MIP for demand side optimisation

by

Gergely Marton

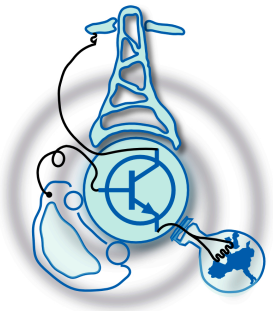

Submitted to the Department of Electrical Engineering, Electronics, Computers and Systems
in partial fulfillment of the requirements for the degree of

Erasmus Mundus Master Course in Sustainable Transportation and Electrical Power Systems

at the

UNIVERSIDAD DE OVIEDO

August 2019

© Universidad de Oviedo 2019. All rights reserved.

Author .....

Certified by .....
Stig Ødegaard Ottesen, PhD
Head of R&D, eSmart Systems
Thesis Supervisor

MIP for demand side optimisation

by

Gergely Marton

Submitted to the Department of Electrical Engineering, Electronics, Computers and Systems
on August 31, 2019, in partial fulfillment of the
requirements for the degree of
Erasmus Mundus Master Course in Sustainable Transportation and Electrical Power Systems

Abstract

As the traditional centralised and vertically integrated power system architecture gives way to distributed generation and open electricity markets, the value of demand side response is increased. Cheap networked sensors give unprecedented insight into electricity consumption, while cloud computing, IoT actuators, and home automation systems unlock the flexibility of consumption.

The goal with demand side response is to optimise the time of use of energy. Successful consumption scheduling allows for high penetration of renewable generation, deferral of infrastructure upgrades, and overall enables better utilisation and thus lower cost for both consumers and producers. The Horizon2020 EU project INVADE is aimed at this opportunity. This thesis will briefly describe a cloud platform for optimal scheduling of flexible loads, such as heating, water boilers, EV chargers, and some industrial loads.

The focus of the thesis is the modelling of the loads in a way that is suitable for mixed integer programming solvers, while considering the imperfection of the predictions and fitting into a rolling horizon approach. Models for residential space heating, residential water boilers, industrial water boilers, and thermal energy storage will be described.

The simulation results presented for each model indicate that the models correctly describe the flexibility of the physical systems, and that there is a significant scope for cost savings, depending on the tariff structure. In one case, the optimal solution leads to reduced energy consumption due to keeping machinery near their most efficient operating point.

The work described in this thesis was performed at eSmart Systems in Halden, Norway, from January to August 2019.

Thesis Supervisor: Stig Ødegaard Ottesen, PhD
Title: Head of R&D, eSmart Systems

Acknowledgments

Big thank you to the team at eSmart for believing in my ideas and providing me with a platform to implement them. Equal gratitude goes to our customers and partners for being demanding and accepting nothing but the very best solutions. This thesis would not be the same without them.

Contents

List of Figures

List of Tables

Chapter 1

Introduction

1.1 eSmart Systems AS

eSmart Systems AS [3] is a Norwegian software company that provides solutions for electrical utilities. It is an expert in energy flexibility optimization and power line inspection automation. eSmart Systems was founded in 2012 by an experienced team previously responsible for large IT service exports including delivery of the world's first power exchange to Nord Pool, the Nordic Power Exchange, in 1991. Since 2012 the company has grown rapidly and currently consists of almost 80 employees with offices in Norway, Denmark, the UK, and the US.

All the development takes place in the Norway office in Halden, a small town in the South-East of Norway, located in the fjord separating Norway from Sweden. I worked for eSmart in 2017 as a System Developer, then being deeply involved in the EMPOWER EU-funded researched project [2], where we developed the foundation for the technology described in this thesis.

For the master thesis project, I was working again at eSmart from end of January, and the thesis covers work done up to the end of August 2019.

1.2 INVADE Horizon 2020

My work at eSmart Systems consisted almost entirely of designing and implementing optimisation models for the INVADE Horizon 2020 [5] European Union research project. This project is about the development and trial of the "Integrated INVADE Platform", a cloud service that provides optimal control to flexible assets, such as batteries, EV chargers, shiftable loads, and curtailable generation resources. The end goal of this project, by the end of year 2019, is to demonstrate the whole platform in operation at five large scale pilot sites:

- Norway - residential heating, water heaters, and home EV-chargers
- Germany - home PV and battery management
- Spain - utility battery management
- The Netherlands - aggregated EV charging management
- Bulgaria - battery and water heater management for hotels

eSmart Systems' role in the INVADE project is to develop and operate the Integrated INVADE Platform (IIP). This consists of:

- Building a database of all assets, their parameters, and their relationship to other assets
- Connecting to various IoT management systems to retrieve meter values and send control signals
- Storing all meter vales and calculate derived time series, such as residual loads or zone level loads
- Developing models to make predictions for base line consumption or generation of assets
- Integrating standalone optimization models for various asset categories into a single global formulation

- Developing optimisation models for some asset categories

- Providing an easy-to-use interface for visualising meter values and manage assets

I worked on designing and implementing the optimisation models described in this thesis, and I was also involved in supporting the pilot in Bulgaria.

1.3 REMA 1000's Vagle distribution centre

The only project outside of INVADE that I participated in at eSmart has been creating a high level control system to take advantage of thermal storage at REMA 1000's Vagle distribution centre. REMA 1000 is the largest Norwegian retailer, with hundreds of supermarkets and various distribution centres spread around Norway. This project focused on one of these distribution centres, located at Vagle, near to the coastal city of Stavanger. The site has huge storage spaces for dry goods, refrigerated warehouse for perishable goods, and freezing warehouse for frozen goods, as well as loading and office spaces. REMA has installed industrial cooling machines for the cold storages, and uses some of the rejected heat to warm the offices in the winter.

The site had a 300.000 litre fire water tank, and the idea for this project has been to utilise the tank as thermal storage. The benefit of storage is that the cooling machines could operate near their peak efficiency, and that they can produce most of the heating/cooling in off-peak periods. This required modelling the cooling machines, the cooling and heating loads, the local PV generation, the on-site battery, the heat exchangers, and the water tank itself.

While the agreement with REMA is a commercial software as a service (SAAS) contract, it shares the underlying architecture with INVADE. It required a specific model for its unique thermal system, but the other electrical components, such as the PV and the battery, were modelled exactly the same way as they are in the INVADE project. The general modelling approach and other platform components are also shared.

1.4 Scope of my work

When I joined the work at eSmart at the end on January 2019, most of the work on the cloud platform itself was done, and the general methodology for joining standalone models has been developed, but models describing various asset categories were missing, and no assets were controlled yet.

1.4.1 H2020 INVADE

I started my work with residential water heaters. I researched the typical sizes and configurations found in Norway and worldwide, and using my experience from the EMPOWER project in 2017, I developed a computationally more favourable, as well as more user friendly formulation. I worked with the assumption that the water heaters are separately metered, and that they can be controlled by switching a switch in series with the thermostat - this is how the water heaters participating in the INVADE project are equipped. The water temperature is not metered. I presented this model to the INVADE project participants, and convinced the committee of researchers from the Universitat Politècnica de Catalunya (UPC) to drop their formulation in favour of mine. I also did the mapping and integration work so that the model can work as part of complex multi-asset optimisation. It is going to be tested in the Norwegian pilot site before the end of year 2019.

I also developed a space heating model, aimed for residential households with electrical heating, very typical in Norway. The model assumes that the electrical consumption of the heating elements are separately metered, and that each temperature zone is controlled by a smart thermostat. The smart thermostats are IoT devices that report measured temperature to the platform, and can accept temperature setpoints from the platform. The model describes the buildings as a thermal system, and therefore establishes a dynamic connection between the temperature setpoints and the electricity consumption. I presented this model to the INVADE project participants, and convinced the committee of researchers to drop the earlier approach developed by the Norwegian University of Science and Technology (NTNU) in favour

of mine.

In the process of integrating this model into the platform, I developed a system identification solution based on SciPy, a Python-based ecosystem of open-source software for mathematics, science, and engineering. [15]

The integration work for this model is still ongoing. The Norwegian pilot site has 11 households where this model is going to be tested.

Finally, I developed a standalone model for the industrial water heaters for the Bulgarian pilot site. The first version of this model went into production in July, and I was involved with evaluating the results and fine tuning the system.

1.4.2 REMA 1000's Vagle distribution centre

I joined this project in April 2019. In a previous phase, a high level strategic control system has already been developed, aimed at avoiding peak power fees and ensuring self consumption of the local PV generation. As a second phase, the site's fire water tank had been renovated with improved isolation and connected to the thermal system via a heat exchanger and a set of valves allowing it to connect to either the hot water or the cold water circuit.

The thermal system had to be modelled on a high level. It had to be detailed enough so that the temperature- and power level dependent coefficient of performance (COP) of the different cooling machines is accounted for, but simple enough that it can be formulated as a mixed integer programming problem using only variables that are already measured. Simplifying assumptions were made in close cooperation with the mechanical engineering team at REMA. I was responsible for finding a reasonable set of assumptions, and formulate a model on top of them.

This work was completed and presented to REMA on June 17th, where I gave a 2-hour presentation on the optimization, using live data gathered from the site through our platform. The results were accepted as a successful factory acceptance test. We are currently waiting for changes to be completed on the physical infrastructure before the control can be tested in production.

1.5 Rolling horizon optimization

The general approach for dealing with uncertainty in the eSmart platform is called a rolling horizon approach. [4] It means that the optimization models do not consider multiple scenarios with different probabilities. They assume that the predictions that parameterise the model are accurate, and find the optimal control plans accordingly.

In each time period, new meter values are logged, new predictions are made, and a new optimisation is run, and new control plans are distributed to the assets. This means that always only the first control action of the control plan is going to be implemented, because the rest will be overwritten by a new control plan before it could be implemented.

This approach makes sure that we always consider all the available information to make the most optimal control plan. It also takes advantage of the nature of predictions, that is to be rather accurate in the near term, and deteriorate going further into the future.

In fact, the new control plans can not include information from the period directly preceding them, because the meter values describing that period only arrive a short time after the new control plan has already been executed. See Figure 1-1. This is especially important to consider with models that take into account the near past, for example because they have limitations on the number of consecutive control actions allowed.

The rolling horizon approach has been in place since before I first joined the company, therefore I had no choice regarding it. I had to be aware of it, and consider it when formulating optimization models.

1.6 Benefits attainable

There are many different benefits one can aim at without significantly changing the modelling approach. Even a simple cost minimization can mean very different objective functions depending on the grid contracts, and it is further complicated if

Figure 1-1: Rolling horizon approach as implemented in the eSmart IIP

depreciation of storage assets is considered.

1.6.1 Operator

For the simplest case of a traditional residential contract where the user only pays for energy consumed per month, the goal for optimisation would be to minimise energy consumption. In the cases discussed in this thesis, there is usually little scope for this. Exceptions are making sure that cooling units operate near their optimal COP-range, minimising battery roundtrips, and ensuring that local generation is locally consumed as much as possible.

All pilot sites discussed in this thesis are equipped with smart meters, and therefore they are typically exposed to the spot price or some sort of time-of-use tariff. This means that there is a benefit in optimal scheduling of the electricity consumption, even if the overall amount remains the same. In many cases, lower energy efficiency is a reasonable sacrifice for a more suitable schedule: this is always the case when batteries are used, since typically the AC-to-AC roundtrip efficiency of a battery system is at most 90%. Preheating a building will similarly reduce energy efficiency, since the heat loss will be higher.

Sites are often equipped with some sort of local generation, typically PV. In these cases, self-consumption is typically encouraged by the grid contract, as the users get a lot less for a unit of surplus energy sold to the grid than what they pay for importing it. This is especially true for industrial scale installations that are not subsidized.

Industrial customers, and increasingly even residential ones, pay a grid tariff proportional to their peak consumption. This fee is often very substantial, and therefore using smart scheduling to minimize peaks is an important use case. Typically, only the highest peak, or the average of the three highest peaks of each month matter, and there is no reward for avoiding smaller peaks or lower-load days. These peaks are to be understood as peek energy usage (kWh) per metering period, typically hour or quarter hour. The real time peaks (kW) are not measured by the smart meters.

In case of some industrial sites, and especially electrical vehicle charging courts, there may be a hard limit on the allowable peaks due to some grid capacity constraint

or fuses. In these cases, a local system must always manage the consumption real time, because even in the best case, the optimisation can only ensure compliance with the limit on an average level.

1.6.2 Other stakeholders

The above listed benefits are to the operator of the consumption sites, that is, the end user in the energy system. It is important to note however, that by increasing the price sensitivity of the demand side, the whole society benefits as the energy system becomes more efficient: the need for peaking plants is reduced, large power plants can operate near their optimal loads, and low marginal cost renewables do not have to suffer export limitations.

Energy retailers can benefit from demand side response to meet their commitments as balance responsible parties. For that to happen, they would communicate their requirements – up or down regulation – to the party that is responsible for the scheduling optimisation. Since the baseline for these regulation requests is the optimal scheduling from the sites' point of view, the new optimal schedule can only be least expensive compared to the other possibilities of achieving that regulation. Allowing this suboptimal control of their assets only makes sense for the site owners if they are compensated for it in some way.

Similarly to energy retailers, the distribution system operator might be interested in requesting deviations from the site-level-optimal schedule: for example to avoid overstressing bottlenecks in the system, or to be able to ensure adequate voltage levels at all nodes in their network. Having this capability might allow for deferring infrastructure investments or avoiding restrictions on residential PV installations.

Both in case of BRP and DSO change request, it is important to know which sites belong to the specific BRP or DSO.

While the eSmart platform is capable of handling such flexibility requests and inferfacing to the BRP or DSO, I was not involved in the development of this functionality, and I could safely ignore it in the development of the asset models. The asset models essentially describe the available set of load curves and their costs, and

this is done regardless of the way that those load curves are evaluated.

Some customers might want to define different goals that cost minimization: minimizing their CO2-footprint, maximizing self-consumption at all costs, etc. These are not discussed in this thesis, but they would be easy to implement in the described system.

This overview of the different benefits provides a peak into the complexity of the objective function. It is a good intuitive indicator that simple isolated operating strategies will not be able to attain or even well approximate a globally optimal consumption schedule, and there is value to be unlocked through centralised optimization. Measuring that value is a goal of the H2020 INVADE project, but this conclusion will only be reached by the end of 2019, and therefore is out of the scope of this thesis.

1.7 Objectives of the thesis

The objective of this thesis is to develop mathematical models for the consumption of the following loads:

- Chained water boiler system
- Chilled warehouse and office complex with thermal energy storage
- Residential heating system
- Residential electric water boiler

The models should be formulated such that they fit into eSmart / INVADE modelling approach. The constraints and objectives specific to the models are discussed in their respective chapters.

Chapter 2

The eSmart optimisation platform

Many systems need to cooperate to optimally schedule the electricity consumption.

- Metering infrastructure
- Database of assets and contracts
- Prediction system
- Optimization system
- Smart assets capable of receiving, storing, and enacting control plans

All of this must be able to work on a tight schedule to minimise the effects of the inherent imperfection of the predictions.

The time resolution is itself a parameter to each optimisation. Typically eSmart employs 15-minute or hourly resolution. The 15-minute resolution is in compliance with the reporting frequency of the smart meters installed in Norway, while the hourly resolution is a good fit for the energy price's time resolution. The platform is capable to use other resolutions as well, such as half-hour, five-minute, or one-minute. The INVADE project used 15-minute resolution, while the REMA one hourly.

Within a single period, the system must be able to collect the most recent meter readings, make new predictions based on the recent data, find the optimal scheduling for the whole planning horizon, and send the control signals to the IoT devices for

execution in the next period. This can be challenging with hundreds or even thousands of assets within the same optimisation problem - indeed one of the goals of the INVADE pilot is to test how well the system can scale up.

While I only worked on the models for assets, it was important for me to keep the rest of the platform in mind: I usually had some freedom in deciding on what measurements to incorporate, what parameters to store, and which variables to predict.

2.1 Metering and smart asset infrastructure

Norway started rolling out smart meters in 2015, and all electricity consumers have received their smart meters by the end of 2018. [12] The metering information is available to energy market participants through a central platform called Elhub [11], which was developed by the Norwegian transmission system operator, Statnett.

More detailed metering requires dedicated infrastructure. Solar installations typically come with some kind of online platform that can be accessed to retrieve generation data. These tend to be proprietary, for example the German manufacturer SMA's Sunny Portal. [17] There is a similar situation with the electrical vehicle smart charger interfaces, although many manufacturers now use the Open Charge Alliance's protocols, Open Charge Point Protocol (OCPP) and Open Smart Charging Protocol (OSCP). [13]

The separate metering of residential water heaters and electrical space heaters must be managed by some home automation system. The same system handles the implementation of the control signals as well. These home automation systems were installed in the pilot sites by the utilities responsible for them.

For more complex sites, such as the hotels in Albena and the Rema 1000 distribution centre, a custom SCADA system is interfacing with the Integrated INVADE Platform.

It is important that the system is resilient to temporary network congestions: this is handled by always sending control signals for the entire planning horizon, so that if there is a delay in delivering the most recent version, the local system has something

to fall back on. Note that if the predictions were accurate, there should be very little difference between the old and the updated control plan - only the effect of the planning horizon being rolled forward in time by one period.

The local system must also bridge the gap between the high granularity of the control plan and the continuous time reality. The control signals from the platform are to be interpreted as average values or "budgets" for the time periods. For example, if the optimal control plan for a battery is to be discharging at 5 kW for an hour long period, it is up to the battery management system what profile to choose, as long as the integral is 5 kWh. This is especially interesting for curtailable assets, like renewable generation or EV chargers.

2.2 Database of assets and contracts

A lot of parameters and timeseries go into each optimisation, and it is important to keep track of them in an orderly and efficient manner.

2.2.1 Assets

The most fundamental units in the optimisation are the assets. An asset is for example a residential water heater, an EV charging point, or a battery. Depending on the type of the asset, they might have different parameters, such as installed capacity, or maximum discharge power. The type also determines the types of timeseries attached to the asset. For example a PV installation will have historic and forecasted generation, while a battery will have state of charge.

Some assets are virtual, and their attached timeseries are calculated instead of measured. Such are the residual load, and the heat demand at the Albena hotels for example. Some assets are not controllable: the above mentioned virtual loads, but also some residential solar or wind generation is operated as non-curtailable.

Assets that are controllable in some way will also have an optimisation model attached to them. An optimisation model is a set of equations that describe the set of possible load curves that the asset can have. The optimisation model's parameters

are the asset parameters and the timeseries attached to the asset - typically some kind of forecast.

2.2.2 Forecasts

Of the various timeseries recorded in the platform, some have forecasts attached to them. These are machine learning models that detect patterns and relationships in historical data and produce the most likely future scenario. Forecasts are typically made 24 hours into the future.

In some instances, for example predicting production from solar generation, the forecasts are quite reliable. In that case, the machine learning algorithm picks up the relationship between the weather forecasts, the time of day, the time of year, similarly to how one would do it on a first principles approach.

The advantage of using machine learning is that it is also applicable when modelling based on principles of physics is hopeless - for example it can pick up patterns in homes' electrical consumption to provide reasonable predictions for residual load, or detect regularities in a family's showering habits on weekdays and weekends to produce a load forecast for a residential water heater.

Great care must be taken to choose which variables to predict. When the underlying reality is fundamentally unpredictable, such as the connection times of electrical vehicles at a charging point, even the most sophisticated prediction approach will fail. In these cases it might be smarter to observe patterns in average connection times, typical energy demand, or average number of cars arriving at different times of the day.

The result of smart control should not skew the forecasts. For example a family might always shower at 5 PM, and as a result the forecast will be some load on the water heater at 5 PM. If then the optimisation decides that is is better to delay that load until 7 PM, the metered data on that they will show load at 7 PM. If this data is then used to train the prediction model, it will be confused regarding the expected time of load. If possible, forecasts should only be on timeseries that are not affected by control, for example the heat demand. If this is not possible, the baseline must be

estimated and stored. For example if the output of a solar installation is routinely curtailed, and the forecast mistakes that for the baseline, it will not even try to move more load to the periods where marginal consumption would be effectively free.

Dealing with uncertainty

Forecasts will always be inaccurate to some level. One approach to deal with this is to calculate a confidence interval along with the best guess forecast, and apply some robust optimisation method. [10] In eSmart's platform uncertainty is not explicitly modelled: we implemented a rolling horizon approach [4] instead. It means that we only implement the first period of each control plan we generate, because we update the control plan to incorporate the latest information as it arrives. (Unless there is a delay.) This approach takes advantage of our greater confidence in near term predictions, and is expected to produce acceptable control quality. Verifying this assumption is one of the key expected outcomes from the INVADE pilots.

Within the optimisation models, the forecasts are treated as if we expected them to be perfectly accurate. However, extra constraints are introduced to limit the potential damage from bad predictions. These are unique to the each and every model and will be discussed with the models themselves.

2.3 Mixed integer programming

Mixed integer programming (MIP) is based on linear programming (LP), with integrality constraints on some variables. There are robust and fast methods for solving linear programming problems, and MIP solvers build on top of that capability. The details of the internal workings of the solvers are not relevant to this thesis, but the restrictions of MIP are briefly discussed. MIP solvers include:

- presolve step for removing redundant constraints and variables
- branch and bound algorithm to deal with integrality constraints
- heuristics for selecting the branch to explore first

- an LP solver to solve the relaxations

There are various solvers available, such as Gurobi, GLPK, CPLEX, etc. There are significant differences in the performance of these solvers as well as their licensing costs. They share a common interface, which makes it relatively easy to switch between them. I used the GLPK solver, because it is free and open source, and its performance was sufficient for the problems described in this thesis - they typically solved in less than a second.

In general, the more variables a model has, the more difficult it is to solve the model. Integer variables are much more costly, and therefore it is desirable to avoid them. Model complexity grows non-linearly with the number of integer variables!

MIPs must still be linear: the objective function and all the constraints must be linear combinations of the decision variables. This is a major restriction when formulating the models. More complex relationships are possible to approximate with piecewise linear functions. Piecewise linear functions employ binary variables to select between the sections, and therefore each piecewise linear relationship will introduce as many binary variables as it has sections. A binary variable is essentially an integer variable that is constrained $0 \leq x \leq 1$. Each binary variable is a branching point for the branch and bound algorithm, and therefore they also significantly slow down the model.

MIP solvers are generally written in C for the highest possible performance, and they take standard file formats as inputs, and produce standard outputs. eSmart uses the python package pyomo to create the formulation in the desired format and simplify the extraction of the results.

Chapter 3

Albena chained water boiler system modelling

3.1 Introduction

The Albena pilot is the Bulgarian site of the INVADE project, consisting of three large five-star hotels in the beach resort of Albena, 30 kilometers north of Varna, on the coast of the Black Sea.

The objective of the Albena pilot was to: "to shift as much as possible of the load from the peak-load hours to the off-peak hours" - that is, prosumer flexibility service. The Albena pilot consist of three hotels: Flamingo Grand, Hotel Borjana, and Paradise Blue. Flexibility comes from an electric battery at Flamingo Grand and water heaters at Hotel Borjana and Paradise Blue.

All installations in the Albena pilot communicate with a SCADA-system delivered by Schneider Electric. All exchange of data, both meter readings and control signals, are between the Integrated INVADE Platform (IIP) and this SCADA-system.

The Albena pilot only includes a battery and water boilers as flexible devices. The PV generation can also be curtailed at the instructions of the grid operator in critical situations, but this is outside of the scope of the INVADE project.

The hotels' water boilers are so significant in energy consumption and large in water volume that the flexibility from them is comparable to the flexibility from the

battery!

The objective for the flexibility algorithms will be to minimize the costs of electricity purchase. In general:

- The battery will be charged in hours with low prices and discharged in hours with high prices
- The water boilers' electricity consumption will be shifted from high- to low-priced hours as much as possible

The working assumption is that meter readings will be received every 15 minutes. This is closely linked to the time resolution of the optimization algorithms, which also are assumed to have 15 minutes time intervals. This means that every decision (output from the algorithms) has the granularity of 15 minutes. Implementation of these decisions in real-time is the responsibility of the local systems. As an example, the optimization algorithm can decide that a battery is going to charge 5.0 kWh in a given 15-minute time slot. How this is implemented inside the time slot is left to the local system, which in this case might be the battery management system (BMS) or the water heating system's SCADA.

3.2 Structure information and meter values

The structure information is data that is rarely changed. The pilot site can be illustrated as in the figure 3-1.

PV solar panels were installed at the beginning of July 2018, metering data for them was available on the SMA Sunny Portal. The solar thermal preheating and the overall water heater infrastructure was installed in its current form in 2017, and 2019 was the first year that they were operated, therefore there was no historical data available.

The IIP receives consumption and/or production from each meter from the Schneider Electric SCADA system as meter readings (i.e. counter values) every 15 minutes.

Outside temperature and cloudiness data is provided by Montel AS.

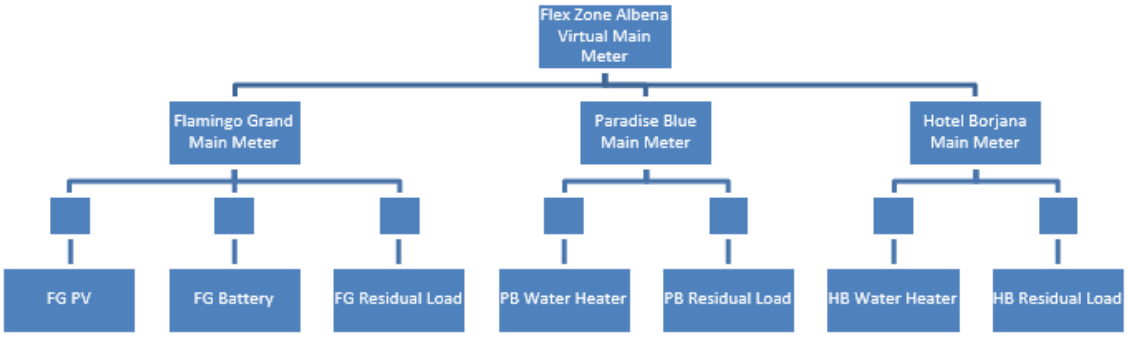

Figure 3-1: Illustration of resource model for the Bulgarian pilot site (squares represent metering points)

3.3 Commercial agreement information

Since the beginning of December 2018 Albena is billed for electrical energy based on the prices on the Day-Ahead Market on the Independent Bulgarian Energy Exchange (IBEX). See figure 3-2 for example. The prices are set and published once a day and are active for the next day.

