distributed over the entire year as they seem to be at this point in time (this could change in the future if prices become more volatile due to an increased share of renewables). Note, that even if we assume that the time intervals with low energy prices are somewhat evenly distributed across the year, i.e. the charging of the buffer can be somewhat evenly distributed across the year, the heat demand is not distributed evenly across the year. As such, the situation in which the time intervals with low energy prices are evenly distributed across the year, does not diminish the benefits of long-term energy storage.

By solving Problem *TG* for given vectors of energy prices **p** and heat demand **D**, we obtain targets $V_{\text{tar}(j)}$ for the energy content of the Ecovat buffer throughout the year. Even though by solving Problem *TG* a target is generated for each time interval, only the targets $V_{\text{tar}(j)}$ at the end of each day are used in the ILP model. This ensures the ILP model has some flexibility to deal with unexpected circumstances, such as prediction errors in the energy prices. In other words, by solving

Problem *TG* we obtain useful energy targets for the buffer for every day in the time horizon, which determine how much energy needs to be charged during each day. These targets subsequently will be used as input for the ILP model described in Chapter 3, which will then determine the charging/discharging strategy of the Ecovat buffer over the year. This charging/discharging strategy gives the specific operational plans for the different devices in the system. For a simulation of a year, Problem *TG* only needs to be solved once to obtain the targets for the entire year.

4.2.2   IMPLEMENTATION

Although Problem *TG* has a much smaller size than the ILP presented in Chapter 3, it still is an ILP problem with around 35,000 integer variables, which leads to long computational times if a general ILP solver is used to solve the problem to optimality. Therefore, we decided to develop a fast heuristic algorithm to solve Problem *TG*.

The heuristic follows a 'greedy' principle to decide which time intervals are used to provide energy to the buffer. To describe this principle, we first rewrite Constraints (4.3b), (4.3d) and (4.3e) as:

$$C_{max} + \sum_{i=1}^{j} D_i \geqslant U_0 + \sum_{i=1}^{j} e_i z_i \geqslant \sum_{i=1}^{j} D_i + C_{min} \quad \forall\, j \in \mathcal{J}. \qquad (4.4)$$

This constraint tells us that the sum of the useful energy we store up to time interval j plus the useful energy we initially start with has a lower bound equal to the sum of the demand up to time interval j plus the minimum allowed target value C_{min}, and an upper bound equal to the sum of the demand up to time interval j plus the maximum allowed target value C_{max}. For simplicity, we rename these as the lower bound (LB_j) and upper bound (UB_j), respectively:

$$UB_j \geqslant U_0 + \sum_{i=1}^{j} e_i z_i \geqslant LB_j \quad \forall\, j \in \mathcal{J}. \qquad (4.5)$$

Similarly, we rewrite Constraint (4.3c) by defining:

$$LB_{N_{int}} := \sum_{i=1}^{N_{int}} D_i + U_0 \,. \qquad (4.6)$$

Equations (4.5) and (4.6) specify the bounds on the solution of Problem *TG*. These bounds are used in our greedy algorithm designed to solve the problem.

Our algorithm starts by setting all decision variables, z_i, to zero. It then checks at which time interval j the first violation (which has to be a violation of the lower
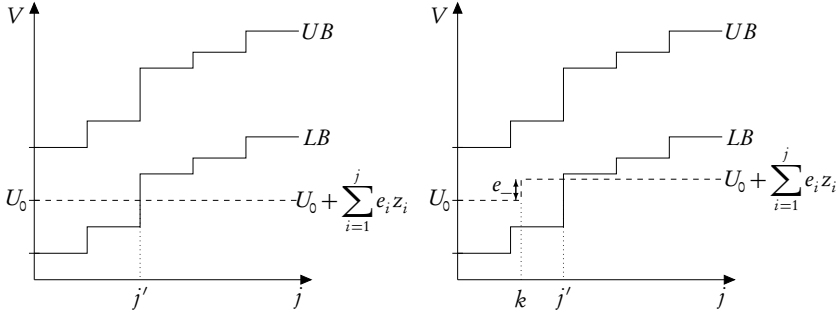
Figure 4.7: On the left, the algorithm starts with an initial (infeasible) solution with $z_i = 0 \; \forall \; i \in \mathcal{I}$. j' is the first interval violating the lower bound (LB) and k is the interval with the lowest energy price in $\{1, ..., j'\}$ (in this case, we have $p_k < 0$). On the right, the change after setting $z_k = 1$ is shown

bound) occurs; we call this time interval j'. The algorithm then determines the time interval in $\{1, ..., j'\}$ with the lowest energy cost, we call this interval k. If setting $z_k = 1$ does not violate the upper bound, we update the solution with $z_k = 1$; this process is depicted in Figure 4.7. The algorithm then repeats this process until the violation of the lower bound at time interval j' has been resolved. Subsequently, the algorithm checks where the next violation (if any) of the lower bound occurs and resolves it in the same way. This is repeated until all violations of the lower bound have been resolved.

To avoid checking the same interval multiple times, every interval is assigned a flag f_i that designates whether adding this interval to the solution does not violate the upper bound. At the start, all the flags f_i are set to true, denoting that all intervals can be added without violating the upper bound. While looking for the interval k with the lowest price, only intervals with $f_i = \text{true}$ are considered. If it is concluded that setting $z_k = 1$ would violate the upper bound, it is instead kept at $z_k = 0$, and the flag f_k is set to false. Furthermore, we know that if adding time interval k to the solution violates the upper bound, that all other time intervals before k, with the same or a larger amount of energy to be stored (i.e. time intervals for which the amount of energy e_i that can be stored is larger or equal to e_k), will also violate the upper bound. This means that we set all flags f_i to false for time intervals $i < k$ for which the energy amount e_i during time interval i is larger than or equal to the energy amount during time interval k, e_k.

When all the violations of the lower bound have been resolved, the algorithm has found a feasible solution. However, it may be possible to find a better solution, since there may still be time intervals with negative energy prices that have not been used and can be added to the solution without violating the upper bound, thus decreasing the objective value given by (4.3a). For this reason, the

algorithm checks if there are any time intervals left that have $p_i \leqslant 0$, $z_i = 0$ and $f_i =$ true. If such time intervals exist they are added to the solution in order of ascending energy price, as long as they do not violate the upper bound, until no such time intervals remain. The entire algorithm is summarized in pseudocode in Algorithm 1. It should be noted that similar problems are discussed in the literature; for example, [87] discusses Problem *TG* with continuous instead of binary decision variables.

Algorithm 1: Determine daily target values.

1 **input**   : $p_i, D_i, e_i, C_{min}, C_{max}, U_0$
2 **output** : U_j

3 set $z_i = 0$ and $f_i =$ true $\forall\ i \in I$
4 **while** solution is infeasible **do**
5   determine first violation of *LB*
6   let j' be the corresponding interval
7   let I' be the set of intervals $i = \{1, ..., j'\}$
8   **while** *LB* is violated at interval j' **do**
9   let k be the interval in I' with the lowest p_i while satisfying $z_k = 0$ and $f_k =$ true
10   **if** setting $z_k = 1$ does not violate *UB* **then**
11   set $z_k = 1$ and update current solution
12   **else**
13   set $f_k =$ false
14   set $f_i =$ false for all intervals $i < k$ that satisfy $e_i \geqslant e_k$

15 **while** there exists an interval $i \in \mathcal{I}$ with $p_i \leqslant 0$, $z_i = 0$ and $f_i =$ true **do**
16   let k be the interval in I with the lowest p_i while satisfying
17    $z_k = 0$ and $f_k =$ true
18   **if** setting $z_k = 1$ does not violate *UB* **then**
19   set $z_k = 1$ and update current solution
20   **else**
21   set $f_k =$ false
22   set $f_i =$ false for all intervals $i < k$ that satisfy $e_i \geqslant e_k$

If this greedy algorithm would be adapted for the LP relaxation of Problem *TG* (i.e. if Constraint (4.3g) is relaxed so that z_i can take a fractional value), it would provide an optimal solution [87]. This implies that the error of our greedy algorithm is somehow related to the integrality gap of Problem *TG*. Furthermore, it is easy to see (by an interchange argument) that the algorithm gives the optimal solution to Problem *TG* if $e_- = e_+$. However, in our case, where $e_- \neq e_+$, there are specific cases imaginable where our algorithm does not give the optimal solution. However, based on the relation with the integrality gap

and the specific structure of the problem, we expect that for the cases considered here, the solution will be very close to optimal.

4.2.3   INCORPORATION OF THE TARGET VALUES INTO ILP MODEL

To incorporate the targets $V_{\text{tar}(j)}$ resulting from the solution of Problem *TG* into the ILP model, an additional term that penalizes the solution for being under the target $V_{\text{tar}(j)}$ at the end of a day is added to the objective function. This term is given by: $-c_j(U_{\text{tar}(j)} - V_{\text{tar}(j)})$, which ensures the further the useful energy content of the buffer $U_{\text{tar}(j)}$ is below the target at the end of the day $V_{\text{tar}(j)}$, the more the solution is penalized. The value of c_j is set at the start of each day j based on whether the target for the previous day is met, i.e. based on the value of $U_{\text{tar}(j-1)} - V_{\text{tar}(j-1)}$. The further away from the target $V_{\text{tar}(j-1)}$ the useful energy content $U_{\text{tar}(j-1)}$ was at the end of the previous day the higher the value of c_j is chosen. The specifics of setting the value for c_j are described in Section 4.2.4. Based on this way of incorporating the targets into the model, the value of c_j determines the energy price the model is willing to accept during day j to charge the buffer. As mentioned, the further the useful energy content of the buffer gets below the target, the higher the value of c_j becomes, and thus the higher the costs are the model is willing to accept to increase the useful energy content of the buffer. The objective function, including this new term then becomes:

$$\min \sum_{i \in \mathcal{I}} \left(K_i - c_1 \sum_{s \in \mathcal{S}} (6-s) \, T_{i,s} - c_2 \, Q_{i,s}^{pvt} \right) - \sum_{j \in \mathcal{J}} c_j (U_{\text{tar}(j)} - V_{\text{tar}(j)}), \quad (4.7)$$

By incorporating long-term planning, via daily useful energy targets, into the model long-term trends such as seasonal effects are taken into consideration. As such, long rolling horizon lengths are no longer required to incorporate such effects, contrary to what we observed in Section 4.1, where no long-term planning was included. Note, that in the approach used in Section 4.1 the optimization considered k days (the rolling horizon length), but only the first of those k days was realized. As mentioned previously, without long-term planning k needs to be on the order of a few months to properly include seasonal effects. However, once long-term planning is included by the introduced targets much lower values of k may be used since the long-term effects are already taken into account by means of these pre-generated targets. In our approach the value of c_j, which determines the energy prices the model is willing to accept to charge the buffer, is set on a daily basis. This means that if we take $k > 1$ we either need to somehow estimate the values of c_j for the days we optimize over, or we take the same value of c_j for all days within the optimization. The problem with the first option is that it is hard to estimate this value for each day within the optimization, since c_j is based on the amount of useful energy content of the buffer at the end of

day $j-1$, which in turn is based on the heat demand and the energy prices, which are hard to predict on beforehand. The problem with the second option is that in this case the realization of the first day is based on a constant value of c_j for the entire optimization period. This means that the decisions made in the first day are based on the information that c_j remains constant for the entire optimization period, while in reality the value for c_j might change after realizing the first day. Note that, due to the fact that the value of c_j is determined in a non-linear way (we refer to Section 4.2.4 for more details) we can integrate the change of this value in a linear optimization over multiple days. Additionally, since future days are already taken into account through the long-term planning, we estimate that the difference in performance between taking a rolling horizon length of $k=1$ or a value k of a few days is small. As such, we reduce the rolling horizon approach to its 'extreme' case, where we optimize the behaviour of the Ecovat system over only one day and subsequently realise that day, i.e. the horizon length is equal to the period realised, i.e. $k=1$.

4.2.4 Simulation setup

The input parameters and profiles for the weather data and heat demand we use to evaluate the impact of adding long-term planning to the ILP model are the same as those in Section 4.1.1, i.e. those given in Figures 4.2 and 4.3 and Table 4.1. In this section we consider the energy prices for 2011, 2013, 2014 and 2015 to investigate the effect of different energy prices on the model. We have omitted 2012 from this analysis, since we wanted to avoid normalization issues because 2012 was a leap year.

If we want to apply the developed model in practice an additional complication arises, namely that the energy prices on the imbalance market are not known on beforehand and are generally hard to predict. Note that, whereas a good estimate for the expected heat demand can be made based on historical data and weather predictions, this is much harder for future energy prices. For this reason, we also investigate the impact of errors in energy price predictions on the achieved results. To do so, we consider two cases. In the first case, we use perfect predictions (PP) to generate the daily target values for the useful energy content of the Ecovat buffer. In this case, we assume that all prices of the entire year are known to us beforehand. In the second case, we consider the other extreme, namely that we have no predictions (NP) available to generate the target values. This implies that the intervals in which energy needs to be stored get distributed equally over the entire year, since we have no reason to favour one period in time over another in this case. Note, that we do not even include seasonal information concerning the energy prices in this case. In the NP case, we therefore do not require any input data for the energy prices. All the other input data used are the same for both the PP and NP cases.

As discussed previously, the weight c_j for deviating from the targets $V_{\text{tar}(j)}$ in the additional term added to the objective function of the ILP model, is varied

from day to day based on the difference between the target and the useful energy content of the buffer at the end of the previous day ($V_{\mathrm{tar}(j-1)}$ and $U_{\mathrm{tar}(j-1)}$ respectively). We distinguish between three different cases; 1) the useful energy content of the buffer is below the target, 2) the useful energy content of the buffer is above the target and 3) the buffer is almost full, i.e. the useful energy content is close to the useful capacity of the buffer. The case for an almost full buffer is added to ensure the model saves some of its useful capacity for any very low energy prices that may come up in the future, instead of filling the entire buffer as soon as it can. The value of c_j is determined by:

$$c_j = a_1 \left(\bar{U} - a_2 - U_{\mathrm{tar}(j-1)} \right) \qquad \text{if } U_{\mathrm{tar}(j-1)} > \bar{U} - a_2 \tag{4.8a}$$

$$c_j = 0 \qquad \text{if } U_{\mathrm{tar}(j-1)} \geqslant V_{\mathrm{tar}(j-1)}, \tag{4.8b}$$

$$c_j = \left[a_3 \left(1 - \frac{U_{\mathrm{tar}(j-1)}}{V_{\mathrm{tar}(j-1)}} \right) \right]^2 + a_4 \qquad \text{if } U_{\mathrm{tar}(j-1)} < V_{\mathrm{tar}(j-1)}, \tag{4.8c}$$

where a_1, a_2, a_3 and a_4 are constants. The range of possible values for c_j, which equates to the maximum energy price the model is willing to accept for charging the Ecovat buffer, is based on the values of these constants. In the case that the useful energy content of the buffer is below the target (Equation 4.8c) the constant c_j varies between $a_4 \frac{€}{kWh}$, when the useful energy content is just below the target, and $a_3 + a_4 \frac{€}{kWh}$ if the useful energy content of the buffer is at zero. This increase in c_j is non-linear, in this way the further the useful energy content of the buffer drops below the target the more quickly the costs increase to ensure that the model stays close to the target, while still allowing some flexibility. When the useful energy content of the buffer is above the target (Equation 4.8b) c_j is set to 0, i.e. a maximum energy price of $0 \frac{€}{kWh}$ is accepted. Finally, when the buffer is almost full (Equation 4.8a) c_j in in the interval $[-a_1 \cdot a_2, 0] \frac{€}{kWh}$, depending on how close the useful energy content of the buffer is to its useful capacity, i.e. maximum energy prices between $-a_1 \cdot a_2 \frac{€}{kWh}$ (0 kWh capacity remaining) and $0 \frac{€}{kWh}$ (a_2 kWh capacity remaining) are accepted by the model. The condition for Equation 4.8a can be true at the same time as one of the conditions for the other two cases. In this case Equation 4.8a takes priority over the other case.

The used values for the constants in Equations (4.8a), (4.8b) and (4.8c) used are given in Table 4.4. These values have been determined experimentally to work well. However, when controlling a real Ecovat system, these values may be changed depending on the circumstances during the year, such as unexpected changes in heat demand or energy prices. For example, if the winter months are expected to be colder than usual, and thus there is a higher heat demand during those months, it could be desirable to accept higher energy prices than

Constant	Value
a_1	$10^{-5} \frac{\text{€}}{kWh^2}$
a_2	$15000 kWh$
a_3	$0.241 \frac{\text{€}}{kWh^2}$
a_4	$0.009 \frac{\text{€}}{kWh}$

Table 4.4: The values of the constants in Equations (4.8a), (4.8b) and (4.8c) used in this work.

one would normally accept during autumn to mitigate the risk of ending up with too little useful energy in the buffer during winter, which could lead to even higher costs.

The final parameters that we need to specify are the minimum and maximum values for the targets C_{min} and C_{max} in Problem TG and the values of e_- and e_+ in Equation 4.2, which give estimates for the amount of energy that can be charged to the Ecovat buffer during a time interval with a negative energy price and a time interval with positive energy price respectively. The most straightforward way would be to choose the minimum and maximum values for the targets, C_{min} and C_{max}, at 0 kWh and the buffers useful capacity, respectively. However, as discussed previously, this is undesirable since targets at these extremes leave very little flexibility when solving the ILP model. As such, we choose slightly different values to allow for a 'safety margin'.

We wish to retain enough flexibility in the solving of the ILP model, such that when a period with a very low number of time intervals with low energy prices occurs, the model is not forced to charge the buffer during time intervals with very high energy prices. As such we choose the minimum target value C_{min} such that two days of heat demand can still be satisfied without any charging, which in the cases considered in this work corresponds to $C_{min} = 5000$ kWh. This value was found to work sufficiently for the cases considered in this work, since throughout the year, there are always some intervals with low energy prices when considering periods of a few days. However, if long periods of time with only a few or no intervals with low energy prices are expected, a higher minimum capacity would have to be used to avoid unnecessarily high costs. This is especially important for the no predictions (NP) case, where such an expectation is not taken into account specifically (the targets are generated by distributing the amount of time intervals during which the buffer is charged, to be able satisfy the heat demand throughout the year, equally over the days in a year in this case). Even though we assume in that case that no predictions are available for the energy prices, we may still have the more general expectation that during winter, a longer period without low energy price time intervals might occur. In that case, taking a higher value for C_{min} may be advisable to avoid potentially high costs.

Table 4.5: Table showing values of e_- and e_+.

	e_- (kWh)	e_+ (kWh)
Low estimate	256	6
Medium estimate	262	12
High estimate	268	18

Similarly, to ensure the model has some flexibility to accept very low energy prices when the buffer is almost full we set C_{max} equal to 95% of the actual maximum useful energy capacity of the Ecovat buffer. Since the maximum useful capacity of the buffer depends on the demand temperature we have chosen for this case to save 5% of the maximum useful capacity for this flexibility instead of a constant value.

The values for the estimates of the useful energy storable during an interval at a negative or positive energy price, e_- and e_+ respectively, depend on the average amount of useful energy generated by the heat pumps and resistance heater. To estimate these values we make the simplifying assumption that the resistance heater is only used to store energy when the energy price is non-positive. Note, that this is not necessarily true since depending on the values of c_j the resistance heater may also be used when the energy price is positive. We evaluate the effect this assumption has in Section 4.2.5. When considering e_-, the total estimate for the useful energy that can be charged during a time interval with negative energy price is a combination of the estimated amount of useful energy produced by the resistance heater, and by the heat pumps in the system during such a time interval. The amount of energy the resistance heater contributes to e_- is easy to determine because of its constant coefficient of performance of one. However, the contribution of the heat pumps is much harder to estimate. Their contribution depends on the demand temperature T_{dem}, the temperatures inside the buffer $T_{i,s}$, the temperature ranges of the heat pumps and their coefficients of performance. This implies that these values are dependent on the current state of the buffer and that therefore e_- and e_+ can only be rough estimates. For this reason, we also investigate the effect of various estimates of these parameters on the results of the model. To this end, we use three estimates for e_- and e_+, a low estimate, a medium estimate and a high estimate. These estimates are given in Table 4.5. For the results presented in Section 4.2.5, the medium estimate is used, unless specified differently. The differences in the estimates for e_- are relatively small because the contribution of the resistance heater, which supplies by far the largest portion of the energy during intervals with negative energy prices, is easily estimated as mentioned before. On the other hand, the estimates for the heat pumps are much more uncertain, however, they supply a much smaller portion of the energy during such times, leading to small relative differences.

4.2.5 RESULTS

By solving the target generation problem *TG* using the method described in Section 4.2.2, we obtain a target for the useful energy content of the Ecovat buffer at the end of every day within the time horizon. First, we investigate whether the approach we applied for the no predictions (NP) case, namely that we distribute the amount of time intervals in which storing energy is required evenly throughout the year, leads to target values that are close to those obtained from the perfect predictions (PP) case. In Figure 4.8, the resulting target values are shown for a demand temperature of 40 °C for the PP case when using real Dutch energy prices from the years 2011, 2013, 2014 and 2015, as well as for the NP case (which is the same for every year, since the heat demand in all considered cases is equal). The targets for the PP case using energy prices from 2013 or 2015 have a similar shape as the targets for the NP case. However, the targets for the PP case using energy prices from 2011 or 2014 are significantly different from the targets for the NP case. This means the choice to distribute the number of intervals in which energy needs to be stored equally over the year seems reasonable for the years 2013 and 2015, but leads to significantly different targets for the years 2011 and 2014 compared to the PP case. This is due to an uneven spread in energy prices in 2011 and 2014, while the NP case assumes a fairly even spread throughout the year.

Figure 4.9 shows the same type of results for the PP and NP cases, however, this time for a demand temperature of 60 °C. Similarly to Figure 4.8, the targets for the PP case using energy prices from 2013 or 2015 have a similar shape compared to the targets for the NP case, while the targets for the PP case using energy prices from 2011 or 2014 are significantly different. The largest difference between Figures 4.8 and 4.9 is that in the case with a demand temperature of 60 °C the buffer is almost completely depleted at the end of winter or during spring, except for the PP case using energy prices from 2014, while in Figure 4.8, this only occurs for the PP case using energy prices from 2011. The buffer being almost completely depleted after winter for most of the years when a demand temperature of 60 °C is used, is due to the fact that the energy demand stays the same as for the 40 °C case, but at the same time the initial useful energy content of the buffer is much lower, due to the higher demand temperature.

From the previous figures, we have seen that the PP and NP cases can lead to target values that differ significantly. Next, we investigate whether this difference in target values also leads to a significant difference in the results of the ILP model. Furthermore, we compare the results obtained by using the model extended with long-term planning, with the previously obtained results for the model that did not incorporate targets for the useful energy content of the buffer.

