

DELFT UNIVERSITY OF TECHNOLOGY

FACULTY OF ELECTRICAL ENGINEERING,
MATHEMATICS AND COMPUTER SCIENCE

SUSTAINABLE ENERGY TECHNOLOGY

MASTER THESIS

Operating power-to-heat in thermal grids

Mitigating wind power intermittency
through the electrification of a heat network

T. Spruit
4713788

supervised by
Prof.dr. Peter Palensky
and
Dr. Milos Cvetkovic

DEPARTMENT OF ELECTRICAL SUSTAINABLE ENERGY
INTELLIGENT ELECTRICAL POWER GRIDS
IEPG GROUP

April 26, 2020

Summary

Natural changes in climates around the world are increasingly affected by carbon emissions from fossil fuel combustion. Negative externalities of these conventional energy sources encourage a transition to a sustainable energy system. Both the intermittency of renewable energy sources and the carbon-intensity of the heat sector challenge this transition. District heating networks allow waste heat to satisfy a heat demand in the built environment and are therefore considered an energy-efficiency measure. Nonetheless, most sources in heat networks are still carbon-based. An increasing share of renewables in the electricity mix could enable heat networks to be powered by renewable sources through power-to-heat (P2H) technologies. In addition, the thermal inertia of heat networks enables a flexible use of P2H as a source in district heating. In the context of heat and renewable power integration, this report focuses on balancing uncertain wind power production using a heat network.

The setting is a wind farm, two greenhouses, a data centre and a district heating network with P2H. The greenhouses have a heat demand and the data centre provides a limited amount of waste heat. Operation of the P2H heat source is aimed at responding to wind power availability, explored here in two different ways. First, power demand of the heat source should follow the predicted wind power supply. Second, power imbalances caused by wind power prediction errors should lead to shifting the demand of P2H to balance these short-term fluctuations.

As such, the problem is split into two parts, day-ahead and intra-day operation. To these ends, a multi-consumer heat network with a centralised electric boiler has been modelled, taking into account heat transportation times and temperature-dependent heat losses. The boiler should be scheduled a day ahead to be active when electricity prices are low, for the boiler to turn off when electricity prices are high. This day-ahead scheduling problem has the objective to minimise electricity costs for operating the boiler and is formulated as a mixed integer linear programming problem. Regarding the intra-day operation, wind power prediction errors should be mitigated by changes to the boiler schedule, thereby decreasing the need for additional intra-day balancing power. This intra-day operation of the boiler is regulated by model predictive control, as this allows the operator to consider the most recent and therefore most accurate prediction for wind power production. Statistical models have been used to simulate wind power forecasts and production profiles. This report is intended to gain insights into the factors and parameters that affect the performance of the heat network regarding its response to the availability of wind power.

It was found that a number of constraints are necessary for a reliable optimal control of the heat network. The optimal solution to the day-ahead scheduling problem proves to be more cost-efficient than the more straightforward rule-based control. In addition, the linear relaxation of the integer programming problem could not provide a proper approximation of the computationally complex integer problem. For the boiler load shifting in response to wind power forecast errors, various constraints are evaluated in terms of their effect on the reliability of operating the heat network with the model predictive control. The flexibility of the boiler load is compared for a varying boiler power capacity, as well as a varying prediction horizon used in the model predictive control.

Within the context of heat and renewable power integration, heat networks are found to help alleviate wind power intermittency. In a broader context, heat network operators are recognised as possible participants in intra-day power trading by providing balancing power through the energy flexibility of the heat network. This report contributes to the growing interest in demand response and flexibility across energy systems, in this particular case power and heat.

Nomenclature

Abbreviations

AR	Autoregressive
CHP	Combined heat and power
EB	Electric boiler
MA	Moving-average
MPC	Model predictive control
P2H	Power-to-heat
WH	Waste heat source

Greek symbols

α	Coefficient in an autoregressive process
β	Coefficient in a moving-average process
γ	Half width at half maximum
ε	Stochastic white noise term
ϵ	Effectiveness
η	Heat loss factor
ϑ	Time index in integrals
κ	Thermal conductivity
λ	Electricity price
ρ	Volumetric mass density
σ	Standard deviation
τ	Time delay from one network node to another
v	Flow velocity
ϕ	Forecast error

Other symbols

\dot{m}	Mass flow rate
\dot{Q}	Heat flow
\dot{Q}_{hx}	Nominal heat flow through a heat exchanger
\dot{Q}_{wh}	Waste heat production
A	Constraint matrix in an optimisation problem
b	Constraint vector in an optimisation problem
C	Cost
c	Cost vector
c_b	Boiler capacity factor

c_i	Heat exchanger capacity factor
c_p	Heat capacity at constant pressure
c_v	Heat capacity at constant volume
C_w	Heat capacity of a water volume
$c_{\dot{q}}$	Heat transfer factor
D	Power delivery day
d	Number of time steps in a day
E	Energy
e	Internal energy
h	Specific enthalpy
i	Index denoting consumer
k	Time step within prediction horizon
L	Characteristic length
m_i	Consumer water mass
N	Number of consumers
N_h	Number of load hours
P	Boiler power use
p	Order of an autoregressive process
P_b	Boiler power capacity
p_k	Number of hours in a prediction horizon
P_w	Wind power
P_{inst}	Installed wind power capacity
Pe	Péclet number
q	Order of a moving-average process
t	Time indicator
T_0	Initial consumer water temperature
T_a	Ambient temperature
$T_{0,s}$	Initial temperature throughout the heat network
T_{avg}	Average temperature throughout the heat network
T_{max}	Maximum consumer water temperature
$T_{min,s}$	Minimum supply pipe temperature
T_{min}	Minimum consumer water temperature
$T_{r,i}$	Return pipe temperature
$T_{s,0}$	Temperature of supply water after leaving the boiler

$T_{s,i}$	Supply pipe temperature
u_b	Boiler control variable
u_i	Heat exchanger control variable
u_k	Control variable for changing the day-ahead boiler schedule
\mathcal{U}	Uniform probability distribution function

Sets

$\dot{\mathcal{Q}}$	Set of heat exchanger control variables
\mathcal{P}	Set of boiler control variables
\mathbf{u}	Set of all control variables
\mathcal{K}	Set of time steps in a prediction horizon
\mathcal{S}	Set of consumers
\mathcal{T}	Set of time steps in a day

Contents

Summary	i
Nomenclature	ii
1 Introduction	1
1.1 Motivation	1
1.1.1 District heating networks	1
1.1.2 Renewable power intermittency	2
1.1.3 Heat and renewable power integration	3
1.1.4 Related work	5
1.2 Problem definition	6
1.2.1 Day-ahead scheduling problem	7
1.2.2 Wind power forecast error mitigation	8
1.3 Report structure	10
2 Modelling a heat network	11
2.1 Background on heat networks	12
2.1.1 Temperature dynamics and the node model	12
2.1.2 Heat transport	13
2.2 Multi-consumer heat network	14
2.2.1 Supply pipe temperature	14
2.2.2 Return pipe temperature	15
2.2.3 Temperature after the boiler	15
2.3 Heat network components	16
2.3.1 Single consumer model	16
2.3.2 Data centre as waste heat source	17
3 Network operation schedule	18
3.1 Background	18
3.1.1 Electricity markets	18
3.1.2 Optimisation techniques	19
3.2 Problem formulation	21
3.2.1 Cost function and control variables	21
3.2.2 Constraints	22
3.2.3 Optimisation problem	23
3.3 Benchmark: rule-based operation	24
4 Smart network control	25
4.1 Background	25
4.1.1 Wind power forecasting	25
4.1.2 Model predictive control	27
4.2 Wind power simulation model	29
4.2.1 Daily wind power production profile	29
4.2.2 Rolling wind power forecast	31
4.3 Problem formulation	34
4.3.1 Objective function and control variables	34
4.3.2 Constraints	35
4.3.3 Optimisation problem	36
5 Results	37
5.1 Standard experimental setup	37
5.2 Day-ahead schedule	38
5.2.1 Feasibility and runtime	38
5.2.2 Rule-based control comparison	42
5.2.3 Integer versus linear programming	45

5.2.4	Doubling heat network time delays	46
5.3	Boiler load shifting	48
5.3.1	Electric boiler capacity	48
5.3.2	Effectiveness and reliability	51
5.3.3	Heat network operator at the intra-day market: a comparison	54
6	Discussion	57
6.1	Network time delays	59
6.2	Boiler shifting performance	59
6.3	Future work	61
7	Conclusion	62
	References	63
A	Background on waste heat	67
B	Derivations	68
B.1	Consumer temperature	68
B.2	Supply temperature	69
B.3	Return temperature	71
C	Optimisation problem implementation	72
C.1	Temperature control variables	72
C.2	MATLAB scripts	73

1 Introduction

Carbon-based fuel-intensive activities have shown a significant link to atmospheric carbon dioxide levels. The physical properties of the carbon dioxide molecule allows for the absorption of infrared radiation, turning the electromagnetic energy into thermal energy. Consequently, the atmosphere rises in temperature. This rise in average global temperatures comes at a cost. The warming and subsequent expansion of ocean water lead to increased sea levels and natural threats in coastal areas. Intensified weather extremes affect the economy and the well-being of people. Climatic change challenges ecosystems around the world to adapt at unusual rates.

These negative externalities of fossil fuel combustion encourage a transition from conventional sources of energy to a sustainable energy system. Numerous challenges are central to this energy transition. Thus far, various ways to harness solar power have been developed and employed; both directly, with photovoltaic cells, and indirectly, in the form of wind power¹ and biomass². One of the challenges is the intermittency problem of renewable energy sources for electricity.

Another problem for the energy transition is that the heat sector has been neglected compared to the production of electricity in the prioritisation of lowering carbon emissions, even though heat is the largest energy end-use [1]. Half of total energy consumption is heat for the industry sector and the buildings sector. Total industrial heat consumption amounts to over 1800 TWh per year in the EU, of which over 500 TWh is attributed to the iron and steel industry [2]. Most of this energy comes from burning fossil fuels [1]. Improvements in the heat sector include a better insulation of buildings [3], total site heat integration [4], combined heat and power (CHP) plants [5] and the use of geothermal energy [6]. Furthermore, district heating networks have proven to be a feasible way of providing heat to communities. Nonetheless, as communities grow continuously, so too does their heat demand. The capital investments for extension of the heat network and expansion of central heat production capacity play a limiting role for district heating networks. Furthermore, current district heating relies heavily on fossil fuels and is therefore limited by environmental constraints challenged by carbon emissions. The integration of industrial waste heat is a viable response, as is the electrification of the heat sector where possible, enabling renewable energy sources to satisfy a heat demand.

1.1 Motivation

Before the main problem in this report can be properly formulated, it is instrumental to provide the relevant context. This background and the provided literature form the motivation for this report. District heating networks, their energy flexibility and their potential to enable renewable power integration through electrification will be discussed presently, to be followed by the description of the problem.

1.1.1 District heating networks

Conventional district heating networks connect a distinct, centralised thermal plant to densely populated residential areas that solely consume heat in the form of domestic hot water and space heating in residential or office buildings. The heat supply is often excess heat from industrial plants, heat from waste incineration or from cogeneration in a combined heat and power plant [7]. Pressurised water of below 373 K is the heat carrier transported through a network of pipes buried in the ground [8]. To separate the main, primary heat network from secondary heat networks operating at varying temperatures, heat networks have substations throughout the grid. These contain circulation pumps, metering systems and at least one heat exchanger, through which consumers can extract heat from the network. A second heat exchanger is usually present to heat up water for domestic hot water use [9].

District heating networks are considered an energy efficiency measure because of the recovery of excess heat from industries³, thereby reducing primary energy demands. Even though heat demand in buildings is expected to be reduced through improved insulation, district heating networks can be expanded cost-effectively in European cities [10]. After all, space heating makes up most of the heat demand in the domestic sector and, along with hot water, is consumed at relatively low temperatures, i.e. below 373 K;

¹Wind is caused by atmospheric pressure differences, which are in turn caused by the differential heating due to differences in solar energy absorption around the globe. Therefore, wind power is solar power converted to wind.

²The energy in biomass is electromagnetic energy from the sun converted to chemical energy through photosynthesis.