The IBEX day-ahead market prices are loaded into the IIP via a third party contract between Albena and Montel SA. The prices are sent to Albena along with the controls signals. Grid contract information is not included.

3.4 Battery

A battery has been installed at the Flamingo Grand Hotel with the following parameters:

- Maximum charging power: 180 kW
- Maximum discharging power: 180 kW
- Maximum SOC: 180 kWh
- Minimum SOC: 20 kWh
- Round-trip efficiency: 98%

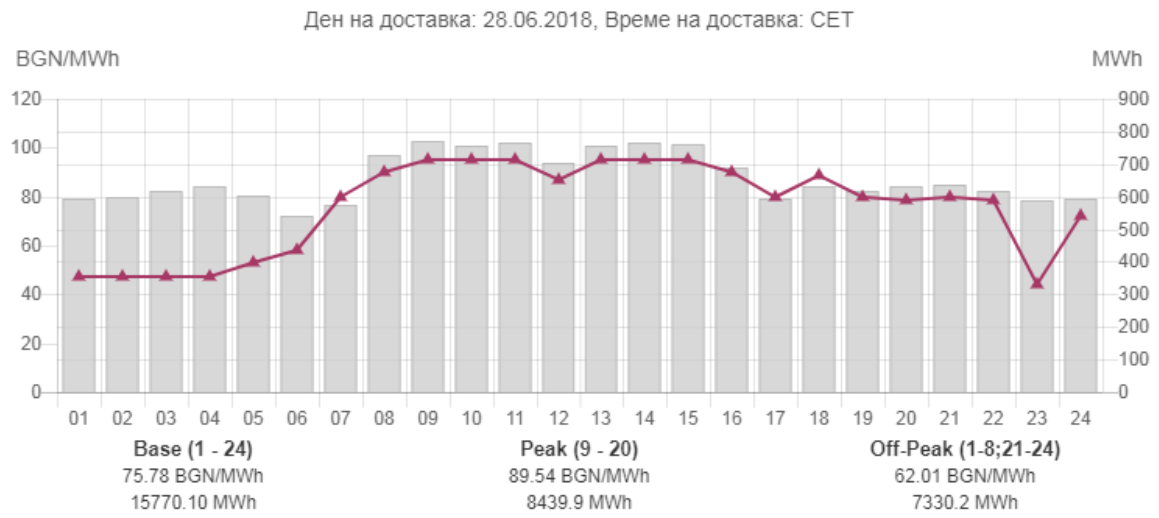

Figure 3-2: Day-ahead prices from the Bulgarian market place for the delivery day 20/06/2018

Meter values for the battery:

- Charging energy (active energy as counter values, metered in kWh)
- Discharging energy (active energy as counter values, metered in kWh)
- Battery state of charge (energy level, metered in kWh)

The control signals (charging and discharging strategy) for the battery from the IIP will be a charging or discharging setpoint for each time slot. The setpoint will either be an average power level (kW) or energy (kWh). Implementation of the setpoint will be performed locally, i.e. by the BMS or the SCADA system. Constraints related to the change of control signals between two periods may be implemented to avoid stress which results in reduced life time.

Under the INVADE pilot phases, the battery will not be used for other purposes, meaning that the Integrated INVADE Platform is the only system that defines control signals.

3.5 Water boilers

The only currently controllable loads are electrical water heaters. The large overall volume and wide temperature range means plenty of potential for storing thermal energy and realizing a Power2Heat energy management approach. [1]

The water heaters are to be controlled indirectly. The IIP sends energy commands for each period and the local SCADA will decide which heating elements to turn on or off and for how long to use exactly the required energy within the period. This control must be transparent towards the IIP, just like the battery management system is.

- In case the energy command would not be enough to maintain a setpoint temperature level in the tanks, the SCADA will override the energy command and use more energy.

- In case consuming the prescribed energy would heat some of the tanks above their maximum rated temperatures, the SCADA will use less energy than commanded.

The SCADA-controlled EWH system at Paradise Blue consists of twelve 2000 litre, 34.5 kW water boilers. The solar thermal system channels heat into these boilers. The heat captured by the solar fields is measured by the local control system. This arrangement is depicted on figure 3-3.

There are three additional buffer tanks with their own heating that is not controlled by the SCADA. They are set up such that if the SCADA-controlled part is in normal operation, the buffer heaters never turn on. All boilers are electric. The whole boiler room electricity consumption is metered together, including lightning, control, and pumps.

The structure of the boiler piping can be changed by opening/closing manual valves. There are no automatically controlled valves in the system: the piping configuration is constant during normal operation.

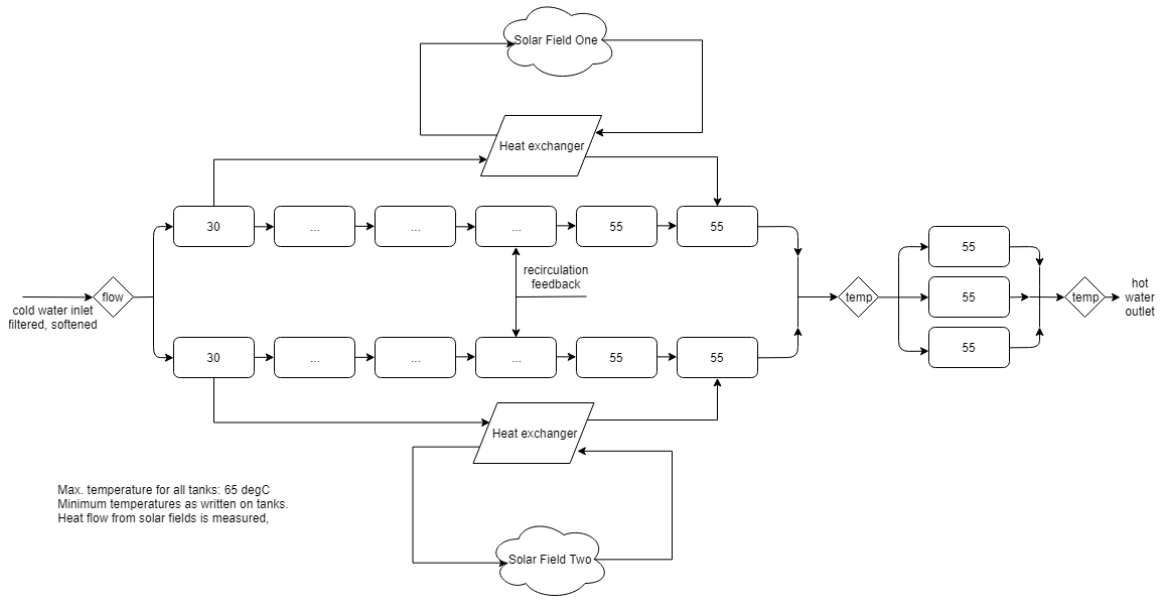

Figure 3-3: Water boiler system at the hotel Paradise Blue

The solar thermal system uses SUNSYSTEM flat plate solar collectors, 2.7m2 each. They are arranged into two solar fields, and each solar field has its own heat exchanger, and feeds heat into a different string of boilers, as shown in the figure below. In the current configuration, the heat exchanger gets water from the first tank, and feeds back to the sixth tank.

Cold water is input to the water boilers, shown on the figure in the upper left corner. The water flow is measured only at this point – while it is still cold and dense! Two boiler strings are connected in parallel. The temperatures are measured in each boiler, and the SCADA can switch each heating element independently.

The site has extensive piping to enable easy reconfiguration. Figure 3-3 only shows the pipes that are currently used, and therefore none of the valves. The hotel's hot water recirculation system takes water from the main hot water output and brings it back in the last boiler of each string.

The boilers are controlled by a Schneider Electric Modicon M251 PLC. Physically, the boilers are switched on and off by dedicated relays that are controlled directly by the controller.

The boilers are set up as a min-max funnel: the minimum temperature setpoints

are increasing, 25 degC for the first boiler, and 55 degC for the last boilers of each string. The maximum is 63 degC for all boilers. It is within that funnel that flexibility may be exercised.

The system at Hotel Borjana is very similar to the one at Hotel Paradise Blue, but much smaller. Compared to Paradise Blue's 414 kW total installed power, Borjana has only 120 kW, and the total flexible storage capacity is just 400 kWh, compared to over 700 kWh at Paradise Blue.

3.5.1 Why is the residential water heater model not applicable?

The residential water heater models used in other pilots in the INVADE project have been designed with 120-200l residential hot water tanks in mind, where temperature information is not available, the only source of heat is the electrical heating element, and control action is limited to blocking the heating element with a series switch.

Albena's complex water heating system can not be viewed as an enlarged version of a residential electrical water heater because it requires electricity consumption values, not a block/unblock type signal. This alone means that it needs its own optimisation model within the platform.

3.5.2 Handling the solar thermal production

The solar thermal heat production is closely linked to the solar irradiation and should be quite well predicted based on the weather forecast. (Errors 5-10%)

The hotel's hot water demand depends on the occupancy and many other factors, and the prediction is expected to be worse. (Errors 20-30%)

The solar thermal production is provided by Albena as an increasing meter reading, in kWh. The heat demand of the building will not be provided but can be calculated from an energy balance equation. This needs to be done in the platform. These historic values for the heat demand will be used for prediction model training and for making new predictions.

$$SOC[i] = SOC[i-1] + SolarThermal[i] + ElectricHeating[i] - HotelHeatDemand[i] \tag{3.1}$$

3.5.3 Calculating the heat demand for past periods

Timing must be considered to ensure consistency.

The solar heat capture is measured by an integrating meter. The heat production in each period is then simply the difference between the last two readings. The heat capture recorded to 0.15 is the amount of heat captured between 0.15 and 0.30.

The electrical consumption is handled the same.

The SOC reading is a snapshot valid at the time it was recorded, the amount of heat exchanged with the tank in the period starting at 0.15 is the difference between the values recorded at 0.15 and 0.30.

3.5.4 Local SOC approach - calculated on the SCADA

The flexible thermal storage capability can be obtained as:

$$(T^1_{max} - T^1_{min}) \times C^1 + (T^2_{max} - T^2_{min}) \times C^2 + ... + (T^{12}_{max} - T^{12}_{min}) \times C^{12} \tag{3.2}$$

With the current configuration, T^x_{max} is always 63 degC, and C^x is always $2000l \times 0.00116\frac{kWh}{kg \times degC} = 2.36\frac{kWh}{degC}$ without bothering about the density variation of the water. Assuming a T^x_{min} funnel of 25, 35, 45, 50, 55, 55 degC:

$$(63 \times 12 - 2 \times (25 + 35 + 45 + 50 + 55 + 55)) \times 2.36 = 533kWh \tag{3.3}$$

The SOC in kWh at any given moment would be given as:

$$(T^1 - T^1_{min} + T^2 - T^2_{min} + .. + T^{12} - T^{12}_{min}) \times 2.36 \tag{3.4}$$

This approach simplifies the communication to:

- Solar energy captured in the period (two fields together)
- SOC
- Electrical energy consumed

The electrical energy consumed should be only that consumed by the 12 boilers – excluding pumps, other heaters, lights, controllers, etc. Since this is not metered separately, it should be calculated in the SCADA as well:

$$W = \int_{period_{start}}^{period_{end}} 34.5kW \times \sum_{i} ON_i(t) \tag{3.5}$$

Where W is the energy consumption for the period for all the controlled boilers together; and ON_i is the relay status for the given heating element.

Thermal and electrical kWhs are used freely in this pilot because the conversion factor with immersion heating elements is 1. This would not be true if we used heat pumps or an electric stove for example.

The consumption of the boiler room is also recorded, in electrical kWh, as a steadily increasing meter reading. This can be used to verify that the control is working as hoped for.

3.6 Optimisation model

3.6.1 Sets

There are three sets indexing constraints and variables for the Albena pilot:

- S is the set of sites
- B is the set of boilers in a chain
- T is the set of periods in the planning horizon

Note that the set of sites in not explicit in the code, because each site is implemented as in independent block in pyomo. The results presented in the next section are for a single site, because that is sufficient to demonstrate the behaviour.

The parallel chains are not modelled separately: the assumption is that the chains behave identically, therefore it can be modelled as a single chain with 'virtual tanks' - in this case double the size of the real tanks, because there are two chains in parallel. Likewise, their maximum heating power is also double of the real tanks.

3.6.2 Decision variables

Five continuous decision variables for each period and each boiler:

- Heating energy $\pi_{s,b,t}^{AWH}$
- Constrained temperature $\tau_{s,b,t}^{AWH,constr}$
- Positive temperature error $\delta_{s,b,t}^{AWH,pos}$
- Negative temperature error $\delta_{s,b,t}^{AWH,neg}$
- Temperature $\tau_{s,b,t}^{AWH}$

The heating power is bound within the range zero to max heating power. The sum of all the heating powers will be treated as the output of the optimisation.

The constrained temperature is bound within the range defined by the funnel settings. The slack variables are included to make the model feasible under strange inputs, for example zero or negative or too high expected heat load. These situations will not occur during normal operation, therefore non-zero slack values are indicative of invalid input and should raise an alarm. They are included because it facilitates testing and troubleshooting.

3.6.3 Constraints

Constraints are used to describe the temperature evolution in the chain of tanks, the relationship between the temperature variables of each tank, and to enforce a final

state of charge.

$$\tau^{AWH}_{s,b,t} = \tau^{AWH,constr}_{s,b,t} + \delta^{AWH,neg}_{s,b,t} - \delta^{AWH,pos}_{s,b,t} : \forall s \in S, t \in T, b \in B \tag{3.6}$$

The temperature evolution of the tanks depends on the water flow in the system. In the case of high solar production and low hot water demand, the direction of the flow might reverse! Indeed that is the only way that the solar collectors can heat all the tanks, not only the final ones.

The water flow in the system can be estimated in a preprocessing step as follows:

$$flow[t] = \frac{HotelHeatDemand[t] - SolarThermal[t]}{0.00116\frac{kWh}{l \times degC} \times SolarOutletTemp[t]} : \forall t \in T \tag{3.7}$$

The temperature difference between the tank and the water flowing into the tank must be established for each tank. This can not be done in preprocessing anymore, because it depends on the tank temperatures, which are decision variables. Therefore $\Delta T_{s,b,t}$ is an expression:

$$\Delta T_{s,b,t} = \begin{cases} T_{inlet} - \tau^{AWH}_{s,b,t}, & flow > 0, b = 1 \\ \tau^{AWH}_{s,b-1,t} - \tau^{AWH}_{s,b,t}, & flow > 0, b > 1 \\ \tau^{AWH}_{s,b+1,t} - \tau^{AWH}_{s,b,t}, & flow <= 0, b < 6 \\ 63 - \tau^{AWH}_{s,b,t}, & flow <= 0, b = 6 \end{cases} \tag{3.8}$$

The actual temperature evolution constraint:

$$C \times \tau^{AWH}_{s,b+1,t} = C \times \tau^{AWH}_{s,b,t} + \pi^{AWH}_{s,b,t} + \Delta T_{s,b,t} \times 0.00116 \times flow_t : \forall s \in S, t \in T \setminus max(T), b \in B \tag{3.9}$$

$$\tau^{AWH}_{s,b,1} = T^{init}_{s,b} : \forall s \in S, b \in B \tag{3.10}$$

Constraining the decision variables:

$$0 \leq \pi_{s,b,t}^{AWH} \leq P_s^{AWH,max} : \forall\ t \in\ T,\ \forall\ s \in\ S \tag{3.11}$$

$$T_{s,b}^{min} \leq \tau_{s,b,t}^{AWH,constr} \leq T_{s,b}^{max} : \forall s \in\ S, b \in B, t \in T \tag{3.12}$$

$$0 \leq \delta_{s,b,t}^{AWH,pos} \leq 100 : \forall s \in\ S, b \in B, t \in T \tag{3.13}$$

$$0 \leq \delta_{s,b,t}^{AWH,neg} \leq 100 : \forall s \in\ S, b \in B, t \in T \tag{3.14}$$

3.6.4 Objective function

The real cost is simply the cost of energy. There is no preference regarding the state of charge. The slack variables must be included in the objective function with a large penalty to ensure that they are only selected non-zero if there is no other way to keep the model feasible.

$$Penalty \times \sum_{t \in T, b \in B} \delta_{s,b,t}^{AWH,pos} + \delta_{s,b,t}^{AWH,neg} \tag{3.15}$$

The penalty should be chosen as sufficiently large, but not too large as to make the model ill-conditioned. 10 euro/kWh would be a reasonable choice.

3.7 Results - model behaviour

Figure 3-4 shows predicted heat load and heat generation values from Paradise Blue on the 23rd of July. The water flow calculated based on this input in shown on Figure 3-5. In some high solar heat production periods the direction of water flow actually reverses, because the heat demand is lower than the production. Finally, Figure 3-6 shows the price input. That is not a real price export, but fictitious one, designed to highlight the behaviour of the model.

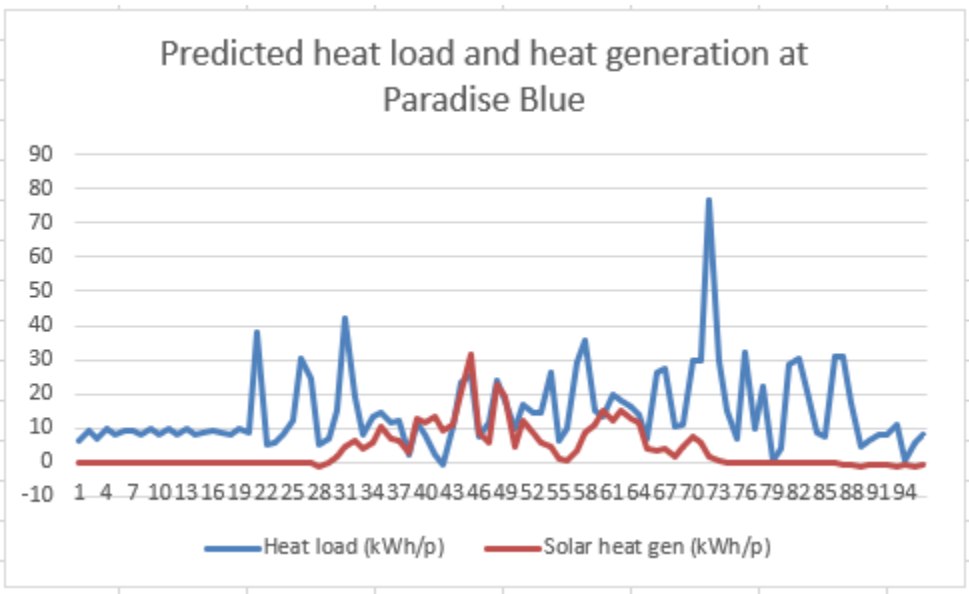

Figure 3-4: Expected heat load and heat production at Hotel Paradise Blue on 23/07/2019, serving as input to the optimisation model

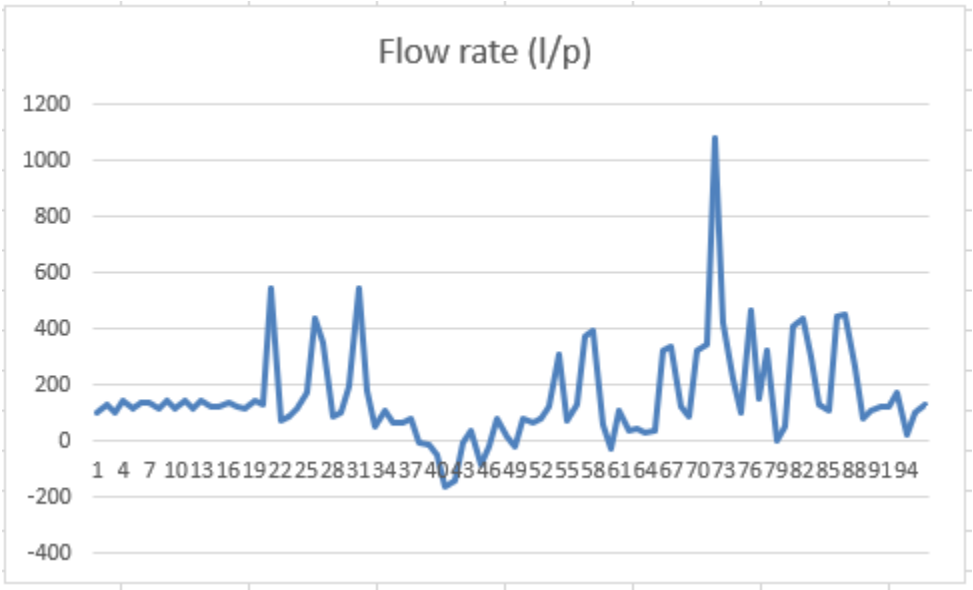

Figure 3-5: Calculated water flow based on the data shown in Figure 3-4

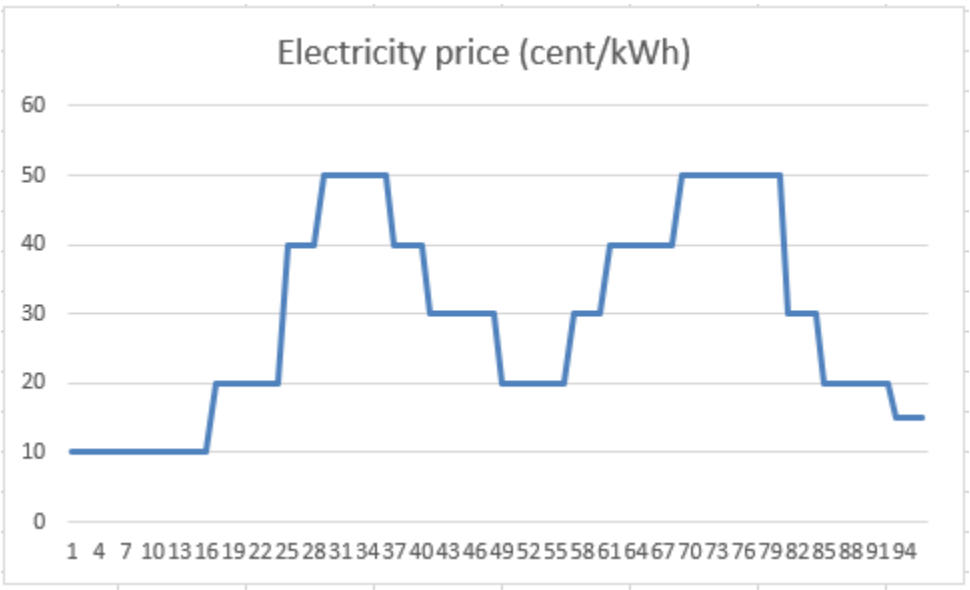

Figure 3-6: Fictitious and simplified electricity price to exercise the model

Figure 3-7 shows the energy budgets decided for each time period. Note that is the sum of all boilers! If compared to the price of energy, it visibly shifts electricity consumption away from the price peaks. The model decides zero heating energy for the last period, because the result of that heating would only show up outside of the planning horizon - therefore the model does not see any benefit to it. This is fine, since the rolling horizon approach ensures that only the first decisions are carried out.

Figure 3-8 shows the temperature evolution in the chain. The temperature funnel is respected, it is most visible in the periods 75 to 89, when T6 is at its lowest allowed value. The tanks higher up tend to have higher temperature, but not necessarily. For example, the model decides to heat up T1 and T2 to prepare for a long period of no-consumption - this way the funnel settings won't be violated as the water flows further in the system. Note that the model decides energy consumption for each boiler, but these are not enacted - only the summed up budget is communicated towards the site, and the local controller decides how to distribute that budget between the boilers. The point in modelling the individual chain components is only to avoid sending non-realisable budgets.

The local system does not have any concept of the future, and it distributes the consumption budget between the boilers such that:

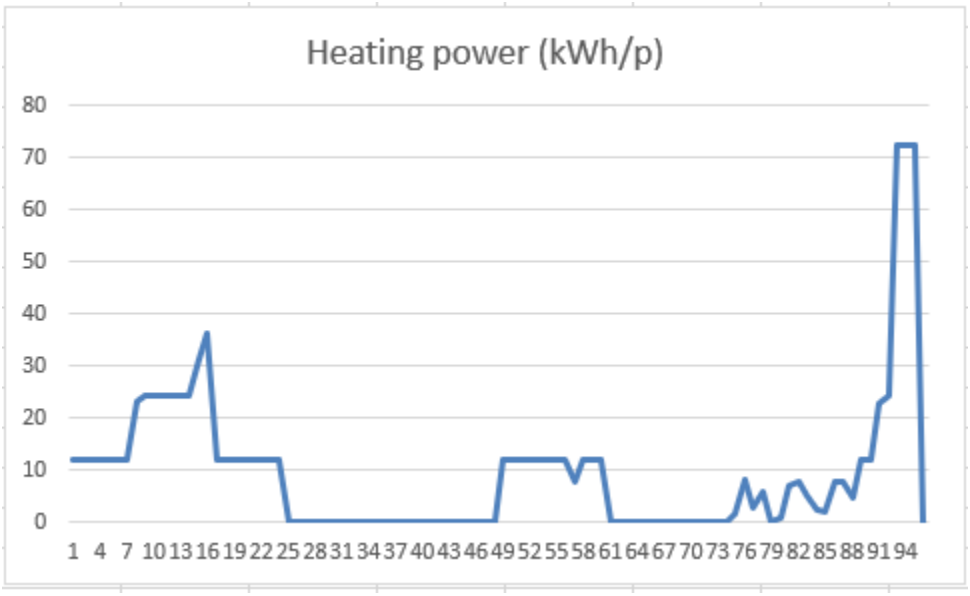

Figure 3-7: Consumption budget as decided by the model based on the above input data. Note that it is for the whole boiler system!