In Figure 4.10, the results of the optimization with a demand temperature of 40 °C without targets are compared with optimizations using the long-term planning. For the case without targets we use the results from the optimization with a horizon length of two days presented in Section 4.1.2. As already observed

Figure 4.8: Targets obtained by solving Problem *TG* for the perfect predictions (PP) case using the energy prices from 2011, 2013, 2014 and 2015 for a demand temperature of 40 °C. Additionaly, the targets obtained for the no predictions (NP) case are shown.

previously, the optimization without long-term planning does quite well in this case, as there is always enough energy to supply the requested heat demand. In this case, adding a long-term planning (both in the PP and NP cases) has the effect of a larger amount of useful energy being present in the buffer throughout the year, which is shown in the figure by the temperatures in the buffer being higher throughout the year. However, this leads to a slight increase in the objective value of the ILP as shown in Table 4.6 as well. The higher objective value for the model that includes targets is compensated by the higher useful energy content of the buffer at the end of the optimization, due to the safer strategy used by adding these targets. Even though in this case, it would be possible to not include long-term planning in the model, and still obtain good results, we observe that the differences are small. The advantage of not including the long-term planning in cases such as this is a slightly lower objective value, i.e. slightly higher profit. However, the disadvantage is that if in such a case a time period without low energy prices may occur at a moment that the buffer is low on energy, meaning high costs may be incurred. Conversely, the advantage of including the long-term planning is that a time period without low energy prices will not lead to such an increase in costs, since the long-term planning ensure the buffer has enough
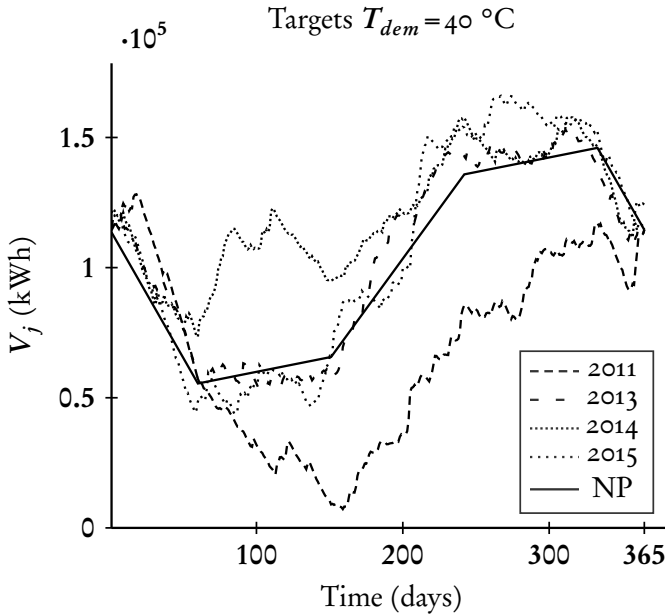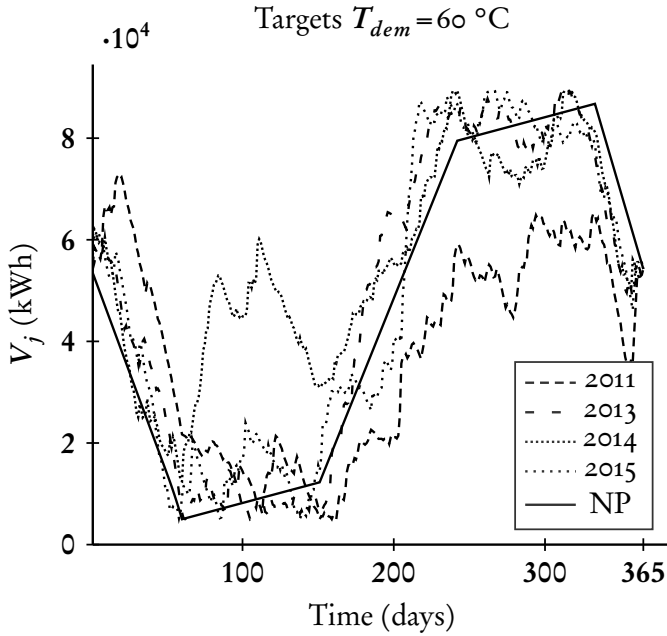
Figure 4.9: Targets obtained by solving Problem *TG* for the perfect predictions (PP) case using the energy prices from 2011, 2013, 2014 and 2015 for a demand temperature of 60 °C. Additionaly, the targets obtained for the no predictions (NP) case are shown.

energy stored before this time period occurs. The disadvantage of including the long-term planning is that in some cases, such as the one presented here, the profit may be slightly lower. This means there is a trade-off between a chance of higher profit at a higher risk. Since in this case the differences in profit are only small, we think it is advisable to incorporate targets in order to employ a safer charging strategy. Furthermore, when the long-term planning is incorporated lead to a higher useful energy content of the buffer at the end of the year. The monetary advantage of this extra useful energy is hard to estimate, since it depends heavily on the energy prices at the start of the next year. If high prices are expected the extra useful energy is very valuable, while if a lot of negative prices are expected the extra useful energy can even be a disadvantage, since this means there is less capacity remaining to charge at a profit. However, in general at the start of the year, i.e. during winter, the heat demand is high and time intervals with low energy prices not very frequent. This means that we consider a higher useful energy content of the buffer at the end of a year a benefit, although we are not able to specify exactly how large the benefit is.

It should be noted that even though there are some differences between the

Case	Obj. value (€)	Useful energy (kWh)	Opt. time (s)
No targets	-28131	74081	6462
PP	-27865	110370	28163
NP	-27776	110145	46632

Table 4.6: Table comparing optimizations using the energy price data from 2014 and $T_{dem} = 40°C$.

Case	Obj. value (€)	Useful energy (kWh)	Opt. time (s)
No targets	-27694	291	853646
PP	-27170	47894	39839
NP	-26409	57100	50325

Table 4.7: Table comparing optimizations using the energy price data from 2014 and $T_{dem} = 60°C$.

PP and NP cases in Figure 4.10, the differences in objective value and useful energy content at the end of the time horizon are very small, as shown in and Table 4.6 (less than 1% in objective value). We finally note that the incorporation of the long-term planning increases the computational time by a factor of 4 to 7 compared to the case that uses no targets and a horizon length of two days.

While the model without long-term planning performs well for the case with a demand temperature of 40 °C, we showed in Section 4.1.2 that this is no longer the case when the demand temperature is increased to 60 °C. In Figure 4.11, the results are shown for a demand temperature of 60 °C. In this case, the optimization without long-term planning performs very poorly as was already observed previously. The addition of long-term planning makes a large difference in this case. As can be seen in the figure, for these two cases there is always enough useful energy in the buffer to be able to supply the heat demand. Furthermore, the amount of useful energy in the buffer at the end of the optimization is a lot higher compared to the case when no long-term planning is used. Table 4.7 shows the objective values and the amount of useful energy in the buffer at the end of the optimization for the considered cases. While the objective value is slightly lower when incorporating the long-term planning, the useful energy content at the end of the year is much higher and the optimization times required are much shorter. When a demand temperature of 60 °C is used, it is clear that the addition of long-term planning to the ILP model based approach leads to a substantial improvement in the results. Again, we see that the differences between the PP and NP cases are small (less than 3% in objective value).

In Figures 4.10 and 4.11, there seems to be little difference in the results when comparing the PP and NP cases. To get some more insight regarding this be-
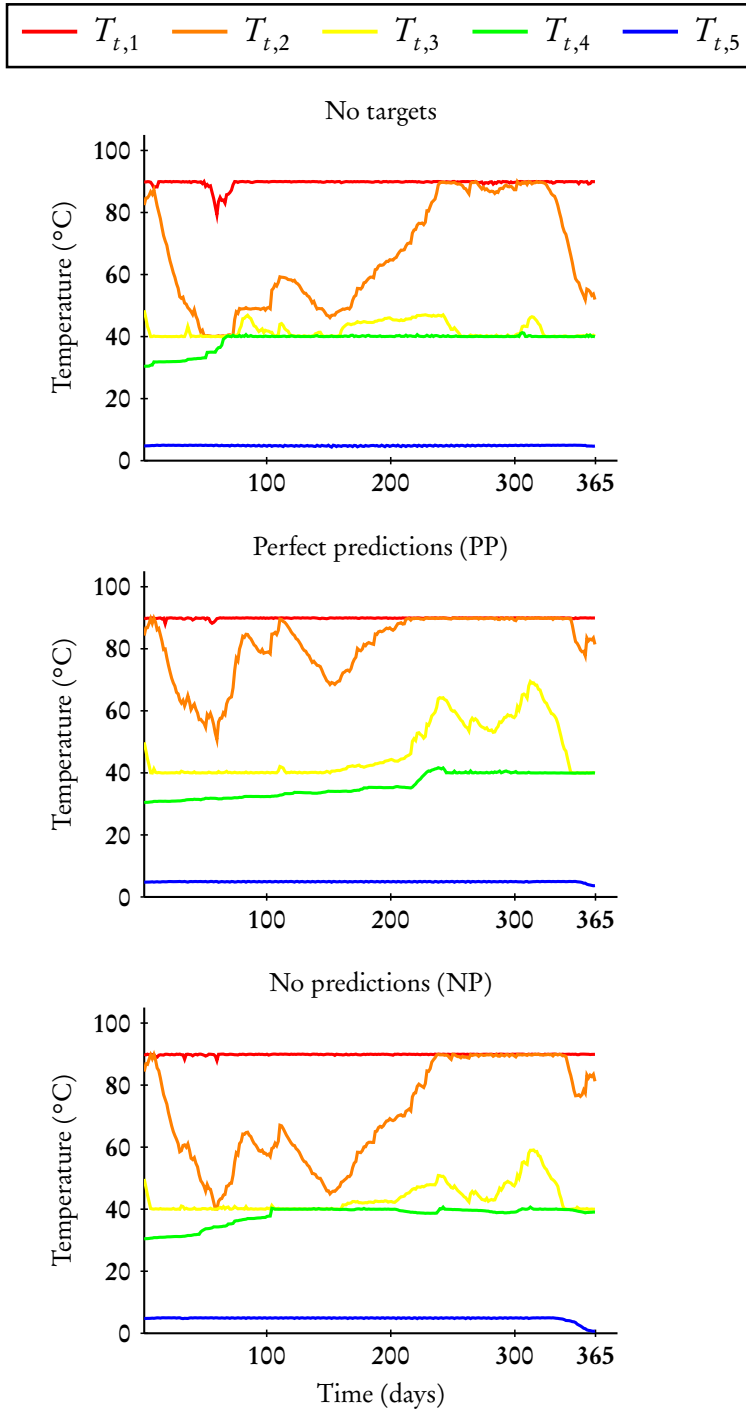
Figure 4.10: Temperature evolution inside the Ecovat buffer throughout the year without using targets, as well as the PP and NP cases when using targets. These results are for a demand temperature of 40 °C and energy prices from the year 2014.
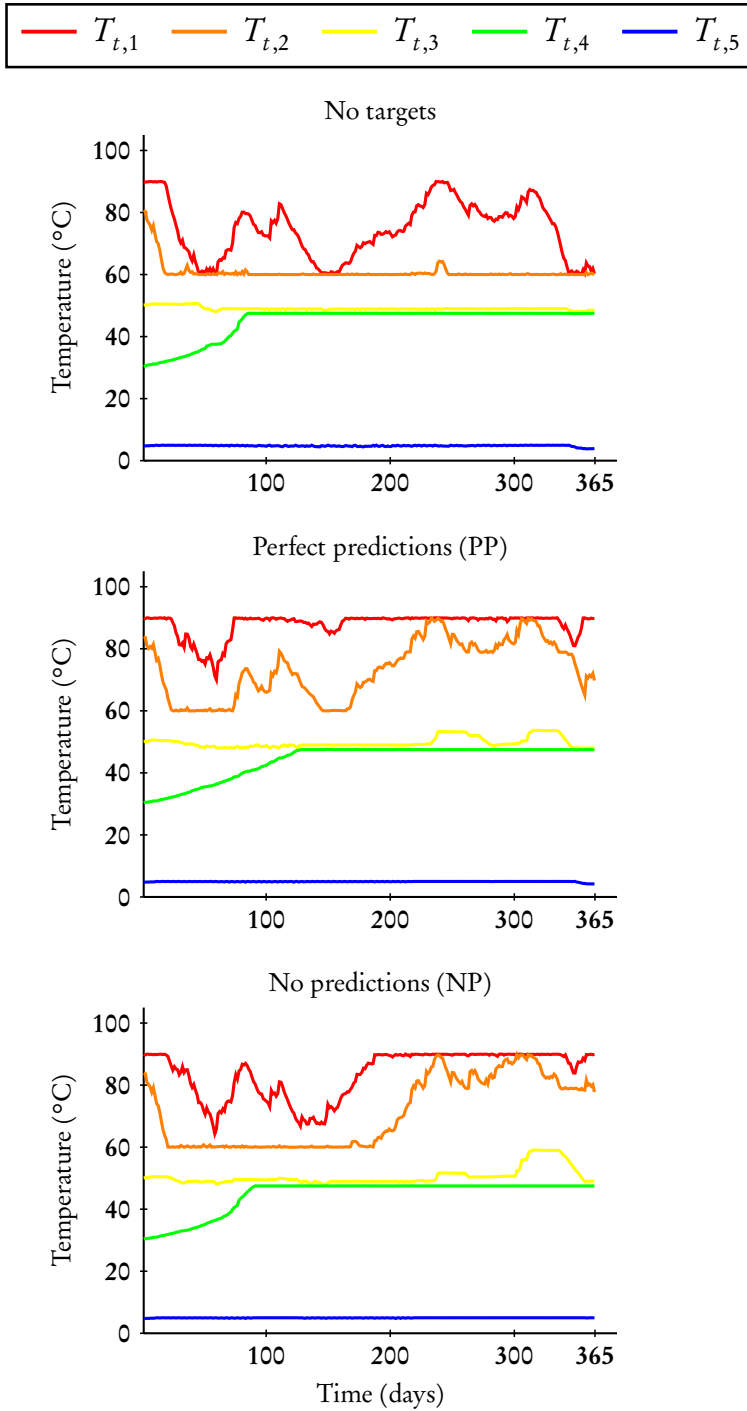
Figure 4.11: Temperature evolution inside the Ecovat buffer throughout the year without using targets, as well as the PP and NP cases when using targets. These results are for a demand temperature of 60 °C and energy prices from the year 2014.

Case	Obj. value (€)	Useful energy (kWh)	Opt. time (s)
$T_{dem} = 40°C$, 2011, PP	-22709	115696	35369
$T_{dem} = 40°C$, 2011, NP	-22097	145635	211282
$T_{dem} = 40°C$, 2013, PP	-41711	114694	28876
$T_{dem} = 40°C$, 2013, NP	-41609	117096	100755
$T_{dem} = 40°C$, 2014, PP	-27865	110370	28163
$T_{dem} = 40°C$, 2014, NP	-27776	110145	46632
$T_{dem} = 40°C$, 2015, PP	-50634	125215	2831
$T_{dem} = 40°C$, 2015, NP	-50528	125739	3018
$T_{dem} = 60°C$, 2011, PP	-22339	56187	79324
$T_{dem} = 60°C$, 2011, NP	-21164	78399	66952
$T_{dem} = 60°C$, 2013, PP	-41343	52309	35715
$T_{dem} = 60°C$, 2013, NP	-40681	59093	45735
$T_{dem} = 60°C$, 2014, PP	-27170	47894	39839
$T_{dem} = 60°C$, 2014, NP	-26409	57100	50325
$T_{dem} = 60°C$, 2015, PP	-52196	55796	39417
$T_{dem} = 60°C$, 2015, NP	-51640	61916	39239

Table 4.8: Table comparing the PP and NP cases for the years 2011, 2013, 2014 and 2015 as well as $T_{dem} = 40°C$ and $T_{dem} = 60°C$.

haviour, we investigate whether this is the same for the other years: 2011, 2013 and 2015. The corresponding results are summarized in Table 4.8. We can see that for all the years and both demand temperatures considered, the PP and NP cases give very similar results. The largest relative difference is observed for 2011, but even there, the difference in the objective value is only slightly higher than 5%. We observe that the NP case generally gives slightly higher objective values for the years and demand temperatures considered, but it also has a higher useful energy content at the end of the optimizations to compensate for this. Whether this trade-off is worthwhile depends on the expected energy prices occurring in the time intervals after the optimization horizon.

Finally, we investigated the influence of the used estimates for the amounts of energy that can be charged by the Ecovat buffer during a time interval with a negative energy price e_- and during a time interval with a positive energy price e_+ on the results. We again take the energy prices from 2014 and a demand temperature of 60 °C and compare the effect of different estimates for e_- and e_+. The results of this comparison are presented in Figure 4.12 and Table 4.9. As can be seen, the results are again very similar. The higher the estimates for e_- and e_+, the lower the objective value, but the lower the amount of useful energy in the buffer at the end of the optimization. We also performed this analysis using energy price data from the other years, yielding similar results to those shown in Figure 4.12 and Table 4.9.

We observe that different values for e_- and e_+, which lead to different sets of
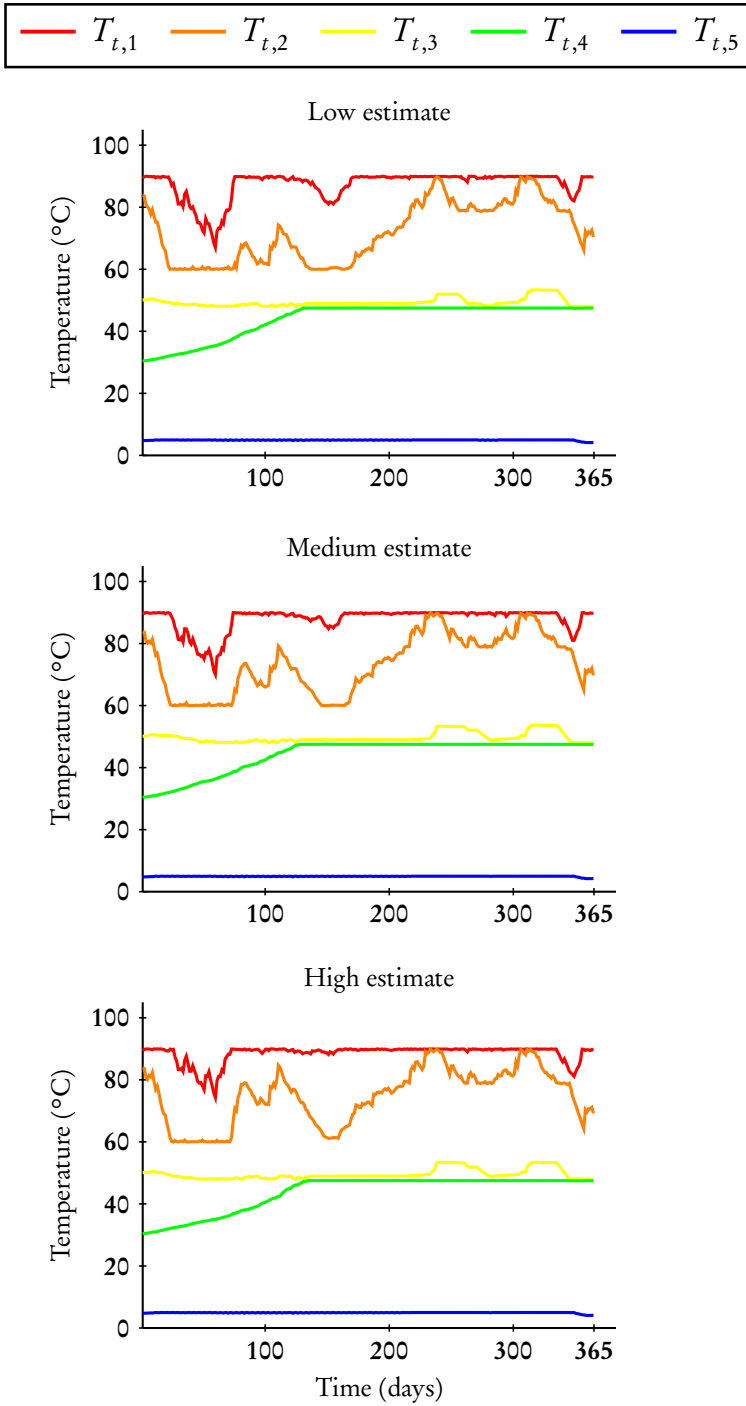
Figure 4.12: Temperature evolution inside the Ecovat buffer throughout the year for different estimates of e_- and e_+, using the energy price data from 2014 and $T_{dem} = 60°C$.

Estimate	Obj. value (€)	Useful energy (kWh)
Low estimate	-27132	48529
Medium estimate	-27170	47894
High estimate	-27249	46890

Table 4.9: Table comparing optimizations using the energy price data from 2014 and $T_{dem} = 60°C$ for different estimates of e_- and e_+.

target values, give very similar results when used as input for the ILP model. And more generally, we find that two quite different sets of target values still give very similar results when used as input for the ILP model. This leads us to believe that the simplifying assumption we made at the end of Section 4.2.4, namely that for determining e_-, we assume the resistance heater only runs when the energy price is non-positive, has a negligible effect on the results of the optimizations presented here.

The comparison between the ILP model with and without targets shows that adding long-term planning to the model improves the results. Furthermore, the results show that the model is only slightly influenced by differences in the prediction of the energy prices, as we observed that the difference between cases where perfect predictions were assumed, and cases where no predictions for the energy prices were assumed to be available, were small. This means that as long as the ILP model is given target values that somewhat follow the actual energy prices, or more generally, if they follow the expected seasonal behaviour (for example, a high target at the end of summer), the results of the ILP model will differ only slightly (by a few percent). However, if as a consequence of the energy transition, the energy prices will get more volatile in the future, this may change.

These results may seem surprising, but can be explained by taking a closer look at the energy prices. In Figure 4.13, the energy prices from 2014 are plotted after sorting them from lowest to highest. The figures shows that there are only a few very low energy prices, which the model will always try to use, but also only a few very high energy prices, which the model will always try to avoid. However, between those extremes, there are many energy prices that differ only slightly from each other. This means that even if the ILP model is not given the optimal targets, and as a consequence some time intervals with a higher energy price are unnecessarily selected, the difference in objective value will be very small.

4.3 CONCLUSION

In this chapter we first introduced a rolling horizon approach to be able to solve the ILP model presented in Chapter 3. We observed that this approach works well if the demand temperature is low, but fails once a higher temperature

Sorted energy prices (2014)
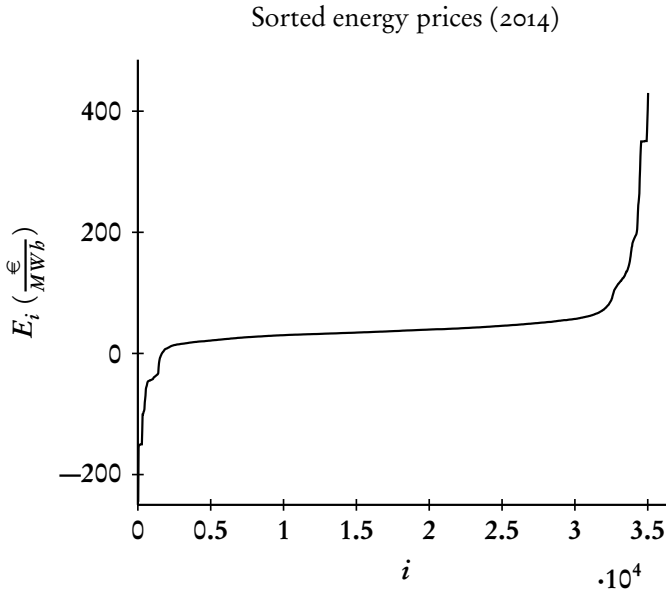
Figure 4.13: Energy prices from 2014 sorted from lowest to highest.

is demanded. Furthermore, we observed that the optimization time increased dramatically in that case. The reason that the developed rolling horizon approach fails (the energy content of the buffer is close to 0 for long periods of time) when considering a higher demand temperature is that the approach is incapable of looking far enough ahead to be able to incorporate seasonal effects. Due to the long optimization times required, even for horizon lengths of only a few days, it is impossible to extend the horizon far enough to be able to include such seasonal effects in that approach.