³Appendix A further discusses waste heat recovery.

the temperatures required for space heating and domestic hot water are considered above 296 K and 323 K, respectively. Supply and return temperatures for low temperature heat networks are envisioned to be 323 K and 293 K. Nonetheless, a return temperature below 303 K has not been recorded thus far [11]. In addition to the domestic sector, industrial heat demand for space heating, hot water and process heat at low temperatures amounts to a quarter of the total industrial heat demand [12]. Therefore, heat networks have a history of presenting opportunities for the use of waste heat in both district heat networks and, more recently, low temperature industrial processes.

Data centres worldwide are estimated to have had a total electrical energy use of 270 TWh in the year 2012 [13]. This global power demand of data centres is expected to increase further [14]. In addition, since most of that electricity is converted to heat eventually, data centres are identified as a reliable source of heat in a district heating network [15]. For further reading, Werner [16] provided various contexts for district heating around the world, including, but not limited to, market, technical and institutional contexts.

A problem with existing district heating networks is that most energy used for heating still comes from fossil fuel combustion. A substantial effort is necessary to decrease the carbon intensity of district heating networks. District heating related carbon emissions remained relatively constant throughout the world in recent years [1]. Especially the reliance of China on coal for district heating is a critical factor behind this trend, as global emissions keep rising because of that (among many other factors).

1.1.2 Renewable power intermittency

Unlike conventional power sources, wind and solar power output is uncertain and therefore never fully predictable [17]. The increasing penetration of these intermittent and variable renewable energy sources into the energy mix has significantly complicated power system balancing for grid operators. Errors in wind and solar power forecasting could result in great difficulties for real-time operations of a power grid. For example, technical stability of power systems is increasingly affected by the intermittency of these renewable energy sources [18]. Furthermore, renewable power production exceeding the regional or national power demand can lead to negative electricity prices and the curtailment of energy. After all, the grid feed-in of wind and solar power does not necessarily respond to electricity demand during operation.

With the current share of renewable, intermittent energy injected into the electricity grid, the necessity of large-scale balancing power capacity is not yet critically urgent. Consider the case of photovoltaic panels on a home. When producing more power than is locally used at that moment, the power is supplied to the grid. Thereby, in effect, a signal is sent to the grid operator through the network frequency imbalance to reduce the output from controllable sources such as coal plants or hydroelectric power dams. When too little power is generated by those photovoltaic panels, electricity is taken from the grid and controllable plants ramp up. With this flexible generation capacity, called positive balancing power, electrical energy can be virtually stored in the grid without a need for any storage technology. Homes with solar panel rooftops have become both consumers and producers of electrical energy, while using the grid as "storage", through which electrical energy exchange is possible when balancing energy supply and demand for a household is necessary. However, this only works as long as the share of renewable electricity in the energy mix is low enough.

Residual load is defined as the total load in the electricity network minus the total power from the renewable energy sources [19], [20]. A negative residual load indicates the power generation exceeds the total load due to an underestimation of renewable power production [21]. These residual loads cause an imbalance to the stability of the system. Balancing power is usually generated by fossil fuel plants, but the increase in renewable sources shows this solution reaches its limit at some point.

Storage

The variability in the energy supply from renewable sources has brought to attention the necessity for energy storage systems [22]. In response to this intermittency problem, batteries and power-to-fuel technologies have been in development to enable the storage of electrical energy. In the case of storage in the form of synthetic fuels, produced with hydrogen from electrolysis, electrical energy can be produced again through combustion in power plants. Negative residual loads can be responded to by charging the available energy storage system, while discharging a storage system balances out the moments of positive residual loads. Nonetheless, costs and energy losses both remain high for storage technologies, while their respective development trajectories remain uncertain [23], [24].

Demand response through flexibility

The energy transition is not merely improving technology to increase efficiency or switching to low-carbon energy sources. It is a paradigm shift that includes the way demand responds to momentary changes in energy supply. Storage was mentioned earlier as a solution to the intermittency problem of renewable energy sources, but there are downsides to energy storage such as high costs. Demand response is a further option. One that is regarded as a (partial) solution to the issues of uncertain and fluctuating power supply [25]. The intermittency problem of renewable energy sources has encouraged the investigation into energy flexibility. The flexibility of electrical loads can support the power system by shifting loads from moments of positive residual load to moments of negative residual loads. Nonetheless, not all electrical loads are flexible enough to solve the intermittency problem altogether. Furthermore, with energy demand responding to energy supply, consumer demands might not always be satisfied.

1.1.3 Heat and renewable power integration

Cogeneration or CHP generates electricity and heat at the same time. Efficiencies in CHP plants are significantly higher than in separate power generation and heat generation. The reason for this is that a CHP plant recovers heat that would otherwise be wasted. In the case of solar power generation, heat at a relatively low temperature is a byproduct and can be recovered to some extent. To wind power, cogeneration does not apply. Nonetheless, after wind power is generated, heat can be produced using this wind energy, for example by powering a heat pump or an electric boiler using electricity from wind turbines. This way, wind power could indirectly be a heat source in a heat network. The various aspects of the electrification of a heat network will be explored here. Ultimately, the heat storage capabilities of a heat network will be coupled to the power production by intermittent sources.

Thermal inertia and flexibility

The subject of flexibility as discussed in Section 1.1.2 further applies to energy use in the form of heat. Buildings as heat consumers store heat and, although their capacity is limited, are the most operable and feasible heat storage systems [26]. After all, a building would not suffer an immediate drop in temperature after switching off its supply of heat. Indeed, it is possible to manipulate the heat demand of a building without affecting the indoor climate. Another straightforward method of storing thermal energy is the installation of thermal storage tanks, but these are to be additionally installed. It is more economical to use the storage capacity of a heat network itself through its pipelines.

In the case of heat networks, flexibility in heat and temperature demand in heat networks are to be exploited to function as a buffer. This so-called negative balancing means that energy demand is lowered during a period of high prices, while a surplus demand at lower prices is essential to balance out this periodic lower demand. In the case of a heat network, the prices are those for operating the heat source, while the demand is heat at a consumer-specific temperature. The thermal inertia of the heat network is the buffer due to which demand response is possible. Using the thermal inertia of a heat network has already been identified as an effective energy-efficiency method [27], [28]. One of the reasons for this is that it enables the integration of waste heat sources, as mentioned earlier.

A newer concept of storing energy in a heat network arises with the integration of alternative heat producing units, e.g. the aforementioned heat pump running on wind power. Similar to the waste heat sources, these heat generators function as additional heat sources in a heat network. This presents further opportunities for handling the problem of time-varying energy production by renewable sources. What this concept entails in more detail, will be explained forthwith.

From electricity to heat

Power-to-heat (P2H), or electricity-to-thermal, is converting electric power to heat. In particular, with the conversion of renewable power into heat, this electrical energy is often energy that would otherwise have been lost through curtailment [29]. Curtailment is one of the manifestations of the intermittency problem, as mentioned in Section 1.1.2. Power-to-heat is an attempt to resolve some of these current issues with the integration of wind and solar power. Overall, power-to-heat technologies are potential solutions for these issues facing the development toward a sustainable energy system, as has been concluded for the European case [18]. This potential of power-to-heat is categorised into a theoretical potential, defined as the total heat demand, because this heat could theoretically come from electric power and therefore renewable power. The technical potential, however, is the amount of negative residual load in the electricity grid that could

be used to satisfy the heat demand, on condition that these moments of negative residual load coincide with the heat demand.

The industrial sector accounts for a large part of heat demand in the Netherlands. Temperature requirements range from under 400 K to over 1300 K, while power-to-heat can be used only for relatively low temperature industrial processes (of below 530 K). In the Netherlands, the theoretical potential of power-to-heat for industrial heat consumption amounts to 36 TWh per year [30].

Horticulture — with a total annual heat demand of 27 TWh — is identified as another sector in which power-to-heat has a considerable potential to satisfy a heat demand, due to the low temperature requirements [30]. The large heat buffers of greenhouses make a case for investigating the technical and economic potential of power-to-heat in the horticulture sector.

For district heating networks, the Netherlands shows a power-to-heat potential of .8 TWh per year, due to the low percentage (5 – 10 %) of district heating in the total heat demand in the residential sector [30]. In contrast, the total annual supply of district heat in Sweden amounts to 57 TWh, while the technical potential is estimated to be between .2 and 8.6 TWh [31]. In Germany, the theoretical potential of power-to-heat in district heating amounts to an annual 98 TWh, while the technical potential amounts to an estimated 8.5 TWh in the year 2030 [21]. To increase the technical potential of P2H in a heat network, some flexibility in heat demand should enable the heat network to shift its load to moments of negative residual load. A practical explanation can help clarify why this condition is necessary. Excess electricity from renewable sources intended to satisfy a heat demand can be stored in the form of heat for later use through a power-to-heat connection to the heat network, i.e. a heat pump or an electric boiler. Temperatures throughout the network are raised by P2H to levels that are sufficiently high to endure a period of low power supply from renewable sources later on. As stated earlier, this form of demand response is made possible by the buffer that is the storage capacity of the heat network itself.

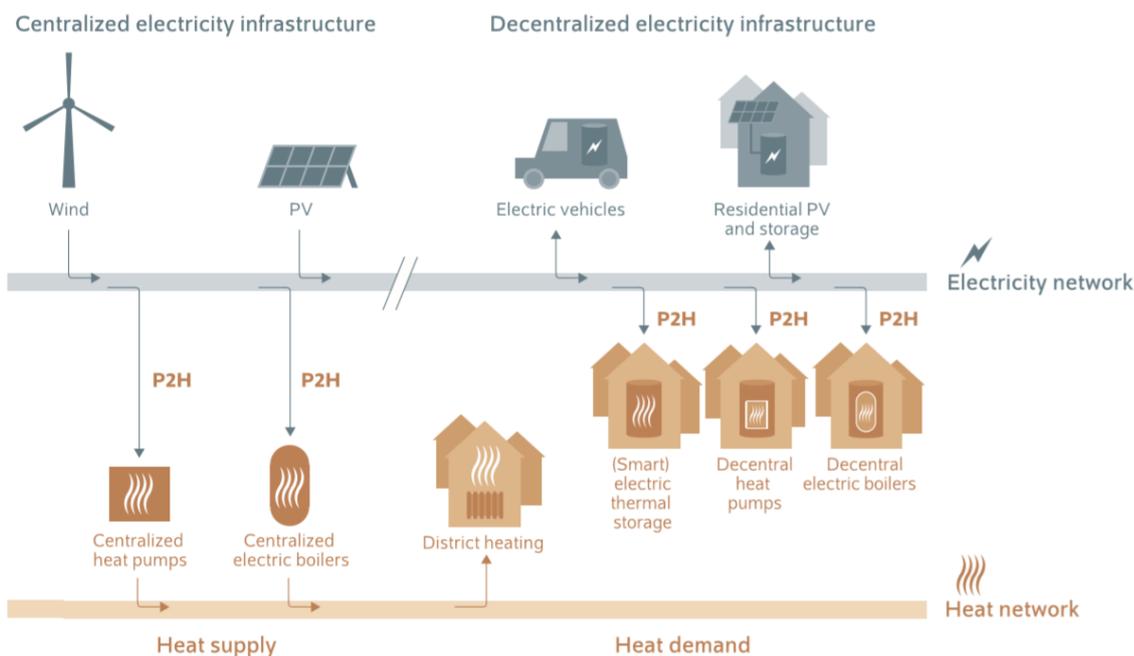


Figure 1: From renewable power to heating in the built environment, adapted from Bloess et al. [32]

Figure 1 shows the two pathways from electricity to residential heating: from the power grid to decentral heating, and from the power grid through the heat network to district heating. The latter pathway will be the main focus here. This use of electricity in heat networks might further decarbonise the energy sector in its supply of heat and enable the integration of a larger share of intermittent power producers into the electricity grid [32]. Indeed, if heat networks use energy at times of high renewable power production, they increase the share of low-carbon power for heat production and they lower the need for renewable power curtailment. With the desire to decarbonise the heat sector on one side and the variability issue of increased renewable electricity on the other, both sides can help each other advance into the desired direction.