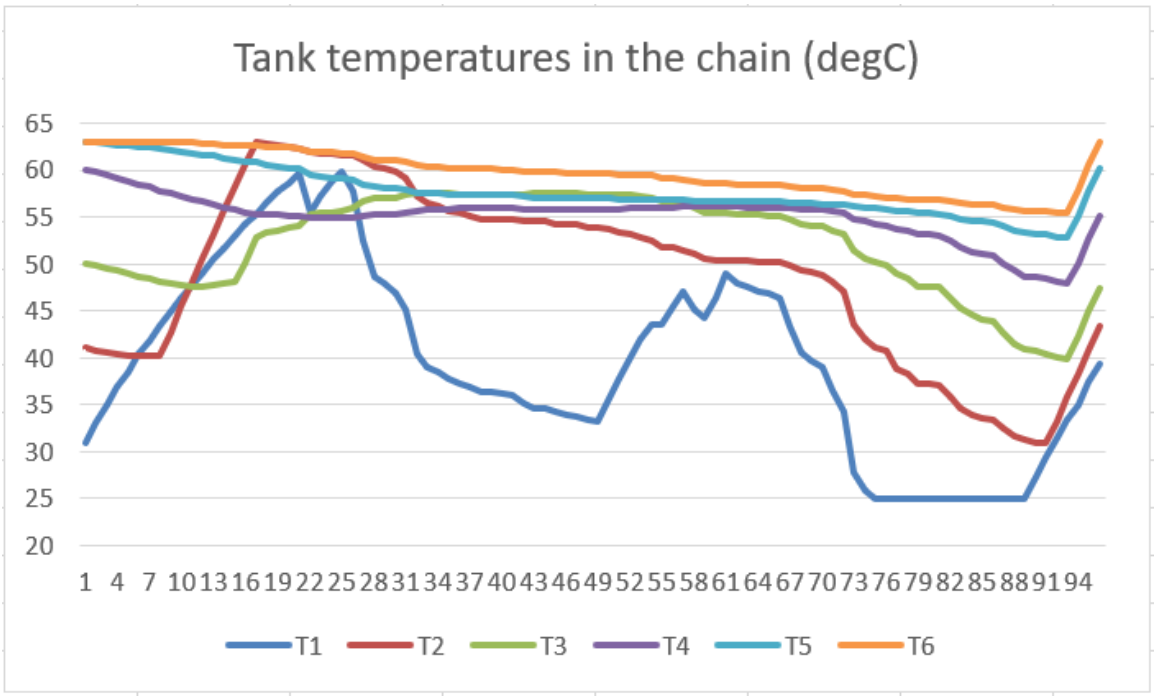

Figure 3-8: Modelled temperatures in each chain element in the boiler system. These are internal variables for the optimisation, and are not considered outputs other than for analysis.

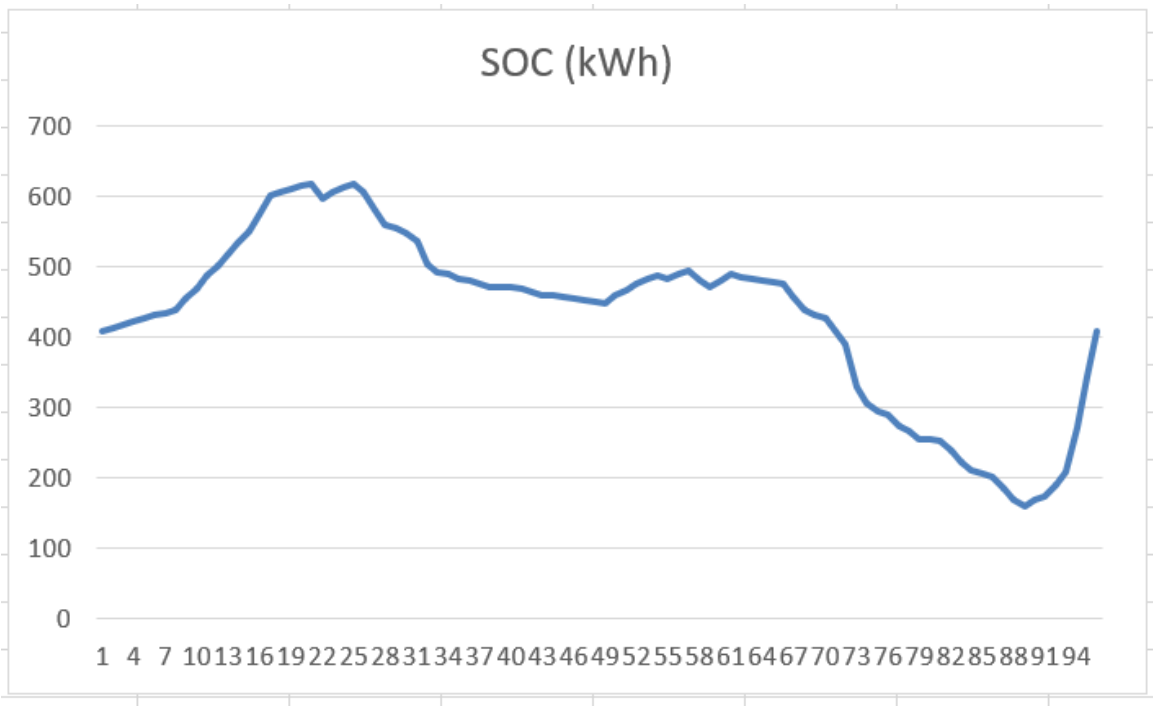

Figure 3-9: State-of-charge in the planning horizon - calculated based on the decided temperatures and the funnel settings.

- Boilers whose temperatures have fallen below the minimum are turned on - whether they fit into the budget or not
- The remaining energy is used to heat the remaining boilers evenly

This nonlinear behaviour is not modelled, because it would slow the model down, it might change in the future as the system is tuned, and because the error introduced by this simplification is insignificant compared to that introduced by the imperfect predictions.

Finally, Figure 3-9 shows the state of charge evolution of the boilers during the planning horizon. Note that the SOC is always high before the price peaks, and falls during the price peaks themselves.

3.8 Results - business case

To provide an upper bound of the saving potential with the model, it was run on an entire week's metered data from Hotel Paradise Blue. Instead of predictions for heat

demand and solar production, the actual metered values were used. There are three reasons why this approach could give an overly rosy picture of the savings:

- The real model is exposed to prediction error, which will lead to suboptimal decisions
- The real model's shorter planning horizon will lead to suboptimal decisions
- The real boiler system might not be able to accurately follow the control signal all the time

To make sure that the comparison is fair, the model was bound to start from, and end with, the same state-of-charge as the real system had at the beginning and the end of the week. This way the model would use exactly as much energy as the real system did, and the only way it can lower costs is by better scheduling.

The real Bulgarian spot price for that particular week was used.

Figure 3-10 shows the inputs and outputs. The top graph shows that the model schedules a lot more fluctuations in the state-of-charge of the boiler system, thereby exercising flexibility. The temperatures are shown to rise and fall in unison with the state-of-charge, as expected. The temperatures stay within their respective limits other than for some brief anomalies when the heat load has large peaks. Indeed the peaks in the load are so strange that they might signify corrupted data.

The optimised heating schedule has more peaks than the original, and the peaks coincide with the low-price periods. This is reasonable, that is exactly how the flexibility is exercised.

The original energy cost was 398.27 euros for that week for just the boiler room. The model's schedule would have cost only 336 euros, a 15.6% reduction. Note that this is just the price of electricity itself, not considering grid tariff, connection fee, etc., since those are all independent of the time-of-use.

Note that the solar heat production typically coincides with the high price periods. This means that the electricity consumption was already concentrated in the low price periods in the base case. When the same test is run for the Hotel Boryana, the

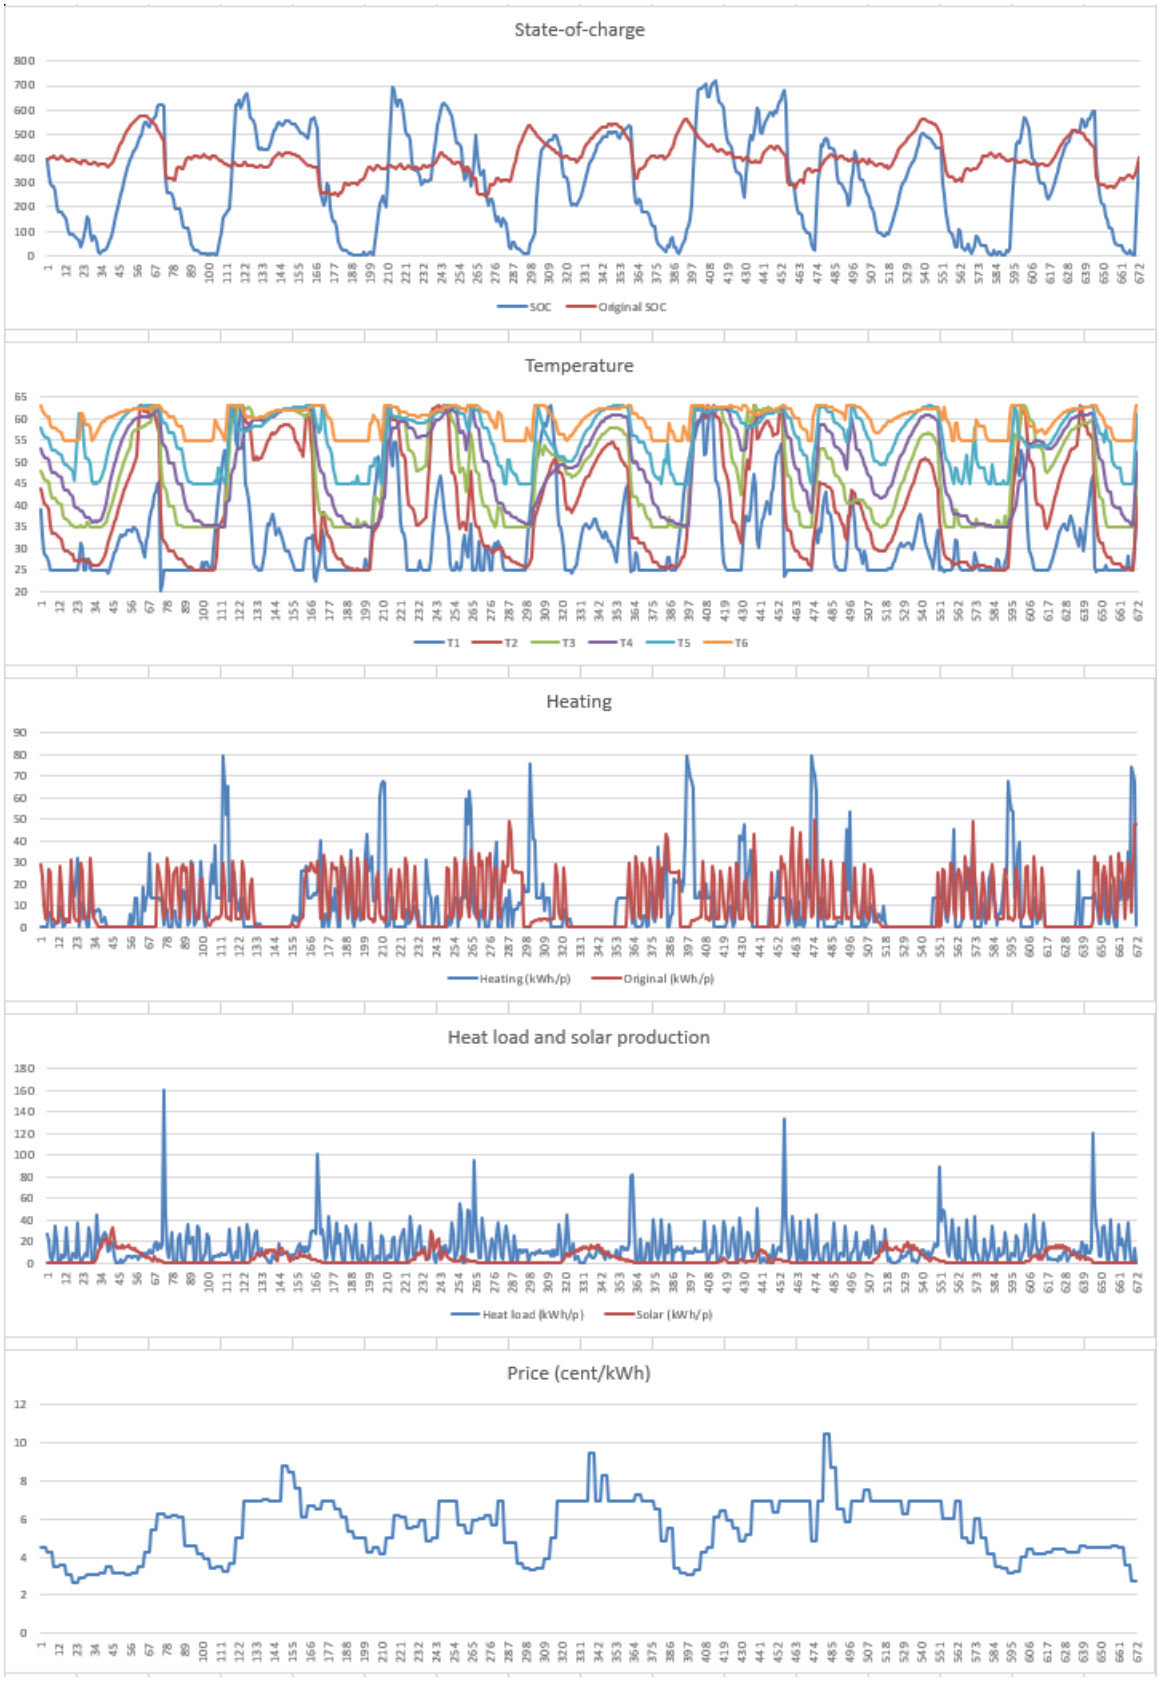

Figure 3-10: The week July 28 to August 3 at the Hotel Paradise Blue

percentage savings are lower (6%), because there the size of the solar collector system is bigger relative to the heat demand. This also means that the savings potential is much larger for hotels without a solar collector system.

The same test was executed on multiple weeks, producing similar results - therefore this week can be accepted as representative.

Chapter 4

REMA 1000'S Vagle distribution centre - thermal system modelling

4.1 Introduction

Earlier in the project, an optimization model has been defined which focuses only on the electrical system: deciding how the battery should be charged and discharged so that the total energy costs are minimized.

Only the second phase is discussed in this thesis, which is the control of the thermal system with the aim of using an on site water tank as thermal energy storage. This must work as an extension of the first phase of the project: In addition to the battery decisions, the model must also decide how the water tank should be filled and emptied of thermal energy.

4.2 General description of the thermal system

The thermal system can be run in two different operating situations:

- Summer operation, where the tank is used for storing cooling energy
- Winter operation, where the tank is used for storing heating energy

It changes from summer to winter operation once a year, and correspondingly from winter to summer operation.

4.2.1 Summer operation

In summer operation, this project focuses on the production of chilled water. This is done in two kuldeaggregats and a kjølemaskin. Both of these words mean 'cooling machine' in Norwegian, but it was convenient to retain the Norwegian name to tell them apart. The kuldeaggregats are also referred to as KA01 and KA02 further in the document. They produce chilled water, which is used to cover cooling loads, i.e. comfort cooling of the office space and cooling of computer room and the UPS. In addition, the kuldeaggregats produce cooled CO2, which is used for cooling the refrigerated warehouse and the cold store. Since the CO2 circuits do not have any connection with the accumulator tank, this is kept outside the project. The kuldeaggregats also produce surplus heat. We also disregard this in summer operation, but it is important when we come to winter operation. A kjølemaskin, referred to as KM06 further in the document, also produces chilled water, but not CO2 or heat. This is started only when KA01 and KA02 fail to meet the need for chilled water.

Figure 4-1 shows how the chilled water flows in different operating situations in summer operation. Top left shows a situation where the accumulator tank is not in use. Chilled chilled water flows from the kuldeaggregats and kjølemaskin to the cooling load. The computer cooling and comfort cooling is combined for convenience. Cooling energy is extracted via heat exchangers, increasing the temperature of the chilled water. The warmer chilled water returns to the kuldeaggregats and the kjølemaskin for re-cooling.

The top right shows what happens when the heat storage tank is charged with cooling energy. The cooled chilled water (from KA / KM) is split into a part that covers the load and another part that goes to the heat storage tank. In the accumulator tank, the cooling energy is extracted, the water in the tank is cooled and the chilled water increases in temperature. This returns to KA / KM together with the return water from the cooling load for re-cooling.

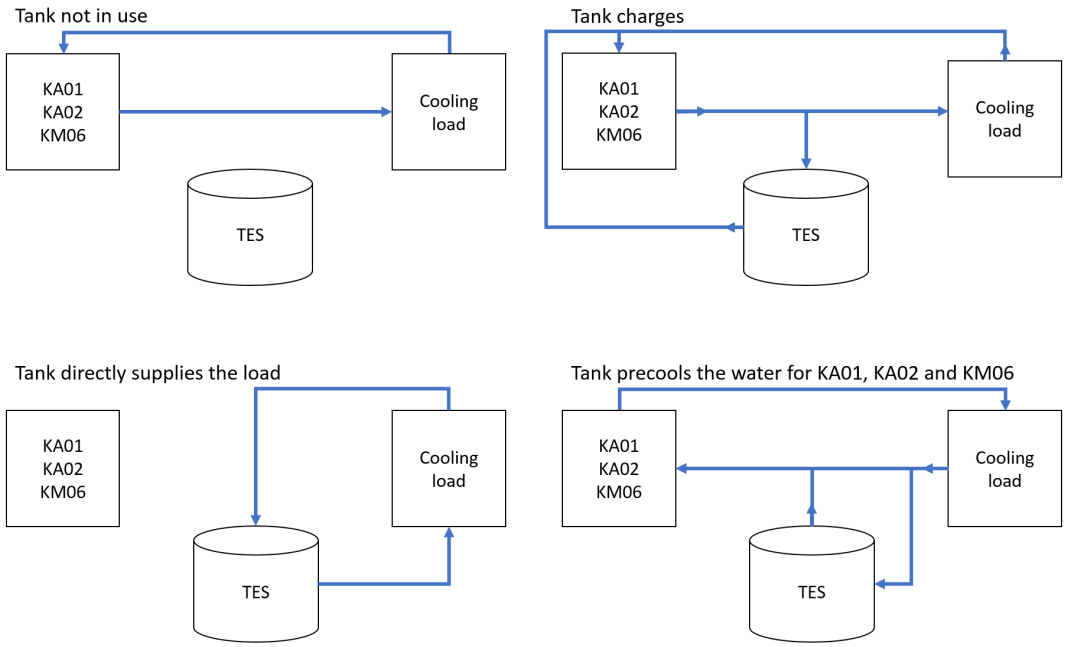

Figure 4-1: Simplified illustration of chilled water flow where the accumulator tank is in different operating situations

The two bottom illustrations show two different situations for discharging the tank. To the left is a situation where cooling energy from the tank is extracted and goes directly to the cooling load, i.e. without KA / KM contributing with the production of chilled water. To the right is a situation in which cold energy from the tank used to precool return water from the cooling load. This allows the KA / KM to produce chilled water using less energy in that period.

When the tank is being charged, the electrical consumption in KA / KM is increased. When discharging the tank, the electrical consumption is reduced. By charging and discharging appropriately, one can:

- use solar power internally instead of selling back to the network
- utilize price variations over the day (store when prices are low and use when prices are high)
- ensure that the cooling units only run in their most efficient operating points
- produce cooling energy at higher efficiency (when outdoor temperatures are low) and thus reduce the total electricity demand.

4.2.2 Winter operation

In winter operation, the focus of this project is on the production and storage of heat energy. The heat energy must cover the heating of the building and has two different sources. One is the cooling units, KA01 and KA02, where surplus heat from the production of chilled water and CO2 can be extracted in the form of hot water. The other source is an electric boiler, which produces hot water in situations where KA01 and KA02 do not produce enough heat to cover the heat load.

It is understood that the amount of heat available to extract from the cooling machines, and the COP at which it might be extracted, depends on how much cooling they are supplying. Therefore heat is available in a different distribution than desirable for the heating load, and the heat storage tank can be used to bridge this gap.

Figure 4-2 illustrates, again simplified, how the hot water flows in different operating situations. The top left shows the flow when the tank is not in use. Water is heated in KA, and optionally further heated in the electric boiler, and flows directly to the heat load. Here, heat energy is extracted, and the water returns to KA for re-heating. At the electric boiler, the lines are dotted to indicate that all, part, or none of the water can flow through the electric boiler.

At the top right we see what happens when the tank is charged. Heated hot water is split between the tank and the heat load (again possibly via the electric boiler). In the tank, heat energy is extracted, the water in the tank is heated, and water is returned mixed with the return the water from the heat load.

The bottom of the figure shows the discharge of the tank. Then water from the tank is mixed with hot water from the cooling units. The return water from the heat load is split into one part to the accumulator tank and one part back to the cooling units.

The purpose of this process is to store heat to avoid using the electric boiler and to take full advantage of the cheaply available heat reject from the cooling machines. In this way one can avoid electric power peaks, reduce the purchase in hours with

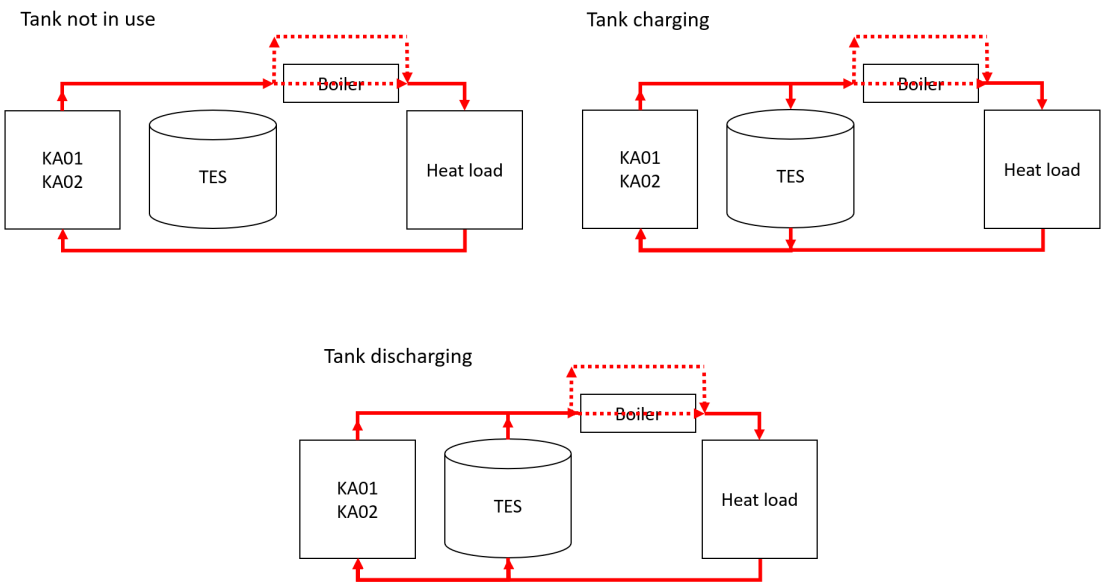

Figure 4-2: Simplified illustration of hot water flow where the accumulator tank is in different operating situations

high electricity prices, and reduce the total electricity consumption.

4.3 Component model in eSmart's platform

4.3.1 Phase 1

In the first phase of the project, only the electrical system was considered, and the battery was controlled. See Figure 4-3. A single component model was established, consisting of the main meter, the PV system, the battery, and the collection of all electrical loads.

- The main meter measures net imports and net exports from / to the network, i.e. buying and selling.
- The PV plant consists of roof-mounted solar panels and inverters.
- Battery bank consists of batteries and inverter.
- The Total Load represents all electrical consumables in the building. The load is not measured directly, and the electrical consumption is calculated in the

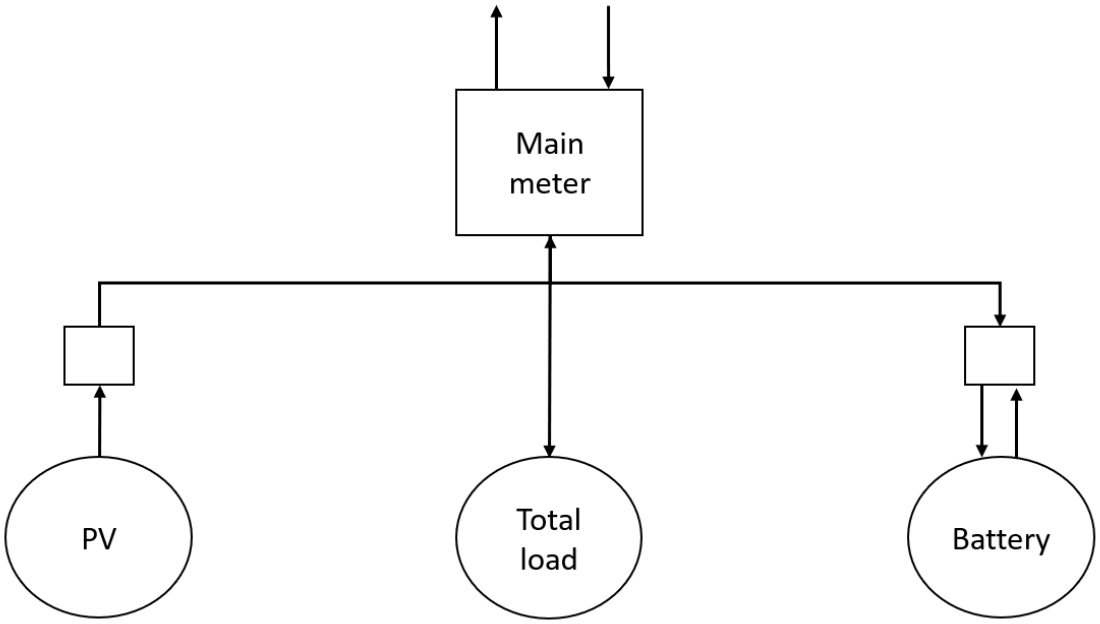

Figure 4-3: Component overview for phase 1

FlexClient as a residual load.