Subsequently, we presented a long-term planning step, to be executed before solving the ILP model. This step determines targets for the useful energy content of the Ecovat buffer for every day throughout the year. Based on this long-term planning the model is then capable of including seasonal effects properly. We showed that including long-term planning greatly improves the results. Furthermore, we showed that by adding long-term planning the model is very robust against prediction errors. The results obtained when assuming perfect predictions, and the results obtained when assuming we do not know anything about the energy prices on beforehand, show only small differences. Finally, we showed that the estimates of the amount of energy that can be stored in one time interval by the Ecovat buffer (i.e. the values for e_- and e_+), only have a small effect on the results of the ILP model based approach as well.

In conclusion, we observe that as long as the targets for the useful energy content of the Ecovat buffer that are used as input to the ILP model, are not too differ-

ent from reality, i.e. if they show the expected seasonal behaviour, the results
obtained from the ILP model are very similar.

Even though the ILP model extended with long-term planning provides good
charging strategies for the Ecovat system it is still computationally too expensive
to be used for the control of an Ecovat system in practice, for which very fast
decisions will be necessary. The ILP model in some cases requires 12 hours or
more to determine a strategy for a year, which means an average of almost 2 min-
utes (and in some extreme cases a much longer time) is required to determine the
strategy for a day, while for the control of an Ecovat system in practice decision
making in less than a second is needed. Additionally, shorter solution times
allow the integration of a model of the Ecovat system into DEM simulations.

For these reasons in the next chapter we present a heuristic method to control
the Ecovat system, which is based on the insights obtained by studying the ILP
model and its output data for different cases. In that chapter we present not only
this method, but also compare the results obtained with it to the results obtained
when using the ILP model.

5
Method for operational control in practice

Abstract – *In this chapter we present a heuristic method to control the operation of an Ecovat system. This heuristic method is based on the insights obtained by using the ILP model presented in the previous chapter. The goal of this heuristic method is to provide charging/discharging strategies that are of comparable quality to the strategies provided by the ILP model, while being fast enough to be usable in a practical setting, i.e. when controlling a real Ecovat system. To determine the quality of the strategies obtained by using the heuristic method we compare them against the strategies obtained with the ILP model.*

In Chapter 3 we presented an ILP model to determine charging/discharging strategies for an Ecovat system. For deriving that approach we did not take into account any restriction on the runtime of this model. The combination of the long time horizon and the short time interval length that are required for modelling the Ecovat system, make that with current solvers the model as it is presented in Chapter 3 can not be solved within a reasonable time. Subsequently, in Chapter 4 we first presented a rolling horizon approach, which makes the ILP model based approach feasible. However, as the results showed, this approach is insufficiently capable of incorporating important seasonal effects. To deal with this, we presented a long-term planning step, to be run before the rolling horizon approach, which determines a target for the useful energy content of the Ecovat buffer for every day. We then slightly modified the previously developed ILP model such that it ensures a useful energy content in the Ecovat buffer close to

---

This chapter is based on [G:3].

those pre-generated targets. We observed that this method is capable of incorporating seasonal effects and leads to an improvement in the results, compared to the rolling horizon approach without long-term planning. However, as observed in Chapter 4 the computational times for yearly simulations based on this model are still in the order of a few days for some cases. This means that the ILP model based approach is still not usable for the real-time control of an Ecovat system in practice, since in that case very fast decision making (around a second or less) is required.

In this chapter we present an approach of controlling the Ecovat system for which the computational times are short enough for use in a practical situation. We set up this approach such that it considers all time intervals iteratively, and only requires input for the current time interval it is considering, i.e. that it does not require input data for future time intervals, such as weather predictions and future energy prices. The exception to this are pre-generated targets for the useful energy content of the Ecovat buffer, which are generated in the same way as for the ILP model based approach in the previous chapter. The approach developed in this chapter is a heuristic method, which is developed based on the insights obtained from the ILP model based approach, as well as the output data from the ILP model based approach for the different considered cases in Chapter 4. More specifically, based on output data from the ILP model based approach we derive some 'rules of thumb' for the control of charging/discharging of the Ecovat system. The heuristic then applies these rules of thumb to each time interval in an iterative manner, to obtain a charging/discharging strategy for the Ecovat system. Clearly, this approach of applying some generalized rules of thumb should lead to results of lower quality compared to an optimization based approach, such as the presented ILP model based approach. Therefore, our goal is to develop a fast heuristic method that gives charging/discharging strategies which are at least of comparable quality to those obtained using the ILP model based approach. Summarizing, our goals for designing this heuristic method are:

1. Produce charging/discharging strategies that are still of good quality, i.e. close to the strategies determined by the ILP model based approach.

2. Produce such charging/discharging strategies in a short computational time (seconds).

3. Produce such charging/discharging strategies without requiring input data, such as energy prices or weather data, beyond the time interval that the heuristic method is considering at that moment.

The remainder of this chapter is structured as follows. In Section 5.1 we present the heuristic method derived from the insights obtained from the ILP model based approach. We describe the rules of thumb derived from the output data of the ILP model based approach, as well as the iterative process the heuristic uses to construct a charging/discharging strategy. In Section 5.2 we compare

the results obtained with the heuristic approach presented in this chapter to the previously presented results from the ILP model based approach. Finally, in Section 5.3 we present our conclusions with regards to the presented heuristic method.

5.1   Heuristic approach

Similar to the ILP model based approach, the heuristic first generates a target $V_{\text{tar}(j)}$ for the useful energy in the buffer at the end of every day j of the year. This is done in the same way as for the ILP model based approach, i.e. Algorithm 1 is used (see Section 4.2.2). After this, the heuristic iteratively considers every time interval in the time horizon and determines the decisions which have to be made for this time interval. This process consists of eight steps. We first summarize these steps below and go into more detail on each individual step after that. During each time interval i the following steps are performed:

1. If the time interval is the first time interval of a new day, determine the maximum accepted energy price for running the devices in the system for the coming day.

2. Calculate the expected output temperature of the PVT panels.

3. Decide if the PVT panels are connected to the bottom buffer segment, and whether or not the lthp and/or the hthp run during time interval i.

4. Decide from which segment the heat demand is supplied and if the resistance heater and/or awhp run during time interval i.

5. Calculate the heat losses of the Ecovat buffer to the surroundings during time interval i.

6. Calculate the changes in energy content in the system due to the decisions made in Steps 2–4 and the calculated heat loss from Step 5.

7. Update the temperatures in the segments of the Ecovat buffer.

8. Calculate the operational costs resulting from the decisions made in Steps 2–4.

After the last time interval in the time horizon the total operational costs incurred over the entire horizon are calculated. The objective of the heuristic is the same as that of the ILP model based approach, i.e. to minimize the total operational cost over the time horizon while satisfying the heat demand of the neighbourhood at all times.

To execute the above steps for a time interval i of day j, the heuristic considers only the current state of the buffer, i.e. the temperatures of the buffer segments $T_{i,s}$, as well as the input data for that time interval, and the energy target $V_{\text{tar}(j)}$ at the end of day j as input. Note, that the heuristic does not take any future data into account, besides the target of day j, $V_{\text{tar}(j)}$. The heuristic uses a set of rules

of thumb to determine, based on the current state and input data, the devices that will be connected to specific buffer segments. Whether a device is connected to a buffer segment is modelled by the binary variable $y_{i,s}^{dev}$. If $y_{i,s}^{dev} = 1$ device dev is connected to segment s during the complete time interval i and if $y_{i,s}^{dev} = 0$ it is not connected. All decisions variables $y_{i,s}^{dev}$ are set to 0 at the start of the heuristic.

It is worth mentioning that within the heuristic all the constraints of the Ecovat system still have to be taken into account, as well as restrictions on the temperatures of the buffer segments, i.e.:

» A device may only be connected to at most one buffer segment during each time interval.

» A buffer segment may have only one device connected to it during each time interval.

» Each buffer segment must have a temperature lower than a specified maximum temperature at all times.

» Each buffer segment must have a lower temperature than the buffer segment above it at all times.

The derived rules of thumb have to ensure that these constraints are not violated by the decisions made. While in Chapter 3, when introducing the ILP model, we ensured the model can accommodate any number of buffer segments, we now focus specifically on the case where the Ecovat buffer has 5 segments when discussing the heuristic method. If a buffer with a different number of segments is considered these rules may have to be adapted.

In the following the steps taken by the heuristic are explained in more detail.

Step 1: If the time interval is the first time interval of a new day, determine the maximum accepted energy price for running the devices in the system for the coming day.
Running any of the devices in the system leads to operational costs, which depend on the energy price during the time interval in which the device runs. Therefore, there is a trade-off between the price one is willing to pay to run a device and the need to charge the buffer, which is related to the targets $V_{\text{tar}(j)}$. To take this trade-off into account we specify a maximum energy price we are willing to accept during day j to charge the buffer using the devices in the system. As such, this maximum accepted energy price is an input for Steps 3 and 4 of the heuristic, in which the heuristic decides which devices will run during a given time interval i. However, note that devices may also be turned on for reasons other than a low energy price, for example to prevent a constraint on the temperatures of the buffer segments from being violated. In such cases the introduced price constraint will be overruled.

The determination of the maximum accepted energy price only takes place if the considered time interval i is the first time interval of day j, i.e. only if $i = (j-1) \cdot 96 + 1$.

The maximum accepted energy price $\bar{p}_{j+1}$ for the next day $j+1$ is determined by comparing the useful energy target $V_{\text{tar}(j)}$ at the end of day j, with the useful energy content of the Ecovat buffer $U_{\text{tar}(j)}$ at the end of day j, where $\text{tar}(j) = j \cdot 96 + 1$ as in Chapter 4. The chosen approach is similar to how c_j (the weight factor in the objective function of the ILP model that ensures the model stays close to the targets for the useful energy content of the buffer) was determined in the ILP model based approach. Here we consider the same three cases; 1) the useful energy content of the buffer is below the target, 2) the useful energy content of the buffer is above the target and 3) the buffer is almost full, i.e. the useful energy content is close to the maximum useful energy content $\bar{U}$ of the buffer. The used formula for the maximum accepted energy price $\bar{p}_{j+1}$ for the next day is then given by:

$$
\bar{p}_{j+1} = \begin{cases} 0.01 \cdot \left(\bar{U} - 15000 - U_{\text{tar}(j)} \right) & \text{if } U_{\text{tar}(j)} > \bar{U} - 15000 \\ 0 & \text{if } U_{\text{tar}(j)} \geqslant V_{\text{tar}(j)} \\ 241 \cdot \left(1 - \frac{U_{\text{tar}(j)}}{V_{\text{tar}(j)}} \right)^2 + 9 & \text{if } U_{\text{tar}(j)} < V_{\text{tar}(j)}. \end{cases} \tag{5.1}
$$

The maximum accepted energy price $\bar{p}_j$ when the useful energy content of the buffer is below the target ranges from $9 \frac{\text{€}}{\text{MWh}}$, when the useful energy content is slightly below the target, to $250 \frac{\text{€}}{\text{MWh}}$, when the useful energy is at 0 kWh. If the useful energy content of the buffer is above the target a maximum accepted energy price of $\bar{p}_j = 0 \frac{\text{€}}{\text{MWh}}$ is chosen unless the buffer is almost full. In that case a maximum accepted energy price ranges from $0 \frac{\text{€}}{\text{MWh}}$, when still 15000 kWh of useful capacity is remaining, to $150 \frac{\text{€}}{\text{MWh}}$, when the buffer is full. These values are the same as those used when determining c_j in the ILP model as discussed in Section 4.2.4. We set the maximum accepted energy price for the first day to $0 \frac{\text{€}}{\text{MWh}}$, since the first target is only defined at the end of day 1.

Step 2: Calculate the expected output temperature of the PVT panels.
To determine whether the PVT panels are used for charging the Ecovat buffer (i.e. are connected to the bottom segment) is based on the output temperature of the PVT panels. Similar to the ILP model, the heuristic only connects the PVT panels to the bottom segment of the buffer when the output temperature of the PVT panels is higher than the temperature of the bottom segment of the buffer. However, even if this requirement is met, the PVT panels may not be connected to the buffer if it is necessary to connect another device to the bottom buffer segment to prevent a constraint violation.

The output temperature T_i^{out} of the PVT panels in the heuristic is determined in the same way as for the ILP model, i.e. using equation:

$$T_i^{out} = \frac{2 \, \dot{m} \, c_p \, T_i^{in} - a_{th} \, A \, T_i^{in} + 2 \, A \, \eta_0^{th} \, G_i + 2 \, a_{th} \, A \, T_i^{amb}}{a_{th} \, A + 2 \, \dot{m} \, c_p},$$ (5.2)

where $\dot{m}$ is the mass flow rate of the heat transfer fluid through the PVT panels, c_p is the specific heat of water, T_i^{in} is the input temperature of the PVT panels, a_{th} is the thermal loss coefficient of the PVT panels, A is the surface area of the PVT panels, η_0^{th} is the thermal efficiency of the PVT panels at a reduced temperature of zero degrees and T_i^{amb} is the ambient temperature. The values of the parameters used in the heuristic approach are the same as those used in the ILP model based approach (see Section 4.1.1). This output temperature T_i^{out} is used in combination with the maximum accepted energy price for the current day $\bar{p}_j$ to determine which devices will run during time interval i.

Step 3: Decide if the PVT panels are connected to the bottom buffer segment, and whether or not the lthp and/or the hthp run during time interval i

In this step of the heuristic, the two water/water heat pumps (the lthp and hthp, covering the low and high part of the total temperature range of the Ecovat buffer respectively) and the PVT panels are considered. We schedule these devices sequentially, whereby aspects of the other devices scheduled later are taken into account. In Step 4 of the heuristic, the air/water heat pump (awhp), the resistance heater and the heat demand are scheduled simultaneously, since the choices for these devices influence each other to a much higher degree than the lthp, hthp and PVT panels do.

From the data obtained by the ILP model we observed that the lthp is mostly used to ensure the bottom segment of the Ecovat buffer stays below its maximum temperature $\bar{T}_5$, which has to ensure the PVT panels are cooled. From this we take as rule of thumb that the lthp is only used by the heuristic to keep the temperature of the bottom segment of the buffer below this maximum temperature. For this we consider three cases, based on two temperature levels $\bar{T}_5 - \delta_1^{lthp}$ and $\bar{T}_5 - \delta_2^{lthp}$, where δ_1^{lthp} and δ_2^{lthp} are positive constants with $\delta_1^{lthp} > \delta_2^{lthp}$. These two temperature levels are used to specify how close the bottom buffer segment is to its maximum temperature. The three considered cases for the lthp are:

» Case 1: The temperature $T_{i,s}$ of the bottom segment of the buffer is low enough, meaning that it is below or equal to $\bar{T}_5 - \delta_1^{lthp}$. In this case the lthp is not used.

» Case 2: The temperature of the bottom segment is between the two temperature levels, i.e. $\bar{T}_5 - \delta_1^{lthp} < T_{i,5} \leqslant \bar{T}_5 - \delta_2^{lthp}$. In this case the

lthp will be turned on when the energy price is at most the maximum accepted energy price $\bar{p}_j$ for the current day j, i.e. if $p_i \leqslant \bar{p}_j$.

» Case 3: The temperature of the bottom segment is higher than $\bar{T}_5 - \delta_2^{lthp}$. In this case the lthp will be turned on no matter what the energy price is.

Note that, this only specifies whether or not the lthp is turned on, but not yet to which segments of the buffer it is connected when it is turned on. In case the lthp is turned on the heat source of the lthp is always the bottom buffer segment, so $y_{i,5}^{lthp-} = 1$. In Case 2, where energy price is at most the maximum accepted energy price, the heat sink of the lthp is taken to be the lowest temperature segment (excluding the bottom segment) that has a temperature lower than the maximum temperature the lthp can generate (i.e. $T_{i,s} < T_{max}^{lthp}$) and that does not break the thermal stratification of the buffer if it is selected as the heat sink (i.e. $T_{i,s} + \Delta T^{lthp} < T_{i,s-1}$, where ΔT^{lthp} is the change of the temperature of the buffer segment when selected as the heat sink for the lthp). The change in the temperature ΔT^{dev} of a buffer segment when it is charged by device dev is given by:

$$\Delta T^{dev} = \frac{C^{dev} COP^{dev} \Delta t}{m_s c_p}, \tag{5.3}$$

where C^{dev} is the capacity of device dev, COP^{dev} is the coefficient of performance of device dev, Δt is the length of a time interval, m_s is the mass of segment s and c_p is the specific heat of water.

The reason to select the lowest temperature segment as heat sink for the lthp is to keep the higher temperature segments available for other devices. Note, that in Case 2 the energy price is below the maximum accepted energy price, and it is beneficial to turn on the largest number of devices possible to minimize the operational costs. Furthermore, leaving the upper segments 'available' gives the potential to connect the awhp and resistance heater to segments with high temperatures leading to a higher energy quality of the buffer, when compared to a situation where those devices are connected to lower temperature buffer segments. As such, it is beneficial to specifically keep the high temperature segments available for those devices.

In Case 3 as heat sink the highest temperature segment is chosen (excluding segments 1) that has a temperature lower than the maximum temperature the lthp can handle and that does not break the thermal stratification of the buffer if it is selected as the heat sink. The two differences compared to the second case are that the highest temperature segment is selected as the heat sink instead of the lowest temperature segment, and that the highest segment (i.e. $s = 1$) is excluded as heat sink as well. The reason to select the highest temperature segment is that in this case the energy price is above the maximum accepted energy price, and

as such most other devices will not run. The advantage of selecting a higher temperature segment in this case is the resulting increase in energy quality of the buffer. The reason to exclude $s = 1$ when selecting a segment as heat sink is to ensure that there is a segment with sufficient temperature available to satisfy the heat demand during the current time interval.

The process of determining the decision variables for the lthp is summarized in Algorithm 2. As mentioned previously, all decision variables are set to 0 at the start of the heuristic, so the heuristic only needs to change any decision variables if a device is turned on.

The results obtained with the ILP model based approach indicate that the hthp has similar behaviour as the lthp. The hthp is mostly used to keep the temperature of segment 4 below its maximum temperature $\bar{T}_4$, such that the lthp will always have a buffer segment to use as heat sink when cooling the bottom buffer segment. As such, we take a similar rule of thumb for the hthp as for the lthp, namely that the hthp is only used to ensure that the temperature of segment 4 does not reach the maximum allowed temperature of that segment.

Similar to the lthp, we define two temperature levels for the hthp, $\bar{T}_4 - \delta_1^{hthp}$ and $\bar{T}_4 - \delta_2^{hthp}$, where δ_1^{hthp} and δ_2^{hthp} are positive constants with $\delta_1^{hthp} > \delta_2^{hthp}$. The difference, compared to the lthp, is that the hthp connects to segments of the buffer that are used by other devices when the energy prices are low, such as the resistance heater. For this reason the hthp often runs when the energy costs are (slightly) higher than the maximum accepted energy cost that is used to determine if the lthp, the awhp and the resistance heater run (see also Step 4). In this way, the other devices can be connected to the buffer segments when the energy prices are low, which outweighs the extra costs made by turning on the hthp when the energy price is higher. As such, the rule of thumb we use for the hthp in the heuristic considers three cases:

Algorithm 2: Set decision variables for the lthp.

1 **if** $p_i \leqslant \bar{p}_j$ & $T_{i,5} > \bar{T}_5 - \delta_1^{lthp}$ **then**

2   set $y_{i,5}^{lthp-} = 1$

3   set $y_{i,s}^{lthp+} = 1$ for lowest temperature segment s, excluding $s = 5$ that satisfies $T_{i,s} < T_{max}^{lthp}$ & $T_{i,s} + \Delta T^{lthp} < T_{i,s-1}$

4 **else if** $T_{i,5} > \bar{T}_5 - \delta_2^{lthp}$ **then**

5   set $y_{i,5}^{lthp-} = 1$

6   set $y_{i,s}^{lthp+} = 1$ for highest temperature segment s, excluding $s = 1$ and $s = 5$, that satisfies $T_{i,s} < T_{max}^{lthp}$ & $T_{i,s} + \Delta T^{lthp} < T_{i,s-1}$

» Case 1: The temperature of the fourth segment of the buffer $T_{i,4}$ is low enough, meaning that it is below or equal to $\bar{T}_4 - \delta_1^{hthp}$. In this case the hthp is not used.

» Case 2: The temperature of the bottom segment is between the two temperature levels, i.e. $\bar{T}_4 - \delta_1^{hthp} < T_{i,4} \leqslant \bar{T}_4 - \delta_2^{hthp}$. In this case the hthp will be turned on when the energy price is between the maximum accepted energy price and a predefined maximum energy price for the hthp $\bar{p}_{hthp}$, i.e. $\hat{p}_j \leqslant p_i \leqslant \bar{p}_{hthp}$, assuming it is possible to turn the hthp on during the current time interval.

» Case 3: The temperature of the bottom segment is higher than $\bar{T}_4 - \delta_2^{hthp}$, in this case the lthp will be turned on no matter what the energy price is, assuming it is possible to turn the hthp on during the current time interval.

Note that, even when the temperature and energy price constraints for Case 2 and 3 are valid during the current time interval, it might not be possible to turn on the hthp because of the decisions made for the lthp, i.e. the lthp may already be connected to the fourth buffer segment during the current time interval.

The three cases describe when the hthp is turned on, but not yet to which segments it is connected when it is turned on. The heat source of the hthp is always taken to be the fourth buffer segment, to ensure the lthp always has a buffer segment to use as heat sink when cooling the bottom buffer segment. The segment chosen as heat sink differs between Case 2 and Case 3. In Case 2, the energy price is above the maximum accepted energy price used to determine whether the lthp, awhp and resistance heater run, meaning those devices are not turned on during any time interval where this case holds. This means the hthp can be connected to any segment without taking into account other devices scheduled after it, since those devices will not run during such a time interval. As such, the highest temperature segment, excluding segment 1, that has a temperature lower than the maximum temperature the hthp can generate (i.e. $T_{i,s} < T_{max}^{hthp}$) and that does not break the thermal stratification of the buffer if it is selected as the heat sink (i.e. $T_{i,s} + \Delta T^{hthp} < T_{i,s-1}$, where ΔT^{hthp} is the change in temperature of the buffer segment when selected as the heat sink for the hthp, obtained using Equation (5.3)). Segment 1 is again excluded to ensure there is at least one segment available to satisfy the heat demand during the current time interval. Again, the allowed segment with the highest temperature is chosen to maximize the energy quality of the Ecovat buffer.

In Case 3, where the hthp is turned on independent of the energy price, the segment chosen as the heat sink for the hthp is always segment 3. The reason for this is that in this case the energy price may be lower than the maximum accepted energy price, which is used to determine if the other devices run. This means that it is beneficial to keep the segments with higher temperatures available for

the other devices that are scheduled after the hthp, to increase the potential profit that can be made during the current time interval.

The process of determining the decision variables for the hthp is summarized in Algorithm 3. In the pseudocode for this algorithm, and the subsequent algorithms, we use the shorthand notation "segment s is free" when there is no previously considered device connected to segment s yet, i.e. $\sum_{dev \in \mathcal{D}} y_{i,s}^{dev} = 0$ for the current time interval i and the considered segment s.

Algorithm 3: Set decision variables for the hthp.

1   **if** $\bar{p}_j \leqslant p_i \leqslant \bar{p}_{hthp}$ & $T_{i,4} > \bar{T}_4 - \delta_1^{hthp}$ & segment 4 is free **then**

2   set $y_{i,4}^{hthp-} = 1$

3   set $y_{i,s}^{htph+} = 1$ for highest temperature segment s, excluding $s = 1$, that satisfies $T_{i,s} < T_{max}^{hthp}$ & is free

4   **else if** $T_{i,4} > \bar{T}_4 - \delta_2^{hthp}$ & segments 3 and 4 are free **then**

5   set $y_{i,4}^{hthp-} = 1$

6   set $y_{i,3}^{hthp+} = 1$

The rule of thumb used for the PVT panels is very straightforward: the PVT panels are connected to the bottom segment of the Ecovat buffer if the output temperature of the PVT panels is higher than the temperature of the bottom segment of the buffer, i.e. $T_i^{out} > T_{i,5}$ and no other device is connected to segment 5 yet. Algorithm 4 summarizes the process of determining the decision variable for the PVT panels.