1.1.4 Related work

The potential of heat pumps in thermal grids, bidding strategies of the heat network operator in the electricity market and uncertainties in wind forecasting are mentioned as ongoing fields of research [33]. The concept of heat pumps has already been connected with heat networks as a potential alternative heat source running on renewable power [34]. Waste heat can be recycled to heat at a higher temperature when a peak in electricity production by renewable sources drives down electricity prices. Such a price response strategy is part of an overall demand response management, which optimises energy demand over time [35]. Research on the P2H potential in China showed electric boilers in combined heat and power systems could reduce wind power curtailment [36]. Compared to a heat pump, the investment cost for an electric boiler, or immersion heater, is low. When operated for relatively few hours annually, it is preferred over the heat pump [31]. Meanwhile, research done in Germany on the profitability of using the power-to-heat technology concluded that, although the excess electricity at peak renewable power production is not yet cost-efficient for use in heat networks, this will change with cost reductions of electric boilers and the expansion of renewable energy sources. This is partly explained by the fact that the current number of hours with negative residual power load is still too low [20], [21]. Other research on the German market shows higher fuel prices favour the use of heat pumps over combined heat and power plants, whilst both technologies provide a significant potential for power flexibility, with a preference for heat pumps [37].

Since the theoretical potential of P2H has been established, it is worth to look into the different approaches to actualise this potential. Case studies in the urban environment have shown power-to-heat systems with thermal storage and load shifting can decrease surplus grid power, through the use of an optimal control approach instead of the often employed rule-based control [38]. The surplus grid power, or negative residual load, is often a result of peak wind power production. Using a signal from the wind as an environmental incentive for load shifting with the intention to reduce system operating costs has been proposed by others in the past [39]. In essence, the wind signal merely functions as an indicator for the electricity price on the market, which shows modestly, but statistically significant, to be inversely correlated with wind power generation throughout the entire year. Taking the paper by Salpakari et al. [38] as an example of optimal control, the scheduling of power-to-heat and load shifting is formulated as a mixed integer linear problem, solved sequentially for each 24-hour time period. It is however pointed out that forecast errors are not included. Furthermore, optimal scheduling considers only the day-ahead market in its analysis and optimisation. It is therefore ill-suited for the fast dynamics of the intra-day market with volatile prices and demand changes, or the uncertainties and errors in wind power forecasting. As a result, the optimal schedule proves to be sub-optimal, as the controller follows its predetermined direction even when the updated conditions would signal it to deviate. In short, the predictions at the base of the optimisation degrade with time, because of changing energy prices and weather forecasts [40]. A heat network operator must be responsive to dynamic pricing through the appropriate use of optimisation techniques [41]. Challenges to the optimisation problem include the temporal variation of price, heat supply and demand profiles [42].

Table 1 lists a few papers highlighting their focus and the limits of their scope. Although the potential of power-to-heat in the context of wind power integration has been investigated, the use of power-to-heat in heat networks is largely unexplored within that context. Furthermore, most studies on the topic include (carbon-based) cogeneration in their analyses. Moreover, the uncertainty in wind power generation is neglected, which motivates the investigation into the (intra-day) operation of power-to-heat in heat networks with uncertain power generation by wind farms.

Table 1: Papers on the potential or control of power-to-heat in buildings and heat networks

Author(s)	Approach	Limitation
Salpakari et al. [38]	Mixed integer linear programming for optimal control of P2H over a 24-hour period.	The heat network itself is not modelled and intra-day load shifting of P2H for power imbalance management is not investigated.
Chen et al. [36]	Linear modelling of a heat network with CHP, electric boilers and heat storage tanks to minimise wind power curtailment.	Wind power production fluctuates in the simulation, but is not regarded as uncertain.
Bianchini et al. [40]	Model predictive control for optimal heating of a building using price-volume demand response request signals.	A simple thermal model of a building is used instead of something as complex as a heat network.

1.2 Problem definition

In the context of heat and renewable power integration, the challenge arises to power a heat network through power-to-heat technologies. The electrification of a heat network would enable renewable power sources to be the heat source in the system. In the same way, power-to-heat may be used to help integrate fluctuating renewable energy sources. The purpose of this report is to present a method to use the electrification of a heat network to alleviate the intermittency problem in wind power generation. The main research question can be formulated as follows:

R. *How to balance uncertain wind power production using a heat network?*

Now, the case to be studied will be described. The system under consideration consists of a wind farm and a heat network with an electric boiler. Two greenhouses with a heat demand are connected to the heat network (Figure 2). Greenhouses are chosen as heat consumers because of their thermal buffer and the power-to-heat potential of horticulture. The heat network must ensure consumer satisfaction by maintaining temperature levels in the two greenhouses. In addition to heat from the electric boiler, waste heat from a data centre is added to the network.

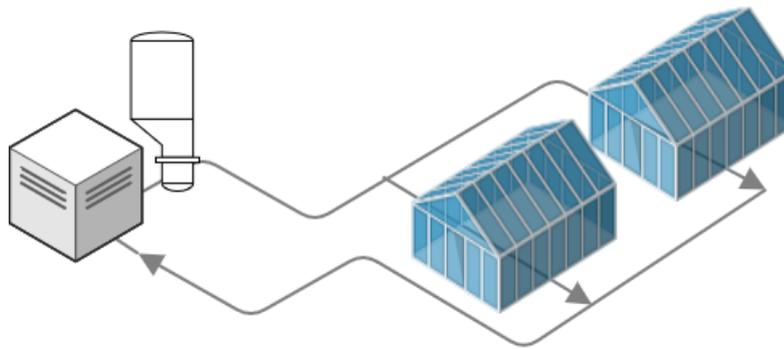


Figure 2: Case setting; data centre, electric boiler and greenhouses

The electric boiler is the central heat source in the network. Operating costs consist of the electricity costs for the boiler and costs coupled to the wind farm's need for intra-day balancing. The main problem can be broken down into two parts, which will be discussed now. First, by operating the centralised electric boiler, greenhouse temperature levels in the heat network must be maintained. Because it takes time for heat to be transported from the source to the consumers, a proper planning of the heat network operation is needed. The first part of the problem therefore concerns the day-ahead scheduling of the heat network:

R1. *How to plan a day ahead the operation of power-to-heat as a centralised source in a heat network?*

The electric boiler is instructed to respond to the availability of wind power, with the intention to reduce operating costs. Nonetheless, even if day-ahead wind forecasts indicate no wind, there would still be a heat demand. Therefore, electricity to power the boiler is bought at day-ahead prices. These electricity prices reflect the predicted wind power production. The first part of the problem ends here, but will be described in more detail in Section 1.2.1.

The second part of the problem is related to the uncertainty in wind power production. After the day-ahead wind power forecast is established, the wind power production is actualised on the day itself, the intra-day. Here, the unpredictability of the wind is reflected by the prediction errors in the day-ahead wind power forecast. These errors must be mitigated with balancing power. This will be explained in more detail in Section 1.2.2. A flexible operation of the electric boiler as heat source in the network may supply (part of) the necessary balancing. The second part of the problem therefore concerns the intra-day operation of the heat network:

R2. *How to mitigate wind power forecast errors during intra-day operation of the heat source?*

By modelling a heat network and wind power forecasts, an optimal control system can be constructed for operating the different components in the heat network. Optimisations to be made throughout the process of mitigating the intermittency of wind power form the core of this research, with the intention to gain insights into the factors and parameters that affect the performance of the heat network in this regard.

1.2.1 Day-ahead scheduling problem

Consider the operation of the heat network. Here, the focus lies on operating the electric boiler. The heat network operator must satisfy a heat demand in the network at all times. For the day-ahead scheduling, the electric boiler is scheduled to be either turned on or off at every hour of the day. Electricity can be bought by the heat network operator to power the boiler. At the day-ahead electricity market, power is bought and sold at hourly spot prices. The buying strategy must be such that the heat demand is satisfied. A constrained amount of flexibility in this demand can be taken into account. The day-ahead electricity price profile is assumed to be deterministic. With this in mind, the first sub-question can now be formulated:

R1. How to plan a day ahead the operation of power-to-heat as a centralised source in a heat network?

- a) *How to minimise electricity costs when scheduling the electric boiler considering heat network time delays?*

Assumption 1. *Day-ahead electricity prices are known beforehand.*

Figure 3 shows a flow chart of this day-ahead scheduling of the heat network operation. The optimisation is aimed at providing an operation schedule. It thereby takes into consideration both the physical constraints of the heat network, as well as certain temperature constraints to satisfy the heat demand of each greenhouse. The resulting schedule for the electric boiler depends further on the cost of the electricity needed to power the boiler.

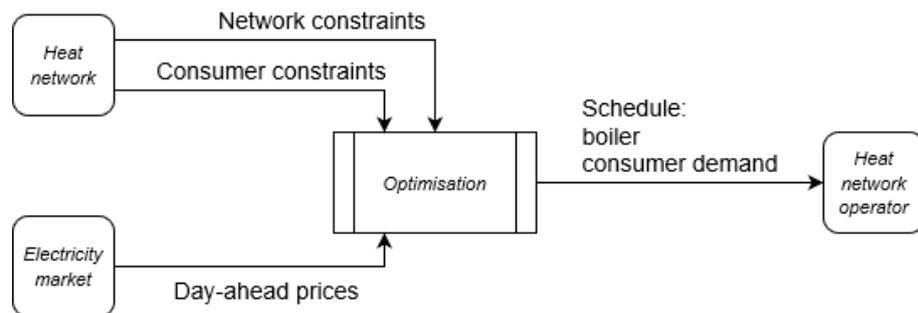


Figure 3: Schematic overview of the day-ahead scheduling problem

In addition to the boiler schedule, each consumer connected to the network can schedule when to draw heat from the network. Heat can be extracted from the network through a heat exchanger to satisfy the demand of the greenhouse. Here too, a constrained amount of flexibility in this demand can be taken into account. Due to the effect the consumer heat demand has on operating the heat network, this heat demand over time is to be optimised along with the electric boiler use. The second sub-question is thus:

R1. b) *How to schedule the heat demand to maintain consumer temperature levels?*

In answering the questions and evaluating the result, making a comparison with a benchmark is a useful step. Operating a heat network consists of decision-making regarding the use of an electric boiler and the exchange of heat at a consumer substation. A simple and straightforward way of decision-making over time is a rule-based control strategy. For example, if the temperature level in a greenhouse falls below a predetermined minimum temperature, the greenhouse will start extracting heat through the exchanger and will only stop doing so when the temperature reaches a maximum temperature. In Section 3.3, a rule-based operation of the heat network will be described. However, such a control system does not take into account price fluctuations and will likely lead to a sub-optimal cost of operating the electric boiler in the heat network. Nonetheless, the elementary nature of a rule-based control system makes it valuable as a benchmark in the performance comparison with other control strategies.

1.2.2 Wind power forecast error mitigation

A power producing wind farm is now to be considered. At the same time as the scheduling, albeit independently from the electric boiler, the wind farm operator submits an electricity bid for each time slot at the day-ahead market. The amount of electricity offered per hour is based on a probabilistic prediction for the wind power production profile, to be discussed in Section 4.2.1. The wind farm operator makes a commitment with the intention of maximising profits and is obliged to deliver the power that has been offered. Naturally, the power produced by the wind farm depends on the wind speeds and the wind farm operator can only produce power using the wind farm.

Next, the following problem is described. Considering a probabilistic prediction is used to sell wind power at the day-ahead market, a problem arises when the prediction turns out to be erroneous. The intra-day electricity market allows for changes in power exchange after the day-ahead market has closed. Nonetheless, prices are usually more volatile and oftentimes higher, because supply is limited and a change in demand is often more urgent. Most power is thus traded through the day-ahead market [43], but with an expansion of intermittent power sources, the need for intra-day balancing increases [44]. The balancing potential of the heat network in support of the wind farm operation is therefore explored.

Avoiding the intra-day market

It is worth explaining further why — especially for a wind farm operator — intra-day trading is disadvantageous. As stated earlier, the day-ahead wind power forecast is submitted at the day-ahead market. An actualised hourly wind power production exceeding the forecast, forces the wind farm operator to sell at the intra-day market. Similarly, a lower actualised power production leads to buying at the intra-day market to be able to satisfy the obligation of supplying the electricity offered at the day-ahead market. This way, a dependency on the intra-day market electricity prices is created for the wind power producer [45].