4.3.2 Phase 2

In the second phase the component model is expanded so that the thermal system can be optimally handled. The main meter, PV system, and battery will be the same as in phase one; but the following changes and additions will be made:

- The electric boiler is introduced.
- The two kuldeaggregats (KA01 and KA02) and the kjølemaskin (KM06) are introduced.
- The heat storage tank is introduced.
- Heat load is introduced.
- Server room cooling and comfort cooling are introduced. These constitute the need for chilled water.

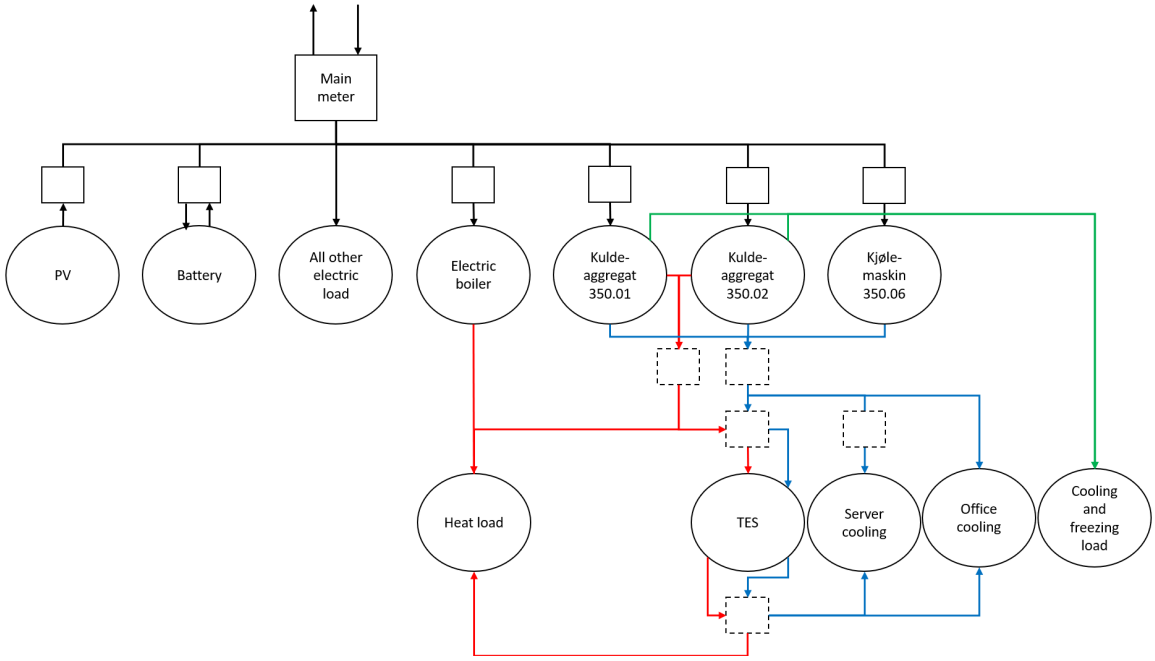

Figure 4-4: Component overview w/ full representation for phase 2

- The load component (Total load in Figure 4-3) from phase 1 is redefined. Now this represents all electric consumption except the boiler, the kuldeaggregats, and the kjølemaskin.

A full overview of all physical components is shown in Figure 4-4, where the circles represent electrical / thermal components that produce, consume or store energy while solid rectangles represent meters of electrical energy and dashed squares represent meters of thermal energy. Black arrows represent electricity, green CO2, blue chilled water, and red hot water. Note that all meters are real. In the figure, there are measures for energy into the accumulator tank and another for out. In reality, this is the same meter with different counters.

4.3.3 Phase 2 in summer operation

Figure 4-5 shows the simplified representation for summer operation. The accumulator tank is used for cooling storage. Note that the kuldeaggregats on the diagram represent only the part related to the production of chilled water! That is why their electricity meters were removed from the figure. In addition, all conditions related

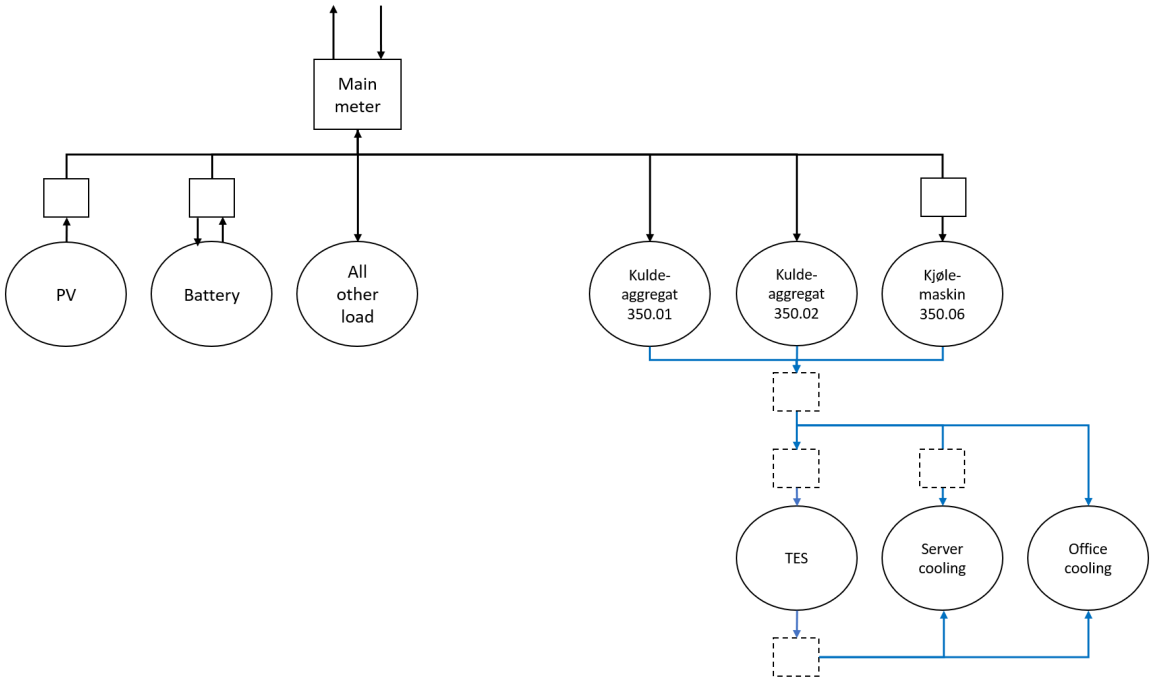

Figure 4-5: Component overview for phase 2 with simplified representation and summer operation

to CO2 and heat have been removed. The arrows with electricity into the kuldeaggregats now represent only the electrical energy used to produce chilled water. This cannot be measured but can be calculated from some simplified assumptions. More on this later.

4.3.4 Phase 2 in winter operation

Figure 4-6 shows the component overview in winter operation, when the accumulator tank is used for heat storage. Here, the CO2 part and cooling load chilled water is removed, as heat energy is stored in the accumulator tank. Note that chilled water is also produced in winter (for computer rooms and UPS), but it is not directly relevant to heat storage.

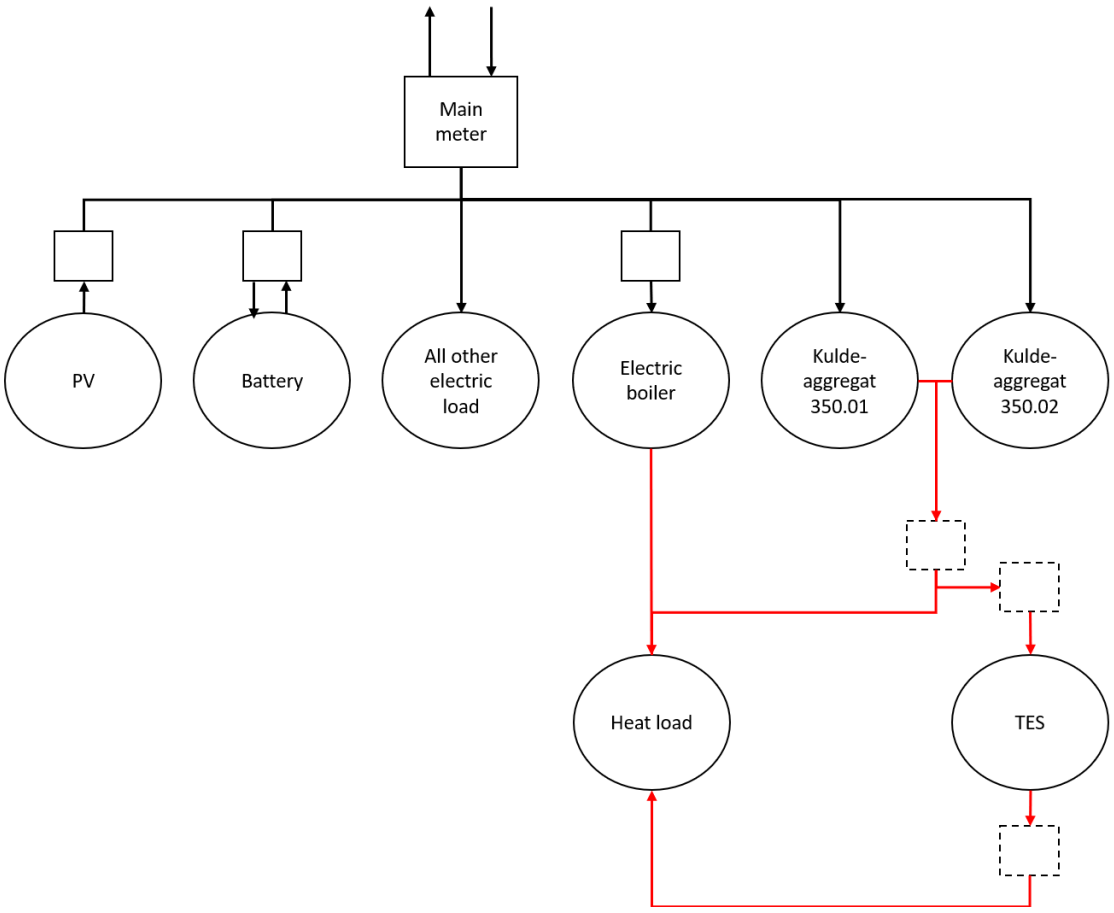

Figure 4-6: Component overview for phase 2 with simplified representation and winter operation

4.4 The operating principle of the thermal components

4.4.1 Kuldeaggregats' and kjølemaskin's chilled water production

Kuldeaggregats convert electric energy to thermal energy in the form of cold CO2, chilled water, and hot water. The kjølemaskin only produces chilled water.

The relationship between the amount of electricity in and the thermal energy out (in the form of the three components) is not unambiguous and depends on many different operating conditions: temperature, pressure etc. We ignore the CO2 part and focus first on the chilled water section and the summer situation, as shown in Figure 4-5

There are five relevant meters on figure 4-5:

- Electricity meter to KM06
- Energy meter from KA01, KA02, KM06
- Energy meter in to and out of acc. tank
- Energy meter for computer room cooling

Note that energy in and out of the accumulator tank is the same meter and same measurement value. If the flow goes in or out must be calculated and depends on the temperature differences and whether the tank is in summer or winter operation.

It is essential to know the relationship between electricity into each of the two kuldeaggregats and the kjølemaskin and thermal energy in the form of chilled water from each of them. This is not known exactly, both because it is not possible to measure specifically the part of electricity that are associated with chilled water production and because there is no separate thermal energy measurement from the three machines. These values must therefore be calculated. Central to this is an

Figure 4-7: Overview of KM and AC compressors in KA01 and KA02

overview of COP values (Coefficient of Performance or efficiency), which was estimated by Rune Grytnes of ECOFrigo. These COP values define the relationship between electrical energy and thermal energy in the form of chilled water from each kuldeaggregat/kjølemaskin.

For KA01 and KA02 there are two different operating situations: With and without heat recovery. When the building has heating needs, heat recovery is activated, and KA01 and KA02 are adjusted to produce waste heat with a higher temperature to supply the heat load. This reduces the COP for chilled water production. The operating mode without heat recovery is also called Max COP.

In operating conditions without heat recovery, there are again two different operating states, which have a great impact on the COP:

- MT, which means that all air-conditioning load is covered by the cooling compressors (KM) and thus the air-conditioner compressors (AC) are not running, see Figure 4-7. This happens when the cooling load is not sufficiently large.

- Parallel, which means that also the AC compressors are running in parallel with the cooling compressors. This gives considerably higher COP than what MT provides.

In both situations, the COP is dependent on the outdoor temperature and the supply temperature of the chilled water circuit, see Figure 4-8. To find the right COP value, one must determine if the system is running in parallel or MT operation, find

the return temperature on the chilled water and the outside temperature. The power (35 kW) is thermal, and that means that for example if the supply temperature is 7 degC, the outside temperature is 10 degC, and the system runs in parallel mode, then the COP is 7.89. This means that an electric power of 4.4 kW is required to produce 35 kW of cooling in the form of chilled water.

In practice the supply temperature is usually very close to 9 degC, and therefore the model will always use that part of the COP table.

The missing values of the tables signify operating conditions that are not allowed. The model assumes a COP of 2 in those conditions. That is sufficiently low that the model will always avoid using those conditions. Damaging the systems is not a concern, because the only output of the optimization model is a recommendation for the amount of energy to be exchanged with the thermal energy storage tank in each period.

While the supply temperature and the outside temperature are the important parameters to extract the correct COP in a situation without heat recovery, it is the supply temperature and pressure which are vital in situations with heat recovery. Measured pressure is obtained and reference pressure can be obtained from the IWMAC system.

An example of the COP table in the heat recovery situation is shown in Figure 4-9. Here we see that the COP curve has 5 breakpoints for different cooling liquid temperatures (væsketemperatur) out of the gas cooler.

The table shown in Figure 4-9 is be used to estimate the COP of heat recovery for past periods. The COP depends on how much of the available heat is being extracted, because the system increases the pressure to extract more heat! eSmart should detect the relationship using machine learning techniques and implement the dependency into the model. Before this happens, a historical average COP might be used.

At the moment, only the sum heat production of KA01 and KA02 is measured. In the near future, a new machine KA04 will be added, and the heat production of the three machines will be metered individually. The distribution of the heat production

System **350.01/350.02**

Uten varmegjenvinning

		COP	
		Parallell	MT
Turtempera-tur	Utetemperatur	35 kW	35 kW
7	10	7,89	4,64
	12	7,86	4,57
	14	7,74	4,49
	16	6,83	4,07
	18	6,03	3,67
	20	5,33	3,31
	22	4,69	2,97
	24	4,43	
	26	4,18	
	28	3,81	
	30	3,35	
	32	2,98	
	34	2,67	
9	10	8,00	4,64
	12	7,87	4,57
	14	7,74	4,49
	16	7,61	4,07
	18	6,68	3,67
	20	5,87	3,31
	22	5,14	2,97
	24	4,72	
	26	4,47	
	28	4,13	
	30	3,63	
	32	3,21	
	34	2,87	

Figure 4-8: COP-table for KA01 and KA02 when run without heat recovery

System **350.01/350.02**

Med varmegjenvinning

		COP				
		Væsketemperatur ut av gasskjøler				
Turtempera-tur	Trykk	14	18	22	26	30
7	84	4,01	3,8	3,56	3,3	2,99
	82	4,14	3,92	3,67	3,4	3,06
	80	4,29	4,05	3,79	3,5	3,12
	78	4,44	4,2	3,92	3,61	3,19
	76	4,61	4,35	4,06	3,72	3,24
	74	4,79	4,52	4,21	3,84	3,26
	72	4,74	4,44	4,1	3,69	3,07
	70	4,97	4,65	4,29	3,87	
	68	5,18	4,86	4,49	4,05	
	66	5,41	5,07	4,68	4,22	
	64	5,76	5,39	4,98		
	62	6,1	5,71	5,59		
	60	6,47	6,05	5,59		

Figure 4-9: COP table for heat recovery for KA01 and KA02

between the machines is not equal. At small heat loads, only KA01 is producing heat, and at a certain threshold KA02 joins in as well.

For KM06 the supply temperature and outdoor temperature are crucial parameters, but here the COP curve has four break points for different thermal power output levels. Again, the missing values of the tables signify operating conditions that are not allowed. The model assumes a COP of 2 in those conditions.

Production of thermal energy in the form of chilled water is only measured for all three plants overall, see Figure 4-5. Distribution of thermal energy from each plant must be calculated based on the cooling system's as-built documentation. Once thermal energy is calculated from the plants, electrical energy can be calculated using the COP tables.

4.4.2 The kuldeaggregats' heat production

In the winter situation, the heat energy (i.e. the surplus heat) from KA01 and KA02 is to be stored in the heat storage tank. This is illustrated in Figure 4-6. Only the

System 350.06

		COP basert på ytelsesbehov [kW]			
Turtempera-tur	Utetempera-tur	30	60	90	120
7	10		6,59	6,92	7,09
	12		6,46	6,82	6,99
	14		6,26	6,63	6,37
	16		5,67	5,95	5,70
	18		5,11	5,34	5,02
	20		4,62	4,81	4,52
	22		4,18	4,32	4,08
	24		3,82	3,88	3,77
	26		3,62	3,71	3,57
	28		3,28	3,12	3,20
	30	2,70	2,93	2,80	2,86
	32	2,42	2,64	2,54	2,57
	34	2,23	2,39	2,31	2,34
9	10				
	12				
	14		6,91	7,31	7,49
	16		6,15	6,52	6,69
	18		5,54	5,83	5,60
	20		4,97	5,23	5,01
	22		4,48	4,68	4,50
	24		4,03	4,18	4,05
	26		3,89	3,96	3,84
	28		3,52	3,58	3,44
	30		3,13	2,98	3,06
	32	2,60	2,81	2,68	2,75
	34	2,34	2,54	2,44	2,49

Figure 4-10: COP-table for KM06

overall heating energy (provided by the two plants together) is measured. This heat is obtained not only from the chilled water production in KA01 and KA02, but also from the production of cold CO2 to the refrigeration and freezing warehouse. The amount of heat available therefore depends on how much cooling the machines are supplying overall!

Similarly to the production of chilled water, there is only a sum measurement of the heat energy from the two refrigeration units.

The exact way of calculating the distribution of heat production between the two cold plants is to be specified by the customer, and it has not been decided yet.

4.4.3 The heat storage tank

The accumulator tank works in principle similarly to the battery: The tank is charged by supplying thermal energy and the tank is discharged by extracting thermal energy. Thermal energy is calculated as volume multiplied by temperature difference and scaled with volumetric heat capacity of course. There are measurements of how much is supplied and extracted. Figure 4-6 which shows the accumulator tank and the

meters in a situation where the tank is used for heat storage. In order to know whether the tank is being charged or discharged, it is necessary to know the temperature difference on the supply and return water and whether the plant is in winter or summer operation.

If we are in refrigerated storage mode (summer operation), the tank is charged if the temperature into the tank is lower than the temperature out of the tank. In the opposite case, the cold energy is taken out.

If we are in heat storage mode (winter operation), the tank is charged if the temperature supplied to the tank is higher than the temperature withdrawn from the tank.

Like a battery, the thermal storage has a balance that is based on energy content in the previous period and the heat exchanged with it in the current period, considering efficiency parameters. However, the accumulator tank as thermal storage differs from the electric battery in that the energy content also changes with temperature conditions in the plant. The amount of usable energy depends on the temperature difference between the water in the tank and the return water, that is, the water passing from the accumulator tank to the kuldeaggregats and the kjølemaskin.

The temperatures in the tank are measured using 12 different sensors (RT20 to RT31), in four different layers in the tank, see the left part of Figure 4-11.

To calculate the energy content of the tank, the average temperature must first be calculated. The tank's capacity is 300.000 litres; therefore, its thermal mass is 348 kWh/degC. The thermal layering effect is considered insignificant.

E2U supplies eSmart with four temperature points, each the average of their respective levels. The eSmart platform averages them to supply the optimisation model with a single temperature.

Heat exchanger

The amount of energy available from or storable in the tank depends on temperatures elsewhere in the system. In both operating modes:

- the return temperature determines the amount of energy available,

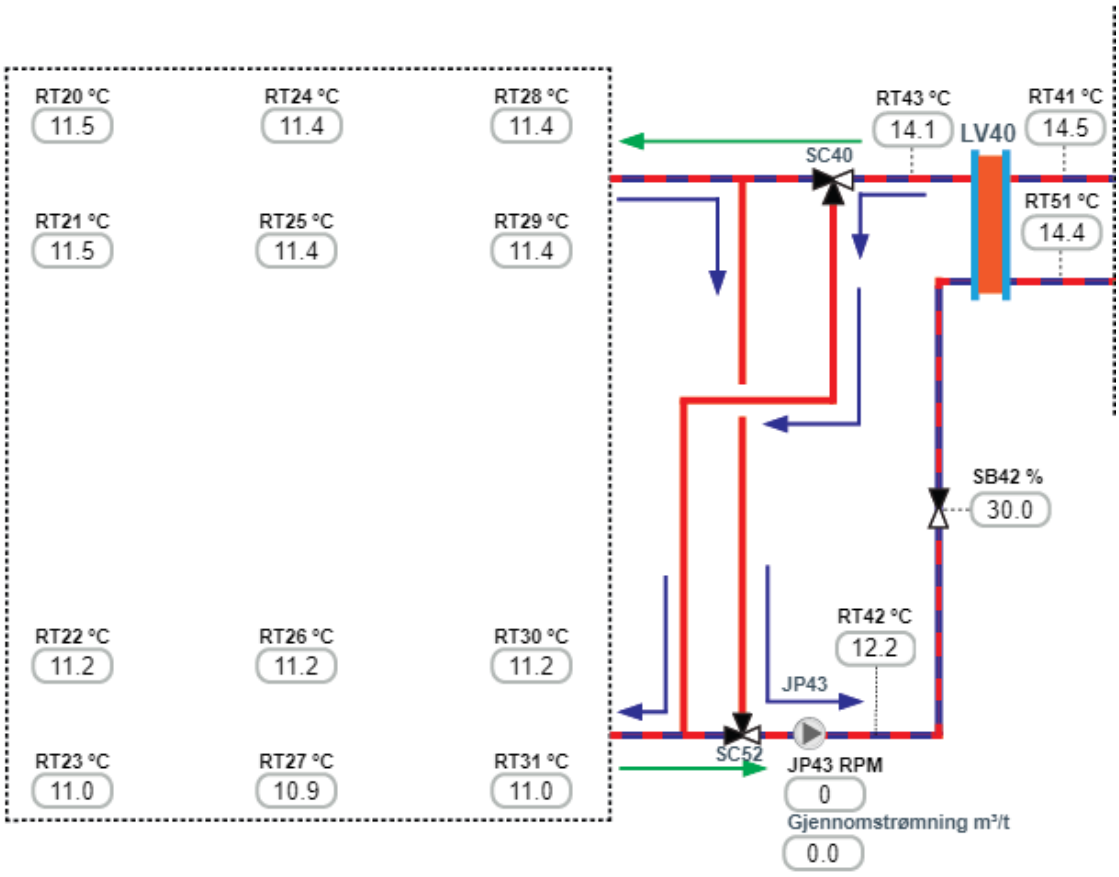

Figure 4-11: Temperature measurements in the tank and on the heat exchanger

- the supply temperature determines the amount that could be stored.

These temperatures can be obtained using a formula from Håkon Person of Multiconsult relating it to the outside temperature, in winter and summer operation:

$$T_{tur}^{winter} = \begin{cases} 52degC, & T_{outside} \leq -10degC \\ -0.1 \times T_{outside} + 51degC, & -10degC < T_{outside} \leq 10degC \\ -1.43 \times T_{outside} + 64.3degC, & 10degC < T_{outside} \leq 17degC \\ -3.33 \times T_{outside} + 96.7degC, & 17degC < T_{outside} < 20degC \\ 30degC, & 20degC \leq T_{outside} \end{cases} \tag{4.1}$$

$$T_{tur}^{summer} = \begin{cases} 12degC, & T_{outside} \leq 5degC \\ -0.2 \times T_{outside} + 13degC, & 5degC < T_{outside} \leq 15degC \\ -0.6 \times T_{outside} + 19degC, & 15degC < T_{outside} \leq 20degC \\ -0.2 \times T_{outside} + 11degC, & 20degC < T_{outside} < 25degC \\ 6degC, & 25degC \leq T_{outside} \end{cases} \tag{4.2}$$

The temperature drop between the supply and return in winter heating mode is 25 degC. The temperature increase in summer cooling mode is 6 degC.

The heat rate is related to the surface area of the heat exchanger, the flow rate of the fluid, the heat capacity of the fluid, the temperature difference between the fluid and the tank, and an additional margin temperature. For example, the maximum heat rate into the tank in winter operation:

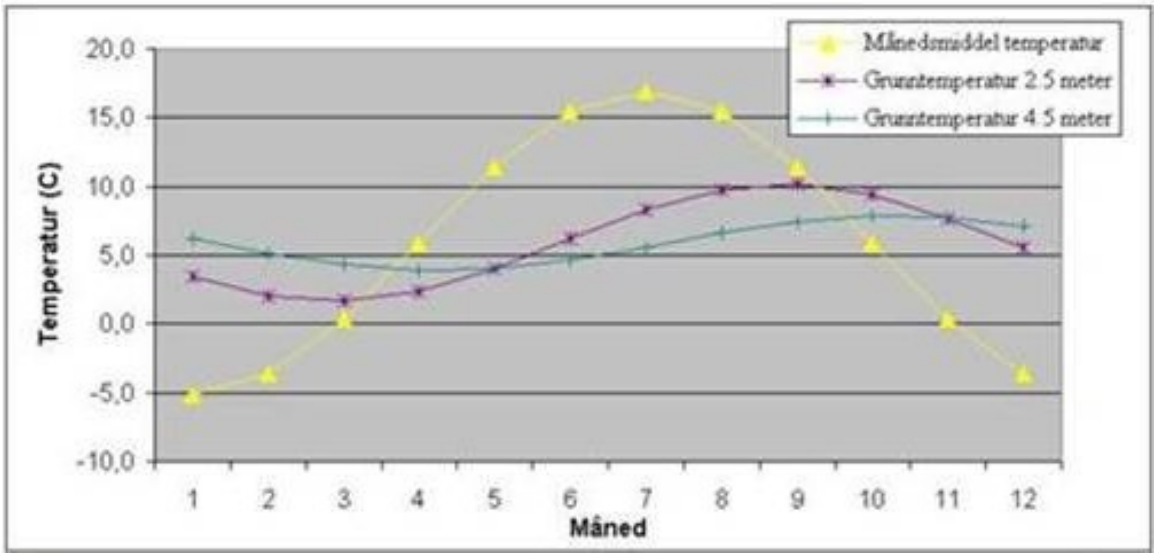

Figure 4-12: Ground temperature from [18]

$$Q_{max,in}^{winter} = V \times \rho \times c_p \times (T_{tur}^{winter} - T_{TES} - K_{tap}) \tag{4.3}$$

$$V = 4.45 l/s \tag{4.4}$$

$$\rho = 0.99 kg/l \tag{4.5}$$

$$c_p = 4.18 \frac{kJ}{kgdegC} \tag{4.6}$$

$$K_{tap} = 2 degC \tag{4.7}$$

Heat loss

The energy content also changes by losing heat energy to the surroundings. The heat loss is calculated from the tank temperature, the surrounding ground temperature, the heat transmission factor, and the surface. The heat loss model is simplified and does not consider the disturbances from some less insulated surfaces in one of the side walls and parts of the roof surface.