Algorithm 4: Set decision variables for the PVT panels.

1   **if** $T_i^{out} > T_{i,5}$ & segment 5 is free **then**

2   set $y_{i,5}^{pvt} = 1$

Step 4: Decide from which segment the heat demand is supplied and if the resistance heater and/or awhp run during time interval i.
In this step, the heat demand, the resistance heater and the awhp are considered. Note that, the heat demand of the neighbourhood must be satisfied at all times, meaning that the heuristic must always select one of the buffer segments to supply the heat demand. Furthermore, to maximize the profit resulting from turning on the devices in the system (and thus minimize the operational costs) the resistance heater and the awhp must be turned on during time intervals with negative energy prices as much as possible. As such, these devices are often turned on in the same time interval. This fact, combined with the restriction that only one device can be connected to a given buffer segment during a time

interval, means the decisions with respect to the operation of the three devices considered in this step strongly influence each other. For this reason, these three devices are considered simultaneously by the heuristic.

Even though the heat demand, the resistance heater and the awhp are often connected to the Ecovat buffer during the same time intervals, this is not always the case. Therefore, we first determine in which specific situations each of these devices is turned on. Since the heat demand of the neighbourhood must be satisfied at all times, one of the buffer segments must be assigned to supplying the heat demand during every time interval. Furthermore, the heuristic will try to turn the resistance heater on when the energy price is at most the maximum accepted energy price. However, in some circumstances this may not be possible due to the temperature distribution over the buffer segments, e.g. when the buffer is already at its capacity. Finally, the heuristic will try to turn on the awhp when the energy price is at most the maximum energy price multiplied by the coefficient of performance of the awhp. The reason for this multiplication is that this leads to the same accepted maximum energy price per unit of thermal energy as for the resistance heater, which has a coefficient of performance of one. However, as with the resistance heater, it might not be possible to turn the awhp on even if the energy price is low enough. To summarize, this leads to three different cases, based on the energy price p_i during a given time interval i:

» Case 1: $p_i \leqslant \bar{p}_j$. In this case, we try to turn on the resistance heater and awhp if possible. Furthermore, a segment must be selected to supply the heat demand.

» Case 2: $\bar{p}_j < p_i \leqslant \bar{p}_j COP^{awhp}$. In this case, we try to turn on the awhp. Furthermore, a segment must be selected to supply the heat demand. (this case does not occur if $\bar{p}_j \leqslant 0$).

» Case 3: $p_i > \bar{p}_j COP^{awhp}$. In this case, we only select a segment that supplies the heat demand.

The general procedure of determining the decision variables for the heat demand, the resistance heater and the awhp is the same in all three cases. In the following, we only discuss the first case, which is the most complex, in detail. The other two cases are very similar, and therefore only summarized in Algorithm 5 at the end of the discussion.

For specifying this step of the heuristic we make use of the following parameters: $\bar{y}_i^{res}$, $\bar{y}_i^{awhp}$ and $\underline{y}_i^{dem}$, which are defined in the following way: $\bar{y}_i^{res}$ is the highest temperature segment that can be charged by the resistance heater, $\bar{y}_i^{awhp}$ is the highest temperature segment that can be charged by the awhp and $\underline{y}_i^{dem}$ is the lowest temperature segment that can supply the heat demand. For all three parameters it holds that selecting that segment should not lead to a violation of the temperature stratification constraint, the maximum temperature constraint

of the buffer or a device specific constraint. Below we discuss in more detail the determination of these parameters for each device. Note, that $\bar{y}_i^{res}$, $\bar{y}_i^{awhp}$ and/or $\underline{y}_i^{dem}$ can take the same value.

When determining $\bar{y}_i^{res}$, the highest temperature segment s that the resistance heater can charge without breaking any constraints, we need to ensure that if the heuristic selects the resistance heater to charge segment s that the maximum temperature constraint of the buffer or the temperature stratification constraint are not violated. To avoid breaking the maximum temperature constraint segment s needs to satisfy $T_{i,s} + \Delta T^{res} \leqslant \bar{T}_s$, where ΔT^{res} is the temperature change caused by the resistance heater charging a buffer segment (ΔT^{res} can be obtained using Equation (5.3)). When checking whether charging segment s with the resistance heater violates the temperature stratification constraint, the heuristic takes into account that segment $s-1$ might be chosen to supply the heat demand, i.e. segment s needs to satisfy $T_{i,s} + \Delta T^{res} \leqslant T_{i,s-1} - \Delta T^{dem}$, where ΔT^{dem} is the temperature change caused by the heat demand being supplied by a buffer segment (ΔT^{dem} can be obtained using Equation (5.3), where the 'capacity' of the heat demand is simply the amount of heat demand during that interval, i.e. $C^{dev} = D_i$). An important observation is that if a segment s can be charged by the resistance heater, any segments below that can be charged by it as well, since those segments have a lower temperature than segment s.

Determining $\bar{y}_i^{awhp}$ is done in a similar way as determining $\bar{y}_i^{res}$. As mentioned before, $\bar{y}_i^{awhp}$ is defined as the highest temperature segment s that the awhp can charge without violating any constraints. Again, we need to ensure that the temperature stratification is kept intact, which is done in a similar way as discussed above, when determining $\bar{y}_i^{res}$. Furthermore, segment s needs to have a temperature below the maximum temperature that the awhp can supply, i.e. $T_{i,s} < T_{max}^{awhp}$. Similar to the resistance heater, we know that if the awhp can charge segment s it can also charge segments with lower temperatures.

Finally, we determine $\underline{y}_i^{dem}$, which is defined as the lowest temperature segment s that can supply the heat demand without breaking any constraints. When determining $\underline{y}_i^{dem}$ the heuristic needs to ensure the temperature stratification stays intact, which is done in a similar way as discussed above when determining $\bar{y}_i^{res}$. Furthermore, segment s needs to have a sufficiently high temperature to satisfy the heat demand from the neighbourhood, i.e. $T_{i,s} > T_{dem}$. In the case of the heat demand, we know that if segment s can supply the heat demand of the neighbourhood, so can any segment of higher temperature.

The procedure to determine the segments to which the resistance heater, awhp and heat demand are connected considers the buffer segments in order of increasing temperature, excluding the bottom segment. This implies that the procedure starts with segment 4. It first checks in the following order if one of the devices can be connected to this segment: resistance heater, awhp, heat demand. If for

one of the segments a connection is possible, this connection is established. The reason the heat demand is checked last, even though it is the only 'device' that always has to run, is that the top segment, which is considered last in this procedure, always has the highest temperature and thus is always able to supply the heat demand of the neighbourhood, assuming the buffer has a sufficiently high useful energy content to satisfy the heat demand. However, the heuristic ensures that this assumption is always the case. Below we go into more detail on this procedure, by specifying the procedure of setting the decision variables for the resistance heater, the awhp and the heat demand for a given segment s.

The entire procedure of determining the decision variables for the resistance heater, the awhp and the heat demand is given in Algorithm 5. First, the heuristic checks whether none of the devices considered in Step 3 (lthp and hthp) are connected to segment s yet (for example in line 5 of Algorithm 5). If this is not the case, it checks whether the resistance heater can be connected to segment s, i.e. if $\bar{y}_i^{res} = s$. If this is possible, the resistance heater is connected to this segment, i.e. $y_{i,s}^{res} = 1$ (e.g. in line 6 of Algorithm 5). Since only one device can be connected to a buffer segment during any time interval the procedure stops considering segment s at this point and proceeds to the next segment, i.e. segment $s - 1$. If the resistance heater is not connected to segment s, the process is repeated for the awhp, i.e. if $\bar{y}_i^{awhp} = s$ the awhp is connected to segment s ($y_{i,s}^{awhp} = 1$) and the procedure proceeds to the next segment (e.g. in line 7 of Algorithm 5). If the awhp is not connected to segment s, the heuristic checks whether this segment can supply the heat demand and the heat demand is not yet supplied by another segment yet, i.e. if both $\underline{y}_i^{dem} \geqslant s$ and $\sum_{s \in \mathcal{S}} y_{i,s}^{dem} = 0$ are satisfied, segment s is assigned to supply the heat demand ($y_{i,s}^{dem} = 1$, e.g. in line 14 of Algorithm 5). The reason to check $\underline{y}_i^{dem} \geqslant s$ instead of $\underline{y}_i^{dem} = s$ is due to the fact that any segment with a higher temperature than segment s can also supply the heat demand (contrary to $\bar{y}_i^{res}$ and $\bar{y}_i^{res}$ which give the highest temperature segments those devices can be connected to).

The above procedure works well if $\bar{y}_i^{res}$, $\bar{y}_i^{awhp}$ and $\underline{y}_i^{dem}$ have different values. However, this is not necessarily the case. To integrate the cases in which some of these parameters are equal as well, some extra checks have to be added to allow more devices to be connected to the Ecovat buffer. To give a specific example, if both $\bar{y}_i^{res}$ and $\bar{y}_i^{awhp}$ have the same value s, the above procedure would connect the resistance heater to segment s and the awhp would not be connected to any segment. However, it might be possible to connect the awhp to segment $s + 1$ in that case, since that segment has a lower temperature. To take this into account the heuristic checks whether $\bar{y}_i^{res} = s$ and $\bar{y}_i^{awhp} = s$ when considering segment $s + 1$. If this is the case then the awhp is connected to segment $s + 1$ and the resistance heater can be connected to segment s as before. Line 8 of Algorithm 5 gives an example of the above mentioned extra checks. Note, that if $\bar{y}_i^{res} = 3$

and $\bar{y}_i^{awhp} = 3$ the awhp gets connected to segment 4 in line 8, i.e. $y_{i,4}^{awhp} = 1$, and the resistance heater gets connected to segment 3 in line 11, i.e. $y_{i,3}^{res} = 1$. Furthermore, note that due to the way the decision variables for the lthp and hthp are set, segment 4 can not be free while segment 3 is not free when $p_i \leqslant \bar{p}_j$. As discussed before, in the cases in which less devices are considered the decision variables are determined in a similar way, as can be seen in Algorithm 5.

Algorithm 5: Set decision variables for the resistance heater, the awhp and the heat demand.

1. set $\bar{y}_i^{res} = s_1$ where s_1 is the highest segment that res can charge without breaking the maximum temperature or temperature stratification constraints
2. set $\bar{y}_i^{awhp} = s_2$ where s_2 is the highest segment that awhp can charge without breaking the temperature stratification constraint and that satisfies $T_{i,s} < T_{max}^{awhp}$
3. set $\underline{y}_i^{dem} = s_3$ where s_3 is the lowest segment that can supply the heat demand without breaking the temperature stratification constraint and that satisfies $T_{i,s} > T_{dem}$
4. **if** $p_i \leqslant \bar{p}_j$ **then**
5.   **if** segment 4 is free **then**
6.   **if** $\bar{y}_i^{res} = 4$ **then** set $y_{i,4}^{res} = 1$
7.   **else if** $\bar{y}_i^{awhp} = 4$ **then** set $y_{i,4}^{awhp} = 1$
8.   **else if** $\bar{y}_i^{res} = 3$ & $\bar{y}_i^{awhp} = 3$ **then** set $y_{i,4}^{awhp} = 1$
9.   **else if** $\underline{y}_i^{dem} \geqslant 4$ **then** set $y_{i,4}^{dem} = 1$
10.   **if** segment 3 is free **then**
11.   **if** $\bar{y}_i^{res} = 3$ **then** set $y_{i,3}^{res} = 1$
12.   **else if** $\bar{y}_i^{awhp} = 3$ **then** set $y_{i,3}^{awhp} = 1$
13.   **else if** $\bar{y}_i^{res} = 2$ & $\bar{y}_i^{awhp} = 2$ **then** set $y_{i,3}^{awhp} = 1$
14.   **else if** $\underline{y}_i^{dem} \geqslant 3$ & demand not satisfied yet **then** set $y_{i,3}^{dem}$
15.   **else if** $\bar{y}_i^{awhp} \leqslant 2$ & $\bar{y}_i^{res} \leqslant 2$ & $\underline{y}_i^{dem} \leqslant 2$ **then** set $y_{i,3}^{awhp} = 1$
16.   **if** segment 2 is free **then**
17.   **if** $\bar{y}_i^{res} = 2$ **then** set $y_{i,2}^{res} = 1$
18.   **else if** $\bar{y}_i^{awhp} \leqslant 2$ & awhp not connected to a segment yet **then** set $y_{i,2}^{awhp} = 1$
19.   **else if** $\underline{y}_i^{dem} \geqslant 2$ & demand not satisfied yet **then** set $y_{i,2}^{dem} = 1$
20.   **else if** $\bar{y}_i^{res} = 1$ & $\underline{y}_i^{dem} = 1$ **then** set $y_{i,2}^{res} = 1$
21.   **if** $\underline{y}_i^{dem} \geqslant 1$ & demand not satisfied yet **then** set $y_{i,1}^{dem} = 1$
22.   **else if** $\bar{y}_i^{res} = 1$ **then** set $y_{i,1}^{res}$

---

23   **else if** $\bar{p}_j < p_i \leqslant \bar{p}_j COP^{awhp}$ **then**

24   **if** segment 4 is free **then**

25   **if** $\bar{y}_i^{awhp} = 4$ **then** set $y_{i,4}^{awhp} = 1$;

26   **else if** $\bar{y}_i^{awhp} = 3$ & segment 3 not free **then** set $y_{i,4}^{awhp} = 1$;

27   **else if** $\underline{y}_i^{dem} \geqslant 3$ **then** set $y_{i,4}^{dem} = 1$;

28   **if** segment 3 is free **then**

29   **if** $\bar{y}_i^{awhp} = 3$ **then** set $y_{i,3}^{awhp} = 1$;

30   **else if** $\bar{y}_i^{awhp} = 2$ & segment 2 not free **then** set $y_{i,3}^{awhp} = 1$;

31   **else if** $\underline{y}_i^{dem} \geqslant 2$ & demand not satisfied yet **then** set $y_{i,3}^{dem} = 1$;

32   **if** segment 2 is free **then**

33   **if** $\bar{y}_i^{awhp} \leqslant 2$ **then** set $y_{i,2}^{awhp} = 1$;

34   **else if** $\underline{y}_i^{dem} \geqslant 2$ & demand not satisfied yet **then** set $y_{i,2}^{dem} = 1$;

35   **if** $\underline{y}_i^{dem} = 1$ **then** set $y_{i,1}^{dem} = 1$;

36   **else**

37   **if** $\underline{y}_i^{dem} \geqslant 4$ & segment 4 is free **then** set $y_{i,4}^{dem} = 1$;

38   **else if** $\underline{y}_i^{dem} \geqslant 3$ & segment 3 is free **then** set $y_{i,3}^{dem} = 1$;

39   **else if** $\underline{y}_i^{dem} \geqslant 2$ & segment 2 is free **then** set $y_{i,2}^{dem} = 1$;

40   **else** set $y_{i,1}^{dem} = 1$;

Step 5: Calculate the heat losses of the Ecovat buffer to the surroundings during time interval i.

The previous steps consider the charging/discharging of the Ecovat buffer by the devices in the system. Next to this, the state of the Ecovat buffer is also influenced by energy losses to the surrounding environment over time. In this step the heuristic determines the amount of heat lost by each of the buffer segments during time interval i. This is done using the same equation as in the ILP model:

$$Q_{i,s}^{loss} = \left(1 - (1-\beta)^{\frac{2}{N_{int}}}\right)(T_{i,s} - T_{gw})\, m_s\, c_p, \qquad (5.4)$$

where $Q_{i,s}^{loss}$ is the heat lost by buffer segment s during time interval i, β is the given heat loss coefficient of the Ecovat buffer over 6 months, m_s is the mass of segment s, c_p is the specific heat of water and the factor $\frac{2}{N_{int}}$ is equal to one over the number of time intervals in half a year (see also Equation 3.30 in Chapter 3).

Step 6: Calculate the changes in energy content in the system due to the decisions made in Steps 2–4 and the calculated heat loss from Step 5.

The decisions in Steps 3 and 4 of the heuristic lead to changes in the energy content of the buffer segments, which in turn lead to temperature changes. In

this step the heuristic determines the changes in the energy content induced by the charging/discharging $Q_{i,s}^{dev}$ done by the devices in the system. For the lthp, hthp, awhp and the resistance heater these changes depend on the capacity C^{dev} of the device and the coefficient of performance of the device COP^{dev}. For the heat demand it depends on the energy demand D_i during that time interval. Finally, for the PVT panels the changes depend on the thermal efficiency of the panels. Of the involved parameters the thermal efficiency of the PVT panels is the only value that is not an input of the model. As in the ILP model, the thermal efficiency is modelled as a linear function of the reduced temperature, which is defined as the difference between the mean temperature in the PVT panels and the ambient temperature divided by the global irradiation [36]. This means the reduced temperature T_i^{red} is given by:

$$T_i^{red} = \begin{cases} \dfrac{\dfrac{T_{i,5} + T_i^{out}}{2} - T_i^{amb}}{G_i}, & \text{if } G_i > 0 \\ 0, & \text{otherwise,} \end{cases} \tag{5.5}$$

where T_i^{amb} is the ambient temperature during time interval i and G_i is the global irradiation during time interval i. Using the reduced temperature, the thermal efficiency of the PVT panels η_i^{th} is given by:

$$\eta_i^{th} = \eta_0^{th} - a_{th}\, T^{red}, \tag{5.6}$$

where η_0^{th} is the thermal efficiency of the PVT panels at a reduced temperature T^{red} of zero and a_{th} is the thermal loss coefficient of the PVT panels. Once the thermal efficiency of the PVT panels is known, the changes in the energy content $Q_{i,s}^{dev}$ of the buffer segments, resulting from the charging/discharging done by the devices in the system, can be calculated using the same equations as in the ILP model:

$$Q_{i,s}^{pvt} = \begin{cases} 0 & \text{if } s = 1,..,4, \\ \eta_i^{th}\, G_i\, A\, y_{i,s}^{pvt} & \text{if } s = 5, \end{cases} \tag{5.7}$$

$$Q_{i,s}^{awhp} = C^{awhp}\, COP^{awhp}\, y_{i,s}^{awhp}, \tag{5.8}$$

$$Q_{i,s}^{lthp-} = -C^{lthp}\, (COP^{lthp} - 1)\, y_{i,s}^{lthp-}, \tag{5.9}$$

$$Q_{i,s}^{lthp+} = C^{lthp}\, COP^{lthp}\, y_{i,s}^{lthp+}, \tag{5.10}$$

$$Q_{i,s}^{hthp-} = -C^{hthp}\, (COP^{hthp} - 1)\, y_{i,s}^{hthp-}, \tag{5.11}$$

$$Q_{i,s}^{hthp+} = C^{hthp}\, COP^{hthp}\, y_{i,s}^{hthp+}, \tag{5.12}$$

$$Q_{i,s}^{res} = C^{res}\, y_{i,s}^{res},\qquad (5.13)$$

$$Q_{i,s}^{dem} = -D_i\, y_{i,s}^{dem}.\qquad (5.14)$$

Additionally, we have to determine the amount of electrical energy produced by the PVT panels which is needed in Step 8 of the heuristic. For this the electrical efficiency of the PVT panels is required, which is modelled in the same way as in the ILP model:

$$\eta_i^{el} = \eta_0^{el} - a_{el}\, T^{red}.\qquad (5.15)$$

Using the electrical efficiency of the PVT panels, the amount of electrical energy produced during time interval i is given by:

$$E_i^{pvt} = \eta_i^{el}\, G_i\, A \sum_{s\in\mathcal{S}} y_{i,s}^{pvt}.\qquad (5.16)$$

Step 7: Update the temperatures in the segments of the Ecovat buffer.
The changes in the energy content of the buffer segments imply also changes in their temperatures. In this step the heuristic calculates those temperatures. The temperature changes in the buffer segments follow from the following equation:

$$T_{i+1,s} = T_{i,s} + \frac{\Delta t}{m_s\, c_p}\Big(\sum_{dev\in\mathcal{D}} Q_{i,s}^{dev} - Q_{i,s}^{loss}\Big),\qquad (5.17)$$

where the temperatures $T_{i+1,s}$ are the temperatures of the buffer segments at the start of time interval $i+1$. These temperatures are used by the heuristic when determining the decision variables for the next time interval $i+1$.

Step 8: Calculate the operational costs resulting from the decisions made in Steps 2–4.
In the final step the heuristic determines the operational costs incurred during the current time interval i, due to the decisions made in Steps 3 and 4. As for the ILP, these costs consist of the costs of running the devices in the system, and are reduced by the amount of electrical energy produced by the PVT panels. As before, we only consider the operational costs and do not consider costs such as maintenance costs, or the profit made from supplying the heat demand of the neighbourhood, since we assume these to be equal for different charging/discharging strategies. The operational costs K_i incurred during time interval i are given by:

$$K_i = p_i\, \Delta t \Big(\sum_{s\in\mathcal{S}} (C^{awhp}\, y_{i,s}^{awhp} + C^{lthp}\, y_{i,s}^{lthp-} + C^{hthp}\, y_{i,s}^{hthp-}$$
$$+ C^{res}\, y_{i,s}^{res}) - E_i^{pvt}\Big).\quad (5.18)$$

Parameter	Value
δ_1^{lthp}	0.3 °C
δ_2^{lthp}	0.1 °C
δ_1^{hthp}	0.3 °C
δ_2^{hthp}	0.1 °C
T_{max}^{hthp}	50 $\frac{€}{\text{MWh}}$

Table 5.1: Parameters specific to the heuristic method for controlling an Ecovat system.

The overall heuristic repeats Steps 1–8 for every time interval in the time horizon, and sums up the operational costs over all time interval to obtain the total operational costs over the entire time horizon, i.e. the total operational costs are given by:

$$K_{tot} = \sum_{i \in \mathcal{I}} K_i. \tag{5.19}$$

5.2   RESULTS

The heuristic method presented in the previous section requires, next to the daily energy targets $V_{\text{tar}(j)}$, only input data for the time interval that the heuristic is considering at that moment. It uses future data only indirectly through the long-term planning introduced in Chapter 4, which depends on historical data for the heat demand of the neighbourhood and may depend on predictions for the energy prices. Similar to the results presented in Section 4.2.5, in this section we compare cases where perfect predictions (PP) are available for the energy prices with cases where no predictions (NP) are available for the energy prices when determining the daily useful energy targets for the Ecovat buffer in the long-term planning step.

For determining the performance of the heuristic method we again use the energy price data from the years 2011, 2013, 2014 and 2015. Additionally, we consider two different demand temperatures $T_{dem} = 40°C$ and $T_{dem} = 60°C$, leading to the same 16 cases we considered in Section 4.2.5 where we presented the results obtained with the ILP model based approach. In this section, we compare results obtained using the developed heuristic with the results obtained using the ILP model based approach. Except for the parameters given in Table 5.1, which are specific to the heuristic method, the inputs used for the heuristic are the same as for the ILP model based approach discussed in Chapter 4.

We first compare the quality of the solutions obtained using the heuristic and the ILP model based approach. For this we consider the same metrics as in Chapter 4,

i.e. the operational costs over the time horizon, and the useful energy content of the Ecovat buffer at the end of the time horizon. In Table 5.2 the operational costs and the useful energy content of the Ecovat buffer at the end of the time horizon are shown for the 16 cases mentioned above, for both the heuristic as well as the ILP model based approach. Furthermore, the relative difference in operational costs between the heuristic and the ILP model based approach is shown in the last column of Table 5.2.