These intra-day prices come with a greater uncertainty in a future energy system with a high share of intermittent renewable energy sources such as wind and solar in the energy mix. As a result, the main problem arises. The wind power producer may have to sell its overproduction at low prices and buy balancing power at high prices. Because an overproduction in wind power is a direct result of prediction-exceeding wind speeds, it is likely that other wind farms suffer an overproduction as well and need to sell on the intra-day market; thus the price is lowered. Similarly, a lower power production is coupled to lower wind speeds. This is expected to force other wind farm operators to buy on the intra-day market, thereby raising the price. In short, selling at low prices and buying at high prices is noticeably cost-inefficient. Additionally, most intra-day markets are highly illiquid, which limits the opportunities of advantageous trades for a wind power producer [46].

Assumption 2. *A wind farm operator benefits from avoiding intra-day trading.*

The role of the boiler

The problems associated with participating on the intra-day market have been discussed in the previous paragraph. Here, the electric boiler and the heat network come into play. Use of the electric boiler has already been scheduled based on day-ahead electricity prices. Meanwhile, the wind farm operator is committed to a day-ahead schedule of its own, to supply a certain amount of energy for every hour of the day.

Intra-day operation of the wind farm and the electric boiler in the heat network are now to be considered. Figure 4 shows a flow chart of this intra-day operation, which takes into account an updated and therefore more accurate wind power forecast. Prediction errors in day-ahead wind power forecasts can be (partially) offset, using the flexibility of the boiler. For example, if more wind power is produced than predicted and the boiler is turned off according to its day-ahead schedule, the boiler could turn on in spite of its day-ahead schedule, powered by the surplus power from the wind farm. In short, a change to the day-ahead boiler schedule aims to counterbalance the prediction error, which is represented in Figure 4 by the new boiler schedule. This option is preferred over selling the surplus electricity on the intra-day market. Conversely, if less power is produced than predicted and the boiler is turned on according to its day-ahead schedule, the boiler could turn off anyway. The power saved by turning off the boiler would mitigate the wind power forecast error.

The wind power forecast error, which plays a role in the intra-day operation of the heat network, will be discussed now. The most accurate forecast is that for the nearest prediction horizon. For the wind farm

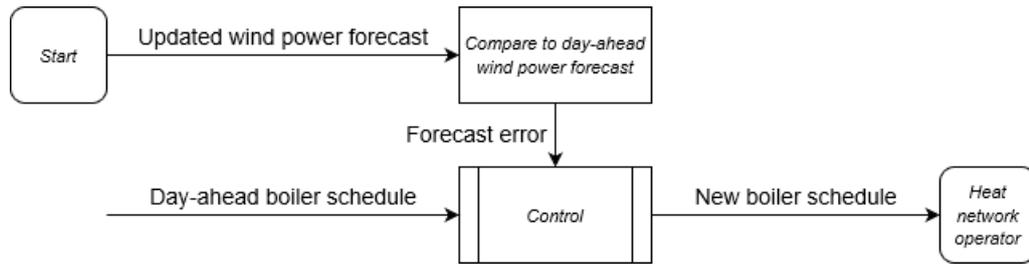


Figure 4: Schematic overview of the control system with continuously updating wind power forecasts

operator at the day-ahead market, the nearness of the prediction horizon is limited by the market closing time. Nonetheless, wind power forecasting continues after the day-ahead market closes. At the start of the day in question, the uncertainty of the forecast for those first few hours has already decreased in comparison with the day-ahead forecast. To use this increased accuracy to the advantage of the operator, the boiler schedule must be updated continuously, following the updates for the wind power forecast. After all, to minimise intra-day market participation, the wind power forecast — used for the boiler control throughout the day — must be as accurate as possible. This way, day-ahead prediction errors are counterbalanced by changes to the boiler schedule most effectively.

Throughout the day, a control system, using a sliding time window with the most updated wind power forecast, optimises the shifting of the boiler operation within the time horizon. Consequently, the choice for this time horizon is expected to be of some relevance to the potential of the boiler to respond to wind power forecast errors for different time scales. After all, a time horizon that is too small might limit the number of time steps available to shift loads, while a time horizon that is too large might suffer from the growing uncertainty in wind power forecasts. The prediction horizons, or control horizons, taken into consideration will be from six hours up to sixteen hours. After describing a method for optimising the day-ahead schedule for the boiler, the following sub-question in this research becomes relevant:

R2. How to mitigate wind power forecast errors during intra-day operation of the heat source?

a) *How does the performance scale with the control horizon?*

Similarly, the capacity of the boiler is expected to affect the extent to which the boiler can be used for demand response purposes. If the boiler is too large, the loads to shift might have a too great of an impact on the heat network, in turn limiting the boiler load shifting capacity. When the boiler power is too low, the boiler would have to be turned on all the time to satisfy the heat demand throughout the network, leaving no possibilities to shift any load.

R2. b) *How does the performance scale with the boiler capacity?*

To conclude, the objective is to minimise operating costs of the heat network and wind farm by means of a flexible electricity demand resulting from the flexibility in heat demand enabled by the thermal inertia of the thermal system. The heat network would thereby respond to wind power forecast signals and should hence alleviate the intermittency problem in the power grid. As stated before, the predictions at the base of the optimal control degrade with time, because of changing energy prices and weather forecasts. Therefore, after the day-ahead schedule is optimised, the formulated problem calls for a receding horizon control strategy.

1.3 Report structure

An overview of the structure of the report will now be presented. The report can be read along two different lines. The first line relates to the first research question (R1). The second line relates to R2. These lines intersect at various points. The methodologies for R1 and R2 are provided in Sections 3 and 4, respectively. Both methodologies rely on a model of the heat network, which is why this model is described first. The description of this model is made into a separate section, Section 2. To mathematically model a heat network, the relevant physics are first to be explored in Section 2.1. Section 2 then continues to describe the different components of the heat network model in Sections 2.2 and 2.3.1. This model will be used for various simulations with the aim of providing answers to the research questions.

Section 3 provides the methodology necessary for answering R1. First, the functioning of electricity markets is explored in Section 3.1.1 to provide a more detailed context for power exchanges. Different optimisation techniques are discussed in Section 3.1.2, which will lead to a proper choice for the optimisation tool in the optimisation problem. The mathematical model of the heat network described in Section 2 is used to formulate the optimisation problem in Section 3.2.1. Section 3.3 describes the benchmark with which the solution to the optimisation problem will be compared. Section 5.2 presents the results of the optimisation that could answer R1, including a comparison with the benchmark (in Section 5.2.2).

Section 4 provides the methodology for answering R2. Section 4.1.1 gives a background on wind power forecasting and discusses a number of statistical models that can be used for wind power forecasting. These statistical models are used in Section 4.2 to simulate wind power production profiles and forecast updates. The background on model predictive control supplied by Section 4.1.2 forms the basis for the formulation of the optimisation problem in Section 4.3, which is used in answering R2.

Section 5.2 shows the results of the experiments related to R1, while Section 5.3 shows the results related to R2. Throughout the results section, Section 5, adjustments to the optimisation algorithms are proposed and the outcomes are explored and compared to each other. Brief deliberations are discussed and some reserved conclusions are drawn. A more detailed discussion of the results and their implications is provided in Section 6. This section ends with a number of ways the experiments could have been conducted differently. The conclusion in Section 7 summarises the main findings and relates them to a broader context.

2 Modelling a heat network

Figure 5 shows a heat network with greenhouses connected in parallel. This set of heat consumers is denoted by $\mathcal{S} = \{1, \dots, N\}$ with N the total number of consumers. The system consists of an uncontrollable heat source, the data centre. This heat source (WH) injects heat into the network using a heat pump. A controllable electric boiler (EB) can further raise the temperature of the water before a mass flow splits off and enters the substation of the first greenhouse. The greenhouse can extract heat from the network through a heat exchanger in the substation. The network pipeline then continues to lead the water to other heat consumers, numbered $i \in \mathcal{S}$. Water is the energy carrier and flows through the network pipeline unidirectionally. The mass flow rate of the water is ensured by a central pump and, depending on the scale and configuration of the network, decentralised pumps. The heat network is operated under a constant mass flow and variable temperature control strategy, resulting in a linear heat network model [47], which will be discussed in this section. This means that throughout the system, the mass flow rate has a constant distribution, while temperatures may vary over time and location.

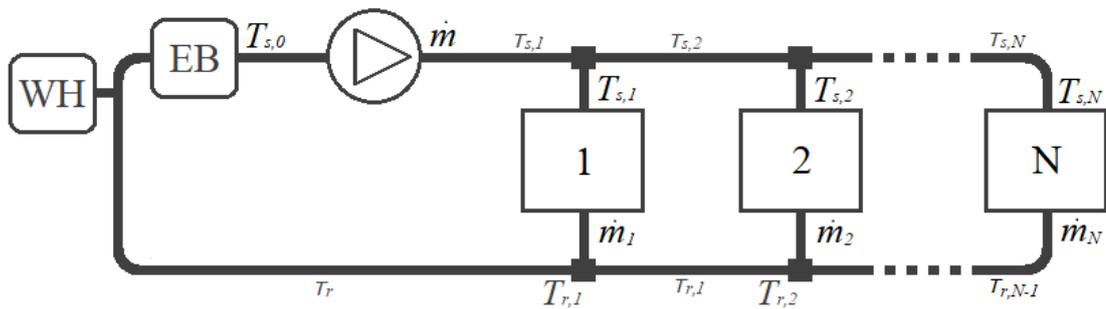


Figure 5: A waste heat source, electric boiler, time delays and temperatures at the consumer nodes

After leaving the substation at the greenhouse, each water flow immediately enters the return pipeline. Where the water flow from the greenhouse joins the already present water flow in the pipeline, a mixing of temperatures occurs. The return pipe leads the low temperature water back to the waste heat source, the data centre, where its temperature is raised again by the waste heat injected through the exchanger. Throughout the heat network, heat losses arise due to a temperature difference between the water in the pipeline and the ambient.

2.1 Background on heat networks

The following steps describe the building blocks for modelling a thermal grid. To simplify the model in terms of mass flow rates and pressure differentials, a constant mass flow and variable temperature control strategy is assumed. Pumps and valves operate independently from other network components to ensure a stable mass flow rate throughout the network. This way, time delays between network nodes are constant.

Assumption 3. *Constant mass flow rates throughout the heat network are maintained.*

2.1.1 Temperature dynamics and the node model

To analyse the temperature dynamics of a mass or volume, one takes a look at the heat balance. This balance considers the heat gains, \dot{Q}_{in} , and heat losses, \dot{Q}_{out} , and calculates the resulting change in temperature per time unit, using the heat capacity, C_w , of the volume (1).

$$C_w \cdot \frac{dT}{dt} = \dot{Q}_{in} - \dot{Q}_{out} \quad (1)$$

When the temperature of the volume, T_w , is higher than the outside temperature, T_{out} , and the outgoing heat is only in the form of heat losses to the environment, the outgoing heat can be calculated using the heat transfer factor, $c_{\dot{q}}$, in Equation (2).

$$\dot{Q}_{out} = c_{\dot{q}} \cdot (T_w - T_{out}) \quad (2)$$

Now, consider a volume of water moving along a line from one node to another (Figure 6). A certain mass flow rate, \dot{m} , is accompanied by a time delay, τ , between the nodes.