The tank is placed in the ground, at an approximate average depth of 2,5 meters. A SINTEF-report [18] estimates the temperature profile at different depths in the ground for Oslo, for the different months of the year in an average climatic year. See

Jan	Feb	Mar	Apr	May	June	July	Aug	Sept	Oct	Nov	Dec
7	6	5	6	7	8	9	10	11	10	9	8

Table 4.1: Suggested monthly ground temperatures in degC by Håkon Person

figure 4-12, the purple profile is for 2,5 meters depth.

The normal ambient temperature average from the closest weather station is 8 degC, which is a couple of degrees higher than the Oslo average. Also, Stavanger has less temperature variations than Oslo. Table 4.1 lists monthly values as suggested by Håkon Person of Multiconsult (representing the customer) for ground temperatures surrounding the tank each month.

The surface area is taken from the 3D model of building to be $(12m+10m+12m+10m)*3m+12m*10m*2 = 372m^2$. Heat transmission factor assuming 200 mm insulation around the tank, is approximately 0,18 W/m2 K. This is approximately 90% of the total surface area. For $(12*3)/372 = 10\%$ of the surface there is approximately 100 mm insulation in average, i.e approximately $0,36W/m^2K$. The average heat transmission value is then calculated to $0,18*0,9+0,36*0,1 = 0,20W/m^2K$. The heat loss is calculated to $0,20*372 = 75W/dT$, of which dT is $(T_{tank} - T_{ground,month})$. A positive value indicates heat loss, and a negative is heat gain. The numbers and calculations in these two paragraphs were provided by Håkon Person.

4.4.4 Electric heater

The electric boiler converts electrical energy to thermal heat energy. Electrical consumption of the electric boiler is separately measured. Heat energy from the electric boiler is not measured but can be calculated from the electric consumption using the efficiency, which has been specified by the customer as 0,9.

Since KA01 and KA02 have COP values of over 2.0, they will produce heat much more efficiently than the electric boiler. The boiler will only be used if the cooling units cannot supply enough heat.

4.4.5 Chilled water supplied cooling load

Two cooling loads are met with chilled water:

- Ventilation cooling of offices and other rooms
- Cooling of computer room and UPS

Only the cooling energy flowing towards the computer room and the UPS are measured, therefore the cooling load for the ventilation must be calculated from Equation 4.8.

$$CoolingLoad_t^{IT} + CoolingLoad_t^{Office} = Cooling_t + Charging_t \tag{4.8}$$

4.4.6 Heat load

The heat load corresponds to heating of rooms in winter. This is metered.

4.5 Changes to the FlexClient

4.5.1 Introduction

This section describes changes that must be made to the eSmart's FlexClient to implement thermal storage management. (The FlexClient is the same as the Integrated INVADE Platform or IIP mentioned in other chapters.) First, it describes which new data must be received from external systems. Then, changes to be implemented by eSmart, divided into two parts: a section describing calculations necessary to produce the input data to the optimization algorithm and a part describing add-ons in the optimization algorithms themselves.

4.5.2 New data that eSmart must receive from external systems

Data for past periods

Data needed for predictions to be performed in the eSmart platform:

- Cooling consumption of the server room and UPS, $CoolingLoad_t^{IT}$
- Consumption for comfort cooling, $CoolingLoad_t^{Office}$
- Total heating load, $HeatLoad_t$
- Available heat, $MaxHeat_t^{KA0x}$
- Heating COP, $COP_t^{Heating,KA0x}$

Data needed to initialize optimisation:

- Average temperature of the heat storage tank, TES_t

Data needed to accurately calculate residual load in the summer:

- Electrical consumption of KM06
- Total cooling energy from KA01 + KA02 + KM06, $Cooling_t$
- $COP_t^{Cooling,KA}$ and COP_t^{KM06}

Winter:

- Heat production from KA01+KA02
- COP for heat energy production from the kuldeaggregats, $COP_t^{Heating,KA01}$ and $COP_t^{Heating,KA02}$
- Electrical consumption for the electric boiler, $Heating_t^{Boiler}$

Data that is nice to have for analysing and verifying behaviour:

- Supply and return chilled water temperature of KA01 and KA02

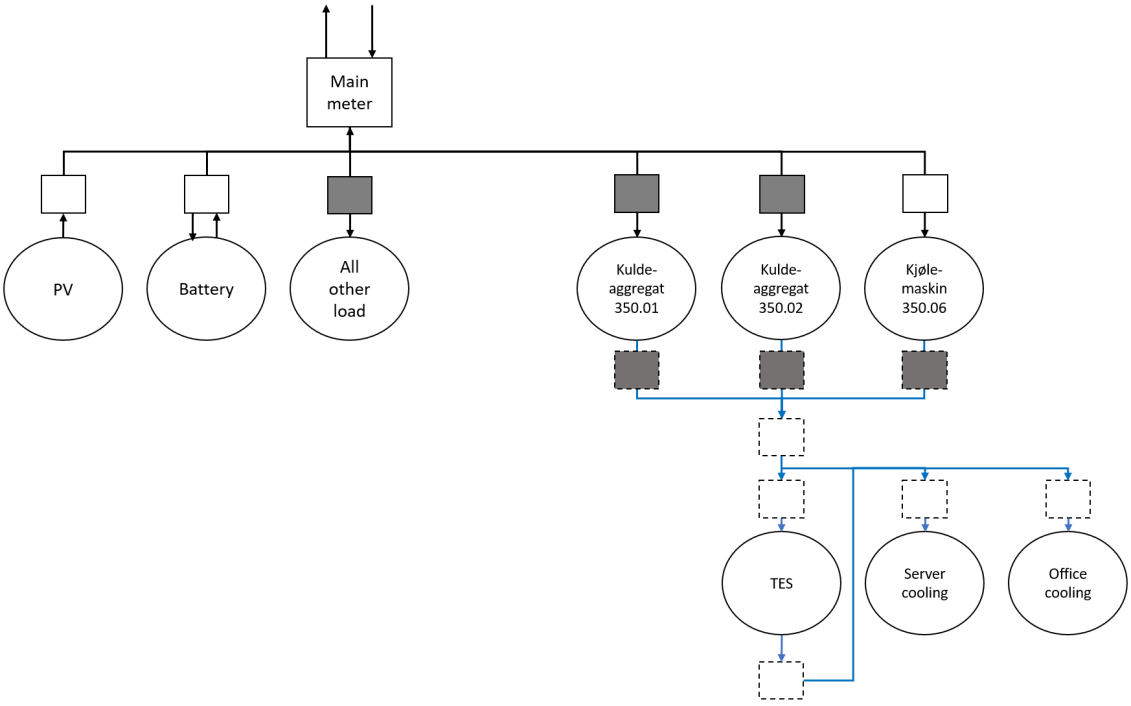

Figure 4-13: Overview of measurement points for which eSmart receives values (squares without fill) and points that eSmart calculates values for (filled squares) in summer operation

- Energy loaded into the TES, $Charging_t$

With recorded values of $Charging_t$, we can calculate the heat loss of the tank and involve it in future improvements in the model. With the chilled water temperatures, we can estimate COP more accurately and improve our residual electric load calculation, which in turn will improve the predictions.

Note that this means that the FlexClient receives measured / calculated data for more points than shown in Figure 4-5 and Figure 4-6. The FlexClient "sees" energy meters as shown in Figure 4-13 for summer operation and Figure 4-14 for winter operation. The white rectangles represent points that eSmart receives data for. Grey rectangles represent points where eSmart calculates historical values. This is specified in the next section.

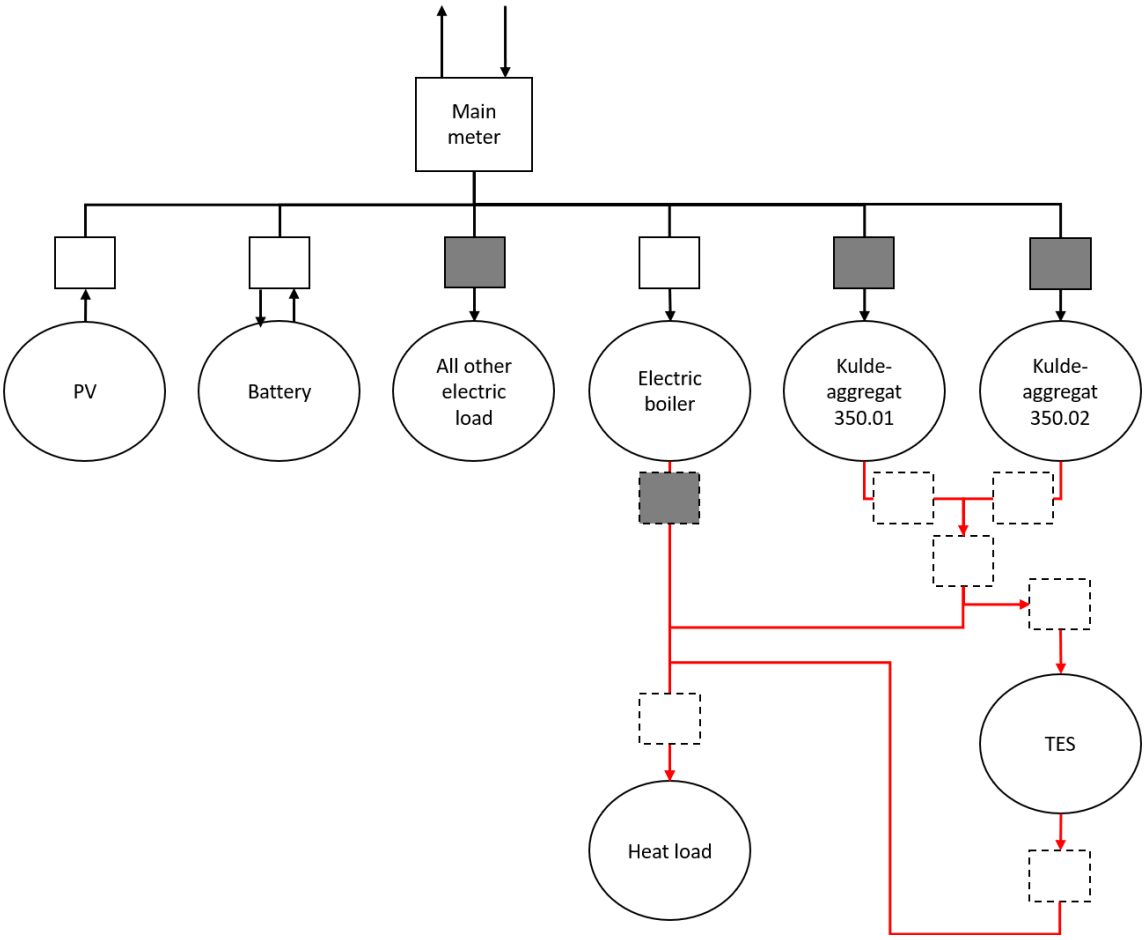

Figure 4-14: Overview of measurement points for which eSmart receives values (squares without fill) and points that eSmart calculates values for (filled squares) in winter operation

4.5.3 New calculations in eSmart's platform

Calculations for past periods

eSmart must calculate the residual electric load for the past periods to be able to train the prediction for the future. This means that the electric load from the kuldeaggregats must be calculated.

In summer:

$$Cooling_t^{KM06} = COP_t^{KM06} \times El_t^{KM06} \tag{4.9}$$

$$Cooling_t^{KA01+KA02} = Coolint_t - Cooling_t^{KM06} \tag{4.10}$$

$$El_t^{KA01+KA02} = \frac{Cooling_t^{KA01+KA02}}{COP_t^{Cooling,KA}} \tag{4.11}$$

In winter, the heating load is not evenly distributed between KA01 and KA02. Their combined heat production is known, it must be divided between them somehow. This is outside of eSmart's scope, and it has not been established by the time of writing. Once the respective heat production values are known, the electrical consumption from each can be calculated using their respective COP values, which are going to be provided by an external system based on the pressure and temperature conditions in the system.

$$El_t^{KA0x} = \frac{Heating_t^{KA0x}}{COP_t^{Heating,KA0x}} \tag{4.12}$$

Calculations for future periods

The following variables need to be predicted:

- $CoolingLoad_t^{IT}$
- $CoolingLoad_t^{Office}$
- $HeatLoad_t$
- $MaxHeat_t^{KA01}$
- $MaxHeat_t^{KA02}$

- $COP_t^{Heating,KA01}$
- $COP_t^{Heating,KA02}$
- Residual load
- PV production

All of these must be predicted at least for the whole planning horizon of the optimisation model.

The water supply and return temperatures must be calculated based on the predicted outside temperature and the formula provided by the client: T_t^{hot}, T_t^{cold}. Note that in the winter, T_t^{hot} is the supply and T_t^{cold} is the return, while is summer it is the other way around.

4.6 Optimisation algorithm for thermal storage

The optimisation's output is a charging/discharging schedule for the thermal energy storage (TES) tank, that is, the $300m^3$ sprinkler water tank. All other variables and expressions are either discarded or recorded for analysis but are not executed in any way.

4.6.1 Summer-winter changeover

The model takes the operating mode as an input for each period – it is that parameter that decides if the cold water energy balance or the hot water energy balance should be enforced.

Since many of the parameters and variables are only interesting in one of these operating modes, the set of periods P is divided into two subsets, the set of summer periods P_{summer} and the set of winter periods P_{winter}.

The model handles smooth changeovers – it will automatically make sure to use up all the heat in the tank before changing from winter mode to summer mode for example. However, because the planning horizon is too short for the model to see the

value in using energy to cool the tank down for operating in the summer, a penalty term is introduced that will force the model to do so: Equation 4.13.

$$Penalty = \begin{cases} max(PEN \times TES_{final} - T_{max}^{cooling}, 0), & Mode_{final} = Summer \\ max(PEN \times T_{min}^{heating} - TES_{final}, 0), & Mode_{final} = Winter \end{cases} \tag{4.13}$$

The max() function in this case is implemented as a piecewise linear function. It therefore introduces binary variables or SOS2 constraints depending on the chosen method and solver. With the glpk solver used for this thesis, only binary variables are possible.

4.6.2 Chilled water production

In order to calculate how the thermal storage will handle the cooling energy in summer operation, the optimization model must calculate the use of the kuldeaggregats and kjølemaskin for future hours. These decisions are not implemented, but which the optimization model internally calculates to find optimal use of the thermal storage.

The chilled water production is described by one variable, $Cooling_t$, which signifies the cooling energy from all three machines combined. As such, it is constrained to the sum of the cooling power of the machines:

$$0 < Cooling_t < 190kW \tag{4.14}$$

This cooling power must be related to the electric power the machines consume. The relationship depends on the ambient temperature and the cooling power.

Let's take the following definition:

$$ElCons(Cooling, AmbientTemp) \tag{4.15}$$

$$Cooling, AmbientTemp, ElCons \in R \tag{4.16}$$

Ambient temp. (degC)	KA - MT	KA - PL	KM30	KM60	KM90	KM120
10	4,64	8	2	6,59	6,92	7,09
12	4,57	7,87	2	6,46	6,82	6,99
14	4,49	7,74	2	6,91	7,31	7,49
16	4,07	7,61	2	6,15	6,52	6,69
18	3,67	6,68	2	5,54	5,83	5,6
20	3,31	5,87	2	4,97	5,23	5,01
22	2,97	5,14	2	4,48	4,68	4,5
24	2	4,72	2	4,03	4,18	4,05
26	2	4,47	2	3,89	3,96	3,84
28	2	4,13	2	3,52	3,58	3,44
30	2	3,63	2,7	3,13	2,98	3,06
32	2	3,21	2,6	2,81	2,68	2,75
34	2	2,87	2,34	2,54	2,44	2,49

Table 4.2: COP values for both machines, with 2 where it was not given

Then the COP values shown in the tables are:

$$COP = \frac{Cooling}{ElCons(Cooling, AmbientTemp)} \tag{4.17}$$

We must use these COP values to reconstruct the underlying two-variable function. Assumptions:

- The supply temperature of the cooling units (Turtemperatur) is 9 degC
- The two kuldeaggregats are always sharing the cooling load 50/50, and are in the same operating mode (MT or parallel) - therefore the two KAs can be modelled as one bigger unit
- The cooling power level where the KA goes from MT to parallel is known – in this document we will assume 30 kW

The Table 4.2 can be understood as samples for this function, and the function can be approximated by using linear interpolation in between the samples.

Since the ambient temperature forecast is known, that dependence can be resolved in a preprocessing step. This way, COP values can be constructed for each period for each operating mode of kuldeaggregats and predefined cooling power levels

for the kjolemaskin: COP_t^{KA-MT}, COP_t^{KA-PL}, COP_t^{KM30}, COP_t^{KM60}, COP_t^{KM90}, COP_t^{KM120} Finally, if we interpolate between these values, we get COP functions that cover the whole cooling power range.

A slice of $ElCons(Cooling, AmbientTemp)$ will be found as:

$$ElCons_t(Cooling) = \begin{cases} \frac{Cooling}{COP_t^{KA-MT}}, & Cooling < 40kW \\ \frac{Cooling}{COP_t KA-PL}, & Cooling \leq 70kW \\ \frac{70kW}{COP_t^{KA-PL}} + ElCons_t^{KM}(Cooling - 70kW), & Cooling > 70kW \end{cases} \tag{4.18}$$

The cooling load is shared between the KA and the KM by the local controller. It will always prioritise using the KA and it will use the KM to cover the slack above 70 kW cooling load. The function $COP_t^{KM}(Cooling)$ will not be directly calculated, because it is easier to just create $ElCons_t^{KM}(Cooling)$

COP's ambient temperature dependence

The tables provide COP values only for a handful of cooling power values. The goal of this step is to select the COP for each cooling power level. Since the ambient temperatures is given in 2-degC steps in the tables, the best approximate COP value will come from linear interpolation. The linear interpolation is justified because the COP depends on the temperature difference in a near-linear fashion, and the output temperature is fixed around 9 degC.

For example, when the outside temperature is 12.5 degC, and we want to know the COP for the KM at 60 kW, we look at the two neighboring data points: COP at 12 degC is 6.46, at 14 degC it is 6.91. Therefore at 12.5 degC:

$$\frac{14 - 12.5}{2} \times 6.46 + \frac{12.5 - 12}{2} \times 6.91 = 6.57$$

The output of this preprocessing step is a list for each power level that contains the COPs for each period.

C^i: Cooling power level (kW)	E_t^i: Electric consumption level (kW)
0	0
30	$30/COP_t^{KA,MT}$
30	$30/COP_t^{KA,PL}$
70	$70/COP_t^{KA,PL}$
100	$70/COP_t^{KA,PL} + 30/COP_t^{KM,30}$
130	$70/COP_t^{KA,PL} + 60/COP_t^{KM,60}$
160	$70/COP_t^{KA,PL} + 90/COP_t^{KM,90}$
190	$70/COP_t^{KA,PL} + 120/COP_t^{KM,120}$

Table 4.3: Relating thermal cooling power to electrical load.

COP's cooling power dependence

The cooling power depends on the decisions of the optimiser, therefore that function must be part of the model itself. It is a piecewise linear function that can be modelled using binary variables or SOS2 constraints. The details of this are not discussed.

From the COP values calculated in the preprocessing, two series are made for each period: see Table 4.3.

Notice that this lookup table is now explicitly f(Cooling)! The piecewise linear function will interpolate between the power levels. Since there are eight points, there will be seven segments. See Figure 4-15 for example.

$ElCons_t$ must be included in the electricity balance.

4.6.3 Hot water production

In winter operation, the heat production of the kuldeaggregats and the electric boiler are decision variables instead of the combined cooling.

Maximum power must be constrained:

$$0 \leq Heating_t^{KA0x} \leq MaxHeat_t^{KA0x} \tag{4.19}$$

$$0 \leq Heating_t^{Boiler} \leq P_{max}^{boiler} \tag{4.20}$$

Note that unlike cooling capacity, the maximum heating power of the kuldeaggregats is not a constant parameter. It makes sense, since only the rejected heat can be

Outside temp (degC)	COP_MT	COP_PL	COP_KM_30	COP_KM_60	COP_KM_90	COP_KM_120		
14	4,49	7,74	1	6,91	7,31	7,49		
Cooling (kW)	0	30	30	70	100	130	160	190
ElCons (kW)	0	6,68	3,88	9,04	39,04	17,73	21,36	25,07

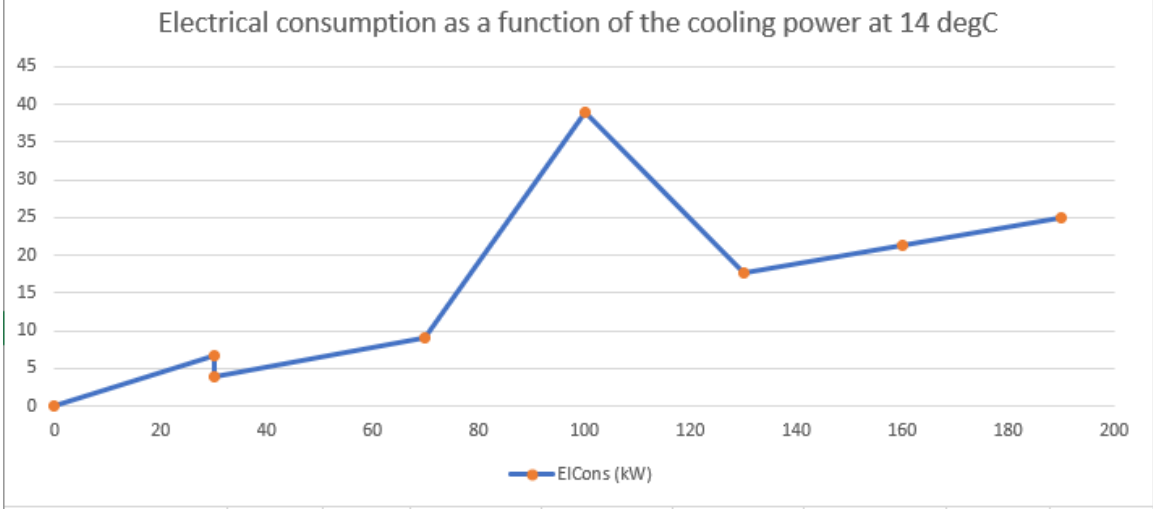

Figure 4-15: Example of piecewise function

captured, and the amount of reject heat depends on how high the cooling load is: the more cooling the machines needs to do, the more heat can be captured from them. Therefore the input to the optimisation model in this case in a predicted MaxHeat timeseries for each machine.

For the electrical boiler the relationship between the electricity consumption and heat production is a simple efficiency term: ε^{EB}.

Electricity consumption must be included in the electricity balance.

$$ElCons_t = \frac{Heating_t^{Boiler}}{\varepsilon^{EB}} + \frac{Heating_t^{KA01}}{COP_t^{Heating,\ KA01}} + \frac{Heating_t^{KA02}}{COP_t^{Heating,\ KA02}} \tag{4.21}$$

4.6.4 Heat storage

The heat storage is characterised in the model by its temperature, represented in the decision variable TES_t.

The heat loss in each period is expressed as:

$$HeatLoss_t = (TES_t - \ T^{Ground}) \times K^{HeatLoss} \tag{4.22}$$

The heat flowing into the tank in each period is expressed as:

$$EnergyIn_t = (TES_t - TES_{t-1}) \times C_{tank} + HeatLoss_t \tag{4.23}$$

The heat storage's energy balance in summer operation:

$$EnergyIn_t = CoolingLoad_t^{IT} + CoolingLoad_t^{Office} - Cooling_t \tag{4.24}$$

Winter operation:

$$EnergyIn_t = Heating_t^{KA01} + Heating_t^{KA02} + Heating_t^{Boiler} - HeatLoad_t \tag{4.25}$$

These energy balances are sufficient to ensure that the thermal loads are met.

The maximum heat exchange rate linearly depends on the temperature difference between the tank and the supply/return temperatures in the system as described above. This can be modelled as an upper and lower bound on the $EnergyIn_t$ variable, that depends on the tank temperature variable. See Figure 4-16 for illustration. Note that the supply and return temperatures are just parameters, not variables. Note also that there is a breakpoint in these bounds at 0 – the exchanges heat is always allowed to be zero. Because of this, they must be modelled with piecewise linear functions as well.