The results presented in Table 5.2 show that for a demand temperature T_{dem} of 40 °C the results obtained using the heuristic presented in this chapter differ only a few percent from the results obtained using the ILP model based approach. To get more insight into the behaviour of the approaches, we have a look at the temperature evolution inside the Ecovat buffer throughout the year for two specific cases. In Figure 5.1 the temperature evolution is given for both the heuristic and the ILP model based approach, for the PP case using energy price data from 2014 and a demand temperature T_{dem} of 40 °C. This is one of the considered cases where the heuristic performs very well, showing only a 1.4% difference compared to the result obtained using the ILP model based approach. The graphs in Figure 5.1 show that the temperature evolution is very similar for both cases.

When considering a demand temperature T_{dem} of 60 °C we observe that the differences between the heuristic and the ILP model based approach become larger. Table 5.2 shows that the largest difference is observed in the NP case using the energy price data from 2011. In this case the difference in operational costs between the heuristic and the ILP model based approach is 14%. In Figure 5.2 the temperature evolution inside the Ecovat buffer for this case is shown for both the heuristic and the ILP model based approach. We can see that the differences between both models are larger compared to Figure 5.1, as expected. However, even in this case we see the same general trends in the temperature evolution throughout the year.

Over all 16 considered cases, we observe an average decrease in performance of 5.2% when comparing the heuristic with the ILP model based approach. Furthermore, we can see that the useful energy content of the buffer is generally very similar for both models. Finally, we observe that the differences between the PP and NP cases are in general very small for the heuristic. This is the same behaviour we observed in Chapter 4 when considering the results obtained with the ILP model based approach. This reinforces the conclusion that as long as the daily targets for the useful energy content of the Ecovat buffer follow the expected seasonal behaviour, the obtained results are very similar.

	Heuristic		ILP		Difference
	Costs(€)	Useful energy (kWh)	Costs(€)	Useful energy(kWh)	in costs (%)
$T_{dem} = 40\ °C$, 2011, PP	-22092	115849	-22709	115696	2.7
$T_{dem} = 40\ °C$, 2011, NP	-21772	140469	-22097	145635	1.5
$T_{dem} = 40\ °C$, 2013, PP	-40998	111530	-41711	114694	1.7
$T_{dem} = 40\ °C$, 2013, NP	-40906	114395	-41609	117096	1.7
$T_{dem} = 40\ °C$, 2014, PP	-27485	104433	-27865	110370	1.4
$T_{dem} = 40\ °C$, 2014, NP	-27375	106434	-27776	110145	1.4
$T_{dem} = 40\ °C$, 2015, PP	-47912	125541	-50634	125215	5.4
$T_{dem} = 40\ °C$, 2015, NP	-47968	122855	-50528	125739	5.1
$T_{dem} = 60\ °C$, 2011, PP	-20056	64464	-22339	56187	10.2
$T_{dem} = 60\ °C$, 2011, NP	-18193	78911	-21164	78399	14.0
$T_{dem} = 60\ °C$, 2013, PP	-39086	52805	-41343	52309	5.5
$T_{dem} = 60\ °C$, 2013, NP	-38463	59794	-40681	59093	5.5
$T_{dem} = 60\ °C$, 2014, PP	-26218	50969	-27170	47894	3.5
$T_{dem} = 60\ °C$, 2014, NP	-25997	56112	-26409	57100	1.6
$T_{dem} = 60\ °C$, 2015, PP	-46338	55099	-52196	55796	11.2
$T_{dem} = 60\ °C$, 2015, NP	-45801	62390	-51640	61916	11.3

Table 5.2: Comparison between the heuristic and the ILP model.
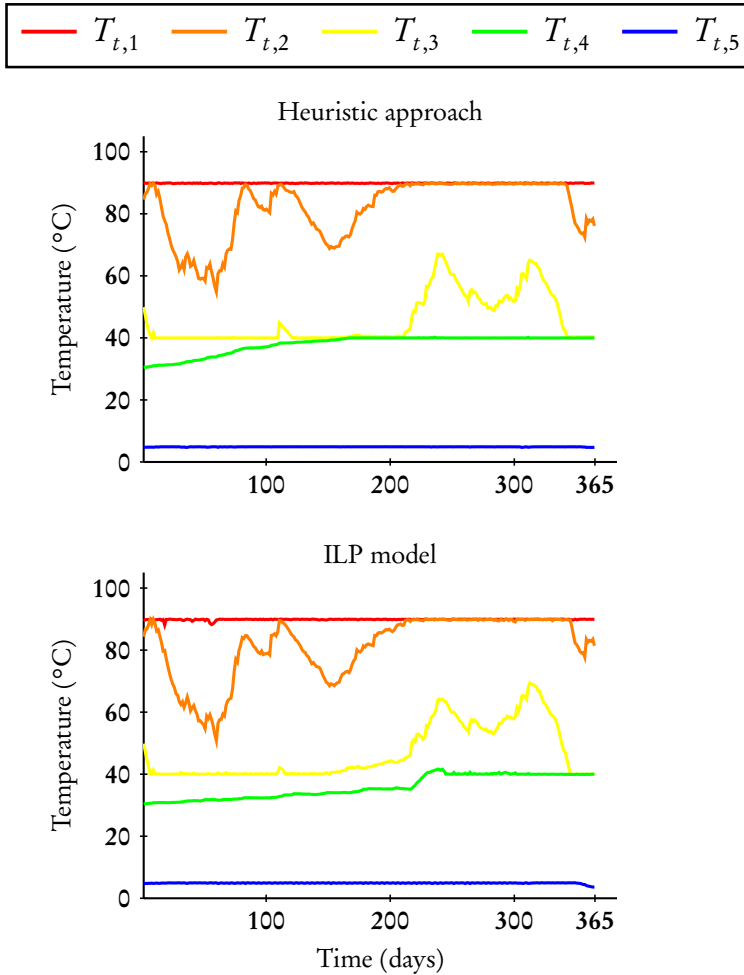
Figure 5.1: Temperature evolution inside the Ecovat buffer throughout the year for the PP case using the energy price data from 2014 and $T_{dem} = 40°C$.

We observe that the cases in which the relative difference between the approaches is largest are the cases in which the buffer is either almost empty (2011) or almost full (2015) at some time during the year. Intuitively it makes sense that the advantage of being able to consider future input data, which is the advantage the ILP model based approach has over the heuristic, is largest in these extreme cases. However, even in these cases the largest observed difference is not larger than 14.0%. Note that, as discussed before, the ILP model based approach only gives an upper bound on the performance of the heuristic as it additionally uses (generally unknown) input data for future time intervals to determine the charging/discharging strategy for the current time interval. In a real world situation
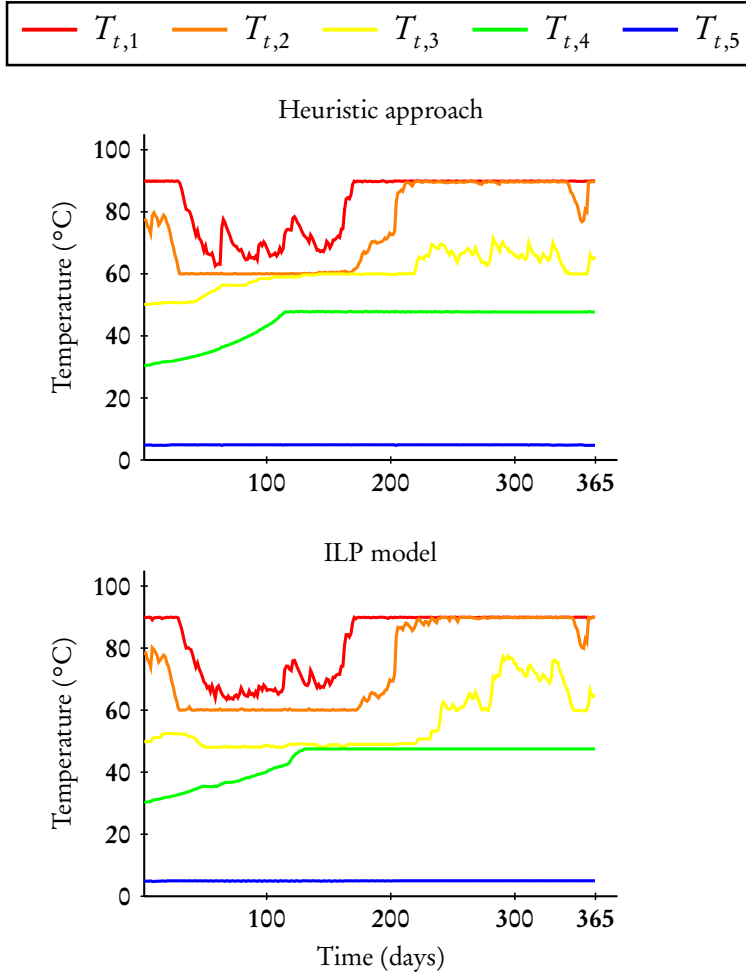
Figure 5.2: Temperature evolution inside the Ecovat buffer throughout the year for the NP case using the energy price data from 2011 and $T_{dem} = 60°C$.

it is unlikely that we have very accurate predictions for these inputs, such as the energy prices, and are instead forced to make operational decisions based only on data available at that point in time or on predictions which may be quite uncertain. As such, a worst case difference of 14.0% between the heuristic, which only uses input data for the time interval it is currently considering, and the ILP model based approach, which uses input data for future time intervals to make better informed decisions, seems to be an acceptable decrease in performance.

We have discussed the quality of the results obtained using the heuristic compared to those obtained with the ILP model based approach, taking into account

the fact that the heuristic only uses input data for the time interval it is considering at that time. The final goal of the heuristic method as described at the start of this chapter is to provide charging/discharging strategies in a short amount of time. While the ILP model has quite a large spread of computational times for the different considered cases, the computational time for the heuristic is a few seconds for all 16 considered cases. This means that the heuristic is computationally a lot cheaper than the ILP model based approach which requires around 8 to 15 hours for most considered cases (with some extreme cases even taking a lot longer). The resulting computational time of a few seconds for the heuristic to simulate an entire year of operation makes the presented heuristic fast enough to be usable for operational (real-time) control of a practical system. This means the heuristic satisfies all three goals outlined at the start of this chapter.

5.3 Conclusion

In the previous chapters we presented an ILP model based approach to control the Ecovat system, which provides good charging/discharging strategies, but which was too slow to be used in a practical situation. Furthermore, that approach required input data for future time intervals (e.g. the energy prices), which are unlikely to be available in practice. Therefore, in this chapter we developed a heuristic to control the Ecovat system, which:

1. produces charging/discharging strategies that are still of good quality, i.e. close to the strategies determined by the ILP model based approach,

2. produces such charging/discharging strategies in a short computational time (seconds), and

3. produces such charging/discharging strategies without requiring input data, such as energy prices or weather data, beyond the time interval that the heuristic method is considering at that moment.

From the results presented in this chapter we observe that the charging/discharging strategies obtained using the heuristic are still of good quality compared to the strategies obtained with the ILP model based approach. We observed a worst case decrease in performance of 14% when comparing the heuristic to the ILP model based approach. However, in most simulated cases the decrease in performance is significantly smaller, with an average decrease in performance of 5.2%. Hereby we have to take into account that the ILP model based approach gives an upper bound on the performance of the heuristic as it considering also input data for time intervals in the future, which generally is not available in a practical situation. The developed heuristic does not require such input data for future time intervals, and only uses available input data for the time interval it is considering at that time. Finally, we observe that for a low demand temperature the differences between the heuristic and the ILP model based approach are especially small. For these reasons we conclude that the decrease in performance

is acceptable, especially considering the decrease in computational time for the heuristic (this computational time is only a few seconds) compared to the ILP model based approach (computation times of generally 8 to 15 hours, but in some extreme cases even longer).

Summarizing, the presented heuristic method satisfies the three goals outlined above, and as such can be used to control an Ecovat system in practical situations. In the next chapter we present a case study of integrating an Ecovat system in a decentralized energy management (DEM) simulation of a neighbourhood. In this case study the Ecovat system is controlled using the heuristic method presented in this chapter.

116

Case study: decentralized energy management simulation including an Ecovat system

Abstract – *In this chapter we present a case study of using an Ecovat system within a neighbourhood of houses. We combine the heuristic method to control an Ecovat system presented in the previous chapter with the Decentralized Energy Management Toolkit (DEMKit), which is a simulator designed to perform decentralized energy management (DEM) simulations. We investigate the sensitivity of the Ecovat system to expected and unexpected increases in the heat demand, by considering different neighbourhood sizes and different heat demand profiles (e.g. due to an unexpectedly cold winter). Furthermore, we determine the benefit of using an Ecovat system to satisfy the heat demand of a neighbourhood of houses as a replacement for using gas boilers, which is currently the most common way of heating houses in the Netherlands.*

In the previous chapter we presented a heuristic method for the control of an Ecovat system. We compared this heuristic method to the ILP model based approach described in Chapters 3 and 4 and found that the heuristic method leads to an acceptable decrease in performance compared to the ILP model based approach, while the heuristic method has much shorter computational times. In this chapter, we consider a case study in which we perform a decentralized energy management (DEM) simulation of a neighbourhood containing an Ecovat system. In this case study the control of the Ecovat system is done by the heuristic method presented in the previous chapter.

This chapter is based on [G:4].

The goals of this case study are: 1) to determine whether the Ecovat system can supply the heat demand of a neighbourhood throughout the year, 2) to determine the sensitivity of the system to major changes in heat demand, and 3) to determine the benefits in terms of CO_2 reduction and self-consumption of using an Ecovat system to supply the heat demand of the neighbourhood, compared to using gas boilers (which is how the majority of Dutch homes is currently heated). These goals are realised by combining the heuristic method to control an Ecovat system, developed in the previous chapter, with the Decentralized Energy Management Toolkit (DEMKit) [47], which is designed to perform DEM simulations (see Section 6.1.1 for more background on DEMKit).

The remainder of this chapter is structured as follows. In Section 6.1 we discuss the way the neighbourhood is simulated by combining the heuristic method for controlling the Ecovat system presented in the previous chapter and the DEMKit simulator. Furthermore, we discuss the input and output of this model and provide some background on the DEMKit simulator. Afterwards, in Section 6.2 we introduce the five different cases we consider during this case study. Subsequently, in Section 6.3 we present the results of the simulations for the considered cases. Finally, in Section 6.4 we conclude this chapter.

6.1 SIMULATION SETUP

The neighbourhood considered in this case study is based on a future scenario, where we expect a higher penetration of emerging technologies such as electric vehicles and PV panels than is currently common. We are interested in the benefits the Ecovat system can provide in such a scenario, where a decentralized approach with regards to the energy management is required.

More specifically, we used data from a neighbourhood with a mix of single worker households, dual worker households, households of families with children and households of retired persons. In this neighbourhood 13% of the households have an electric vehicle and a further 32% have a plug-in hybrid electric vehicle. Furthermore, 50% of the households have PV panels and 10% have a battery. Additionally, 25% of households uses induction cooking and between 20% to 60% of the households have a dishwasher, based on the type of household, whereby households consisting of more persons having a higher chance of owning a dishwasher. Finally, we assume that every household has common devices such as a washing machine, oven and refrigerator. For satisfying the heat demand, we assume every house is heated by a gas boiler, as this allows us to quantify the benefits an Ecovat system has over gas boilers for satisfying the heat demand of this neighbourhood.

To start, we give a general description of the two parts required for the simulation of this neighbourhood, as well as the input and output of both these parts. Subsequently, we discuss the input data used and provide some more details on specific parts of the simulation. In the first part of the simulation we use the

DEMKit simulator to optimize the electricity consumption within the neighbourhood. In our case the objective of this optimization is the flattening of the electricity profile of the neighbourhood. In this part of the simulation the control of the heat consumption is not yet optimized, since we model a neighbourhood in which the heating is done using only gas boilers, which can simply be turned on whenever there is a heat demand. The input required for this first part of the simulation is the weather data for the location of the neighbourhood, as well as a data set containing the loads of the houses in the neighbourhood over the time horizon. This data set includes the electrical and thermal loads in the houses during every time interval, as well as the flexibility of these loads, i.e. whether the loads can be shifted and if so by how much (for more details about this data set and its generation, see Section 6.1.2, from here on we refer to this data set as the load data of the neighbourhood). The output of the first part of the simulation consists of the total heat demand of the neighbourhood as well as any excess of electricity from the neighbourhood (i.e. the amount of generated electricity the neighbourhood could not consume itself).

In the second part of the simulation the heuristic for controlling the Ecovat system is used to satisfy the heat demand of the neighbourhood. The objective of this second part is the same as in the previous chapter, i.e. to minimize the operational costs of the Ecovat system. The input for this part are the outputs of the first part, e.g. the total heat demand of the neighbourhood and the excess electricity, which the system can use to charge the Ecovat buffer. Additionally, like in the previous chapter the heuristic needs weather data and energy prices as input. The output of this part of the simulation consist of the operational costs of the system, the temperature evolution inside the Ecovat buffer and the amount of electricity that is fed back into the grid. Electricity is only fed back if during some time interval the neighbourhood had electricity it could not consume itself and that also could not be used to charge the Ecovat buffer. Figure 6.1 gives an overview of the different components of the simulation.

To perform such a simulation we would in general consider all time intervals in the time horizon iteratively, where we first use DEMKit to optimize the electricity consumption in the neighbourhood, and subsequently use the heuristic for the Ecovat system to satisfy the heat demand of the neighbourhood for the current time interval. However, since in our case there is no feedback from the Ecovat heuristic back to DEMKit, we can simplify the simulation process by splitting it in two distinct parts. First we optimize the electricity consumption of the neighbourhood over the entire year using DEMKit. And second, we use the heuristic for the Ecovat system for satisfying the heat demand of the neighbourhood, where for every time interval we use the output of DEMKit for that time interval.

As discussed above, the input for the first part of the simulation consist of weather data for the location of the neighbourhood, as well as the load data of the houses in the neighbourhood. For the location of the neighbourhood we
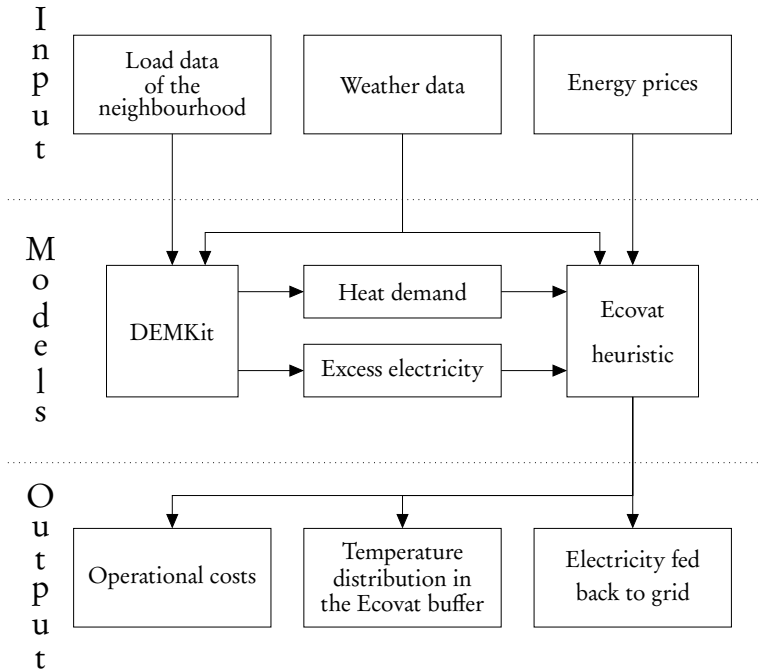
Figure 6.1: Overview of the used models and their inputs and outputs.

choose the municipality of Enschede, as such we use the ambient temperature and solar irradiation profiles measured by weather station Twenthe in Enschede, the Netherlands. For the simulations in this chapter we use the historical data from the the year 2014. The load data of the houses in the neighbourhood are generated using the Artificial Load Profile Generator (ALPG) described in [48]. We give some more background on the ALPG in Section 6.1.2.

Based on the load data and the weather data DEMKit optimizes the electricity profile of the neighbourhood, in our case by flattening it as much as possible, given the flexibility provided by smart devices (which is included in the load data of the houses). DEMKit outputs the heat demand profile of the neighbourhood, as well as the excess local electricity produced within the neighbourhood, which are subsequently used as input for the second part of the simulation. In our simulation we assume that this excess local electricity can be consumed by the Ecovat system at zero cost. This could be in exchange for a lower heating bill, as well as from a desire to improve the degree of self-consumption of the neighbourhood.

In the second part of the simulation the output of the first part, the heat demand of the neighbourhood and the excess electricity, are used as input. Furthermore, as input weather data for the location of the neighbourhood as well as the energy

prices are needed. For the weather data we use the same profiles as in the first step. For the energy prices we use the prices on the Dutch imbalance market of the year 2014.

In general the heuristic method tries to use as much of the local excess electricity as possible to charge the Ecovat buffer, as this improves the self-consumption of the neighbourhood (and decreases the operational costs, assuming the energy would otherwise be bought on the energy market when the energy price is positive). This means that the heuristic method needs to ensure that the local excess electricity is consumed by the Ecovat system if possible. This is done in the following way: if the lthp and/or hthp are scheduled to run during a time interval with an excess of local electricity, this excess local electricity is used to run those heat pumps. After that, if there is any excess local electricity remaining the heuristic first tries to turn on the awhp to consume this excess electricity, and finally the resistance heater if necessary. The awhp is prioritized here because of its higher efficiency compared to the resistance heater. The scheduling of these devices is done using the algorithms presented in Chapter 5. Note, that if the energy price p_i is less than the maximum accepted energy price $\bar{p}_j$ during that day the Ecovat system simply purchases less electricity on the energy market compared to the case considered in Chapter 5, while the scheduling of the devices in the system is not changed. If $p_i \geqslant \bar{p}_j$ some extra devices may be turned on, compared to the case considered in Chapter 5, to consume the excess local electricity. While devices that are turned on based on the energy price during a time interval always run at maximum capacity, devices that are turned on to consume the excess local electricity only run at a power equal to the amount of this excess electricity. This means we are making the simplifying assumption that these devices can run at any power, i.e. they do not require a minimum power to run and are not limited to a number of discrete power values at which they may run.

The values for the parameters used in the heuristic are the same as those given in the previous chapter, with one exception. We use a different initial temperature distribution in the Ecovat buffer to better simulate a real world system. From the results in the previous chapter we observe that after operating for some time the bottom two buffer segments are generally close to their maximum temperatures, while for a demand temperature $T_{dem} = 60°C$ the temperatures of the top three segments do not decrease much below 60°C when their useful energy is depleted. For this reason we take an initial temperature distribution in the Ecovat buffer of $T_{1,s} = \{90, 75, 59.5, 47.5, 4.5\}$ instead of $T_{1,s} = \{90, 75, 50, 30, 5\}$ we used in the previous chapters, to better simulate a real Ecovat system that has been in operation for some time. Note, that for a demand temperature of $T_{dem} = 60°C$ both initial temperature distributions give the same initial useful energy content in the buffer. It is worth mentioning that the Ecovat buffer we are simulating has a diameter of 20 meter, which is relatively small compared to the largest planned buffers which have more than ten times as large volume.