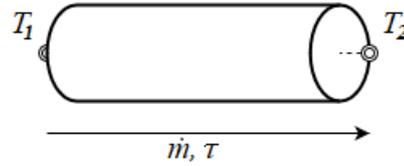


Figure 6: A mass flow from node 1 to node 2

The temperature at node 2, T_2 , depends on the influx of water from node 1. Neglecting heat losses, this temperature at time t can be calculated by Equation (3).

$$T_2(t) = T_1(t - \tau) \quad (3)$$

Taking heat losses into account means adding a temperature loss term (4).

$$\begin{aligned} T_2(t) &= T_1(t - \tau) - \Delta T \\ \Delta T &= \frac{\dot{Q}_{out}}{C_w} \\ \dot{Q}_{out} &= c_{\dot{q}} \cdot (T_1(t - \tau) - T_{out}) \end{aligned} \quad (4)$$

Rewriting Equation (4) results in Equation (5) for $\tau=1$.

$$T_2(t) = T_1(t - 1) \cdot \left(1 - \frac{c_{\dot{q}}}{C_w}\right) + \frac{c_{\dot{q}}}{C_w} \cdot T_{out} \quad (5)$$

Additionally, when multiple flows enter the same node, the temperature of that node is the mass flow weighted average of the different flow temperatures (6).

$$T_{node} = \frac{\sum_i \dot{m}_i T_i}{\sum_i \dot{m}_i} \quad (6)$$

2.1.2 Heat transport

Any axial heat transmission is negligible compared to the heat transport by the axial fluid transport [48, p. 77]. This is indicated by the Péclet number, Pe . The Péclet number is defined by the ratio between the advective heat transfer and the diffusive heat transfer. The first law of thermodynamics and the resulting energy equation (7) will be used for clarification [49].

$$\rho \left[\frac{\partial h}{\partial t} + \nabla \cdot (h\vec{v}) \right] = -\frac{Dp}{Dt} + \nabla \cdot (\kappa \nabla T) + (\vec{\tau} \cdot \nabla) \vec{v} \quad (7)$$

Enthalpy, h , temperature, T , and internal energy, e , are related to each other according to Equations (8) to (10).

$$h = e + \frac{p}{\rho} \quad (8)$$

$$dh = c_p dT \quad (9)$$

$$de = c_v dT \quad (10)$$

For incompressible flows, fluid density is constant and the heat capacity for an isobaric process, c_p , equals the heat capacity for an isochoric process, c_v . Given the relation between enthalpy and temperature, Equation (7) can be reduced to Equation (11). Figure 7 is a representation of this energy balance.

$$\underbrace{\rho c_p \frac{\partial T}{\partial t}}_{\text{internal energy gains}} + \underbrace{\rho c_p (\vec{v} \cdot \nabla) T}_{\text{advection}} = \underbrace{\kappa \nabla^2 T}_{\text{diffusion}} + \underbrace{(\vec{\tau} \cdot \nabla) \vec{v}}_{\text{dissipation}} \quad (11)$$

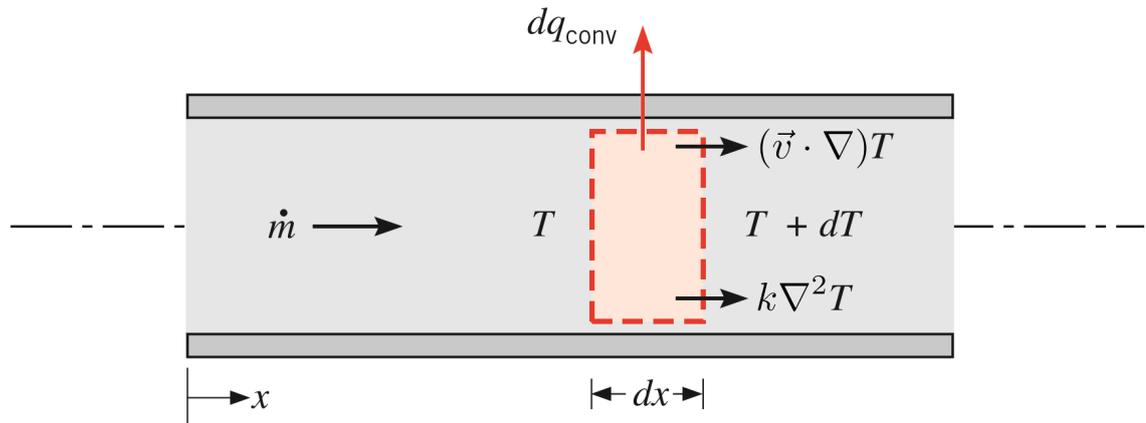


Figure 7: Energy balance in a pipe flow control volume, adapted from Bergman et al. [49, p. 529]

Calculating the temperature in a flow through a pipe, only one dimension of the partial differential equation is considered. Therefore, the dissipation term disappears.

$$\rho c_p \left[\frac{\partial T}{\partial t} + \vec{v} \cdot \frac{\partial T}{\partial x} \right] = \kappa \frac{\partial^2 T}{\partial x^2}$$

Equation (12) can calculate the Péclet number for a pipe diameter of .5 m, a flow velocity of 2 m·s⁻¹, a density of 990 kg·m⁻³, a specific heat capacity of 4180 J·kg⁻¹·K⁻¹ and a thermal conductivity of .6 W·m⁻¹·K⁻¹. Because $Pe \gg 1$, heat convection and conduction in the dimension along which the fluid flows, will be ignored. Nonetheless, a radial heat transmission through the pipe does lead to energy losses.

$$Pe = \frac{Lv\rho c_p}{\kappa} \quad (12)$$

Assumption 4. Axial heat transmission can be ignored.

2.2 Multi-consumer heat network

With Figure 5 in mind, the general case of a number of consumers, N , in a heat network can now be considered to derive a set of equations describing the model. At all times, the principle of mass conservation must be satisfied, which implies the distribution of the mass flows through the substations follows Equation (13).

$$\dot{m} = \sum_{i \in \mathcal{S}} \dot{m}_i \quad (13)$$

The different components in the heat network are located at a distance away from each other, represented by a number of time steps, τ . This number of time steps translates into a certain time delay and depends on the water flow velocity and the length of the pipeline between the two nodes in the network. Table 2 contains an overview of the relevant time delay symbols as used in the formulation of the physical model.

Table 2: Network time delays

From	To	
Electric boiler	Consumer 1	$\tau_{s,1}$
Separation node at consumer $i-1$	Consumer i	$\tau_{s,i}$
Consumer $i+1$	Connection node at consumer i	$\tau_{r,i}$
Connection node at consumer 1	Waste heat source	τ_r

Some distances in the heat network are much greater relative to others. For example, a time delay between consumers is much larger compared to the time it takes the water to flow from the supply pipe, through the consumer substation and into the return pipe. Furthermore, each value for a time delay, τ , is rounded to the nearest integer. Relatively small distances can therefore be neglected.

Assumption 5. *Each time delay is a multiple of the time step size.*

Assumption 6. *The pipeline length toward, through and from a consumer substation is zero.*

Assumption 7. *There is no distance between the data centre and the electric boiler.*

Sections 2.2.1 to 2.2.3 discuss the discretised formulas describing the system dynamics for the heat network in Figure 5 and will be discussed one by one. Equation (14) is an expression for the supply temperature at consumer i , $T_{s,i}$. Equation (15) shows an expression for the temperature at different points in the return pipe, $T_{r,i}$. Equation (17) describes the temperature of the water in the supply pipe right after leaving the electric boiler, $T_{s,0}$.

2.2.1 Supply pipe temperature

The supply temperature at a certain consumer i at time t , $T_{s,i}(t)$, does not depend on the supply temperature at that exact location of one time step ago, at time $t-1$, because that water has already left the substation. Rather, to calculate the supply temperature, one must turn upstream; the supply temperature depends on the upstream temperature of a few time steps ago. To be precise, the previously established temperature that is of influence on the current temperature, is the temperature measured lastly at the upstream agent in the system, which is located at a distance represented by the number of time steps it takes the water in the pipelines to flow from that agent to the consumer. Furthermore, the supply temperature at a consumer, $T_{s,i}$ in Equation (14), is dependent on the heat loss in the pipeline, which depends on the supply temperature for the upstream consumer, $T_{s,i-1}$, and the ambient temperature, T_a . However, for the first consumer, $i=1$, the supply temperature at the consumer depends on the temperature at which the water left the upstream heat sources, $T_{s,0}$. As can be deduced from Figure 5, the upstream heat sources are located at a number of time steps away from the consumer. To calculate the supply temperature at the consumer, $T_{s,i}$, the appropriate heat loss factor, η , and time delay, τ , are $\eta_{s,i}$ and $\tau_{s,i}$ from Table 2. A derivation for Equation (14) can be found in Appendix B.2 and is useful to gain an understanding of the pipeline heat loss calculation method used in the following equations.

$$T_{s,i}(t) = (1 - \eta)^\tau \cdot T_{s,i-1}(t - \tau) + \eta \int_{t-\tau}^t (1 - \eta)^{t-\vartheta} \cdot T_a(\vartheta) d\vartheta \quad (14)$$

As explained before, in a multi-consumer network the water flow in the heat supply pipeline splits off at every consumer, only to converge again in the return pipe. The points at which the mass flow rate in the supply pipe decreases due to the separation of a water flow to a consumer, are shown in Figure 5, along with the points at which the mass flow from a consumer joins the mass flow in the return pipe, which are called the return pipe connection nodes. The division of the mass flow over the consumers has an implication for the calculation of the supply temperatures in Equation (14). After all, if the remaining mass flow in the supply pipe suffers the same heat loss over the same time period, it would lead to a larger temperature drop, because a smaller mass leads to a smaller heat capacity.

2.2.2 Return pipe temperature

As can be seen in Figure 5, the temperature in the return pipe at consumer 1, $T_{r,1}$, depends not only on the water temperature of the mass flow of water through through substation 1, \dot{m}_1 , after leaving the substation, but also on the temperature of the water already in the return pipe with mass flow rate $\dot{m}_{r,1}$. The latter temperature is in turn dependent on $T_{r,2}$ and so forth. Ultimately, the temperature at any return pipe connection node $T_{r,i}$ is calculated as would any other temperature mixing of two water flows under the conservation of energy, by a summation of the product of temperature and mass flow, divided by the combined mass flow [50]:

$$T_{x+y} = \frac{\dot{m}_x T_x + \dot{m}_y T_y}{\dot{m}_x + \dot{m}_y}$$

In general, $T_{r,i}$ is then a value taking into account the temperature mixing by the different temperatures and mass flow rates leaving the various substations, as shown in Equation (15). A glance at Figure 5 would remind that, for the consumer furthest from the boiler ($i=N$), temperature mixing is irrelevant and thus Equation (15) does not apply. Appendix B.2 provides a derivation, taking $\eta_{r,i}$ and $\tau_{r,i}$ as appropriate indicators of the heat loss factor and the time delay, respectively.

$$T_{r,i}(t) = \frac{\overbrace{\dot{m}_i \cdot T_{s,i}(t) - \frac{\dot{Q}_i(t)}{c_w}}^{\text{substation } i \text{ component}} + \left[(1-\eta)^\tau \cdot T_{r,i+1}(t-\tau) + \eta \int_{t-\tau_{r,i}}^t (1-\eta)^{t-\vartheta} \cdot T_a(\vartheta) d\vartheta \right] \cdot \dot{m}_{r,i}}{\dot{m}_i + \dot{m}_{r,i}} \quad (15)$$

For the mass flow rate \dot{m}_i from substation i no further calculation is necessary. The summation of the mass flow rates upstream from i — relevant for the temperature mixing in the return pipe at i — is indicated by $\dot{m}_{r,i}$ and calculated by Equation (16).

$$\dot{m}_{r,i} = \sum_{j=i+1}^N \dot{m}_j \quad \forall i \in \{1, \dots, N-1\} \quad (16)$$

2.2.3 Temperature after the boiler

The temperature leaving the electric boiler, $T_{s,0}$ (see Figure 5), is calculated in Equation (17). Here, the temperature in the final return pipe connection node, $T_{r,1}$ — where the mass flow of water through substation 1, \dot{m}_1 , enters the return pipe — represents the sum of temperatures leaving the various consumer substations. The data centre adds a time-varying amount of waste heat, $\dot{Q}_{wh}(t)$, to the network water in the pipeline. Furthermore, when the boiler is turned on, it will raise the temperature at that location instantaneously, with heat of a power P_b added to the water in the supply pipe. However, the water, and thus the heat, does not stay at that location. It takes one complete circulation for the water to return. Again in Appendix B.2, a derivation for Equation (17) is provided, which is greatly similar to the derivation for Equation (14), except for the inclusion of waste heat from the data centre and the electric boiler. It further includes the appropriate heat loss factor, η_r , and time delay, τ_r .

$$T_{s,0}(t) = \frac{\overbrace{\dot{Q}_{wh}(t) + P(t)}^{\text{heat gain from sources}}}{c_w \dot{m}} + (1-\eta)^\tau \cdot T_{r,1}(t-\tau) + \eta \int_{t-\tau}^t (1-\eta)^{t-\vartheta} \cdot T_a(\vartheta) d\vartheta \quad (17)$$

2.3 Heat network components

The greenhouses and the data centre are the components of the heat network that remain to be discussed. Section 2.3.1 describes the model of a greenhouse as a consumer of heat and Section 2.3.2 describes the model of the data centre as waste heat source in the heat network.