This means a pair of piecewise linear bounds for each period of the planning horizon.

$$MaxHeatOut \leq EnergyIn_t \leq MaxHeatIn_t \tag{4.26}$$

$$MaxHeatIn_t = max(K_{HX} \times (T_t^{Hot} - \Delta T_{HX} - TES_t), 0) \tag{4.27}$$

$$MaxHeatOut_t = max(-K_{HX} \times (TES_t - T_t^{Cold} - \Delta T_{HX}), 0) \tag{4.28}$$

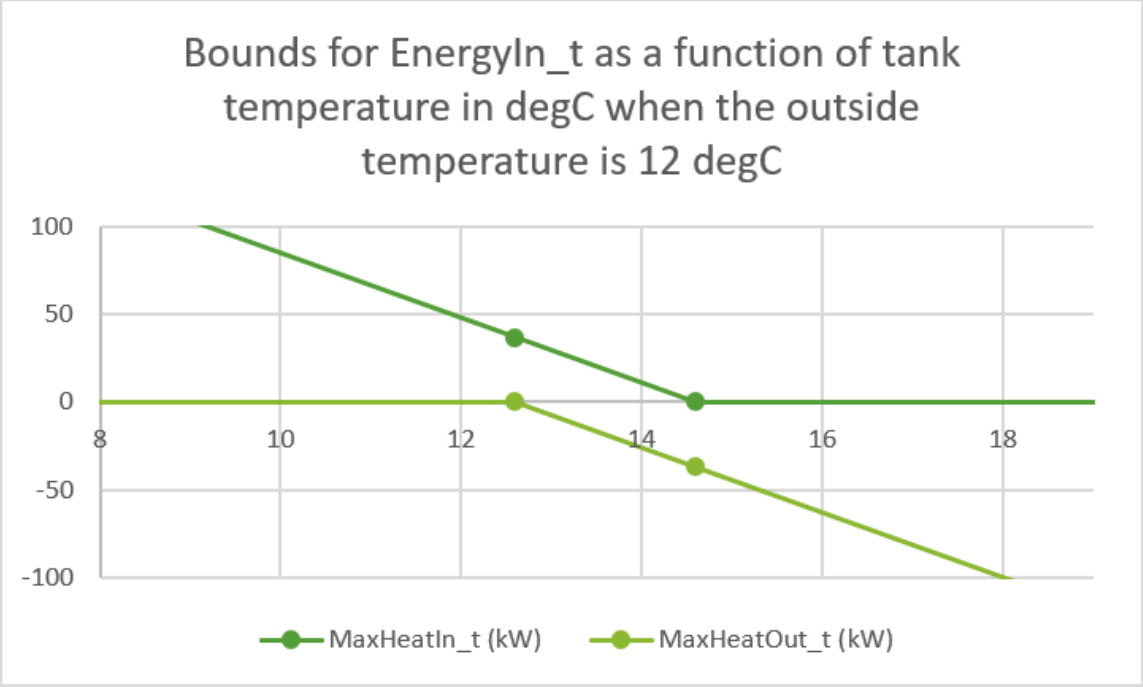

Figure 4-16: Piecewise linear bounds for $EnergyIn_t$

Note that the model has many more variables, many of them binary, due to the piecewise linear constraints. They have a large impact on solver performance.

4.7 Results

The system can be analysed in the two separate operating modes of summer and winter, and the changeover between them is also interesting to observe.

Figure 4-17 shows a base case scenario in the summer mode, using actual data form the site with the exception of the electricity price, which is fictitious. This base case was created by restricting the energy exchange with the energy storage tank to zero. The resulting cost was 101.3 NOK.

4.7.1 Summer operation

The cooling load and the dispatched cooling is equal in all periods, as expected: since the energy exchange with the storage tank is zero, there is no slack in the system.

Figure 4-17: Base case in summer operation - the energy exchange variable is fixed to zero

Notice that the electricity consumption is not only a scaled down version of the cooling production, because the COP is not constant: it depends on both the power and the outside temperature. The bottom left figure shows that the system often runs the kuldeaggregats in MT mode, which has poor efficiency. The right hand side figures show that the tank temperature is near constant, as the only energy exchange is towards the ground.

Figure 4-18 shows the exact same scenario with the storage tank unblocked. This time the kuldaggregats never run in the inefficient MT mode, because the tank can provide the necessary cooling in the periods with low cooling load. This increases overall energy efficiency.

The tank temperature is not constant anymore, indeed the model extracts as much cooling energy from it as possible, to further minimize the energy cost. The amount of cooling available from the tank is restricted by the heat exchanger as explained in a previous section. In the hours 11 to 21, the tank temperature is higher than the supply temperature of the cooling circuit - it means that that is the only time when it is possible to cool the tank further. In the hours 13 to 18, the tank temperature is very near the return temperature of the cooling circuit, therefore it is impossible to heat the tank in those hours.

The model even schedules some cooling of the tank in the hours 16-17-18 to prepare for supplying cooling in the following hours to avoid using the kuldeaggregats in MT mode

The heat loss is proportional to the temperature. Note that in this case, the heat loss is a good thing, because it essentially charges the cooling storage system.

The total cost of operating in this case was only 74.7 NOK, a significant reduction from the 101.3 NOK before. A major component of this is that the tank itself supplied 340 kWh worth of net cooling as it increased in temperature from 10 degC to 10.98 degC. This is worth about 10 NOK assuming a COP of 5 and an energy price of 0.15 NOK/kWh. The rest is due to running the system in a more efficient mode and avoiding the high price periods.

Figure 4-18: Example of summer operation

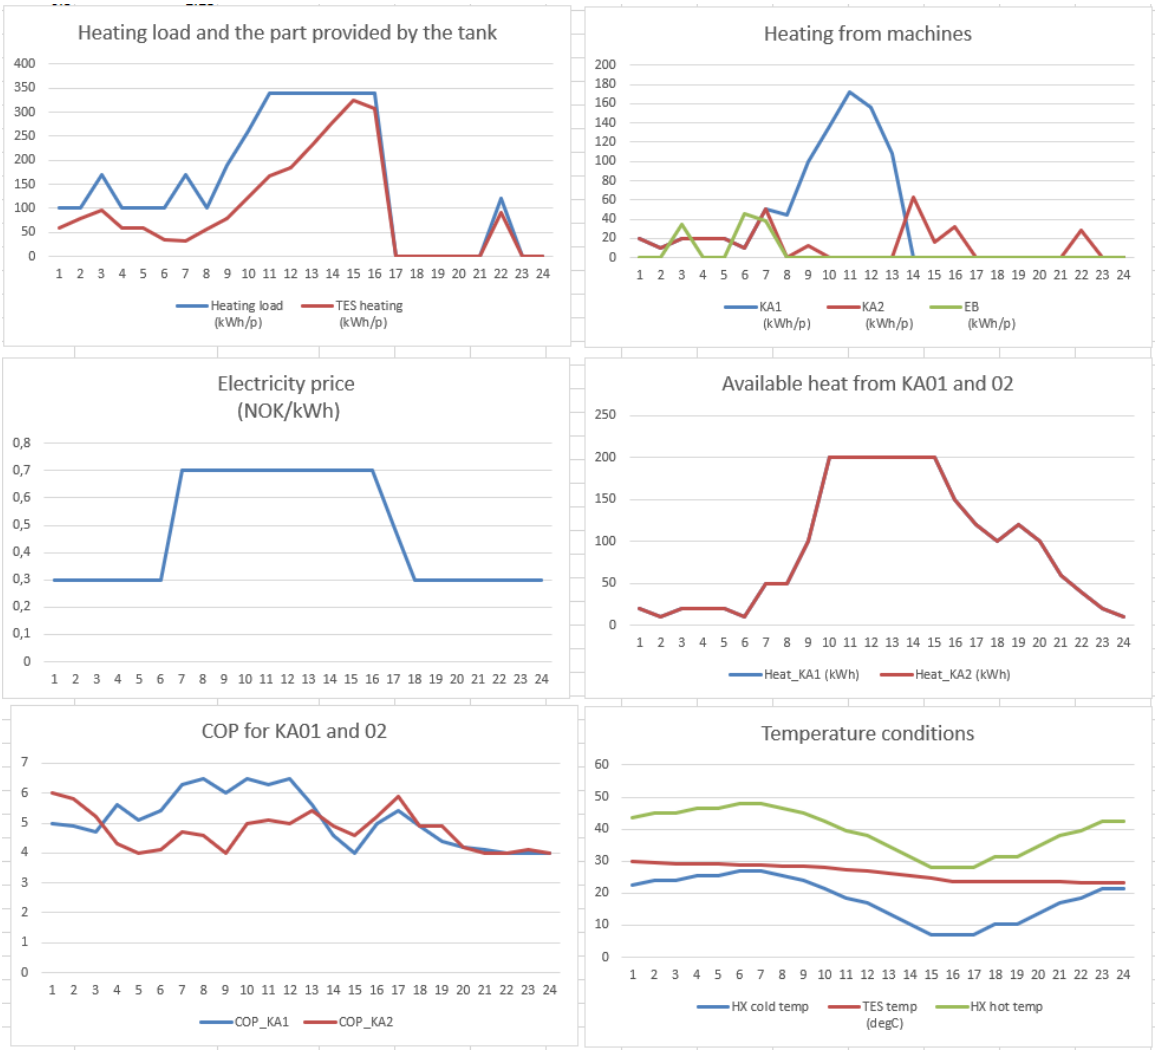

Figure 4-19: Example of winter operation

4.7.2 Winter operation

Figure 4-19 shows an example from winter operation. In principle the heat from the kuldeaggregats is available much cheaper than the heat from the boiler, but the kuldeaggregats can't always provide heat, and the amount available is also variable. In this case the available heat is assumed the same from both, but at slightly different COPs.

The tank temperature is scheduled to decrease from 30 to around 24 degC throughout the planning horizon, using as much heat from it as possible. Note that the amount of heat available from the tank is again restricted by the temperature condi-

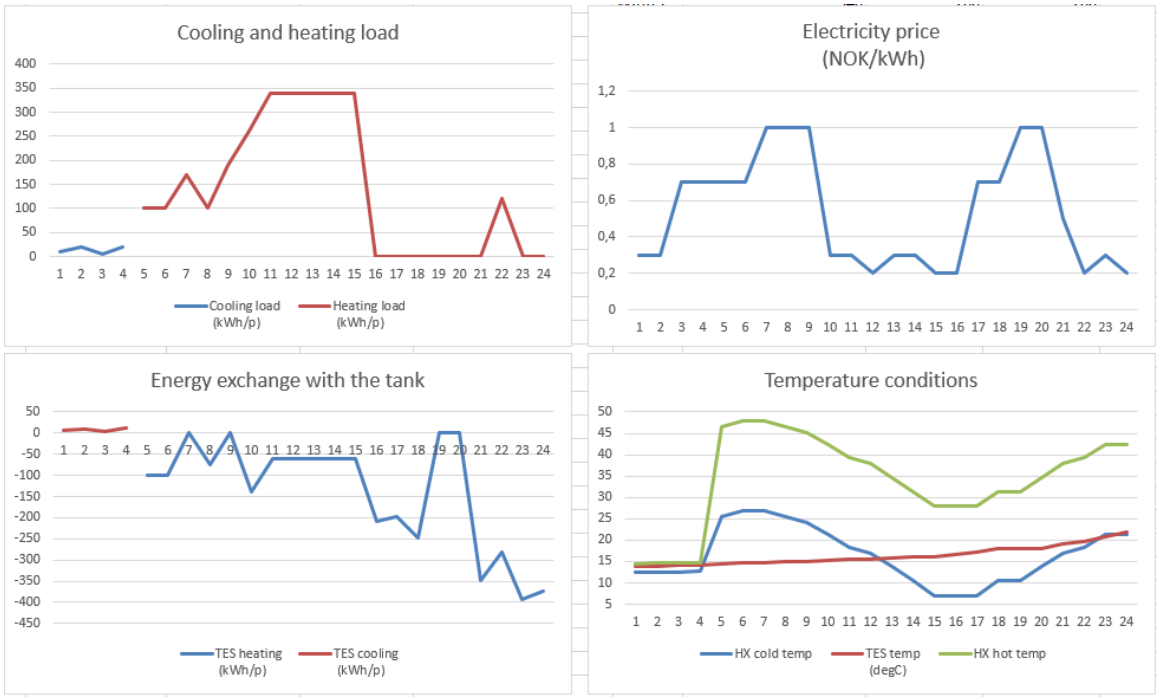

Figure 4-20: Example of changeover from summer to winter operation

tions on the heat exchanger, which in turn depends on the weather. There is a lot more control over the TES in winter operation than in summer, because the temperature difference between the supply and return temperature is much higher: 25 degC vs. 6 degC in the summer.

In the first periods, the heat load is significant, and there is not sufficient heat available from the kuldeaggregats, nor from the storage tank, due to the heating system's return temperature being too close to the tank temperature. Therefore using the boiler is unavoidable. In the following periods where a lot of heat is available from the kuldeaggregats, the optimization schedules the one with higher COP. As the outside temperature rises during the day, the return temperature from the heating system drops, therefore more heat becomes available from the storage tank! The system takes full advantage of that as well.

4.7.3 Changeover

Figure 4-20 shows a changeover from summer operation to winter operation. At changeovers, a large amount of energy is needed to raise or lower the tank temperature so that it can be used in the other operating mode. This does not pay back within the short 24 hours planning horizon, therefore the model is coerced into it by a penalty. Therefore the objective here is to manage the changeover in the lowest cost way possible while avoiding the penalty.

In the case on Figure 4-20 the model achieves the minimum-cost changeover by heating the tank as much as possible during the summer operation, then using the low-load and low-price periods to heat the tank further. The minimum temperature to avoid the penalty is 22 degC, and the model reaches exactly that by the end of the planning horizon.

The available heat during the winter periods is 400 kWh/h from the two kuldeaggregats together. Note that the model managed to never exceed that, and therefore the boiler needs not be used at all. This explains the low heating in the periods 11 to 15 - the heat load is high in those periods, therefore there is little extra heating power left for the tank in those periods. Note that the graph shows 'heating provided by the heat storage tank', therefore when it is negative, it means that the tank itself is being heated.

The temperature conditions figure shows how much the controllability of the tank increases in winter operation compared to summer operation.

4.7.4 Factory acceptance test

The REMA Vagle project finished on June 17th with a factory acceptance test. The optimization behaviour had been scrutinised and accepted as part of the test.

It had been proposed to include two additional operating modes, 'preparing for summer' and 'preparing for winter', in which the thermal storage can only export or import heat, respectively. This is because the 24 hours planning horizon is not sufficiently long to perform an optimal changeover.

The optimization has not yet been tested on the actual system, because the local control system is still lacking some modules.

Chapter 5

State space modelling of residential electrical heating

5.1 Introduction

The goal with the space heating model is providing demand side response utilising the thermal inertia of households through controlling only the temperature setpoint of the thermostat. The opening assumption is that only the following information is available from the building:

- Maximum power of the heating system (kW)
- Electricity consumption of heating system measured separately (kWh/period)
- Inside temperature at thermostat (degC)
- Outside temperature in the weather zone of the building (degC)
- Temperature band available for control (min degC, max degC)

An example of the temperature information is shown on Figure 5-1.

This model was developed in response to a previous model developed by NTNU. [6] I identified that it:

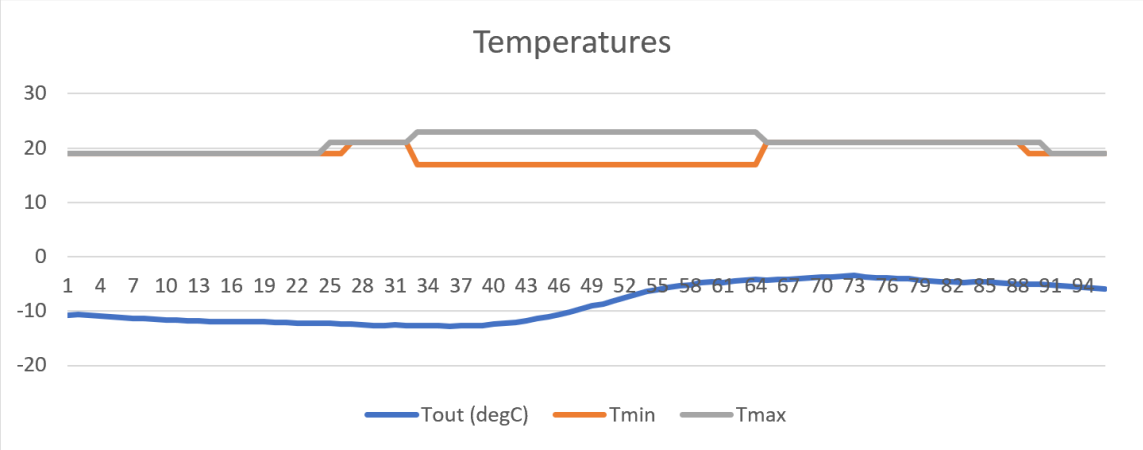

Figure 5-1: Example showing some possible input parameters

- Does not model the thermal system behaviour with sufficient accuracy for demand side response
- Cannot model the relationship between thermally coupled rooms - not generalised for multizone buildings
- Only selects between preset thermostat setpoints, thereby limiting the choice and flexibility
- Assumes machine learning predictions that we do not have and do not know how to make
- Has not been tested on real data
- Runs relatively slowly due to binary variables

5.2 eSmart space heating modelling approach

5.2.1 Thermal models

The use of an accurate model minimises the error both at the baseline prediction and at the optimisation levels. An inaccurate model will result in sub-optimal performance: higher energy cost, higher peaks, or non-conforming load curve.

The model must be linear so that the linear programming optimiser toolkits can handle it. A balance must be found between simplicity and accuracy. The accuracy is limited by the amount of information available of the buildings. Thermal behaviour is fundamentally linear, and it is common practice to create thermal models for buildings using a concentrated parameter model, akin to the typical representation of electrical circuits. [19] [16]

My approach is to try to find the best concentrated parameter model structure for each building and estimate those parameters. A limited set of model structures (4-5) will be used. The approach for finding the best concentrated parameter model is called system identification and is common practice in thermal, acoustic, mechanical, and electrical modelling since the 1980's. [8]

If the structure is well selected, the identified model describes the building's unchanging fundamental characteristics. The expectation is that with our simple models and limited data this condition will not be met, and we must regularly re-identify the model, for example once a month.

A basic assumption in the modelling approach is that the temperature setpoint is always reached by the end of the period - masking the nonlinear behaviour of the thermostat! The model itself ensures that the temperature setpoint will never be too far from the last measured value, because it is one of the state variables, and therefore it cannot have jumps that are bigger than what the real system can produce. This is conditional on the model being sufficiently accurate.

5.2.2 Why not machine learning?

Machine learning is ideal for providing some predictive power over systems that we do not understand. However, buildings and thermodynamics are well understood, and that domain expertise can be used to provide faster, simpler, more accurate, and more dependable models. Machine learning models are called 'black box' models and system identification is a 'grey box' approach. In a later stage, machine learning could provide insight into occupant behaviour for example.

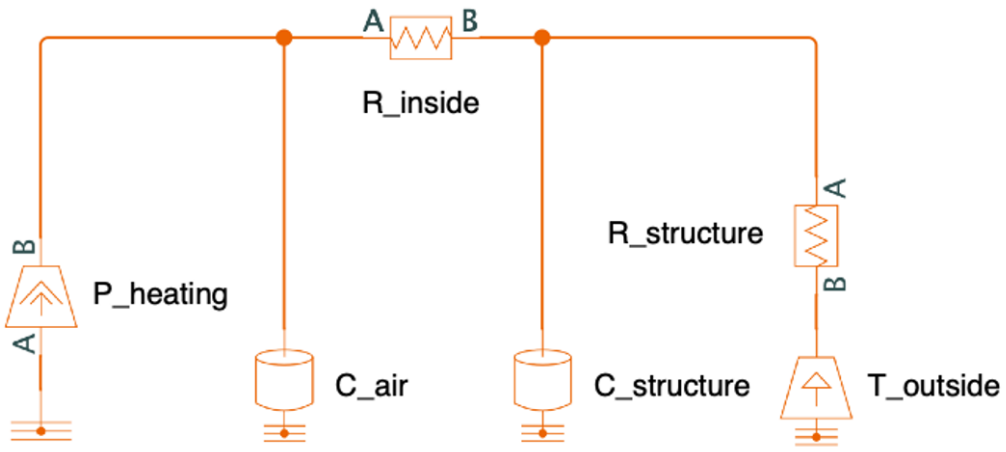

Figure 5-2: Example of second order model for single zone building

5.2.3 Concentrated parameter thermal models

Concentrated parameter thermal models are akin to concentrated parameter electrical models, and the Kirchhoff (and Ohm) laws apply to them in unchanged form. The analogy is as follows:

- Temperature [degC] – Voltage [V]
- Heat flow [kW] – Current [A]
- Thermal coupling [degC/ kW] – Resistance [V/A or Ohm]
- Thermal mass [kWh/degC] – Capacitance [Farad or A*s/V]
- Heat source [kW] – Current source [A]
- Temperature source [degC] – Voltage source [V]

As described above, the coherent unit system used is: kW – kWh – h – degC. The reference can be freely chosen, 0 degC is a convenient choice.

See Figure 5-2 for an example of a concentrated parameter thermal model represented similarly to an electrical network.

5.2.4 State space system description

The temperatures of the thermal masses in the model are akin to the voltages of the capacitances in an electrical network – they cannot change abruptly, as their rate of change is proportional to the heat rate across them. (Just like $I_c = C \times \frac{d}{dt}U_c$!) This makes the temperatures state variables, and therefore the state of the system can be fully described with just the temperatures across the thermal masses.

State space equations are traditionally written using the following notations:

$$\frac{d}{dt}x = A \times x + B \times u \tag{5.1}$$

$$y = C \times x + D \times u \tag{5.2}$$

If 'n' is the number of states (temperatures) to describe the thermal system and it has 'm' zones:

- State vector will be $n \times 1$ column vector (usually denoted with x)
- Input vector will be $(m + 1) \times 1$ column vector (usually denoted with u)
- A matrix will be an $n \times n$ matrix
- B matrix will be an $n \times (m + 1)$ matrix
- C matrix will be an $m \times n$ matrix that exposes the measured temperatures and hides the rest
- D matrix will be an $m \times (m + 1)$ null matrix

$$A_s = \begin{bmatrix} A_{s,11} & \cdots & A_{s,1n} \\ \vdots & \ddots & \vdots \\ A_{s,n1} & \cdots & A_{s,nn} \end{bmatrix} \tag{5.3}$$

$$B_s = \begin{bmatrix} B_{s,11} & \cdots & B_{s,1(m+1)} \\ \vdots & \ddots & \vdots \\ B_{s,n1} & \cdots & B_{s,n(m+1)} \end{bmatrix} \tag{5.4}$$

$$C_s = [I_{m\times m}\ 0_{m\times(n-m)}] \tag{5.5}$$

$$D_s = 0_{m\times(m+1)} \tag{5.6}$$

$$x_{s,p} = \begin{bmatrix} \tau_{s,1,p}^{inside} \\ \tau_{s,2,p}^{inside} \\ \vdots \\ \tau_{s,m,p}^{inside} \\ \tau_{s,1,p}^{struc} \\ \vdots \\ \tau_{s,n-m,p}^{struc} \end{bmatrix} \tag{5.7}$$

$$u_{s,p} = \begin{bmatrix} \pi_{s,1,p}^{heating} \\ \pi_{s,2,p}^{heating} \\ \vdots \\ \pi_{s,m,p}^{heating} \\ T_{s,p}^{outside} \end{bmatrix} \tag{5.8}$$

Having this representation is beneficial because it is concise and generalised for any number of inputs and outputs – in this case, for any number of temperature zones.

The inputs will always be either heat flows of heat sources or temperature of temperature sources. The outputs will always be some of the states, therefore temperatures of temperature masses.

Using Figure 5-2 as an example, the heat flow from the heating system and the

outside temperature will be the two inputs, and the inside temperature will be treated as the output.

$$x = \begin{bmatrix} \tau_{air} \\ \tau_{struc} \end{bmatrix} \tag{5.9}$$

$$u = \begin{bmatrix} \pi_{heating} \\ T_{outside} \end{bmatrix} \tag{5.10}$$

$$C = [1\ 0] \tag{5.11}$$

$$D = [0\ 0] \tag{5.12}$$

$$y = [\tau_{air}] \tag{5.13}$$

In this case, A and B are 2×2, C and D are 1×2 matrices.

Notice that the state matrix A is independent of the choice of inputs and outputs.

5.2.5 From concentrated parameter thermal model to state space representation

This section shows how to obtain the matrices A and B from the concentrated parameter representation. In the example shown in Figure 5-2, the following two Kirchhoff node equations can be written for the two unknown temperatures:

$$P_{heating} + \frac{T_{air} - T_{struc}}{R_{inside}} + C_{air} \times \frac{d}{dt} T_{air} = 0 \tag{5.14}$$

$$\frac{T_{struc} - T_{outside}}{R_{struc}} + \frac{T_{struc} - T_{air}}{R_{inside}} + C_{struc} \times \frac{d}{dt} T_{struc} = 0 \tag{5.15}$$

They can be rearranged to fit the form $\frac{d}{dt}x = A \times x + B \times u$ so that the coefficients that make up A and B can be identified.

$$\frac{d}{dt}T_{air} = -\frac{T_{air}}{R_{inside} \times C_{air}} + \frac{T_{struc}}{R_{inside} \times C_{air}} - \frac{P_{heating}}{C_{air}} \tag{5.16}$$

$$\frac{d}{dt}T_{struc} = \frac{T_{air}}{R_{inside} \times C_{struc}} - \left(\frac{R_{struc} + R_{inside}}{R_{inside} \times R_{struc}}\right) \times \frac{T_{struc}}{C_{struc}} + \frac{T_{outside}}{R_{struc} \times C_{struc}} \tag{5.17}$$

$$A_{11} = -\frac{1}{R_{inside} \times C_{air}} \tag{5.18}$$

$$A_{12} = \frac{1}{R_{inside} \times C_{air}} \tag{5.19}$$

$$A_{21} = \frac{1}{R_{inside} \times C_{struc}} \tag{5.20}$$

$$A_{22} = -\left(\frac{R_{struc} + R_{inside}}{R_{inside} \times R_{struc}}\right) \times \frac{1}{C_{struc}} \tag{5.21}$$

$$B_{11} = -\frac{1}{C_{air}} \tag{5.22}$$

$$B_{12} = 0 \tag{5.23}$$

$$B_{21} = 0 \tag{5.24}$$

$$B_{22} = \frac{1}{R_{struc} \times C_{struc}} \tag{5.25}$$

$$A = \begin{bmatrix} -\frac{1}{R_{inside} \times C_{air}} & \frac{1}{R_{inside} \times C_{air}} \\ \frac{1}{R_{inside} \times C_{struc}} & -\left(\frac{R_{struc} + R_{inside}}{R_{inside} \times R_{struc}}\right) \times \frac{1}{C_{struc}} \end{bmatrix} \tag{5.26}$$

$$B = \begin{bmatrix} -\frac{1}{C_{air}} & 0 \\ 0 & \frac{1}{R_{struc} \times C_{struc}} \end{bmatrix} \tag{5.27}$$

The disadvantage of using the matrix notation is hiding the thermal properties – instead of storing C_{struc} for example, the names will be less descriptive, like A_{11}. The descriptive properties can be retrieved for analysis.