In this case study we only consider a demand temperature of $T_{dem} = 60°C$, as opposed to previous chapters were we also considered a demand temperature of $T_{dem} = 40°C$. The reason for this is that from the results presented in the previous chapters we can see that the performance of the system is always lower for a demand temperature of $T_{dem} = 60°C$, since this is the 'harder' case of the two. As such, if the results from this case study show that the Ecovat system performs well for a demand temperature of $T_{dem} = 60°C$, it should also perform at least as well for a lower demand temperature.

6.1.1   DECENTRALIZED ENERGY MANAGEMENT TOOLKIT

To perform DEM simulations of a neighbourhood of houses the DEMKit simulator, developed at the University of Twente [47], is used. This simulator considers integrated energy systems combining various energy carriers such as e.g. heat and electricity. This open-source simulator allows the user to define a model of a neighbourhood, build up of houses with individually modelled devices. To this end, DEMKit includes a library of device components, such as; washing machines, dishwashers, batteries, heat pumps, PV panels and electric vehicles. Furthermore, DEMKit includes components to control these devices by means of various optimization algorithms.

As mentioned, DEMKit can optimize the consumption of multiple commodities, such as electricity and heat. However, in this case study we only use DEMKit to optimize the electricity consumption within the neighbourhood, while taking the requirements of the heat demand into account. The used optimization approach in this work is the profile steering algorithm [43], which uses a desired power profile (in this case a flat profile) as steering signal instead of electricity prices, which are used in many other tools as steering signal. More specifically, DEMKit uses the extended version of this algorithm presented in [47], which uses a two-phase approach. In the first phase based on predictions a synchronous planning is made for the neighbourhood, while in the second phase profile steering is used to schedule the individual devices. This second phase happens asynchronously and is event driven.

The output of the DEMKit simulator consist of, among others, the consumption and generation of energy on both device and neighbourhood level, and the amount of electricity imported from and exported back to the grid.

6.1.2   ARTIFICIAL LOAD PROFILE GENERATOR

For the generation of the load data of the neighbourhood, which are one of the inputs required by DEMKit, we use the ALPG [48]. This open-source tool uses a bottom-up approach to create an occupancy profile for a given house, by modelling the presence and behaviour of its occupants. Taking into account when the occupants are at home, the usage patterns of devices are created (e.g. a device that requires input of an occupant never runs when no one is at home).
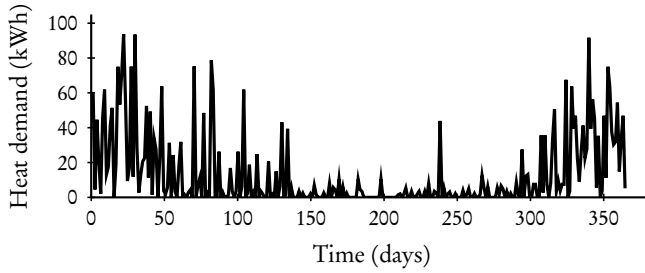
Figure 6.2: Heat demand throughout the year for the h50 case.

The generated profiles include electrical devices, but also consumption of heat used for space heating and hot water. The profiles that the ALPG gives as output consist of one electrical load profile for all uncontrollable devices together (this contains e.g. lighting, televisions and computers), vectors specifying constraints for any controllable smart devices (such as the time range in which a washing machine is required to run) a vector with temperature setpoints for the thermostat, which specify the desired temperature inside the house, a ventilation profile of the house and a hot water consumption profile. This means that the load data also includes the flexibility provided by the smart devices present in the neighbourhood.

To create the load data, the ALPG requires some parameters of the neighbourhood, such as the types of households in the neighbourhood and the devices available in those households. In our case we use the parameters for the neighbourhood specified before, i.e we consider a neighbourhood with a higher penetration of emerging technologies than is currently common, and with a mix of single worker households, dual worker households, households of families with children and households of retired persons.

6.2 Case descriptions

To investigate the benefits of using an Ecovat system instead of gas boilers for satisfying the heat demand of the neighbourhood, as well as determining the influence of a higher heat demand on the performance of the system we consider five different cases. The base case consists of a neighbourhood of 50 houses which aims to maximize its self-consumption of energy (we call this case **h50**). In the base case we assume any excess of electrical energy from the neighbourhood is offered for free to the Ecovat system, where it can be converted to heat and stored for later use by the neighbourhood. In this case self-consumption is preferred even if the energy price is negative. The resulting heat demand profile of the neighbourhood for this case is given in Figure 6.2.

The second case differs from the base case only in its objective, which switches

from maximizing self-consumption to maximizing profit (we call this case **h50-profit**). In this case electricity might be fed back to the grid by the neighbourhood, while simultaneously electricity is bought on the energy market for a negative price by the Ecovat system. This assumes that the neighbourhood is not penalized for feeding back electricity to the grid even though the energy price is negative, or that the neighbourhood and the Ecovat system operate on different markets. In other words the consumers would still receive the normal feed-in tariff from the supplier in this case. It should be noted that while maximizing profit in this way is sound from an economical perspective, it may not be preferable from a social standpoint.

The third case is again similar to the base case with the exception that the heat demand during the winter months is increased by an average of 30%. This will give us an indication how robust the Ecovat system is in handling heat demands which are higher than was expected, for example due to a colder winter than expected (we call this case **h50-winter**). To simulate this case, the previously discussed heat demand profile is taken and the value for the heat demand in every time interval in the months December, January and February is multiplied by a random number between 1 and 1.6 taken from a uniform distribution. To model that the increased heat demand is unexpected, and thus not known on beforehand, this case uses the same targets for the useful energy content of the Ecovat buffer at the end of every day, as those used for case **h50**.

The fourth case is the same as the base case but instead of a 50 house neighbourhood we consider a 65 house neighbourhood, we call this case **h65**. The fifth case combines the third and fourth cases, meaning that we consider a neighbourhood of 65 houses with increased demand in winter (we call this case **h65-winter**). The case **h65-winter** uses the same targets for the energy content of the Ecovat buffer at the end of every day as case **h65**, for the same reasons as for the **h50-winter** case.

Because of the randomness introduced in the heat demands of the **h50-winter** and **h65-winter** cases we have simulated each of these cases five times and determined the average results obtained from these five simulations for both cases. For each of these two cases the simulation that gives the results closest to the average result is presented in the next section. It should be noted that even though some randomness was introduced in the heat demand, the results for all the simulations were close to the average (all within 5% from the average).

6.3 RESULTS

The first goal of this case study is to determine to which extend the Ecovat system can supply the heat demand of a neighbourhood throughout the year. In this case study we consider an Ecovat buffer with a diameter of 20 meter, as in the rest of this thesis. However, larger buffers are planned for construction in the future, which may serve larger neighbourhoods than those considered in this case study.

Case	Operational costs (€)	Useful energy (kWh)
h50	-18894	49569
h50p	-18968	49569
h50w	-18022	43660
h65	-19339	44422
h65w	-15921	42025

Table 6.1: Operational costs and useful energy content of the Ecovat buffer at the end of a year of operation, for the five considered cases.

Figure 6.3 shows the temperature distribution of the Ecovat buffer throughout the year for the cases **h50**, **h50-winter**, **h65** and **h65-winter**. The **h50-profit** case is not included since the temperature distribution is the same for that case as for the **h50** case. This is due to the fact that the decisions whether and how to charge the buffer do not change between these two cases, merely the preferred energy source (excess neighbourhood electricity versus electricity bought on the energy market) changes. The temperature evolutions in this figure show that the Ecovat system is capable of supplying the heat demand of the neighbourhood in all considered cases, as there is always at least one segment with a temperature higher than the demand temperature of $T_{dem} = 60°C$. However, we can see from Figure 6.3 that already in the base case, **h50**, the Ecovat buffer is low on useful energy ($T_2 = 60°C$ and T_1 is around 70°C) at the end of February (around day 60) and almost full during summer (all segments are close to their maximum temperatures), as one would expect. As can be seen in Figure 6.3 the temperature evolution inside the Ecovat buffer is very similar for the different considered cases. While the exact temperature values differ between cases, for example in cases with higher heat demand the temperature of the top segment of the buffer comes closer to the demand temperature of 60°C at the end of the winter, the 'pattern' of the temperature evolution is very similar between all cases.

The second goal of this case study is to investigate the sensitivity of the Ecovat system to changes in the heat demand of the neighbourhood. To do so, we look at the temperature evolution inside the Ecovat buffer throughout the year, the operational costs of the Ecovat system and the useful energy of the buffer at the end of the year for the different cases described in Section 6.2. This gives a good indication about how well the Ecovat system is capable of supplying the heat demand of the neighbourhood for different circumstances. Table 6.1 shows the operational costs of the Ecovat system as well as the useful energy at the end of the simulated year. For reference, the initial useful energy content of the buffer and thus the target at the end of the year is 54246 kWh. As before, we only consider the operational costs and do not include for example the costs for maintenance or the profit made from supplying heat to customers in the neighbourhood.
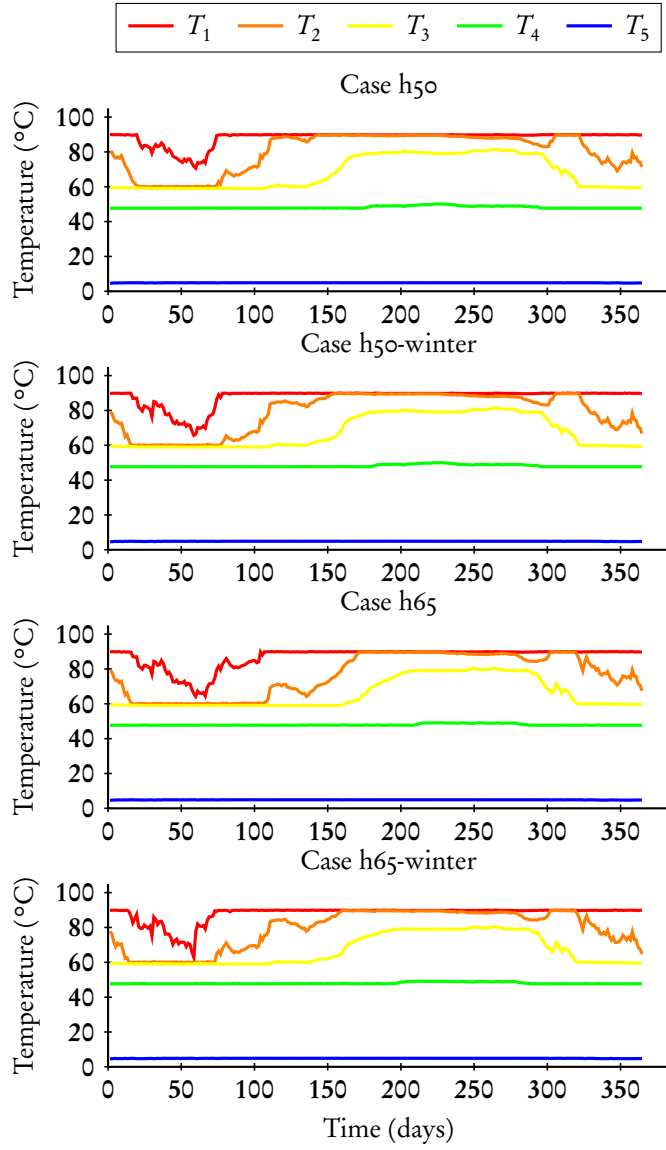
Figure 6.3: Temperature evolution inside the Ecovat buffer for the different simulation cases.

From Table 6.1 we can see that for the **h50** case the useful energy content at the end of the year is about 10% lower than the aforementioned target of 54246 kWh. As mentioned in Section 4.2.5, the cost of a lower useful energy content of the buffer is hard to estimate. However, we concluded that a lower useful energy content of the buffer leads to higher costs during the next year, as the difference in useful energy has to be made up by extra charging during time intervals with positive energy prices, leading to extra costs. When operating a real Ecovat system this deficit would cause the costs to increase in the following year due to starting the year at a lower useful energy content. The costs for the **h50-profit** case are only slightly lower than for the **h50** case. This is due to the fact that the amount of extra electricity that is bought on the energy market (at negative prices) in this case is only small compared to the base case, **h50**. For the **h50-winter** case we see that the buffer is a bit closer to empty at the end of February, but the overall temperature evolution inside the buffer is very similar. As expected, the operational costs are higher than for the **h50** case due to the increased heat demand during winter. Finally, the useful energy content of the buffer at the end of the year is lower than for the **h50** case, again due to the increased demand during winter.

When looking at the **h65** case we see that the operational costs are lower than for the base case, which may be surprising. However, since the energy prices on the market can be negative, a higher energy demand does not necessarily mean higher operational costs. The increase of the energy demand during summer, when the buffer is almost full in all considered cases, actually leads to more opportunities to charge the buffer while the energy price is negative, leading to lower operational costs. However, this does come at the cost of a lower useful energy content compared to the **h50** case. Finally, for the **h65-winter** case we observe higher operational costs and lower useful energy content at the end of the year, compared to both the **h65** and **h50** cases. In this case the extra heat demand, compared to the **h65** case, is concentrated in the winter months. During the winter months the buffer was already getting close to empty in the **h65** case, by only charging during time intervals with low energy prices. This means, that the extra demand in the **h65-winter** case needs to be compensated for by charging the buffer while the energy prices are higher, leading to higher operational costs. Additionally, we observe a lower useful energy content of the buffer at the end of the year. In this case the costs to compensate for the lack of useful energy during the next year of operation would be even larger than for the previous cases.

To investigate the effect of the lower useful energy at the end of the year in the five considered cases, we simulate a second year of operation of the system. For this second year of operation we take the same input data as for the first year, except that we take the temperature distribution (and thus the useful energy content of the buffer) at the end of the first year as the initial temperature distribution in the Ecovat buffer at the start of the second year. The results for these simulations are summarized in Table 6.2, which gives the operational costs during the second

Case	Operational costs (€)	U. energy start (kWh)	U. energy end(kWh)
h50	-18836	49569	49760
h50p	-18909	49569	49760
h50w	-17916	43660	43662
h65	-19139	44422	44565
h65w	-15760	42025	42024

Table 6.2: Operational costs and useful energy content of the Ecovat buffer at the end of the second consecutive year of operation, for the five considered cases.

year of operation as well as the useful energy content of the buffer at start and the end of the second year. We observe that the useful energy content at the end of the second year is very similar to the useful energy content at the start of the year. This means that amount of energy used to charge the buffer throughout the year is almost equal to the sum of the heat demand throughout the year, which gives a better indication of the actual operational costs during a year of operation. From Table 6.2 we observe slightly higher operational costs compared to the first year for all considered cases. This is expected, since the Ecovat buffer starts with a lower useful energy content in this case, which needs to be compensated for by extra charging of the buffer at positive energy prices during winter. Furthermore, we see the same trends as in the first year, i.e. the winter cases have higher operational costs due to the extra demand during the winter months, when the buffer is almost empty already, while extra demand throughout the entire year does not necessarily lead to higher operational costs, due to extra charging opportunities during the summer when the buffer would be full otherwise.

We observe that for all cases considered the Ecovat system is capable of supplying the heat demand of the connected neighbourhood of houses throughout the entire year. The results show that a higher heat demand due to an unexpectedly cold winter leads to higher operational costs as expected. Furthermore, we observed that a higher heat demand due to a larger neighbourhood does not always lead to higher operational costs. However, as seen from Table 6.2 the difference between **h50** and **h50-winter** cases is much smaller than the difference between **h65** and **h65-winter** cases. The reason this difference is larger is that in the **h65** case a larger number of time intervals with low energy prices is already used to charge the buffer, since that case has a higher heat demand than the **h50** case because of the extra houses. This means that an extra increase in the heat demand on top of that, due to a cold winter, forces the model to charge during time intervals with higher energy prices in the **h65** case compared to the **h50** case, leading to a larger difference in operational costs between the **h65** and **h65-winter** cases than the **h50** and **h50-winter** cases. This means that even though a larger neighbourhood may lead to lower operational costs when there is no unexpected increase in heat demand, the risk of much higher operational costs during a cold

Case	Electricity fed back (kWh)	Decrease (%)
No Ecovat, 50 houses	44596	-
No Ecovat, 65 houses	55778	-
h50	4546	89.8
h50p	5454	87.8
h50w	4283	90.4
h65	9148	83.6
h65w	7308	86.9

Table 6.3: The amount of electricity fed back to the grid for the five considered cases. Also shown is the relative decrease in the amount of electricity fed back compared to using gas boilers to satisfy the heat demand of the neighbourhood.

winter are also higher. In other words, for a larger neighbourhood the Ecovat system is more sensitive to unexpected changes in the heat demand. A question for further research is which neighbourhood size is optimal for an Ecovat system, especially when taking into account that a larger neighbourhood will also lead to higher profit for satisfying the heat demand of the neighbourhood. However, this question is outside the scope of this thesis.

The third goal of this case study is to determine the benefits of using an Ecovat system to supply the heat demand of a neighbourhood of houses instead of using gas boilers, in terms of CO_2 reduction and self-consumption. One of the benefits of using an Ecovat system instead of gas boilers is that part of the excess electricity produced in the neighbourhood, for example from PV panels, can be used to charge the Ecovat system instead of being fed back to the electrical grid. Table 6.3 shows the amount of electricity fed back to the grid over the entire year for the considered cases. The 'No Ecovat' cases give the amount of electricity fed back to the grid if gas boilers are used instead of the Ecovat system. The second column gives the percentage decrease of electricity fed back to the grid compared to the 'No Ecovat' case for the same number of houses. We can see that for all the considered cases the amount of electricity fed back to the grid is decreased significantly if an Ecovat system is used instead of gas boilers to satisfy the heat demand. Even for the **h50-profit** case we see a large decrease in electricity being fed back, the decrease being only slightly smaller than for the base case **h50**. In the **h50-profit** case extra electricity, compared to the **h50** case, is fed back to the grid during times where there is both excess local production as well as negative energy prices on the energy market. This happens only for a small number of time intervals. As such, the majority of the excess local production is used to charge the Ecovat system. This extra feedback of electricity would translate to some extra profit in the **h50-profit** case compared to the **h50** case. However, this is only a small amount (about 41 € if we assume a feedback price of 0.04 €/kWh). Combined with the small decrease in operational costs in the **h50-profit** case compared to the **h50** case the total increase in profit is relatively

Case	Total heat demand (kWh)	Reduction CO_2 (ton)
h50	$4.5 \cdot 10^5$	91.4
h50-winter	$5.3 \cdot 10^5$	107.2
h65	$5.9 \cdot 10^5$	118.6
h65-winter	$6.9 \cdot 10^5$	138.8

Table 6.4: CO_2 emission reduction when using an Ecovat system to satisfy the heat demand of the simulated neighbourhood, compared to satisfying the heat demand using gas boilers.

small, which means that even a small incentive to increase self-consumption is enough to disregard the **h50-profit** option. Finally, we note that most of the electricity that is still fed back to the grid in all the considered cases is fed back during summer when the Ecovat buffer is at its capacity.

Another benefit of using the Ecovat system to supply the heat demand instead of using gas boilers is the reduction in CO_2 emissions. The Ecovat system produces zero emission during the charging and discharging of the buffer. Which implies that, assuming the energy used to charge the buffer is generated in a sustainable way, the Ecovat system is emmision free. This means that using an Ecovat system instead of gas boilers causes a CO_2 emission reduction equal to the total emission of the gas boilers. Table 6.4 lists the total heat demand during the year for each case, as well as the amount of CO_2 reduction that can be achieved by replacing the gas boilers by an Ecovat system. The CO_2 reduction is calculated using the emission factor of 56 ton CO_2/TJ for natural gas in the Netherlands [88].

Finally, we compare a scenario with an Ecovat system with an 'all electric' neighbourhood, where the heat demand of the neighbourhood is supplied by heat pumps. Using heat pumps to satisfy the heat demand has the same benefit as the Ecovat, namely no CO_2 emissions. As it is not in the scope of this work to carry out a detailed comparison of such a neighbourhood using an elaborated DEM method, we present a simplified analysis to give a first impression of the advantages and disadvantages of using the Ecovat system to satisfy the heat demand of the considered neighbourhood compared to using heat pumps. To this end, we make some simplifying assumptions. First, we assume the heat pumps have a constant coefficient of performance, and we chose the value 4 for this coefficient. And second, we assume that during any time interval in which the neighbourhood has a non-zero heat demand, this heat demand is directly supplied by the heat pumps during that time interval. This implies that the heat pumps are not combined with a buffer, and we do not apply any smart control to the heat pumps, they simply run at the exact time interval during which heat is demanded. For the comparison presented here there is no difference between considering a large central heat pump for the entire neighbourhood or smaller heat pumps for each individual house. However, note that when considered in
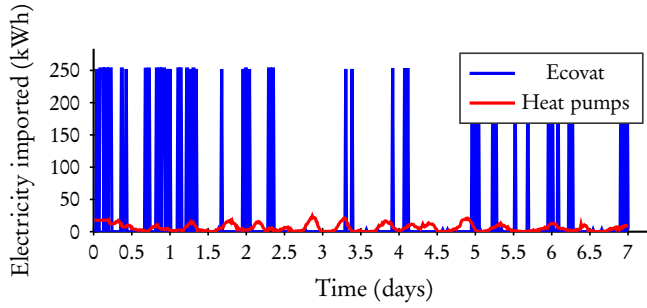
Figure 6.4: Imported electricity to satisfy the heat demand of the neighbourhood.

Case	Total electricity imported (kWh)	Costs (€)
Ecovat	$3.7 \cdot 10^5$	-18894
Heat pumps	$1.3 \cdot 10^5$	5069

Table 6.5: Amount of electricity imported to satisfy the heat demand of the neighbourhood and the associated costs.

a real implementation the option of individual heat pumps does not require a district heating network to be present. This in contrast to a central heat pump or the Ecovat system, where such a system is needed to distribute the heat to the houses. Finally, we have chosen to perform the comparison for the case **h50** described in Section 6.2.

The advantage of using heat pumps compared to the Ecovat system is the fact that a lot less electricity needs to be imported to satisfy the heat demand of the neighbourhood, due to the high coefficient of performance of the heat pumps (compared to the coefficient of performance of the resistance heater, which consumes by far the most electricity in the Ecovat system). We can see this difference in Figure 6.4, where the imported electricity is shown for the first week of January for both the Ecovat system as well as the heat pumps. Furthermore, we can see that the Ecovat system has much higher peaks in the imported electricity, but also imports electricity in a much lower number of time intervals.

The total imported electricity over the year is given in Table 6.5 for both cases. As mentioned before, the total import of electricity is significantly higher for the Ecovat system compared to using heat pumps for satisfying the heat demand. However, we see that the associated costs of these imports are much lower for the Ecovat system. The reason for this is that the Ecovat system can import electricity during time intervals with low energy prices, while in the case where heat pumps are used the electricity has to be imported during the time interval in which there is a heat demand.

While this comparison shows that using the Ecovat system leads to a large financial advantage compared to using heat pumps to satisfy the heat demand, we need to take into account that we have used smart control for the Ecovat system while no smart control is used in the heat pumps case. To make a fair comparison we would have to consider a case where smart control is used for controlling the heat pumps as well, which we expect will lower the costs of the imported electricity for the case using heat pumps. However, unless a buffer is used in combination with the heat pumps, there is only limited flexibility in when to turn on the heat pumps. While some pre-heating of houses may be done during time intervals of low energy prices, this is limited to only a short time before the expected heat demand. Furthermore, to do so predictions of the energy prices for the next time intervals are required, which the heuristic method used to control the Ecovat system does not require. As such, we expect that when using smart control for the heat pumps the costs will decrease, but will not be as low as those for the Ecovat system. Even though the Ecovat system requires substantially more imported electricity, the large advantage of the Ecovat system is the ability to charge during time intervals with low energy prices, even if those are already months in advance of the actual heat demand. Note, that this assumes that time intervals with negative energy prices occur (in which case requiring more electricity is not always purely a disadvantage since it gives more opportunities to make a profit). If time intervals with negative energy prices would not occur any more, or would become much rarer, the costs for the Ecovat system would increase significantly, in which case heat pumps might become a more attractive alternative.