2.3.1 Single consumer model

The consumer model consists of a water volume at greenhouse $i \in \mathcal{S}$ connected through a heat exchanger to a heat network. The time period over which the dynamics of the consumer is modelled, is denoted by $\mathcal{T} = \{1, \dots, d\}$. The greenhouse extracts heat to regulate the temperature of its volume. It experiences heat losses due to a temperature difference between the ambient, T_a , and the volume the greenhouse needs to keep between certain temperature limits. These heat losses cause the greenhouse to have a heat demand. The amount of heat extracted from the heat network, or the heat flow through the heat exchanger, at time step $t \in \mathcal{T}$ is either zero or a fixed quantity, $\dot{Q}_{hx,i}$.

Assumption 8. *The heat flow through an exchanger at a substation is either nominal or zero.*

Figure 8 depicts the substation at greenhouse i . The counter current configuration of the heat exchanger is aimed at optimising system efficiency by maintaining a temperature difference along the entire heat exchanger [51]. Hot water entering the heat exchanger first heats water that has already been warmed up and is about to leave the heat exchanger. The source side of the heat exchanger cools down as the water progresses through the heat exchanger, while still transferring heat, as the temperature on the receiving side is lower closer to the cold water inlet. Deriving an expression for the temperature of the water volume on the greenhouse side of the heat exchanger, T_i , demands looking at the temperature dynamics of the system in Equation (18).

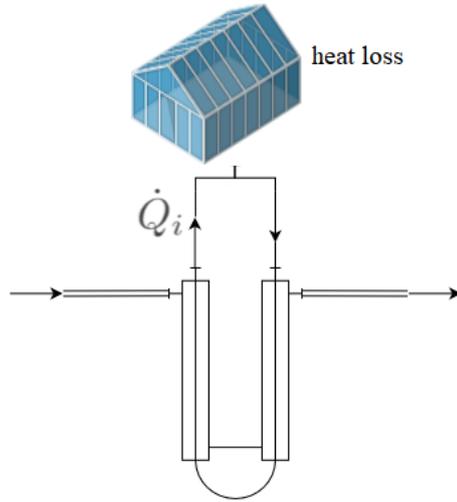


Figure 8: A heat flow through an exchanger with a loss on the side of the greenhouse

$$c_w m_i \frac{dT_i}{dt} = \dot{Q}_i - (T_i - T_a) c_{q,i} \quad (18)$$

The left side of Equation (18) represents the temperature change in terms of the heat gain; the right side shows the heat balance. When the heat gain, \dot{Q}_i , outweighs the heat loss to the ambient, the heat balance is positive and the temperature rises. A negative heat balance leads to a decrease in temperature. A discretisation of Equation (18) leads to Equation (19). This derivation is found in Appendix B.1.

$$T_i(t) = T_i(t-1) + \dot{Q}_i(t) \frac{\Delta t}{c_w m_i} - [T_i(t-1) - T_a(t-1)] \frac{c_{q,i} \Delta t}{c_w m_i} \quad (19)$$

Lastly, it must be noted that the heat flow through the exchanger in the greenhouse substation results in a decrease in the network pipeline temperature in Equation (17), while resulting in a heat gain for the consumer in Equation (19).

2.3.2 Data centre as waste heat source

Up to 40 percent of the energy consumption in data centres constitute the use of cooling systems, as virtually all electric power used in a data centre is sooner or later converted into heat [52]. High temperature build-up in data centres causes the most frequent form of component failure. Typically, the upper temperature limit is about 358 K. Since a number of system types are liquid-cooled, data centres are potential contributors of waste heat to thermal grids. In fact, the growing power densities in data centres support the argument for liquid-cooled systems instead of air-cooled ones, due to the higher per unit heat transfer capacity of water. Furthermore, heat can be captured closer to the data centre's processing units, resulting in waste heat captured at 323 to 333 K in liquid-cooled systems, compared to a 298 to 308 K temperature range for air-cooled systems [53]. Still, the heat carried by the return water could not be directly transferred to a heat network, of which the supply temperature is at least 323 K. A heat pump would be required for effective contribution of heat to the supply water of a heat network.

As described in Section 1.2, the case to be studied in this report includes a data centre as a waste heat source. Directly upstream from the electric boiler, the data centre injects its excess heat into the pipeline. It is assumed the data centre uses a constant amount of electric power, P_{dc} .

Assumption 9. *A constant amount of heat is produced inside the data centre.*

The difference between the ambient temperature, T_a , and the temperature desired inside the data centre, T_{dc} , causes an exchange of heat between the data centre and the ambient. This heat flow depends on a heat transfer factor, $c_{\dot{q},dc}$. Equation (20) calculates the amount of waste heat injected into the heat network, \dot{Q}_{wh} , per time step t .

$$\dot{Q}_{wh}(t) = P_{dc} - c_{\dot{q},dc} \cdot (T_{dc} - T_a(t)) \quad (20)$$

3 Network operation schedule

This section of the report presents the background and methodology for answering the first research question (R1). Section 3.1 provides a background on electricity markets and optimisation techniques. Section 3.2 formulates the optimisation problem relating to R1. In Section 3.3, a benchmark is described to which the result of the optimised schedule will be compared in Section 5.2.2.

3.1 Background

Because the present problem concerns a cost-minimisation, the relevant markets and prices are discussed in Section 3.1.1. For solving the optimisation problem, the different types of problems and their respective solution techniques are described in Section 3.1.2.

3.1.1 Electricity markets

Conventionally, electricity is generated centrally in power systems, while the aggregate load is fairly predictable. Therefore, the generator unit commitment problem is approached as deterministic [41]. In a future sustainable energy system, both electricity demand and supply depend on weather characteristics. For example, temperature influences demand while wind speed and solar irradiance affect the supply of electricity. Furthermore, a balanced demand and supply is essential for power system stability. With an increasing share of renewable energy sources, operation of the power system becomes more stochastic. The consequent price volatility incentivises research into electricity price forecasting. Weron [54] gives an elaborate overview of forecasting techniques, with the main focus on day-ahead electricity prices.

Electricity price characteristics are mainly determined by — from larger to smaller time scale — the season, a reversion to the mean (short-term extremes are levelled out toward long-term trends), price-dependent volatility (e.g. fuel prices, or peaks due to congestion) and correlation to electricity load [41].

Day-ahead market

At the day-ahead market, electricity is sold by power producers, to be delivered at each delivery time the next day. Participants submit offers and bids for electricity for each load period of the next day. The day-ahead market does not allow for continuous trading. With continuous trading, an offer is accepted as soon as there is a high enough bid and a bid is accepted as soon as there is a low enough offer. Instead, on the day-ahead market the bids and offers are collected until all trades are executed in a batch. For the Nord Pool day-ahead market, buyers and sellers have until noon to submit their final bids to Nord Pool for the auction for delivery hours the next day [55]. At noon, the market closes. After the market closing time, when the offers and bids are known, buyers' bids are accepted in descending order, while sellers' offers are accepted in ascending order, until demand and supply meet in terms of quantity and price. This way, the market clearing price is established. All energy is then traded at the market clearing price. As a result, no party sells energy at a price lower than its offer and no party buys energy at a price higher than its bid. In contrast to this uniform-price auction, a pay-as-bid auction pays the energy supplier the price offered by the supplier.⁴

Intra-day market

It happens that on the delivery day, electricity buyers realise their power demand is higher than the power they bought at the day-ahead market. Similarly, it can happen that power producers do not manage to supply the power they sold on the day-ahead market, e.g. due to unexpected failure in the power plant. Changes in electricity demand or supply can be supported by power exchanges at the intra-day market. With an increasing amount of intermittent renewable energy sources, it becomes increasingly challenging for market participants to be in balance in terms of their supply or demand after the day-ahead market closing time. As a result, interest in intra-day trading is growing and the intra-day market is becoming essential in allowing power producers and consumers to handle unexpected outages and changes in demand. In contrast to the day-ahead market, the intra-day market does allow for continuous trading, which means an offer is accepted as soon as there is a high enough bid and a bid is accepted as soon as there is a low enough offer.

⁴In case of transmission congestion, the location of a load or power production unit plays a role in the costs. These costs must be paid. Locational marginal prices thus differ from the market clearing price described before. These locational prices may increase the complexity of the market.

In real-time pricing, electricity prices for the load periods adjust continuously to external conditions or supply and demand dynamics. Here, real-time does not mean instantaneous. It means as quickly as possible; the flow of information determining prices is not purposefully slowed down but instead guaranteed within a certain time delay. Therefore, real-time pricing applies not merely to the intra-day market, but just as much to the day-ahead market. Oppositely, with time-of-use pricing or peak-pricing, electricity consumers know the price they pay for electricity for a certain load period a day or more in advance.

The need for price forecasting

The electricity price forecasting considered in day-ahead markets, is that from hours to a few days ahead. Additionally, this forecasting horizon is of a time scale similar to the flexibility domain of various energy loads, e.g. the heat demand of buildings. The effectiveness of demand response management depends greatly on the electricity price forecasting at these intervals by the network operators and power market participants. Hong discusses in a number of papers the various domains of forecasting in the energy industry [56], [57]: renewable power generation, electric load forecasting, and electricity price forecasting. The first two domains greatly affect the third domain, along with e.g. transmission congestion and market participant behaviour, which are hard to model. That is why electricity price forecasting has thus far been largely data-driven and performed using hybrid models, because, although wind power forecasting and short-term load forecasting are experienced in their respective use of probabilistic and point forecasting techniques, no such maturity exists for any electricity price forecasting approach.

3.1.2 Optimisation techniques

A static optimisation is an allocation of resources at only one instant in time, with the possibility of allocating resources for multiple time periods during that one-time decision. The allocation consists of choosing values for the decision variables, the number of which is called the number of degrees of freedom, from a set of values in the mathematically feasible region, to optimise a predetermined objective function.

Within the class of static optimisation a categorisation can be made, based on the objective, constraints and variables. Table 3 shows an overview of the different types of optimisation problems. The objective and constraints can be linear or non-linear and variables can be either continuous or discrete. These factors ultimately determine the method used for solving an optimisation problem.

Table 3: Types of optimisation problems

Objective	Constraints	Variables	Type
Linear	Linear	Continuous	Linear Programming
Linear	Linear	Continuous and discrete	Mixed Integer Linear Programming
Quadratic	Linear	Continuous	Quadratic Programming
Non-linear	Linear	Continuous	Non-Linear Programming
Non-linear	Linear	Continuous and discrete	Mixed Integer Non-Linear Programming
Non-linear	Non-linear	Continuous	Non-Linear Programming with Non-linear constraints
Non-linear	Non-linear	Continuous and discrete	Mixed Integer Non-Linear Programming with Non-linear constraints

A linear programming problem is an optimisation problem in which the minimum value of a linear expression, the objective function, subject to a number of linear constraints has to be found, along with the corresponding values for the decision variables. To solve a linear programming problem, the so-called feasible region must be found, which is the set of feasible combinations of variables that satisfy all constraints. Figure 9 shows the graphical method for finding the optimal solution. This optimal solution is then a point in the feasible region that gives the optimal value for the objective function. The corners of the feasible region in Figure 9 are called extreme points. The optimal solution for a linear programming problem can be found in one of these extreme points. It must be noted that the graphical method is only practical for problems with two decision variables. The simplex method solves linear programming problems, including those with more than two decision variables. Firstly, an initial extreme point is identified. Then, the algorithm looks from this initial extreme point to the next extreme point along each edge. For all neighbouring

extreme points the value of the objective function is calculated. The algorithm repeats this step using the extreme point with the optimal function value as the new initial extreme point. If the function values do not improve any longer, the extreme point with the optimal value thus far is the optimal solution.