5.2.6 Thermal model identification

The model is identified using MATLAB's System Identification Toolbox based on a few weeks' historical data. The identified model needs to be stored in the INVADE platform.

The input to the system identification process are the parameterised ABCD matri-

ces. These matrices describe the relationship between the parameters and the system, and that guides the identification process to come up with robust models. System identification with well selected structures avoids overfitting the model and allows making sense of the model.

Note: the ABCD matrices describe the whole building and heating system. They must be stored as a parameter of the site/asset group, not the individual heater assets.

5.2.7 Identifying continuous time system based on sampled inputs and outputs

MATLAB's System Identification Toolbox assumes that the samples were taken synchronously. The model described in the earlier chapters would require synchronously sampled temperatures and heat flows. [9]

In the platform:

- Measured temperatures (both in- and outside) are accurate for the time they were recorded
- Temperature setpoint for period 0.00 to 0.15 is recorded to 0.00
- Electricity consumption from 0.00 to 0.15 is recorded to 0.00

The problem is that heat flows are not synchronously sampled. The electricity consumption is known as a 15-minute total, from which the average heat flow can be calculated. The best sample time to attach for the average of the period 0.00 to 0.15 is 0.07:30.

This however would violate STI's restriction that "Output data must be in the same domain as input data." For now, the average heat flow is being considered at 0.00, because this ensures causality. This introduces an error. A more elegant solution shall be found for this.

5.2.8 System identification using SciPy

As MatLab was not part of the eSmart stack, using the System Identification Toolbox would have meant significant changes in the platform, as well as extra licensing costs. Instead, python-based options were explored. SciPy is a free and open-source Python library used for scientific computing and technical computing, that contains modules for optimization and other common engineering tasks. It's minimize() function provides a common interface to constrained minimization algorithms for multivariate scalar functions in scipy.optimize.

The main difference between this and the optimization discussed elsewhere in this thesis is that all other cases are (mixed integer) linear programming cases. The benefit of linear programming is that it scales to huge problems, but at the cost of being restricted to linear formulation. That restriction can not be satisfied for parameter identification, because the parameters are often multiplied together or raised to a power, therefore they can not be decision variables in a LP/MIP. However the performance limitation is not an issue in this case, since each building only has to be identified a few times, and it does not have to happen on a tight schedule.

The approach was the following:

1. Make a simulation of each model that takes the parameters as inputs
2. Take the square error of the simulation results compared to the meter values
3. Use the square error as the scalar function to guide the minimize() function
4. Use the simulation parameters as the variables to the function

The MEEB [7] dataset was used, with the temperature measurements averaged to give a single temperature zone. Figure 5-3 shows the inside and outside temperature movements and the heating power. The output temperature and the heating power are considered inputs, and the inside temperature output. The best results were achieved with the data sliced into 4 days, handled as separate experiments. The first three days were used for the parameter identification, and the last for verification.

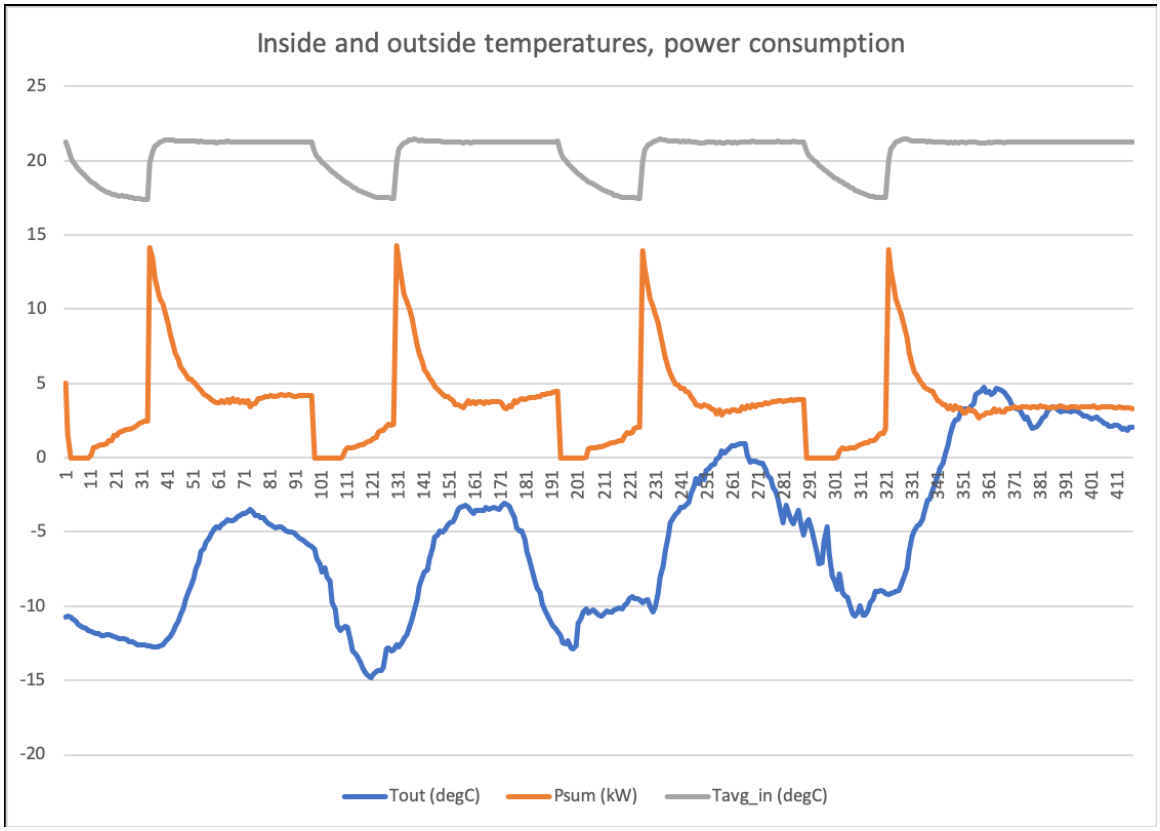

Figure 5-3: Input example recorded at the MEEB test site in Quebec [7]

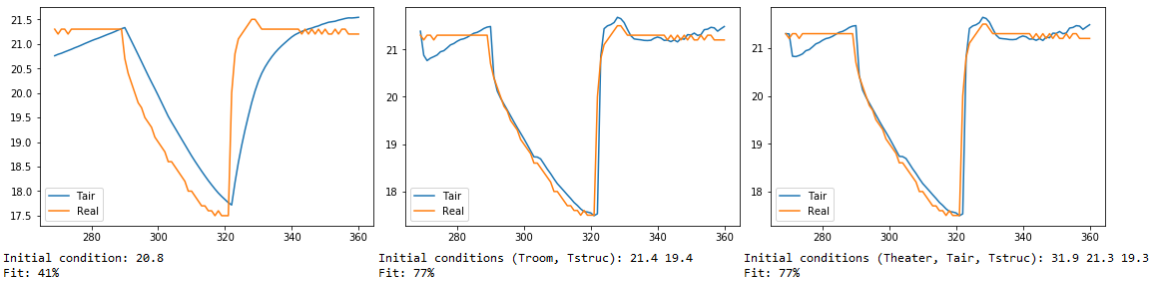

Figure 5-4: Verification of identified models: first, second, and third order

Separate initial conditions had to be identified for all days - the fact that the days were consecutive was not taken into consideration.

Three models were tried: a first, a second, and a third order model. The first order model only has the measured state. The second order model introduces a hidden state to separate the heat capacity of the structure from the heat capacity of the measured air. The third order model introduces a second hidden state to separate the heat capacity of the heating elements.

The models can be evaluated for fit by calculating the normalised root mean square error on the verification data. MatLab uses the same error metric, and it was used to verify the behaviour of the SciPy implementation. The achieved square error from the minimization and the visual comparison of the simulated and metered values provides additional insight. The results were:

- First order model: 41% fit, 3 seconds
- Second order model: 77% fit, 35 seconds
- Third order model: 77% fit, 3 minutes

While the second order model behaves markedly better than the first order one, there is no noticeable benefit in introducing a second hidden state. The sum square error decreased from 30.8 to 29.9, a difference too small to justify the added complexity and significantly increased computation time. The calculations were run on an Intel i7-6500U, 8GB RAM computer. Figure 5-4 shows the similarity in the waveforms produced with the second and third order model.

Note that the initialization is also done by solving a parameter fitting problem. This approach can be used also when the model is integrated into the eSmart platform.

5.2.9 Discretisation

The system described in the previous sections is a continuous time model of the physical system. However, the model needs to be used for simulation at discrete 15 minute time steps, and to obtain from it control signals that are going to be valid for a period of 15 minutes.

For this, the identified continuous time state space representation is discretised with the zero order hold technique with a 15-minute time step. The ABCD matrices resulting from the discretisation are stored in the platform and used in the model. The discrete ABCD representation shown in Equation 5.28 is convenient for modelling, because each line of the matrices can be directly written as a constraint.

$$x[k+1] = A_d \times x[k] + B_d \times u[k] \tag{5.28}$$

$$y[k] = C_d \times x[k] + D_d \times u[k] \tag{5.29}$$

5.2.10 Defining structures

The previous sections used a simple example case with four parameters to demonstrate the process of describing a heating system with ABCD matrices. The question of how to come up with the structure itself was not covered.

One approach is to manually come up with many reasonable structures, identify the parameters for all, and choose the best fit structure. This approach is scalable despite the manual step, because 5-6 model structures would cover all residential buildings with 1-2 zones.

For special buildings, an experienced engineer should be able to draw up a few potentially good structures based on information available on the building, and choose the best fit based on the data available.

If the customer is not worth the dedicated engineering hours, a general model might be identified. The automation comes at a cost of robustness. In this case, all elements of the A and B matrices are parameters.

The following restrictions can still be taken:

- All measured temperatures should be a state (number x)
- The order of the models identified should be in the range $[x + 1, \frac{3x}{2} + 1]$
- The output matrix C must contain ones for the observed states and zero for the rest

5.2.11 Baseline prediction

The baseline electrical energy consumption profile can be predicted based on the outside temperature and the desired indoors temperature using the building's thermal model. This is the prediction application of the model, and it has the same function as the machine learning prediction models for other asset types.

It calculates the energy demand to maintain the minimum temperature profile using a simplified version of the optimisation code. It uses parameter identification to initialise itself.

To provide a more accurate baseline, the occupants' habits, solar irradiation, and wind speed could be considered. This is not in the scope of the current proposal.

It is not strictly necessary to use this model for the predictions – an independent ML model could just as well perform the baseline predictions. It seems reasonable to use this model for predictions since it takes no significant extra effort.

5.3 Model implementation

5.3.1 Optimisation

The optimisation uses the thermal model to establish the connection between the controlled and dependent variables.

S	Sites (buildings)
Z	Temperature zones within buildings – there are m of them!
H	Structural (hidden) temperatures
P	Periods

Table 5.1: Sets used in the formulation

A daily minimum-maximum temperature profile (one profile for each zone) is enforced as a constraint so that the heating system always performs its primary goal of providing comfortable indoors temperature. It is assumed that the occupant chooses a comfortable minimum temperature profile. It could also be called setpoint temperature profile.

The flexibility comes from the model having the freedom to preheat when the energy price or flexibility requests warrant it. Notice that the model would never preheat without reason because it will mean larger overall energy consumption due to the escaping heat.

Of course, the model can only preheat when the constraint profile allows it. The occupant might choose identical min-max temperature for some periods when he wants a particular temperature no matter what.

The standalone optimisation model is implemented in python with Pyomo and uses the free glpk solver. It is to be integrated into the INVADE optimizer framework. The output of the model is the temperature setpoint timeseries for the thermostat: $\tau_{s,z,t}^{inside,limited}$

Note that the temperature setpoint timeseries needs to be shifted forward by one period for causality: In the optimisation model, heating causes temperature increase, in the real world, temperature setpoint causes heating.

5.3.2 Avoiding temperature ramping related infeasibilities

Every heating system's temperature ramping ability is limited in the positive direction due to the finite power of the heating element; and in the negative direction due to the finite heat conduction towards the environment. We assume no air conditioner. If the temperature bounds of the optimisation do not allow for adequate freedom for

$T^{min}_{s,z,p}$	Zone minimum temperature (degC)
$T^{max}_{s,z,p}$	Zone maximum temperature (degC)
$T^{inside}_{s,z,p}$	Zone inside temperature (past) (degC)
$W^{max}_{s,z}$	Zone maximum heating power (kW)
$T^{outside}_{s,p}$	Outside temperature (degC)
$K^{tempError}_{s}$	Temperature error penalty (NOK/degC/period)
A_s	State matrix
B_s	Input matrix
C_s	Output matrix
D_s	Feedthrough matrix

Table 5.2: Model parameters

the temperature between two fixed points, the model will reflect this reality by being unable to satisfy the constraints and turning infeasible.

In a practical system this must be avoided by relaxing the bounds if we suspect that the model will be infeasible.

Estimating the temperature ramping capability

The ramping ability depends on the outside temperature, the temperatures of the different thermal masses, and the parameters of the model. Consequently, it is also time dependent and varies even within the same ramping event. It is therefore impossible to handle maximum ramping rates for the whole planning horizon in a preprocessing step, and the model itself must be changed.

Adding error variables

While enforcing the lower temperature limit, deviations above $T^{max}_{s,z,p}$ could be allowed. That deviation is a non-negative decision variable that would be heavily penalised: $\varepsilon^{temp}_{s,z,p}$. Adding these error terms should ensure that the model is always feasible – since there are no more constraints!

The error variables also avoid infeasibility when a building with multiple temperature zones would want to maintain impossibly high temperature difference between two zones.

$\tau_{s,z,p}^{inside,limited}$	Zone temperature
$\varepsilon_{s,z,p}^{temp}$	Zone temperature error
$\tau_{s,h,p}^{struc}$	Hidden temperature
$\pi_{s,z,p}^{heating}$	Zone heating energy

Table 5.3: Decision variables in the space heating model

5.3.3 Decision variables and bounds

The decision variables are the state variables in the concentrated parameter state space model, and the heating powers. All of these are continuous variables, and there is one of them for each period.

For the structure showed on the drawing of the second order model for single zone building (Figure 5-2), they would be: $\tau_{s,z,p}^{inside}, \tau_{s,c,p}^{struc}, \pi_{s,z,p}^{heating}$

Temperatures that are not measured have no bounds. The heating powers are bound between zero and the maximum output of the heating appliance as described in the asset parameter.

To avoid possible ramping related infeasibilities, the controlled temperatures are not implemented directly as decision variables, but instead split into a bounded variable and an unbounded error variable. The temperature evolution equations are using $\tau_{s,z,t}^{inside}$, which is an expression defined as:

$$\tau_{s,z,p}^{inside} = \tau_{s,z,p}^{inside,limited} + \varepsilon_{s,z,p}^{temp} \tag{5.30}$$

Both $\tau_{s,z,p}^{inside,limited}$ and $\varepsilon_{s,z,p}^{temp}$ are one decision variable for each period. $\tau_{s,z,p}^{inside,limited}$ is bound between the minimum and maximum temperature profiles. $\varepsilon_{s,z,p}^{temp}$ is nonnegative.

With this setup, the actual inside temperature can always deviate from the specified range, but only in the positive direction. This way infeasibilities due to the structure not cooling off fast enough are avoided. (Negative ramping events.) Positive ramping infeasibilities are avoided by preheating as necessary, even if it means that the temperature needs to go above the maximum allowed temperature in some previous periods.

5.3.4 Constraints

The only complex constraints are the state equations. These equations describe the temperature evolutions.

An example for a second order structure:

$$\tau^{inside}_{s,z=1,p+1} = A_{s,11} \times \ \tau^{inside}_{s,z=1,p} + A_{s,12} \times \ \tau^{struc}_{s,h=1,p} + B_{s,11} \times T^{outside}_{s,p} + B_{s,12} \times \pi^{heating}_{s,z=1,p} \tag{5.31}$$

$$\tau^{struc}_{s,h=1,p+1} = A_{s,21} \times \tau^{inside}_{s,z=1,p} + A_{s,22} \times \tau^{struc}_{s,h=1,p} + B_{s,21} \times T^{outside}_{s,p} + B_{s,22} \times \pi^{heating}_{s,z=1,p} \tag{5.32}$$

The number of state equations is equal to the number of states. In residential single-zone buildings it will be typically two, but in general it can be any number. Matrix notation can describe any number or zones and hidden states.

Note that the model knows how many zones and hidden states there are by looking at the C matrix's dimensions.

All state equation constraints are then expressed as:

$$x_{s,p+1} = A_s \times x_{s,p} + B_s \times u_{s,p} : \forall s \in S, p \in P \setminus (last) \tag{5.33}$$

Bounds:

$$T^{min}_{s,z,p} < \tau^{inside,limited}_{s,z,p} < T^{max}_{s,z,p} : \forall s \in S, z \in Z, p \in P_f \tag{5.34}$$

$$0 \leq \varepsilon^{temp}_{s,z,p} : \forall s \in S, z \in Z, p \in P \tag{5.35}$$

$$0 \leq \ \pi^{heating}_{s,z,p} < W^{max}_{s,z} : \forall s \in S, z \in Z, p \in P_f \tag{5.36}$$

$$\tag{5.37}$$

5.3.5 Objective function

The cost function is the cost of energy plus the penalty terms for ramping-related errors. The penalty is a parameter, but it should not be selectable - some value large enough should be used, for example 100 NOK/degC/period for a household is

Name	Source	Frequency	Unit
Temperature	Smartly	15-min	degC
Temperature setpoint	Smartly	15-min	degC
Outside temperature	Weather provider	15-min	degC
Heating power	Smartly	15-min	kWh/period

Table 5.4: Timeseries in the platform

a reasonable number.

The error penalisation is calculated as:

$$K_s^{tempError} \times \sum_{z \in Z} \left(\sum_{p \in P} \varepsilon_{s,z,p}^{temp} \right) : \ \forall s \in S \tag{5.38}$$

Note that there is no explicit penalty for preheating.

5.4 Verification of causality-reversal and discretisation

The above thinking uses a lot of assumptions:

- By forward-shifting the temperature setpoint we can control the heating power
- The interaction of the 15-minute setpoint update frequency and the real time thermostat controller will produce the expected heating power distribution
- The 15-minute discretisation is not too coarse to make the model instable or poorly controllable

To verify these completely will be the role of the pilot, but a Simulink simulation provides good reassurance. The physical system was explicitly modelled in continuous time, a simple hysteresis thermostat controller runs on a high time resolution of 10 seconds, and new setpoints arrive every 15 minutes. The temperature and the electrical load tracks the expectations closely, see Figure 5-5.

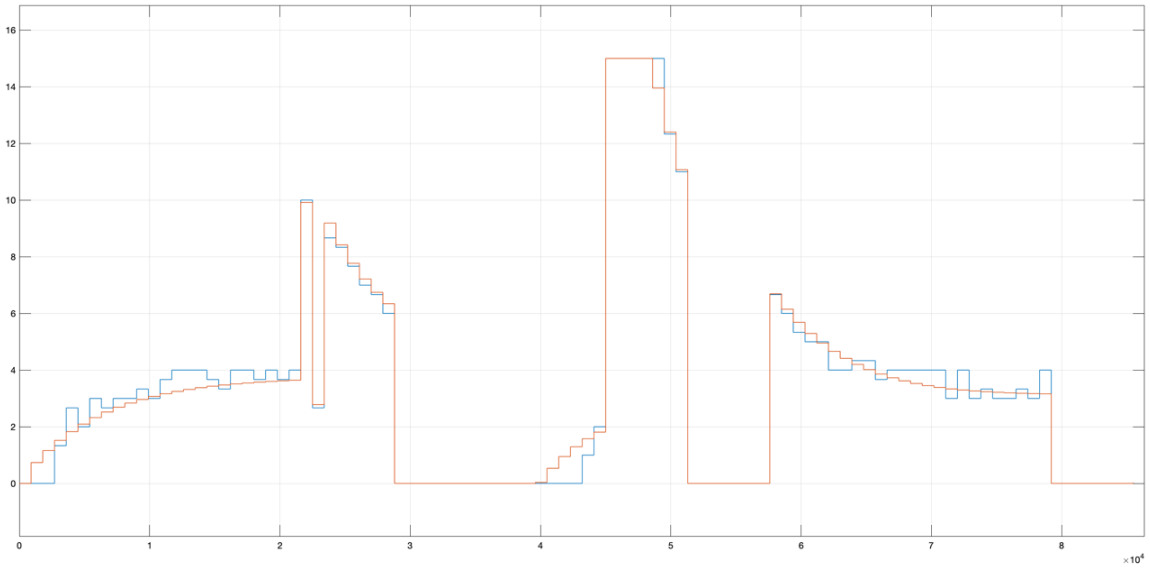

Figure 5-5: Electrical heating load: Red is model expectation, blue is simulated.

5.5 Risk of model failures

If some major error occurred – for example the model received wildly incorrect outside temperature forecast or the house's parameters are significantly changed by opening all doors and windows – the electrical load could be very be different from the one expected. This could result in suboptimal costs, but the occupant comfort will never be compromised because of the model because the setpoint can not leave the specified range.

5.6 Results: Single zone example

A well instrumented experimental house in Quebec has published a dataset [7] that was used to obtain a reasonable set of parameters the model. The second order model

discussed in the identification section was used with the identified parameters:

$$R_{in} = 0.3457 degC/kW$$

$$C_{air} = 0.4827 kWh/degC$$

$$R_{env} = 7.7423 degC/kW$$

$$C_{struc} = 6.3230 kWh/degC$$

A hypothetical, but realistically chosen temperature profile was selected: The house is kept at 19 degC during the sleeping hours, 21 degC when the occupant is home, and it allows for flexibility when the occupant is away for work for example: between 17 and 23 degC.

The example does time of use optimisation for this individual asset only. Figure 5-6 shows the fictitious electricity price and the selected inside temperature and consequent heating power.

The inside temperature found by the model will be set as reference on the thermostat of the house. The heating power is calculated but it is not stored and will not be used in any way whatsoever. Not to be confused with the case when the same model is used for base case estimation, in which case a constant electricity price is considered, and the heating power is the estimation output.

Looking at the heating power graph, a near constant power schedule is seen in the morning hours, which is reasonable, because the outside temperature is steady, the temperature setpoint is also steady, and there is no flexibility allowed, because the minimum and maximum temperature setpoints are equal. As that temperature setpoint rises, so does the heating power, with a peak to start with. This is still constrained behaviour.

During the daytime hours when the allowed temperature range opens up, the model lets the temperature fall all the way to the minimum, then maintains that minimum until the periods with cheap electricity. Then it schedules a lot of heating to take advantage of the cheap energy, heating the building all the way to the top of the allowed range. This preheating is cheaper than waiting for the following constrained

period to regain the temperature. This behaviour showcases the purpose of the space heating model.

The second order nature of the heating system can be seen in the step response of the temperature. When the temperature range opens up and the model lets the temperature fall, it first falls quickly then keeps falling at a lower pace. In the opposite direction, when the model preheats, it expects that the temperature will be first fast to rise, and then it will keep rising at a lower pace. These 'knee points' are characteristic of second order models and occur because the air's thermal mass is much smaller than the structure's.

Another reality of heating systems that also show up in the results is that the temperatures cannot jump, therefore time must be allowed for the transitions. There is such a transition window in the morning and in the evening as well. Sometimes the temperature can not stay within the window, for example the evening window is too small for the inside temperature to naturally fall to the lower setpoint. This is fine, the model allows for higher temperatures than the maximum setpoint, exactly for such cases.

In the very first periods, the temperatures are also too high. This is because the model in the example has been initialised with too high temperatures. This is handled exactly the same way as the insufficient transition window in the evening.

5.7 Multi zone example

This example shows how the notations for a two-zone house, where the bedrooms and the living area are two separate zones. One possible model for it would be the one depicted on Figure 5-7.