To conclude, the advantages of using an Ecovat system compared to gas boilers to satisfy the heat demand of a neighbourhood include a higher self-consumption of electricity in the neighbourhood, as well as a reduction in the CO_2 emissions. The advantage of using an Ecovat system compared to heat pumps to satisfy the heat demand of a neighbourhood is a lower cost according to the simple comparison presented in this chapter. However, the exact difference is hard to specify without a more in depth comparison, which is outside the scope of this thesis.

6.4 CONCLUSION

In this chapter we presented a case study in which we simulate a neighbourhood of houses including an Ecovat system. The goals of this case study are: 1) to determine whether the Ecovat system can supply the heat demand of a neighbourhood throughout the year, 2) to determine the sensitivity of the system to major changes in heat demand, and 3) to determine the benefits in terms of CO_2 reduction and self-consumption of using an Ecovat system to supply the heat demand of the neighbourhood, compared to using gas boilers. To achieve these goals we use the DEMKit simulator developed at the University of Twente [47] to optimize the electricity consumption in this neighbourhood, and the heuristic

method for controlling the Ecovat system, presented in Chapter 5, to determine the charging/discharging strategy for the Ecovat system, while satisfying the heat demand of the neighbourhood.

The results show that the considered Ecovat system is capable of supplying the heat demand of a neighbourhood of 50 or 65 houses throughout the year, even if the heat demand is unexpectedly higher, such as due to a colder winter than expected. For a neighbourhood of 50 houses, such a cold winter leads to a slight increase in operational costs, while the increase in operational costs for a neighbourhood of 65 houses during a cold winter is much larger. This means that the Ecovat system is more sensitive to unexpected changes in the heat demand when the size of the neighbourhood increases. The reason for this is that extra energy is demanded during winter, at times when the Ecovat buffer is almost empty and all time intervals with low energy prices are already used to charge the buffer. This extra energy then needs to be charged during time intervals with higher energy prices, which leads to larger differences in operational costs when considering a larger neighbourhood. The reason for this is that a larger neighbourhood has a higher energy demand initially due to the increased number of houses, which leads to a larger increase in costs when even more energy is demanded.

The presented case study shows that the benefits of using an Ecovat system instead of gas boilers for satisfying the heat demand of the neighbourhood include a reduction in the electricity fed back to the grid by 83.6 to 90.4% (increased self-consumption), as well as a reduction in CO_2 emission by 91.4 to 138.8 ton for the considered cases.

7

CONCLUSION

In this thesis an approach to determine charging/discharging strategies for an Ecovat system, a seasonal thermal storage technology, throughout an entire year is presented. The requirements for the approach were that the resulting charging/discharging strategies for the Ecovat system should be of good quality and that the needed computational time should be low enough such that the approach is usable in practice. The latter means that the computational time required for such an approach needs to be short. Therefore, a heuristic method that satisfies these requirements was developed. This heuristic method was evaluated in a case study, to determine the benefits of using an Ecovat system instead of gas boilers to satisfy the heat demand of a neighbourhood. In this chapter we summarize and conclude the results of the thesis.

The remainder of this chapter is structured as follows. In Section 7.1 we give a short summary of the preceding chapters. Subsequently, in Section 7.2 we answer the research question and the sub-questions presented in Chapter 1. Finally, in Section 7.3 we give some recommendations for future work.

7.1   SUMMARY

The goal of the work presented in this thesis was to develop an approach for the operational control of an Ecovat system that gives good charging strategies, i.e. strategies that have a low operational cost while satisfying the heat demand of the neighbourhood, and does so in a computational time that is short enough to be used in practice, i.e. when controlling an Ecovat system in real-time. To determine an approach that satisfies these constraints we started by developing a control method based on an ILP model for the control of the operation of an Ecovat system, without taking into account the restriction on the computational time. In Chapter 3 we presented such an ILP model, in which the restrictions and constraints of the Ecovat system are taken into account. At the end of Chapter 3 we concluded that even though an approach based on the presented ILP model would give good charging/discharging strategies, it is not useful without

modifications. The reason for this is that to model the operation of an Ecovat system, which is a seasonal thermal storage technology, we need to consider optimizing over a long time horizon (which due to seasonal effects should be at least 1 year), while for the energy market that we wish to incorporate time intervals of 15 minutes are used. This combination of a optimization horizon, short time intervals and the complexity of the system leads to an approach that is by far not solvable to optimality for the entire time horizon at once.

Subsequently, in Chapter 4 we introduce a rolling horizon approach to solve the ILP model presented in Chapter 3. This allows us to solve the ILP model based approach in reasonable computational time. However, we observed that solving the ILP model based approach in this way fails to incorporate seasonal effects, since the rolling horizon length could only be a few days, while to incorporate these seasonal effects the approach would need to be able to look ahead for a few months. To overcome this problem we introduced a long-term planning step, which is executed before solving the ILP model. In this long-term planning step we generate energy targets for the Ecovat buffer for each day of the year. The ILP model is then altered slightly, by introducing an extra term in the objective function that penalizes the model if the useful energy content of the buffer drops below these targets. This ensures that the approach incorporates seasonal effects, for example by making sure that the Ecovat buffer is full at the end of the summer, to ensure that the heat demand of the neighbourhood can be supplied throughout the winter without incurring unnecessarily high operational costs. We observed that with this extension the ILP model based approach performs well, but that it still does not satisfy the constraint of having a computational time short enough to be usable for the real-time control of an Ecovat system.

Therefore, in Chapter 5 we presented a heuristic method to control the Ecovat system. The goal of this heuristic method is to provide charging/discharging strategies for the Ecovat system that are of comparable quality to the strategies obtained with the ILP model based approach, while requiring a computational time that is short enough to be used in practice. To accomplish this we derived a set of rules of thumb based on the insights gained from the ILP model based approach. The heuristic method simply follows these rules of thumb to construct charging/discharging strategies, leading to very short computational times. In contrast to the ILP model based approach the heuristic method does not require input data for time intervals in the future, such as predictions of weather data or energy prices, but instead only uses input data that are known at the time interval it is considering. When comparing the charging/discharging strategies obtained with the heuristic method to the strategies obtained with the ILP model based approach we observe that for the cases considered the heuristic on average only performs 5.2% worse and in the worst case shows a 14 % decrease in performance (i.e. increased operational costs). The differences are especially small when considering a lower demand temperature. Furthermore, the computational time of the heuristic method is only a few seconds, compared to the 8-15 hours generally required by the ILP model based approach (with some extreme

cases taking more than a day).

Finally, in Chapter 6 we presented a case study of a neighbourhood including an Ecovat system. In this case study we used the previously developed heuristic to control the Ecovat system in this neighbourhood. In this case study a detailed simulation of the considered neighbourhood is executed to achieve a proper reflection of the energy needs of a (future) neighbourhood. For the simulation of the neighbourhood the Decentralized Energy Management Toolkit (DEMKit) simulator is used. In detail, we first optimize the electricity consumption and production in the neighbourhood using DEMKit, with the objective of flattening the electricity profile of the neighbourhood. In this first step we assume the heat demand of the neighbourhood is supplied by gas boilers. Subsequently, we use the heuristic method for controlling the Ecovat system to determine the way the Ecovat system would satisfy the heat demand of the neighbourhood instead. In this way we can determine the benefits of using the Ecovat system instead of using gas boilers to satisfy the heat demand. We observe that for the considered cases the amount of electricity fed back to the grid is reduced by 83.6 to 90.4%, increasing the self-consumption of the neighbourhood. Furthermore, we observe a reduction in CO_2 emission by 91.4 to 138.8 ton for the considered cases when using the Ecovat system instead of gas boilers for satisfying the heat demand of the neighbourhood.

7.2 Conclusions

In Chapter 1 we presented the main research question of this thesis, namely:

> *How can we model and determine the operational control of the Ecovat system such that the resulting method provides good charging/discharging strategies, with a computational time short enough to be usable in a real world situation?*

In this section we conclude the thesis by answering this question. We do so by considering the sub-questions introduced in Chapter 1. The first sub-question is:

1. *What is a possible method for determining the control of the Ecovat system, which is able to provide a good charging strategy for the Ecovat, assuming computational time is not a limitation?*

When computational time is not a limiting factor, a global optimization model is a good choice to determine charging/discharging strategies for the Ecovat system. We presented such a model in Chapter 3. In this ILP model we integrated the components and restrictions of the Ecovat system as constraints. If computational time is not a limitation we can simply solve this model for a time horizon of any length (assuming we have access to a computer with sufficiently large memory), and obtain a charging/discharging strategy for the complete time horizon. This would give good charging/discharging strategies for the Ecovat

system, since all time intervals are considered at once when optimizing the operation of the system. However, in practice the computational time of solving the presented ILP model is already very long for only relatively short time horizons of a few days. This is due to the fact that the model is complex, and we require short time intervals of 15 minutes to incorporate the energy market of interest into the model. Furthermore, since we are also interested in the seasonal behaviour of the Ecovat system we require a time horizon of a year, which leads to computational times that are much too long for a practical application. This leads to the second sub-question, namely:

2. *How can we adapt the developed control method to be solvable given limited computational time and resources, while maintaining a good charging strategy?*

To ensure the ILP model presented in Chapter 3 becomes solvable given limited computational time and resources we started by introducing a rolling horizon approach to solve the ILP model in Chapter 4. In this approach we iteratively determined the charging/discharging strategy for of the Ecovat system for every day in the time horizon. We did so by solving the ILP model based approach for a few days at a time (2,3 or 5 days in Chapter 4) and realizing just the first day of these few days. While this approach allows us to create solutions for the ILP model based approach in a computational time of a few days/weeks, we observed that it is insufficiently capable of incorporating important seasonal effects.

To ensure the ILP model based approach can incorporate such seasonal effects we introduced a long-term planning, to be executed before creating a solution using the ILP model. The goal of this long-term planning is to generate a daily useful energy target for the Ecovat buffer. We then integrated these targets into the ILP model. The results show that with the addition of long-term planning the ILP model based approach does incorporate seasonal effects correctly and the computational time is further reduced to a few hours/days, depending on the considered case.

Creating a solution for the ILP model in this way leads to an increase of the objective value, i.e. higher operational costs, compared to solving the model for the entire time horizon at once. However, it also simulates a more realistic use case in which predictions for the input data, especially the energy prices, are not available for the entire time horizon on beforehand.

Even though predictions for the input data are in general not available, the ILP model based approach (which does require such predictions) provides a good charging/discharging strategy to use as a benchmark for a simpler model (which will not require such predictions) that can be used for the real time control of an Ecovat system. This leads to the third sub-question, namely:

3. *How can we use the insights obtained from the approach developed for the second sub-question to derive a heuristic method capable of controlling an Ecovat system in real time, while maintaining a good charging strategy?*

We used the insights obtained from the ILP model based approach to develop a heuristic method for controlling the Ecovat system presented in Chapter 5. Based on these insights and the output data generated by the ILP model based approach for the cases described in Chapter 4 we derive a set of rules of thumb for determining a charging/discharging strategy for the Ecovat system. The heuristic method then simply applies these rules of thumb to every time interval in the time horizon in an iterative way. Furthermore, in this approach only data for this time interval is taken into account and no predictions for future time intervals are needed. This means that the heuristic method is not only a lot simpler than the ILP model based approach, which we see reflected in the computational times, which are only a few seconds for the heuristic method compared to the hours/days for the ILP model based approach, but it is also not dependent on predictions.

From the results presented in Chapter 5 we observe that the heuristic method shows an average decrease in performance of 5.2% compared to the ILP model based approach for the considered cases. We conclude that the charging/discharging strategies given by the heuristic method are still good enough, since we are comparing them to strategies which can only be obtained in a situation where perfect predictions are available for some of the input data a day in advance, which in general will not be the case in a real world situation. In practice, some of the input data for a given time interval will only be available when that time interval starts and obtaining good predictions of these inputs is hard, especially for the energy prices which are very volatile. As such, a decrease in performance of 5.2% seems acceptable compared to an approach which does assume perfect predictions a day in advance.

The developed heuristic method only indirectly uses predictions, namely for obtaining the daily useful energy targets for the Ecovat buffer, which may be based on historical data. Since such predictions may be unreliable, ideally the influence of errors in these predictions on the charging/discharging strategies generated by the heuristic method is only small. This leads to the final subquestion, namely:

4. *How robust is the developed approach of the Ecovat system to input data differing from predictions?*

When generating the daily useful energy targets for the Ecovat buffer, predictions for the energy prices and the heat demand of the neighbourhood are required. In Chapters 4 and 5 we investigate the effect of errors in the predictions for the energy prices on the results of the developed approaches. We consider two cases for the required predictions of the energy prices. In the first case we assume that perfect predictions are available when generating the targets, and thus we know how to distribute the amount of charging (which the Ecovat system has to do during the year to supply the heat demand of the neighbourhood) over the days in the year to achieve the lowest costs for this charging. In the second case

we assume there are no predictions available, and simply distribute the amount of charging that needs to be done throughout a year over all the days. These two cases generate two different sets of targets that can be used as input for the developed approaches. Comparing the results using these two target sets as input we observe that the difference in the results between the two considered cases is very small. As such, we conclude that the approaches are robust against changes in the predictions of the energy prices. As long as the targets somewhat follow the expected seasonal behaviour the results differ only slightly.

In the case study presented in Chapter 6 we investigate the effect of a heat demand that is higher than predicted when generating the daily useful energy targets. For this we compare cases where the heat demand is as predicted, with cases where winter is colder than was predicted when generating these targets. The results presented in Chapter 6 show that the Ecovat system is capable of dealing with a heat demand that is higher than was predicted. As expected, the operational costs increase when the demand is unexpectedly high. We see that this increase is small for a neighbourhood of 50 houses, but becomes significantly larger when considering a neighbourhood of 65 houses with the same size of the Ecovat system. In other words, when an Ecovat system of the same size is connected to a larger neighbourhood the influence of an increased heat demand compared to the predictions leads to a larger increase in the operational costs, i.e. the robustness of the approach to changes in the predicted heat demand decreases when the neighbourhood size increases. A likely reason for this is that in case of a larger number of houses the Ecovat system is operated more at its limits with respect to the capacity of the Ecovat buffer.

From the discussion above we conclude that by using the heuristic method presented in this thesis, we control the Ecovat system in a way that provides good charging/discharging strategies, in a time frame short enough to be usable in a real world situation, thus answering the main research question. We observe that the heuristic method requires only a few seconds to simulate a year of operation of the Ecovat system, which makes it fast enough to be used for the real-time control of an Ecovat system, where quick decision making is important. Furthermore, we show that the heuristic method does not require predictions for input data of future time intervals to still obtain charging/discharging strategies that are on average only 5.2% worse than those resulting from the ILP model based approach, which does require predictions for the input data to be available for an entire day in advance, and as such gives an upper bound on the performance of the heuristic method. Summarizing, the heuristic method gives good charging/discharging strategies for the Ecovat system in a time frame short enough to be usable for the real-time control of an Ecovat system.

7.3   Recommendations for future work

In this final section of the thesis we discuss some potential avenues for future work on the modelling of the Ecovat system, in no particular order.

  » In this thesis we only considered one device configuration for the Ecovat system, namely a configuration consisting of PVT panels, two water/water heat pumps, an air/water heat pump and a resistance heater. However, an interesting research direction for future work would be to consider different device configurations and determine their advantages and disadvantages. An example which seems promising is to consider removing or replacing the PVT panels in the system. In the configuration considered in this thesis, the contribution of the PVT panels is relatively small compared to that of the other devices, while it places some severe restrictions on the system. Most importantly, to keep the panels cooled, and thus keep their efficiency high, the temperature in the bottom buffer segment has to be kept very low. By removing or replacing the PVT panels this restriction can be dropped and gives the system a lot of extra flexibility. Another option would be to connect the output of the PVT panels to a small buffer tank, and use a heat pump to increase the temperature before charging the Ecovat buffer. This would also allow us to drop the low maximum temperature restriction on the bottom segment of the buffer. Alternatively, heat pipes could be used instead of PVT panels. Heat pipes produce heat of a much higher temperature than PVT panels, which make them a more interesting energy source for the Ecovat system.

  » It may be interesting to explore cheap locally available heat sources to charge the Ecovat buffer, instead of just relying on the energy market. An example of this is that a lot of industrial processes produce heat which is useless to the process itself and is thus considered a waste product. This waste heat, which is often of high enough temperature to be useful for the Ecovat system, could be a good cheap energy source for charging the system. While this will not be a feasible option for every residential area, since it requires closeby industry producing waste heat, it is an interesting option to investigate for the areas where it is available.

  » For the work presented in this thesis a number of assumptions concerning the physical properties of the Ecovat system have been made, since there was no measurement data available from an Ecovat system operating in a real world situation at the time of writing. In the future it would be beneficial to verify whether the made assumptions, for example on the effects of charging/discharging the buffer on the temperatures inside the buffer, hold in practice or whether they should be modified, which can be done by analysing the measurement data from an Ecovat system that is operating in a real world situation.

  » As mentioned in Chapter 6 it would be of interest to perform an economical analysis on the ideal neighbourhood size in relation to the size

of an Ecovat system. As we observed in Chapter 6 connecting a larger neighbourhood to an Ecovat system of the same size leads to larger risks of increased operational costs if there is an unexpected increase in the heat demand of the neighbourhood, for example due to an unexpectedly cold winter. However, a larger neighbourhood would of course lead to a larger profit for supplying the heat demand of the neighbourhood. This leads to the question which neighbourhood size is a good fit for a given size of the Ecovat system. Additionally, such an analysis can determine whether an Ecovat system is economically feasible for a given neighbourhood. In that case the analysis should also include other costs and profits which we have ignored in this work, such as the maintenance costs and the revenue from supplying heat. We assumed they are equal for different charging/discharging strategies, however, they are important when determining the economic feasibility of an Ecovat system for a given neighbourhood.

Acronyms

A   ALPG   Artificial Load Profile Generator
    awhp   air-water heat pump

C   CFD   computational fluid dynamics
    CHP   combined heat and power
    COP   coefficient of performance

D   DEM   decentralized energy management
    DEMKit   Decentralized Energy Management Toolkit
    DSM   demand side management

H   hthp   high temperature heat pump

I   ILP   integer linear programming

L   lthp   low temperature heat pump

P   PCM   phase change materials
    PV   photovoltaic
    PVT   photovoltaic thermal

S   STES   seasonal thermal energy storage

T   TES   thermal energy storage

Notation	Description
A	Total size of the PVT panels.
a_{el}	Electrical loss coefficient of the PVT panels.
a_{th}	Thermal loss coefficient of the PVT panels.
β	Heat loss factor of the Ecovat buffer over 6 months.
C^{dev}	Capacity of device dev.
C_{max}	Maximum allowed useful energy content of the Ecovat buffer when generating energy targets.
C_{min}	Minimum allowed useful energy content of the Ecovat buffer when generating energy targets.
c_p	Specific heat of water.
COP^{dev}	Coefficient of performance of device dev.
D_i	Heat demand during interval i.
Δt	Length of a time interval.
e_i	Amount of energy that can be stored in the Ecovat buffer during interval i.
E_i^{pvt}	Electrical energy produced by the PVT panels during interval i.
f_i	Flag that indicates whether adding time interval i to the solution of the algorithm to determine the daily energy targets violates the upper bound.
G_i	Global radiation during interval i.
K_i	Cost of the system during interval i.
K_{tot}	Total cost of the system summed over all intervals.
LB_j	Lower bound up to time interval j in algorithm to determine daily energy targets.
m_s	Mass of segment s.
$\dot{m}$	Mass flow rate of the heat transfer fluid through the PVT panels.

Notation	Description
η_0^{el}	Electrical efficiency of the PVT panels at a reduced temperature of zero degrees.
η_i^{el}	Electrical efficiency of the PVT panels during interval i.
η_{max}^{el}	Maximum electrical efficiency of the PVT panels.
N_{int}	Number of intervals in the time horizon.
N_{seg}	Number of segments in the Ecovat buffer.
η_0^{th}	Thermal efficiency of the PVT panels at a reduced temperature of zero degrees.
η_i^{th}	Thermal efficiency of the PVT panels during interval i.
η_{max}^{th}	Maximum thermal efficiency of the PVT panels.
p_i	Energy price during interval i.
$\bar{p}_j$	Maximum energy price accepted during day j.
$Q_{i,s}^{dev}$	Thermal energy change in segment s during interval i due to device dev.
$Q_{i,s}^{loss}$	Thermal energy lost to the surroundings from segment s during interval i.
T_i^{amb}	Ambient temperature.
T_{dem}	Demand temperature.
T_{max}^{dev}	Maximum temperature of device dev.
T_{min}^{dev}	Minimum temperature of device dev.
T_{gw}	Temperature of the ground water.
T_i^{in}	Input temperature of the PVT panels.
$T_{i,s}$	Temperature of segment s at the start of interval i.
$\bar{T}_s$	Maximum temperature of segment s.
T_i^{out}	Output temperature of the PVT panels.
T_i^{red}	Reduced temperature.
U_i	Useful energy content of the buffer at the start of time interval i.
$\bar{U}$	Maximum useful energy content of the buffer.
UB_j	Upper bound up to time interval j in algorithm to determine daily energy targets.
$V_{\text{tar}(j)}$	Target energy content of the Ecovat buffer for day j.
$x_{i,s}^{dev}$	Decision variable in the ILP model indicating whether device dev is connected to segment s during interval i.
$y_{i,s}^{dev}$	Decision variable in the heuristic indicating whether device dev is connected to segment s during interval i.

Notation	Description
z_i	Indicator variable that indicates whether energy is stored in the Ecovat during interval i.

BIBLIOGRAPHY

[1] Technology Roadmap - Smart Grids. technical report, International Energy Agency (IEA). Last accessed on March 7, 2019. (Cited on page 29).

[2] Aandeel hernieuwbare energie naar 6,6 procent, CBS. [Online] https://www.cbs.nl/nl-nl/nieuws/2018/22/aandeel-hernieuwbare-energie-naar-6-6-procent, *Dutch*. Last accessed on November 1, 2018. (Cited on page 3).

[3] Energie in Nederland 2018, Energie Beheer Nederland. [Online] https://www.ebn.nl/wp-content/uploads/2018/01/EBN-Infographic-2018-pdf.pdf, *Dutch*. Last accessed on April 26, 2019. (Cited on page 6).

[4] Ecovat Renewable Energy Technologies BV. [Online] www.ecovat.eu. Last accessed on March 1, 2019. (Cited on page 7).

[5] Directive 2009/28/EC of the European Parliament and of the Council of 23 April 2009 on the promotion of the use of energy from renewable sources and amending and subsequently repealing Directives 2001/77/EC and 2003/30/EC. [Online] http://data.europa.eu/eli/dir/2009/28/oj. Last accessed on November 1, 2018. (Cited on page 3).

[6] Our World in Data. [Online] https://ourworldindata.org/fossil-fuels. Last accessed on October 11, 2018. (Cited on page 2).

[7] Energieagenda: Naar een CO_2-arme energievoorziening, Ministerie van Economische Zaken. [Online] https://www.rijksoverheid.nl/documenten/rapporten/2016/12/07/ea, *Dutch*. Last accessed on November 1, 2018. (Cited on page 6).

[8] IRENA (2018), Renewable capacity statistics 2018, International Renewable Energy Agency (IRENA), Abu Dhabi. [Online] http://www.irena.org/-/media/Files/IRENA/Agency/Publication/2018/Mar/IRENA_RE_Capacity_Statistics_2018.pdf. Last accessed on October 12, 2018. (Cited on page 4).