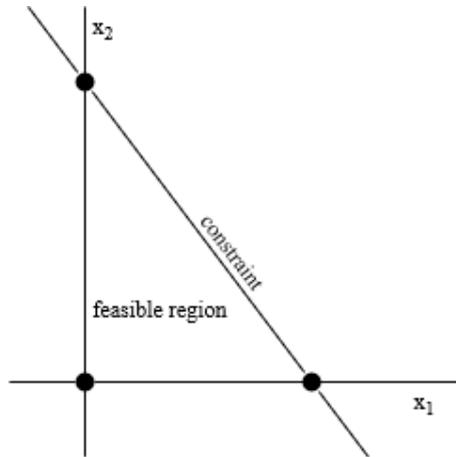


Figure 9: An example of a feasible region in a linear programming problem with $x \geq 0$

A simplified way to solve an integer linear programming problem is to solve it as a linear programming problem — this is called relaxation — and to round the decision values of the solution to the nearest integer value. Nonetheless, this solution might not be optimal and it might not even be feasible; it could fall outside the feasible region. The cutting-plane method repetitively refines the set of constraints by adding linear inequalities, thereby cutting out the non-integer solution to the linear relaxation of the integer problem.

The branch-and-bound approach is a class of methods for linear and nonlinear (mixed) integer programming. The central idea is to draw up a decision tree. As it is extremely time consuming to go through all the possibilities, the branches which clearly show at any early stage that they do not contain the optimal solution will be eliminated (bound). The more branches that are bound, the quicker the solution will crystallize. Branch and bound starts by solving the linear relaxation problem (see Figure 10). If all discrete variables have an integer value, the problem is solved. If any variable has a fractional value, the algorithm chooses for each fractional value a branch; first zero, then one. If the objective function value is better than the optimal value found thus far, it replaces this previously best value. If either sub-problem is unfeasible, the branch is bound and will not be further investigated. Otherwise, another fractional variable is to be found and the steps are to be repeated. Branch and bound has the advantage over the cutting-plane method that it can be stopped early as long as it has found a feasible integer solution without it necessarily being the optimal solution. Furthermore, each feasible solution can be directly compared to the linear optimum.

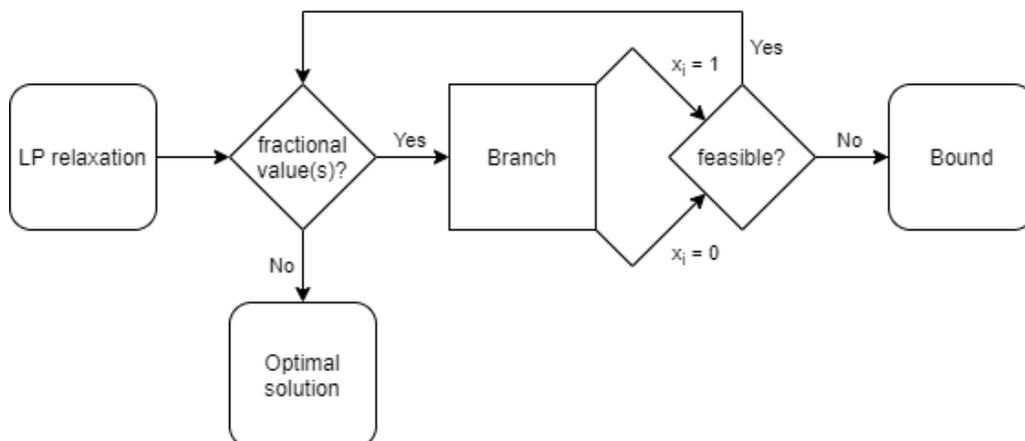


Figure 10: Schematic representation of the branch-and-bound algorithm

3.2 Problem formulation

The scheduling problem is approached as a static optimisation problem, because buying electricity at day-ahead price levels is a one-time decision. In this case, the decision variables are those for the boiler and substations in the heat network in Figure 11 (shown earlier by Figure 5). Because the cost of operating the boiler for an hour is linearly dependent on the power used by the boiler and on the hourly electricity price, the objective function is linear. The model describing the heat network is a linearisation; the resulting constraints to the optimisation problem are therefore linear. The decisions to be made are assumed to be of a Boolean or binary nature. For example, the electric boiler is either on or off. This can be included in the optimisation problem using (mixed) integer linear programming, thereby excluding the linear solutions that do not coincide with an integer solution.

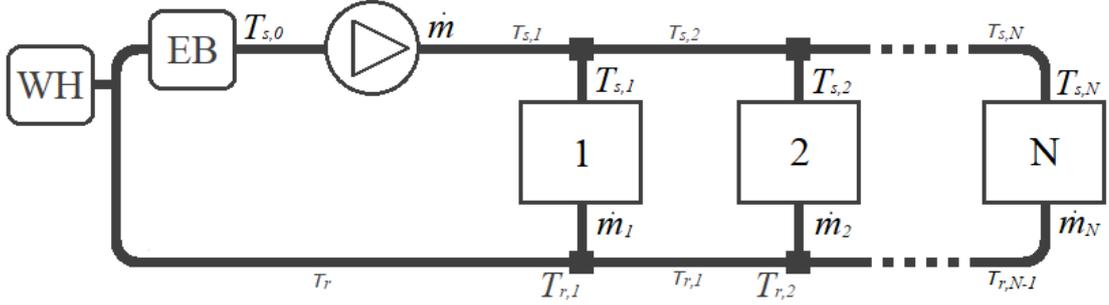


Figure 11: A data centre, electric boiler, time delays and temperatures at the greenhouse nodes

3.2.1 Cost function and control variables

Operating the boiler and the heat exchangers is to be optimised in a cost-minimising way. The electric boiler can be found upstream from the two greenhouses, which is either turned on or off at time t , represented by the unit commitment variable u_b , resulting in a power P_b added to the pipeline water flow when the boiler is turned on. For the boiler power at time t a price $\lambda(t)$ is paid, while the cost for heat extraction through the heat exchangers is neglected. The resulting cost function to be minimised can be seen in Equation (21).

Assumption 10. *The electric boiler adds either a fixed amount of power to the heat network supply water or no power at all.*

$$C(\mathbf{P}, \dot{\mathbf{Q}}) = \sum_{t \in \mathcal{T}} \left[P(t) \cdot \lambda(t) + \sum_{i \in \mathcal{S}} \dot{Q}_i(t) \cdot 0 \right] \quad (21)$$

The degrees of freedom in this optimisation problem are represented by matrices \mathbf{P} and $\dot{\mathbf{Q}}$, and consist of the control variables for the boiler, u_b , and heat exchanger i , u_i , at every $t \in \mathcal{T}$, as shown in Equations (22) and (23).

$$P(t) = [u_b(1) \quad \cdots \quad u_b(d)] \cdot P_b \quad (22)$$

$$\dot{Q}_i(t) = \begin{bmatrix} u_i(1) & \cdots & u_i(d) \\ \vdots & \ddots & \vdots \\ u_N(1) & \cdots & u_N(d) \end{bmatrix} [\dot{Q}_{hx,1} \quad \cdots \quad \dot{Q}_{hx,N}] \quad (23)$$

Each control variable can take a value of either zero or one (24), which represents a true or false statement of whether heat is added by the boiler or extracted by a greenhouse through the heat exchanger.

$$u_b(t), u_i(t) \in \{0, 1\} \quad \forall i \in \mathcal{S}, \quad \forall t \in \mathcal{T} \quad (24)$$

3.2.2 Constraints

The system as a whole is subject to a number of constraints, by which the boiler operation is to be optimised. These constraints relate to the greenhouse temperature demands and the physical conditions necessary for a proper functioning of the network. The constraints presented here coincide with the conditions used in the rule-based operation presented later on in Section 3.3. They will now be discussed in more detail.

Consumer constraints

The system is subject to a number of temperature inequality constraints. Equation (25) ensures the temperature of the water volume on the greenhouse side of heat exchanger i is within the predefined temperature limits, T_{min} and T_{max} , at all times.

$$T_{min,i} \leq T_i(t) \leq T_{max,i} \quad \forall i \in \mathcal{S}, \quad \forall t \in \mathcal{T} \quad (25)$$

The constraint in Equation (26) is present to raise the water volume temperature of greenhouse i at the end of period \mathcal{T} to at least its initial temperature, T_0 . This is to prevent the temperature from falling to the defined minimum temperature at the end of the time period. After all, that might not be the optimal path to take long-term, considering the start of the following time period.

$$T_{0,i} \leq T_i(d) \quad \forall i \in \mathcal{S} \quad (26)$$

Network constraints

Another temperature inequality constraint of the system is a minimum temperature requirement, $T_{min,s}$, for the supply water in the pipeline before entering the substation at greenhouse i to ensure a sufficient flow of heat through the exchanger when desired, shown in Equation (27). Equation (14) — or (71) in Appendix B.2 — calculates $T_{s,i}$.

Assumption 11. *Whenever a consumer wants to extract heat from the network through a heat exchanger, the heat network supply water must be at least at a predefined minimum temperature upon entering the heat exchanger.*

$$T_{s,i}(t) \geq u_i(t) \cdot T_{min,s,i} \quad \forall i \in \mathcal{S}, \quad t \in \mathcal{T} \quad (27)$$

To illustrate, Figure 5 shows the electric boiler with power P_b that operates to satisfy the minimum supply temperature constraint. By inserting u_i into Equation (27), a minimum temperature is only necessary whenever heat is about to be extracted. Namely, u_i equals zero when no heat is extracted and the minimum temperature then becomes zero, too. Naturally, one would be quick to think that defining the constraint in such a way would result in a boiler being activated only at times when heat is actually extracted. Nonetheless, this will not always be true. The boiler activity is optimised to minimise costs. When the electricity price is low, but the consumer does not intend to extract heat from the network, the boiler can still be turned on for the heat from the boiler to be stored in the pipeline of the network for later use, when electricity prices will be higher.

3.2.3 Optimisation problem

The constraints and all control variables have been established. The optimisation problem can now be formulated in full (28). The cost-minimising objective function consists of the cost vector, c , and the control vector, \mathbf{u} . As stated beforehand, the only relevant cost here is the cost for boiler operation at each time step. Therefore, all values in the cost vector relating to control variables other than those for the electric boiler, are set to zero. The equality constraint describes the system dynamics using matrix A and vector b , which, together with the control vector \mathbf{u} , are constructed by Equations (14), (15), (17) and (19). The inequality constraints have already been discussed in Section 3.2.2. The day-ahead schedule optimisation problem can be implemented in MATLAB with Scripts 1 to 4 in Appendix C.2 and, for scheduling consecutive days using optimal control, Script 5 should be added. The schedule for each day is obtained using a branch-and-bound solver for the mixed integer linear programming, because it is preferred to obtain at least a feasible solution within a predetermined maximum runtime rather than obtaining the optimal solution using the cutting-plane method.

$$\begin{aligned}
& \min_{\mathbf{u}} && c^T \mathbf{u} \\
& \text{s.t.} && A\mathbf{u} = b \\
& && u_i(t) \cdot T_{min,s,i} \leq T_{s,i}(t) \\
& && T_{0,i} \leq T_i(d) \\
& && T_{min,i} \leq T_i(t) \leq T_{max,i} \\
& && u_b(t), u_i(t) \in \{0, 1\} \quad \forall i \in \mathcal{S}, \quad \forall t \in \mathcal{T}
\end{aligned} \tag{28}$$

$$c = [P_b \lambda(\mathcal{T}), 0, \dots, 0]$$

$$\mathbf{u} = [u_b(\mathcal{T}), u_I(\mathcal{T}), \dots, u_N(\mathcal{T})]$$

$T_i(t)$ is directly related to $\dot{\mathbf{Q}}$, as shown by Equation (19) in Section 2.3.1:

$$T_i(t) = T_i(t-1) + \dot{Q}_i(t) \frac{\Delta t}{c_w m_i} - [T_i(t-1) - T_a(t-1)] \frac{c_{q,i} \Delta t}{c_w m_i}$$

Although no cost is associated with the heat flow through the exchanger, \dot{Q}_i , the use of the heat exchanger can still be optimised due to its effect on the network temperatures, which are partially controlled by the electric boiler. This way, the greenhouse temperature T_i relates indirectly to the electric boiler unit commitment variable u_b .