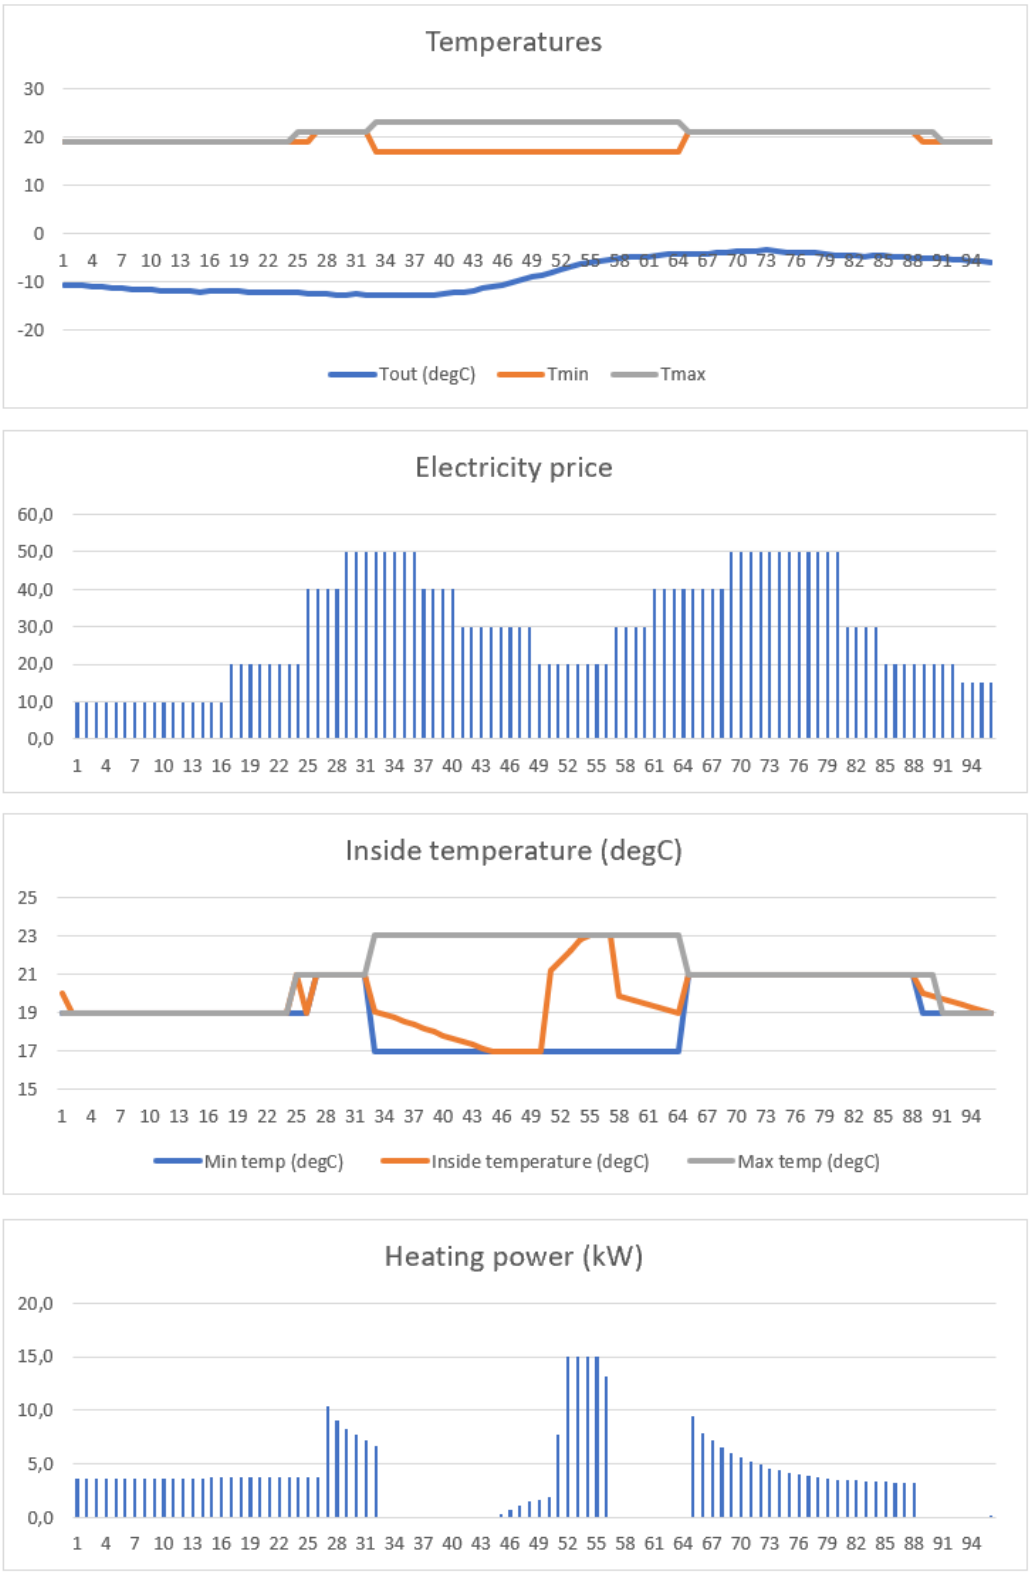

Figure 5-6: Ideal temperature curve solution from the optimiser with parameters as described above

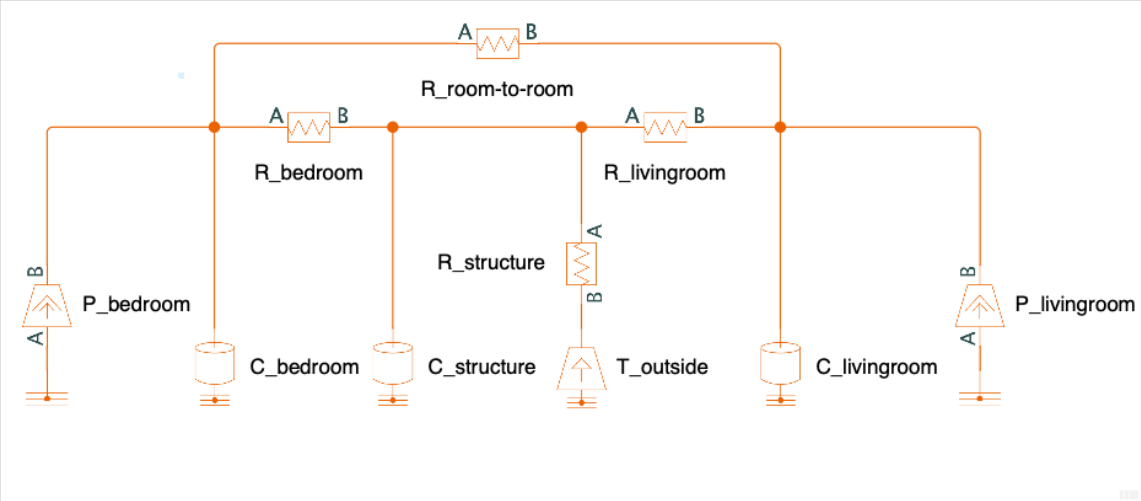

Figure 5-7: Third order two-zone model example

$$x = \begin{bmatrix} \tau_{bedroom} \\ \tau_{livingroom} \\ \tau_{structure} \end{bmatrix} \tag{5.39}$$

$$u = \begin{bmatrix} \pi_{bedroom} \\ \pi_{livingroom} \\ T_{outside} \end{bmatrix} \tag{5.40}$$

A and B are 3×3, C and D are 2×3 matrices.

This model was not tested for lack of data.

Chapter 6

Residential electrical water heaters

6.1 Introduction

The residential electrical water heater model aims at providing demand side response utilising the thermal inertia of the water tanks by controlling only a switch in series with the thermostat and metering only the electrical consumption.

The model is based loosely on the shiftable loads described in [20] and is described as a changeover from the previously used model described in the electric water heater model chapter of [14].

Available information

- The rated power of the heating element (kW)
- Electricity consumption of EWH measured separately (kWh/period)

Available control input

- Switch status (on/off or unblocked/blocked)

Consumption is measured every 15 minutes, and switch status is controlled in according blocks.

Problems with the previous version

- Past decisions must be considered when making new control plan
- Does not scale well due to binary variables (semi-continuous energy consumption!)
- Provides less flexibility on the same dataset
- Constraining parameters are too abstract for the end user

6.2 Thermal model – thinking about heat storage

The model explicitly models the water tank behaviour as a thermal storage, which yields a natural, straightforward formulation that is already linear. The key insight is that the purpose of a water heater's water tank is to store heat. This document describes heat demands, heat storage, and heat production, because this approach helps to understand the energy balance.

In this language, when the occupants take a shower, they are taking heat from the tank. Of course, they are also taking water, but that is immediately replenished by cold water in the tank. Overall, the amount of heat was reduced in the tank. This means the tank temperature is now lower. If someone wants to take a shower again, then more water will be used to carry the same amount of heat. Reminding ourselves of the mixing valve helps to further accept the truth that people need heat from the tank, not water. (In case of no mixing valve, the person in the shower will turn the faucet to mix more hot water and less cold, in the end regulating for heat demand the same way as the mixing valve would.)

In this heat-focused thinking, the electrical heating element provides heat equal to its electrical consumption, because its conversion efficiency is one. The heat loss of the storage tank is minimal and is not practically separable from the heat demand of the occupants – it can be treated as part of the heat demand. The heat demand on the average is equal to the electrical consumption, but the heat demand has higher

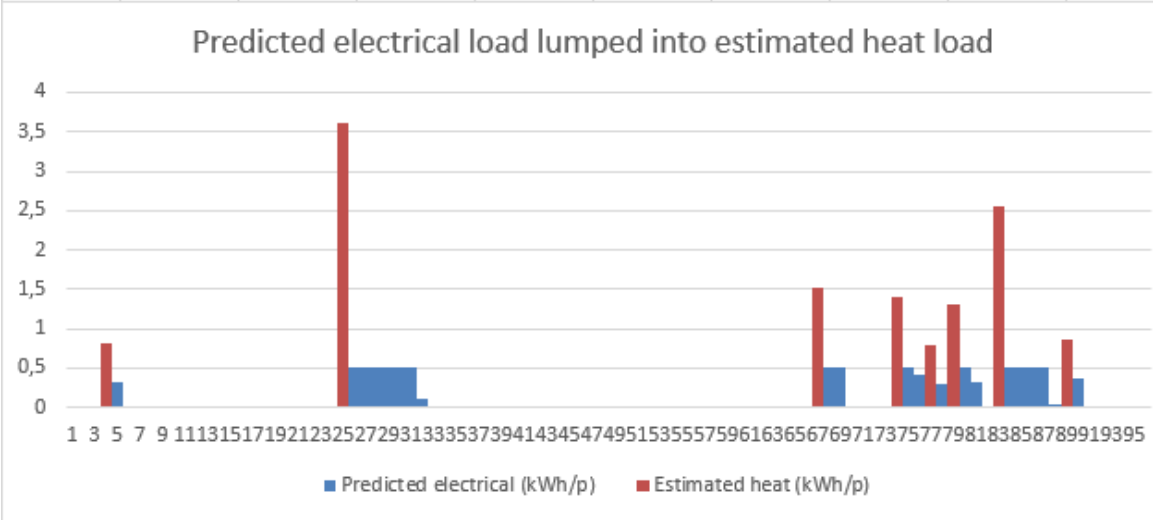

Figure 6-1: Estimation of the underlying heat load based on the electrical consumption of a water boiler.

peaks and shorter durations. (You may take a shower for 5 minutes and your water heater will take 20 minutes to make up for the heat you used.)

If the amount of heat in the tank is tracked, and its volume is known, this is in fact equivalent to keeping track of the temperature in the tank. This makes it very convenient to place constraints on the temperature, which is in the end the only parameter that affects comfort levels; and is also well understood by lay people as well. A typical 200l residential electrical water heater will have a heat capacity of $0.00116\frac{kWh}{degC \times l} \times 200l = 0.23\frac{kWh}{degC}$! If the nominal temperature is 70 degrees, and the occupant allows us to control the temperature down to 50 degrees, this allows a flexibility of 4.6 kWh, the size of a typical home battery.

6.3 Estimating heat load

Only the electrical consumption is metered and predicted. To reconstruct the underlying heat load, the electrical load must be lumped forward. This is best shown in a bar chart as on Figure 6-1.

Of course the lumping is limited to a reasonable heat consumption value – for example if a person is taking a long, hot, high-flow shower at 35 degC at a flow

rate of 8l/min, and the water inlet temperature is 5 degC, he is using $0.00116\frac{kWh}{degC\times l}\times 8\frac{l}{min}\times 15\frac{min}{period}\times 30degC = 4.18\frac{kWh}{period}$. We could accept that as the lumping limit, that is, the highest heat consumption rate we deem possible. (An average water heater will take over 2 hours to reheat the water after his shower if it lasts 15 minutes!)

This lumping also ensures that we are being conservative and look at a worst-case scenario from a temperature-drop point of view. The shower in the example is an extreme case.

The code for lumping algorithm can be found in the appendix.

6.4 Optimisation

Given a price profile and a predicted electrical consumption, the optimisation model first estimates the heat load as described in the previous section, then decides on an optimal heating schedule that will never violate the temperature limits.

The blocking switch commands are decided in a post processing step: in periods when the water temperature is below the nominal but the heating power is zero, the EWH is considered blocked. This trick allows us to do away with binary variables.

$$\tau_{l,t}^{EWH} < T_l^{EWH,max} \; and \; \omega_{l,t}^{EWH} = 0 \;\leftrightarrow B_{l,t} = 1 \;\; \forall \; l \in L^{ewh}, \; \forall \; t \in T \tag{6.1}$$

Where B is 1 when the water heater is blocked and 0 when the switch is on. This is the output of the optimisation which will be stored in the platform and sent to the assets.

6.4.1 Decision variables and bounds

- Heating energy $\omega_{l,t}^{EWH}$
- Tank temperature $\tau_{l,t}^{EWH}$

Both are continuous variables, and there is one of them for each period. The

tank temperature is bound between the minimum and maximum limits. The heating power is bound between zero and the maximum output of the heating element as described in the asset parameter.

6.4.2 Constraints

The only complex constraint is the state equation. This equation describes the temperature evolution.

$$C_l^{EWH} \times 0.00116 \times \tau_{l,t+1}^{EWH} = C_l^{EWH} \times 0.00116 \times \tau_{l,t}^{EWH} + \omega_{l,t}^{EWH} - D_{l,t} \forall\, t \in T,\ \forall\, l \in L^{ewh} \tag{6.2}$$

Where D is the estimated heat load, as described above. This is calculated inside the model and it not stored. It is not used for anything outside of the optimisation model, and it is not worth recording either – it is easy to reconstruct based on the recorded load.

We may wish to forbid blocking in certain periods to protect from prediction errors. This can be done by setting the forecasted electrical consumption as the lower limit for $\omega_{l,t}^{EWH}$ in the periods where we wish to forbid blocking.

The decision variable bounds:

$$T_l^{EWH,min} < \tau_{l,t}^{EWH} < T_l^{EWH,max} \forall\, t \in T,\ \forall\, l \in L^{ewh} \tag{6.3}$$

$$(1 - C_{l,t}^{EWH,\ allow}) \times W_{l,t}^{EWH} < \omega_{l,t}^{EWH} < \frac{Q_l^{EWH,max}}{N^{hour}} \forall\, t \in T,\ \forall\, l \in L^{ewh} \tag{6.4}$$

6.4.3 Temperature penalisation

It must be ensured that the model only maintains a lower-than-nominal temperature if it has good reason to do so – high energy prices that will be followed by low ones before large heat demand is expected. Since the optimised temperature curve can at no point be above the baseline temperature curve, their difference is a trivial linear

expression to penalise. A penalty factor with the dimension $\frac{Euro}{degC \times hour}$ could express how much flexibility-exercise the occupant is willing to tolerate.

Setting a high value would mean zero blocks – therefore gradually decreasing this value would be a good way to tune it to the user preferences and prediction quality. Note that even zero penalty would not compromise comfort if the predictions were perfectly accurate, because of the constraint on the minimum allowed temperature.

6.4.4 Avoiding infeasibilities

If the forecast predicts high enough heat demand the temperature will fall below the setpoint, no matter what. This means infeasibility. To avoid it, the model first calculates what the temperature profile would look like if the water heater was never blocked, and if it has a minimum below the prescribed minimum, it will use the lowest of the two. This way the model is ensured always feasible.

6.4.5 Asset parameters

- $Q_l^{EWH,max}$: Nominal power of the heating element in kW (typically 1.95 kW)
- C_l^{EWH}: Tank volume in litre (typically 120-200 l)
- $T_l^{EWH,max}$: Nominal temperature in degC (typically 60, 65, or 70 degC)
- $T_l^{EWH,min}$: Minimum allowed temperature (typically 45 or 50 degC)
- $T_l^{EWH,inlet}$: Typical inlet temperature (typically 5 degC)
- $W_l^{EWH,water}$: Maximum reasonable water consumption at 35 degC (typically 480l/hour)
- $P_l^{EWH,shift}$: Penalty factor

6.4.6 Cost function

The addition to the cost function is the penalty term:

$$\sum_{t \in T} P_l^{EWH,shift} / N^{hour} \times (T_{s,t}^{base} - \tau_{l,t}^{EWH}) \tag{6.5}$$

Where $T_{l,t}^{base}$ is defined as the estimated temperature for period t in EWH l if it weren't blocked in any period. This is calculated inside the model and it not stored. It is not used for anything outside of the optimisation model, and it is not worth recording either.

In the single-asset simplified example shown on the figures used the following cost function:

$$\sum_{t \in T} \omega_{l,t}^{EWH} \times P_t + P_l^{EWH,shift} / N^{hour} \times (T_{l,t}^{base} - \tau_{l,t}^{EWH}) \tag{6.6}$$

P_t is the electricity price in period t.

6.4.7 Handling wrong predictions

If heat is used in a period when none was expected, there is a chance that the model is blocking in that period. The water temperature will then drop below the estimate, and it might be that it will drop below the setpoint even though the heat consumption was not so extreme that it would be unavoidable.

- The penalty factor ensures that the model only commands long blocking periods if there is good reason to do so.
- The minimum allowed temperature can be set with some safety factor – for example if the true comfort limit is 45 degC, setting 50 degC might shield from the inconvenient consequences of prediction errors.

Once real meter data is in, the prediction-optimisation assembly may be simulated in a loop, and the extent to which wrong predictions cause a problem can be quantified. Additionally, we could consider a different kind of prediction – an upper bound for the load with a 90% confidence for example. Feeding this into the model instead of the original load curve would result in less blocking, but most importantly also reduced risk of cold showers.

6.4.8 Limiting consecutive number of blocks by preprocessing

To eliminate the possibility of cold showers, the number of successive blocks might be limited, and a number of no-blocks enforced. This cannot be done in an optimal way without introducing binary variables. Since it is meant to be a safeguard only during the initial operation when we expect the predictions to be very inaccurate, it is reasonable to move this to preprocessing, even at the cost of somewhat suboptimal solutions.

The preprocessing step would look at the already sent control signals, and if the last $D_l^{EWH,max}$ signals sent are all blocks, it will change the first $D_l^{EWH,min}$ elements of $C_{l,t}^{EWH,allow}$ to zero.

In postprocessing, the code must look at making sure that periods where $C_{l,t}^{EWH,allow}$ is zero are never blocked. This is important when the predictions are zero for the no-block periods, but load is shifted onto them.

6.5 Model performance and parameter tuning

If the predictions were perfect, simply limiting the temperature range would be enough to ensure that the users never experience too cold water, and the model would be able to extract all the possible flexibility from the water heaters in the most optimal way.

Because the predictions will never be perfect, a penalty parameter discourages the model from shifting too much, blocks are forbidden in certain periods, and hard limits are placed on the number of consecutive blocks. It is desirable to loosen these constraints as much as possible to gain access to more flexibility. How loose they can be depends on how good the predictions are, so it is important to be able to measure the prediction quality.

One way to do it is to measure the rebound after blocks are lifted and compare it to the expected rebound. Based on this, we can tell when the water temperature

was lower than it should have been, and we can tune the parameters accordingly – detecting weak predictions before the end users would notice.

It is desirable not to expose the parameters $D_l^{EWH,min}$, $D_l^{EWH,max}$, $T_l^{EWH,inlet}$, $W_l^{EWH,water}$, $P_l^{EWH,shift}$ to the end user, these are for the flexibility operator or eSmart to tune.

6.6 Results

Continuing the example scenario from Figure 6-1, given the prices, the model finds the optimal heating schedule shown on Figure 6-2.

The model attempts to keep track of the temperature inside the tank. The lowermost graph on Figure 6-2 shows the estimated original temperature profile, and the temperature profile resulting from the occasional blocks. By nature, the temperature with the blocks can always only be lower than the original temperature would have been - there is no way to preheat.

The top graph shows the decided master switch positions, and the expected resulting electricity consumption in each period. Indeed it was calculated the other way around: the linear programming model decides only heating powers to avoid having binary variables, and the fitting switch positions are calculated in post processing.

The model behaviour is reasonable as it nicely schedules the electricity consumption to the lowest price periods, while making sure that the temperature in the tank remains acceptable and tries to minimize the difference between the original and the scheduled temperature profile. The priority between the two goals can be set via the penalty property.

Figure 6-3 show the effect of raising the penalty parameter to 1 NOK / (degC period). The overall difference between the original and scheduled temperature is decreased, and the number of block periods as well. The electricity consumption still concentrates in the low-price regions, but the overall cost is higher. Tuning this penalty parameter will be a task of the pilot.

Figure 6-2: Decided heating energy based on the prices and the expected heat load with a penalty parameter of 0.5 NOK / (degC period)

Figure 6-3: Decided heating energy based on the prices and the expected heat load with a penalty parameter of 0.5 NOK / (degC period)

Chapter 7

Conclusions

7.1 Results

This thesis set out to find simple but accurate asset models to work as part of the eSmart Integrated INVADE Platform.

Fast and robust models were developed for residential electrical water heaters and space heating to replace existing models. They have been presented to and accepted by the INVADE partner universities UPC and NTNU. They are going to be validated in a pilot site with 11 households in Western Norway.

An architecture has been developed from the ground up to extract flexibility from two large hotels in Albena, Bulgaria. The testing is underway, and it is the first pilot site to go online as part of INVADE. The simulation results promise a 10-15% cost saving potential.

An architecture has been developed for REMA's Vagle distribution centre to extend the existing battery control capability to controlling a sprinkle tank as a thermal energy storage. It has been successfully presented for the representatives of REMA and accepted for live testing later this year. The simulation results forecast significant energy and cost savings due to keeping the cooling machines in their high-COP regions.

The results for each model have been presented in their respective chapters. Each of them proved sufficiently fast to solve, and provided reasonable control signals.

Work is underway to test them integrated into the INVADE platform, the results of which will be published by the end of the year.

The work done in this thesis proves that mixed integer programming modelling of flexible load and generation assets is possible, and suitable control performance can be achieved with limited information.

7.2 Further work and areas left unexplored

The work remaining is to integrate and test the models in real world conditions and prove the business case for them in real life conditions. Work will continue with all four models. They are going to be used to generate control signals within the eSmart / INVADE platform, and those control signals are going to be implemented by the actual devices on site. Should any significant inaccuracies be detected, the models will be modified to better describe the physical systems.

Work is ongoing to estimate the business potential of each model, as well as to estimate the cost of running the optimisation service. Based on the promising results, eSmart is continuing to developing similar models for various other sites.

As the system scales to handling hundreds or thousands of assets in a single optimisation problem, solver parameters can have huge influence on solve time, which can become a critical bottleneck. Finding the optimal solver parameters, and benchmarking different solvers will become crucial.

Aggregated models could be beneficial to take advantage of the reduced variance of summed variables, as well as to reduce computational complexity. Work is underway for such a model for EV charging.

This thesis started with the general modelling approach already decided. Scenario-based robust optimisation might provide better results with poor predictions.

For models that are not going to be part of a large aggregated optimisation, a more accurate nonlinear model could be computationally manageable. Interfacing such a subsystem to the rest of the system could provide challenging.

Bibliography

[1] Alexander Zerrahnb Andreas Bloessa, Wolf-Peter Schillb. Power-to-heat for renewable energy integration: A review of technologies, modeling approaches, and flexibility potentials. *IEEE Xplore*, May 2015.

[2] H2020 EMPOWER, July 2019. URL: `http://empowerh2020.eu/`.

[3] eSmart Systems, July 2019. URL: `https://www.esmartsystems.com/`.

[4] Miguel Zamarripa et al. Rolling horizon approach for production–distribution coordination of industrial gases supply chains. *Ind. Eng. Chem. Res.20165592646-2660*, February 2016.

[5] H2020 INVADE, July 2019. URL: `https://h2020invade.eu/`.

[6] Venkatachalam Lakshmanan. D5.4 advanced optimal battery operation and control algorithm. *H2020 INVADE Deliverables*, pages 43+, December 2018. URL: `https://h2020invade.eu/wp-content/uploads/2017/05/D5.4-Advanced-Optimal-Battery-operation-and-control-algorithm.pdf`.

[7] Gélinas S. Le Bel C. All-electric experimental twin houses: The ultimate demand management testing tool. *IASTED International Conference Power and Energy*, 2013. URL: `https://meeb.ca/`.

[8] L Ljung. *System Identification: Theory for the User*. PTR Prentice Hall, Upper Saddle River, NJ, second edition, 1999.

[9] Time- or frequency-domain data, July 2019. URL: `https://se.mathworks.com/help/ident/ref/iddata.html`.

[10] Mohammad Fattahi Morteza Shabanzadeh. Generation maintenance scheduling via robust optimization. *Applied Energy*, 2018.

[11] Elhub on NVE, Norwegian regulator, July 2019. URL: `https://www.nve.no/reguleringsmyndigheten-for-energi-rme-marked-og-monopol/sluttbrukermarkedet/elhub/`.

[12] Smart metering on NVE, Norwegian regulator, July 2019. URL: `https://www.nve.no/energy-market-and-regulation/retail-market/smart-metering-ams/`.

[13] Open Charge Alliance, July 2019. URL: https://www.openchargealliance.org/.

[14] Leon Haupt Pol Olivella-Rosell, Pau Lloret and Sara Barja. D5.4 advanced optimal battery operation and control algorithm. *H2020 INVADE Deliverables*, pages 34+, December 2018. URL: https://h2020invade.eu/wp-content/uploads/2017/05/D5.4-Advanced-Optimal-Battery-operation-and-control-algorithm.pdf.

[15] SciPy, July 2019. URL: https://www.scipy.org/.

[16] Xinhua Xu Shengwei Wang. Simplified building model for transient thermal performance estimation using ga-based parameter identification. *International Journal of Thermal Sciences*, August 2005.

[17] SMA Sunny Portal, July 2019. URL: https://www.sunnyportal.com.

[18] Marit Thyholt Tor Helge Dokka. Passiv klimatisering, June 2002. URL: https://www.sintef.no/globalassets/upload/a02513.pdf.

[19] C. P. Underwood. An improved lumped parameter method for building thermal modelling. *Energy and Buildings*, August 2014.

[20] Stig Ødegaard Ottesen. Techno-economic models in smart grids. *Doctoral theses at NTNU, 2017:44*, pages 111+, 2017. URL: https://pdfs.semanticscholar.org/ce2f/6f80162968c63b5aaffffa19e60700ac536d.pdf.