[9] United Nations Convention on Climate Change. Adoption of the Paris Agreement, 2015. [Online] https://unfccc.int/resource/docs/2015/cop21/eng/l09r01.pdf. Last accessed on October 11, 2018. (Cited on page 1).

[10] Ondiepe bodemenergie - open systemen. Rijksdienst voor Ondernemend Nederland (RVO), [Online] https://www.rvo.nl/onderwerpen/duurzaam-ondernemen/duurzame-energie-opwekken/nationaal-expertisecentrum-warmte/bodemenergie-en-aardwarmte-geothermie/bodemenergie, *Dutch*. Last accessed on February 19, 2019. (Cited on page 17).

[11] Energieverbruik 2016, Energie Beheer Nederland. [Online] `http://www.energieinnederland.nl/2016/energieverbruik`, *Dutch*. Last accessed on April 26, 2019. (Cited on page 6).

[12] Beleidsregel redelijke terugleververgoedingen vergunninghouders elektriciteit. Staatscourant, [Online] `https://zoek.officielebekendmakingen.nl/stcrt-2006-82-p28-SC74915.html`, *Dutch*, 2006. Last accessed on December 7, 2018. (Cited on page 51).

[13] Kamerbrief Stimuleringsbeleid lokale hernieuwbare elektriciteitsproductie. [Online] `https://www.rijksoverheid.nl/documenten/kamerstukken/2018/06/15/kamerbrief-stimuleringsbeleid-lokale-hernieuwbare-elektriciteitsproductie`, *Dutch*, 2018. Last accessed on March 8, 2019. (Cited on page 51).

[14] Warmtebaan Enschede. [Online] `https://ennatuurlijk.nl/warmtebaan-enschede`, *Dutch*, 2018. Last accessed on December 13, 2018. (Cited on page 27).

[15] D. Akinyele and R. Rayudu. Review of energy storage technologies for sustainable power networks. *Sustainable Energy Technologies and Assessments*, 8:74–91, 2014. (Cited on page 15).

[16] I. Aldaouab and M. Daniels. Microgrid battery and thermal storage for improved renewable penetration and curtailment. In *2017 International Energy and Sustainability Conference (IESC)*, 2017. (Cited on page 30).

[17] K. Allen, T. von Backström, and D. Kröger. Packed rock bed thermal storage in power plants: design considerations. *Energy Procedia*, 49:666–675, 2014. (Cited on page 19).

[18] G. Alva, L. Liu, X. Huang, and G. Fang. Thermal energy storage materials and systems for solar energy applications. *Renewable and Sustainable Energy Reviews*, 68:693–706, 2017. (Cited on pages 16 and 17).

[19] G. Alva, Y. Lin, and G. Fang. An overview of thermal energy storage systems. *Energy*, 144:241–378, 2018. (Cited on pages 16, 19, 20, and 21).

[20] N. Amaro, J. Murta Pina, J. Martins, and J. Maria Ceballos. Superconducting magnetic energy storage: A technological contribute to smart grid concept implementation. In *Proceedings of the 1st International Conference on Smart Grids and Green IT Systems (SMARTGREENS-2012)*, pages pages 113–120, 2012. (Cited on page 15).

[21] R. Anderson, L. Bates, E. Johnson, and J. Morris. Packed bed thermal energy storage: A simplified experimentally validated model. *Journal of Energy Storage*, 4:14–23, 2015. (Cited on page 19).

[22] A. Arteconi, N. Hewitt, and F. Polonara. State of the art of thermal storage for demand-side management. *Applied Energy*, 93:371–389, 2012. (Cited on page 30).

[23] J. Barron. Steam Blast Jolts Midtown, Killing One. The New York Times, [Online] `https://www.nytimes.com/2007/07/19/nyregion/19explode.html`, 2007. Last accessed on February 8, 2019. (Cited on page 28).

[24] I. Best, J. Orozaliev, and K. Vajen. Economic comparison of low-temperature and ultra-lowtemperature district heating for new building developments with low heat demand densities in Germany. *International Journal of Sustainable Energy Planning and Management*, 16:45–60, 2018. (Cited on page 28).

[25] A. Bland, M. Khzouz, T. Statheros, and E. Gkanas. PCMs for Residential Building Applications: A Short Review Focused on Disadvantages and Proposals for Future Development. *Buildings*, 7:78, 2017. (Cited on page 20).

[26] V. Boicea. Energy Storage Technologies: The Past and the Present. *Proceedings of the IEEE*, 102:1777 – 1794, 2014. (Cited on page 15).

[27] R. Bosch. Modelling the effects of different renovation scenarios of apartments on the configuration of the Ecovat energy storage system. Master's thesis, Technische Universiteit Eindhoven, the Netherlands, 2015. (Cited on pages 56, 57, and 60).

[28] M. Budt, D. Wolf, R. Span, and J. Yan. A review on compressed air energy storage: Basic principles, past milestones and recent developments. *Applied Energy*, 170: 250–268, 2016. (Cited on page 15).

[29] M. Cascetta, G. Cau, P. Puddu, and F. Serra. Experimental investigation of a packed bed thermal energy storage system. *Journal of Physics: Conference Series*, 655:012018, 2015. (Cited on page 19).

[30] Y. Çengel and J. Cimbala. *Fluid Mechanics*. McGraw-Hill, 2nd edition, 2010. (Cited on page 34).

[31] J. Cook, N. Oreskes, P. T. Doran, W. R. L. Anderegg, B. Verheggen, E. W. Maibach, J. S. Carlton, S. Lewandowsky, A. G. Skuce, S. A. Green, D. Nuccitelli, P. Jacobs, M. Richardson, B. Winkler, R. Painting, and K. Rice. Consensus on consensus: a synthesis of consensus estimates on human-caused global warming. *Environmental Research Letters*, 11(4):048002, 2016. (Cited on page 1).

[32] J. Dickinson, N. Buik, M. Matthews, and A. Snijders. Aquifer thermal energy storage: theoretical and operational analysis. *Géotechnique*, 59:249–260, 2009. (Cited on page 17).

[33] I. Dincer and M. Rosen. *Thermal Energy Storage. Systems and Applications, second edition*. John Wiley and Sons, 2011. (Cited on pages 17, 18, and 19).

[34] J. Duffie and W. Beckman. *Solar Engineering of Thermal Processes, fourth edition*. John Wiley and Sons, 2013. (Cited on page 18).

[35] B. Dunn, H. Kamath, and J.-M. Tarascon. Electrical Energy Storage for the Grid: A Battery of Choices. *Science*, 334:928–935, 2011. (Cited on page 15).

[36] R. Dupeyrat, C. Ménézo, and S. Fortuin. Study of the thermal and electrical performances of PVT solar hot water system. *Energy and Buildings*, 68:751–755, 2014. (Cited on pages 41, 60, and 106).

[37] A. El-Sebaii, M. Ramadan, S. Aboul-Enein, and A. Khallaf. History of the solar ponds: A review study. *Renewable and Sustainable Energy Reviews*, 15:3319–3325, 2011. (Cited on page 18).

[38] S. Ephrati and H. Jonker. A numerical study of the Ecovat with IFISS, 2017. BSc thesis, University of Twente, unpublished. (Cited on pages 24 and 36).

[39] A. Fallahi, G. Guldentops, S. Tao, M. Granados-Focil, and S. Van Dessel. Review on solid-solid phase change materials for thermal energy storage: Molecular structure and thermal properties. *Applied Thermal Engineering*, 127:1427–1441, 2017. (Cited on page 19).

[40] S. Fanshawe, A. Al-Habaibeh, and B. Shakmak. An Innovative Design and Evaluation of a Stratified Hot Water Storage System - The Water Snake. In *Proceedings of 2015 IEEE Jordan Conference on Applied Electrical Engineering and Computing Technologies (AEECT)*, 2015. (Cited on page 23).

[41] P. Fleuchaus, B. Godschalk, I. Stober, and P. Blum. Worldwide application of aquifer thermal energy storage – A review. *Renewable and Sustainable Energy Reviews*, 94:861–876, 2018. (Cited on page 17).

[42] S. Furbo, E. Andersen, A. Thür, L. Shah, and K. Andersen. Performance improvement by discharge from different levels in solar storage tanks. *Solar Energy*, 79: 431–439, 2005. (Cited on page 23).

[43] M. Gerards, H. Toersche, G. Hoogsteen, T. van der Klauw, J. Hurink, and G. Smit. Demand Side Management using Profile Steering. In *PowerTech, 2015 IEEE Eindhoven*, 2015. doi: 10.1109/PTC.2015.7232328. (Cited on page 122).

[44] J. Hadorn. Advanced storage concepts for active solar energy - iea-shc task 32 2003-2007. In *Proceedings of Eurosun - 1st International Congress on Heating, Cooling and Buildings*, 2008. (Cited on page 21).

[45] Y. Han, R. Wang, and Y. Dai. Thermal stratification within the water tank. *Renewable and Sustainable Energy Reviews*, 13:1014–1026, 2009. (Cited on page 21).

[46] N. Hoogervorst. Toekomstbeeld klimaatneutrale warmtenetten in Nederland. [Online] `https://www.pbl.nl/sites/default/files/cms/publicaties/pbl-2017-toekomstbeeld-klimaatneutrale-warmtenetten-in-nederland-1926.pdf`, *Dutch*, 2017. Policy study, Planbureau voor de Leefomgeving, PBL-publicatienummer: 1926, Last accessed on December 13, 2018. (Cited on page 27).

[47] G. Hoogsteen. *A Cyber-Physical Systems Perspective on Decentralized Energy Management*. PhD thesis, University of Twente, 12 2017. CTIT Ph.D. thesis series no. 17-449. (Cited on pages 118, 122, and 132).

[48] G. Hoogsteen, A. Molderink, J. Hurink, and G. Smit. Generation of flexible domestic load profiles to evaluate demand side management approaches. In *2016 IEEE International Energy Conference (ENERGYCON)*. IEEE PES, 4 2016. (Cited on pages 120 and 122).

[49] A. Hussain, V. Bui, H. Kim, Y. Im, and J. Lee. Optimal Energy Management of Combined Cooling, Heat and Power in Different Demand Type Buildings Considering Seasonal Demand Variations. *Energies*, 10:789, 2017. (Cited on page 30).

[50] H. Jarimi, D. Aydin, Z. Yanan, G. Ozankaya, X. Chen, and S. Riffat. Review on the recent progress of thermochemical materials and processes for solar thermal energy storage and industrial waste heat recovery. *International Journal of Low-Carbon Technologies*, 14:44–69, 2019. (Cited on page 21).

[51] S. Kalnæs and B. Jelle. Phase change materials and products for building applications: A state-of-the-art review and future research opportunities. *Energy and Buildings*, 94:150–176, 2015. (Cited on page 20).

[52] A. Karim, A. Burnett, and S. Fawzia. Investigation of Stratified Thermal Storage Tank Performance for Heating and Cooling Applications. *Energies*, 11:1049, 2018. (Cited on pages 23 and 24).

[53] L. Kenjo, C. Inard, and D. Caccavelli. Experimental and numerical study of thermal stratification in a mantle tank of a solar domestic hot water system. *Applied Thermal Engineering*, 27:1986–1995, 2007. (Cited on page 26).

[54] A. Khodadoost Arania, H. Karamia, G. Gharehpetiana, and M. Hejazi. Review of Flywheel Energy Storage Systems structures and applications in power systems and microgrids. *Renewable and Sustainable Energy Reviews*, 69:9–18, 2017. (Cited on page 15).

[55] J.-H. Kim and J.-T. Kim. Comparison of electrical and thermal performances of glazed and unglazed PVT collectors. *International journal of photoenergy*, 2012. article ID:957847. (Cited on page 41).

[56] E. Kleinbach, W. Beckman, and S. Klein. Performance study of one-dimensional models for stratified thermal storage tanks. *Solar Energy*, 50(2):155–166, 1993. (Cited on pages 25 and 26).

[57] S. Knudsen and S. Furbo. Thermal stratification in vertical mantle heatexchangers with application to solar domestic hot-water systems. *Applied Energy*, 78:257–272, 2004. (Cited on page 24).

[58] S. Knudsen, G. Morrison, M. Behnia, and S. Furbo. Analysis of the flow structure and heat transfer in a vertical mantle heat exchanger. *Solar Energy*, 78:281–289, 2005. (Cited on page 24).

[59] P. O. Kriett and M. Salani. Optimal control of a residential microgrid. *Energy*, 42:321–330, 2012. (Cited on page 30).

[60] A. Lake, B. Rezaie, and S. Beyerlein. Review of district heating and cooling systems for a sustainable future. *Renewable and Sustainable Energy Reviews*, 67:417–425, 2017. (Cited on page 27).

[61] M. Lanahan and P. Tabares-Velasco. Seasonal Thermal-Energy Storage: A Critical Review on BTES Systems, Modeling, and System Design for Higher System Efficiency. *Energies*, 10:743, 2017. (Cited on pages 17 and 18).

[62] Z. Lavan and J. Thompson. Experimental study of thermally stratified hot water storage tanks. *Solar Energy*, 19:519–524, 1977. (Cited on pages 8, 23, and 24).

[63] G. Li, Y. Hwang, R. Radermacher, and H. Chun. Review of cold storage materials for subzero applications. *Energy*, 51:1–17, 2013. (Cited on page 16).

[64] H. Lund. *Renewable Energy Systems - A Smart Energy Systems Approach to the Choice and Modeling of 100% Renewable Solutions, second edition*. Academic Press, 2014. (Cited on page 29).

[65] H. Lund. Renewable heating strategies and their consequences for storage and grid infrastructures comparing a smart grid to a smart energy systems approach. *Energy*, 151:94–102, 2018. (Cited on page 29).

[66] H. Lund, S. Werner, R. Wiltshire, S. Svendsen, J. Thorsen, F. Hvelplund, and B. Mathiesen. 4th Generation District Heating (4GDH) integrating smart thermal grids into future sustainable energy systems. *Energy*, 68:1–11, 2014. (Cited on pages 27 and 28).

[67] H. Lund, P. Østergaard, D. Connolly, and B. Mathiesen. Smart energy and smart energy systems. *Energy*, 137:556–565, 2017. (Cited on page 29).

[68] R. Oliveski, A. Krenzinger, and H. Vielmo. Comparison between models for the simulation of hot water storage tanks. *Solar Energy*, 75:121–134, 2003. (Cited on pages 26 and 48).

[69] H. Öztürk. Experimental evaluation of energy and exergy efficiency of a seasonal latent heat storage system for greenhouse heating. *Energy Conversion and Management*, 46:1523–1542, 2005. (Cited on page 20).

[70] S. Paardekooper, R. Lund, B. Mathiesen, M. Chang, U. Petersen, L. Grundahl, A. David, J. Dahlbæk, I. Kapetanakis, H. Lund, N. Bertelsen, K. Hansen, D. Drysdale, and U. Persson. Heat Roadmap Europe 4 : Quantifying the Impact of Low-Carbon Heating and Cooling Roadmaps. Aalborg University, Denmark, 2018, [Online] `http://vbn.aau.dk/files/288075507/Heat_Roadmap_Europe_4_Quantifying_the_Impact_of_Low_Carbon_Heating_and_Cooling_Roadmaps..pdf`. Last accessed on 13 December 2018. (Cited on page 27).

[71] P. Palensky and D. Dietrich. Demand Side Management: Demand Response, Intelligent Energy Systems, and Smart Loads. *IEEE transactions on industrial informatics*, 7:381–388, 2011. (Cited on page 29).

[72] M. Perez and R. Perez. A fundamental look at supply side energy reserves for the planet. *Solar Update: IEA SHC Newsletter*, 62, 2015. (Cited on pages 2 and 3).

[73] P. Pinel, C. Cruickshank, I. Beausoleil-Morrison, and A. Wills. A review of available methods for seasonal storage of solar thermal energy in residential applications. *Renewable and Sustainable Energy Reviews*, 15:3341–3359, 2011. (Cited on pages 16, 19, and 21).

[74] S. Rehman, L. Al-Hadhrami, and M. Alam. Pumped hydro energy storage system: A technological review. *Renewable and Sustainable Energy Reviews*, 44:586–598, 2015. (Cited on page 15).

[75] B. Rezaie and M. Rosen. District heating and cooling: Review of technology and potential enhancements. *Applied Energy*, 93:2–10, 2012. (Cited on page 27).

[76] M. Rosen. The exergy of stratified thermal storages. *Solar Energy*, 71:173–185, 2001. (Cited on pages 8, 23, and 24).

[77] I. Sarbu and C. Sebarchievici. A Comprehensive Review of Thermal Energy Storage. *Sustainability*, 10(1):191, 2018. (Cited on pages 16, 19, 20, and 21).

[78] B. Schepers and M. van Valkengoed. Warmtenetten in Nederland: Overzicht van grootschalige en kleinschalige warmtenetten in Nederland. [Online] `https://www.acm.nl/sites/default/files/old_download/documenten/nma/Onderzoek_Warmtenetten_in_Nederland.pdf`, *Dutch*, 2009. CE Delft, Publicatienummer: 09.3031.45, Last accessed on December 13, 2018. (Cited on page 27).

[79] L. Shah and S. Furbo. Correlation of Experimental and Theoretical Heat Transfer in Mantle Tanks used in Low Flow SDHW Systems. *Solar Energy*, 64:245–256, 1998. (Cited on page 26).

[80] L. Shah, G. Morrison, and M. Behnia. Characteristics of Vertical Mantle Heat Exchangers for Solar Water Heaters. *Solar Energy*, 67:79–91, 1999. (Cited on page 22).

[81] A. Sharma, V. Tyagi, C. Chen, and D. Buddhi. Review on thermal energy storage with phase change materials and applications. *Renewable and Sustainable Energy Reviews*, 13:318–345, 2009. (Cited on page 20).

[82] P. Siano. Demand response and smart grids: A survey. *Renewable and Sustainable Energy Reviews*, 30:461–478, 2014. (Cited on page 29).

[83] B. Sibbitt, D. McClenahan, R. Djebbar, J. Thornton, B. Wong, J. Carriere, and J. Kokko. The performance of a high solar fraction seasonal storage district heating system – five years of operation. *Energy Procedia*, 30:856–865, 2012. (Cited on pages 18 and 20).

[84] H. Tabor. Solar Ponds. *Solar Energy*, 27(3):181–194, 1981. (Cited on page 18).

[85] S. Ushak, A. G. Fernández, and M. Grageda. Using molten salts and other liquid sensible storage media in thermal energy storage (TES) systems. In L. Cabeza, editor, *Advances in Thermal Energy Storage Systems*, chapter 3, pages 49–63. Woodhead Publishing, 2015. (Cited on page 16).

[86] J. Van Berkel. Ecovat thermodynamics. concept evaluation report (internal report, unpublished), 2014. (Cited on pages 9, 48, and 60).

[87] T. van der Klauw, M. Gerards, and J. Hurink. Resource allocation problems in decentralized energy management. *OR Spectrum*, 39:749–773, 2017. (Cited on page 72).

[88] A. van Harmelen and W. Koch. TNO-report R2002/174: CO_2 emission factors for fuels in the Netherlands. Available online at: https://www.rvo.nl/sites/default/files/2013/10/Harmelen2002. (Cited on page 130).

[89] V. Velmurugan and K. Srithar. Prospects and scopes of solar pond: A detailed review. *Renewable and Sustainable Energy Reviews*, 12:2253–2263, 2008. (Cited on page 18).

[90] F. Verrilli, S. Srinivasan, G. Gambino, M. Canelli, M. Himanka, C. Del Vecchio, M. Sasso, and L. Glielmo. Model Predictive Control-Based Optimal Operations of District Heating System With Thermal Energy Storage and Flexible Loads. *IEEE Transactions on Automation Science and Engineering*, 14(2):547–557, 2017. (Cited on page 28).

[91] H. Versteeg and W. Malalasekera. *An Introduction to Computational Fluid Dynamics - The Finite Volume Method, second edition*. Pearson Education Limited, 2007. (Cited on page 26).

[92] H. Williams. *Model Building in Mathematical Programming*. John Wiley & Sons Ltd, 5th edition, 2013. (Cited on page 45).

[93] J. Xu, R. Wang, and Y. Li. A review of available technologies for seasonal thermal energy storage. *Solar Energy*, 103:610–638, 2014. (Cited on pages 17, 18, 19, and 21).

[94] W. Yaïci, M. Ghorab, E. Entchev, and S. Hayden. Three-dimensional unsteady CFD simulations of a thermal storage tank performance for optimum design. *Applied Thermal Engineering*, 60:152–163, 2013. (Cited on pages 23, 24, 26, and 34).

[95] T. Yan, R. Wang, T. Li, L. Wang, and I. Fred. A review of promising candidate reactions for chemical heat storage. *Renewable and Sustainable Energy Reviews*, 43:13–31, 2015. (Cited on pages 20 and 21).

[96] J. Yau and B. Rismanchi. A review on cool thermal storage technologies and operating strategies. *Renewable and Sustainable Energy Reviews*, 16:787–797, 2012. (Cited on page 16).

[97] B. Zalba, J. Marín, L. Cabeza, and H. Mehling. Review on thermal energy storage with phase change: materials, heat transfer analysis and applications. *Applied Thermal Engineering*, 23:251–283, 2003. (Cited on pages 19 and 20).

[98] G. Zhang, Z. Shen, and L. Wang. Online Energy Management for Microgrids With CHP Co-Generation and Energy Storage. *IEEE transactions on control systems technology*, IEEE early access articles, 2018. (Cited on page 30).

[99] D. Zhao, Y. Li, Y. Dai, and R. Wang. Optimal study of a solar air heating system with pebble bed energy storage. *Energy Conversion and Management*, 52:2392–2400, 2011. (Cited on page 19).

[100] H. Zondag. Flat-plate PV-Thermal collectors and systems: A review. *Renewable and Sustainable Energy Reviews*, 12:891–959, 2008. (Cited on pages 39 and 41).

[101] Y. Zurigat, A. Ghajar, and E. Moretti. Stratified Thermal Storage Tank Inlet Mixing Characterization. *Applied Energy*, 30:99–111, 1988. (Cited on pages 23 and 24).

LIST OF PUBLICATIONS

[G:1] G.J.H. de Goeijen, G.J.M. Smit, and J.L. Hurink. An integer linear programming model for an ecovat buffer. *Energies*, 9:592, 2016.

[G:2] G.J.H. de Goeijen, G.J.M. Smit, and J.L. Hurink. Improving an integer linear programming model of an ecovat buffer by adding long-term planning. *Energies*, 10:2039, 2017.

[G:3] G.J.H. de Goeijen, J.L. Hurink, and G.J.M. Smit. A heuristic approach to control the ecovat system. In *2018 IEEE PES Innovative Smart Grid Technologies Europe (ISGT Europe)*, 2018.

[G:4] G.J.H. de Goeijen, G. Hoogsteen, J.L. Hurink, and G.J.M. Smit. Using the ecovat system to supply the heat demand of a neighbourhood. In *2019 IEEE Powertech, Milano*, 2019.

THIS THESIS

```
@phdthesis{goeijen2019:thesis,
    author={Goeijen, Gijs J. H. de},
    title={Developing a Method for the Operational Control of an Ecovat System},
    school={University of Twente},
    address={PO Box 217, 7500 AE  Enschede, The Netherlands},
    year={2019},
    month=dec,
    day={20},
    number={IDS Ph.D. Thesis Series No. 19-022},
    issn={2589-7721},
    isbn={978-90-365-4909-7},
    doi={10.3990/1.9789036549097}
}
```

BIBTEX of this thesis