Section 2.2.1 explains the equation to calculate the supply temperature at each greenhouse, $T_{s,i}$ in Equation (14):

$$T_{s,i}(t) = (1-\eta)^\tau \cdot T_{s,i-1}(t-\tau) + \eta \int_{t-\tau}^t (1-\eta)^{t-\vartheta} \cdot T_a(\vartheta) d\vartheta$$

Section 2.2.2 explains the equation to calculate the return temperature at each greenhouse, $T_{r,i}$ in Equation (15):

$$T_{r,i}(t) = \frac{\dot{m}_i \cdot T_{s,i}(t) - \frac{\dot{Q}_i(t)}{c_w} + \left[(1-\eta)^\tau \cdot T_{r,i+1}(t-\tau) + \eta \int_{t-\tau}^t (1-\eta)^{t-\vartheta} \cdot T_a(\vartheta) d\vartheta \right] \cdot \dot{m}_{r,i}}{\dot{m}_i + \dot{m}_{r,i}}$$

Section 2.2.3 explains the equation to calculate the supply temperature right after the boiler, $T_{s,0}$ in Equation (17):

$$T_{s,0}(t) = \frac{\dot{Q}_{wh}(t) + P(t)}{c_w \dot{m}} + (1-\eta)^\tau \cdot T_{r,I}(t-\tau) + \eta \int_{t-\tau}^t (1-\eta)^{t-\vartheta} \cdot T_a(\vartheta) d\vartheta$$

3.3 Benchmark: rule-based operation

In Section 1.2.1, a rule-based control strategy was introduced. A more detailed description will be given here.

Operating a heat network with power-to-heat entails some form of decision-making, such as when to turn on an electric boiler or whether to extract heat from the network at a substation. A rule-based control strategy defines certain rules or conditions. If a condition is met, an action is triggered. In the heat network, actions include turning on the electric boiler and extracting heat through a heat exchanger. The rule-based control algorithm is described by Algorithm 1. It can be implemented in MATLAB with Script 3 in Appendix C.2, while Scripts 1 and 2 function as parameter setting and calculation.

The conditions to be met consist of an upper and a lower temperature boundary for the temperature in the greenhouse. If the greenhouse temperature would exceed the upper temperature limit when extracting heat from the network at that time step, then no heat will be extracted from the network at that time step to prevent the temperature from crossing the upper boundary. Likewise, if the greenhouse temperature would drop below the lower temperature limit without extracting heat from the network at that time step, then heat will be extracted from the network at that time step to prevent the temperature from crossing the lower boundary. Otherwise, i.e. if the temperature would remain within the boundaries either way, the amount of heat extracted from the network is the same as the amount of one time step ago.

Algorithm 1: Rule-based control

```

for  $t \in \mathcal{T}$ 
  for  $i \in \mathcal{S}$ 
    if  $T_{max,i} < (1 - \eta) \cdot T_i(t - 1) + \dot{Q}_{hx,i} \frac{\Delta t}{c_w m_i} + \eta \cdot T_a(t)$ 
       $\dot{Q}_i(t) = 0$ 
    else if  $T_{min,i} > (1 - \eta) \cdot T_i(t - 1) + \eta \cdot T_a(t)$ 
       $\dot{Q}_i(t) = \dot{Q}_{hx,i}$ 
    else
       $\dot{Q}_i(t) = \dot{Q}_i(t - 1)$ 
    end
  end
  if  $T_{min,s} > \frac{\dot{Q}_{wh}(t)}{c_w \dot{m}} + (1 - \eta)^\tau \cdot T_{r,i}(t - \tau) + \eta \int_{t-\tau}^t (1 - \eta)^{t-\vartheta} \cdot T_a(\vartheta) d\vartheta$ 
     $P(t) = P_b$ 
  else
     $P(t) = 0$ 
  end
end

```

A similar condition is applied to deciding whether or not to turn on the electric boiler at time step t . The temperature of the supply water is calculated in case the boiler does not add heat to the water. This temperature is then compared against a predetermined minimum supply temperature. If the supply temperature without the boiler would be below the minimum supply temperature, the boiler is turned on. Otherwise, the boiler is turned off. The comparison of the supply temperature without the boiler against the minimum supply temperature occurs for every time step.

4 Smart network control

This section of the report provides a background and a methodology with the aim of answering the remaining research questions — R2, a) and b). Section 4.1 presents a background on statistical models used in wind power forecasting and a background on the concept of a receding horizon control. In Section 4.2, the model that simulates wind power production profiles and the model that simulates continuously updating wind power forecasts are described. Lastly, Section 4.3 formulates the optimisation problem relating to R2.

4.1 Background

In Section 4.1.1, a number of wind power forecasting methods are discussed, statistical models in particular. These models form the basis of the model described in Section 4.2, which is a simulation model simulating wind power forecasts. The concept of model predictive control with a receding horizon is introduced in Section 4.1.2.

4.1.1 Wind power forecasting

Wind power generation is stochastic in nature due to the uncertainty of wind [58]. Therefore, an increasing share of wind energy in the power system poses multiple problems to maintaining stability and reliability in the power system. Proper wind forecasts can alleviate these problems as they enable an improved planning of power system operation.

Classification of the time-scales related to the various wind forecasting methods is still ambiguous. The very-short term ranges from a few seconds to half an hour ahead. Short-term forecasting ranges from half an hour to six hours ahead, medium-term is defined as ranging from six hours to one day ahead and from one day to one week ahead fills the category of long-term forecasting [59]. According to Foley, Leahy, Marvuglia, *et al.* [58], short-term forecasts range from an hour to a few days and are used in the unit commitment and economic dispatch. Medium-term forecasts range from a few days to weeks and are used in maintenance planning and energy storage operations.

The relevant time horizon dictates the decision for a wind power prediction method. The accuracy of the prediction method determines the cost of increased system balancing due to wind power penetration into the power system. After all, wind power is traded at the day-ahead market along with the power from other electricity generation units. An error in the wind power prediction creates the need for balancing power to be purchased at the intra-day market, because the power offered the day before must be delivered one way or another. Interestingly, forecasting wind power generation during the intra-day market occurs using the same data, available at the same time, as for the day-ahead market.

Hereafter, a number of prediction methods will be discussed, with an emphasis on statistical methods. Artificial or computer intelligence methods are mentioned briefly now and will not be looked into any further, because they are not the main focus of this report while they require a greater amount of explaining. These machine learning approaches used in wind power forecasting include artificial neural networks, fuzzy logic systems and support vector machines.

Persistence method

The persistence method is also known by its rather descriptive name, ‘naïve predictor’. The method states that the forecast value at time $t+k$ equals the measured value at time t ; it makes use of the assumption that the net change in value equals zero (29). Surprisingly, the forecast errors remain on average significantly smaller for short-term to very short-term forecasts in comparison with physical and statistical methods, which are perceived to be more complex [59]. Hence, every forecast method is compared to the benchmark of the persistence method [58].

$$\underbrace{\hat{x}(t+k)}_{\text{forecast}} = \underbrace{x(t)}_{\text{known}} \quad (29)$$

Comparing a wind power data set of a specific location to the persistence model, a forecast error distribution for the persistence model can be obtained for that specific location in particular. From that error distribution, a wind power production profile can be constructed that has comparable characteristics as the measured wind power production from the data set. This will be discussed in more detail in Section 4.2.1.

Time-series based autoregressive moving-average models

Statistical methods are considered as ‘black box’. Whereas physical models often consider how certain atmospheric and geographic properties relate to wind speeds and calculate wind power production from wind speeds, statistical models involve only one step from input data to output [58]. The data used as input is often the weather forecast and historical wind data [60]. These time-series based models are not based on any physical model — as in a numerical weather prediction — but are rather based on patterns in the historical data [59].

The persistence method has already been discussed as a way to relate (forecast) values to measured values from past time steps. Now, time-series models with increasing complexity will be explored, which have the persistence model at their foundation.

A random walk process is a time-series model where the presently measured value equals the previously measured value with a random step up or down (30). Note that a random walk is the same as the persistence model, but with an additional stochastic white noise term, ε .

$$\begin{aligned} x(t) &= x(t-1) + \varepsilon(t) \\ x(t-1) &= x(t-2) + \varepsilon(t-1) \end{aligned} \quad (30)$$

It is clear to see that the value for $x(t)$ in the random walk model is simply the sum of the white noise terms that have occurred up until time step t from the known value for x at $t=0$, as shown by Equation (31).

$$\begin{aligned} x(t) &= \varepsilon(t) + \varepsilon(t-1) + \varepsilon(t-2) + \cdots + x(0) \\ &= x(0) + \sum_{i=0}^t \varepsilon(i) \end{aligned} \quad (31)$$

The random walk will now be used to provide the basis for the autoregressive model, since the autoregressive model is an extension of the random walk. An autoregressive model can include terms that go further back in time than just one previous time step. If the terms go back p number of time steps, then the model is called a p -order autoregressive process. The autoregressive model depends linearly on these multiple previous time step terms, but with a different coefficient α_i for each term i ($\alpha_p \neq 0$ for a p -order process). Equation (32) shows the p -order autoregressive model, AR(p). The influence of the persistence model is still noticeable, as previous terms are the predictors for the current value.

$$\begin{aligned} x(t) &= \alpha_1 x(t-1) + \alpha_2 x(t-1) + \cdots + \alpha_p x(t-p) + \varepsilon(t) \\ &= \sum_{i=1}^p \alpha_i x(t-i) + \varepsilon(t) \end{aligned} \quad (32)$$

Furthermore, the random walk would be exactly the same as the autoregressive model if the AR model would be AR(1) with $\alpha_1=1$.

Now, a model similar to the AR(p) model will be discussed, the moving-average model. Instead of a summation of linear dependencies to past time-series values, the moving-average model consists of linear dependencies on past white noise term values. A q -order moving-average model, MA(q), depends linearly on each past white noise term up to q number of time steps ago, with a coefficient of β_i for each term Equation (33).

$$\begin{aligned} x(t) &= \varepsilon(t) + \beta_1 \varepsilon(t-1) + \beta_2 \varepsilon(t-2) + \cdots + \beta_q \varepsilon(t-q) \\ x(t) &= \varepsilon(t) + \sum_{i=1}^q \beta_i \varepsilon(t-i) \end{aligned} \quad (33)$$

For the MA(q) model, each white noise ‘jump’ from the past q time steps is represented directly in the value for $x(t)$. In contrast, in the AR(p) model these ‘jumps’ are only represented indirectly, in the form of the past time-series values into which these past ‘jumps’ are integrated. Nonetheless, whereas the value

for $x(t)$ in the MA(q) model is only affected by the past q number of white noise terms, the AR(p) model takes all prior values for the white noise terms into account, although to a decreasing extent the further the 'jumps' go back in time.

The so-called Box-Jenkins models consist of a composition of an autoregressive model and a moving-average model. A general ARMA(p, q) model for data x is presented by Equation (34), with parameters α and β for the autoregressive and moving-average terms, respectively. The white noise term, ε , has a Gaussian probability distribution.

$$x(t) = \sum_{i=1}^p \alpha_i x(t-i) + \sum_{i=1}^q \beta_i \varepsilon(t-i) + \varepsilon(t) \quad (34)$$

The values for the model parameters in Equation (34) — with given p and q — are obtained by fitting the model to a data set. Variations to this model include the autoregressive integrated moving average (ARIMA), the seasonal ARIMA (SARIMA) and an autoregressive model with exogenous input (ARX/ARMAX).

4.1.2 Model predictive control

Section 3.1.2 discussed static optimisation problems. Optimal control is an allocation of resources over an interval of time by choosing an optimal time-dependent path for a control variable. Decisions are made for the control variables as time progresses, instead of there being only a one-time decision. Model predictive control (MPC) uses a dynamic model of a system to predict the behaviour of the system [61]. This will now be discussed in more detail.

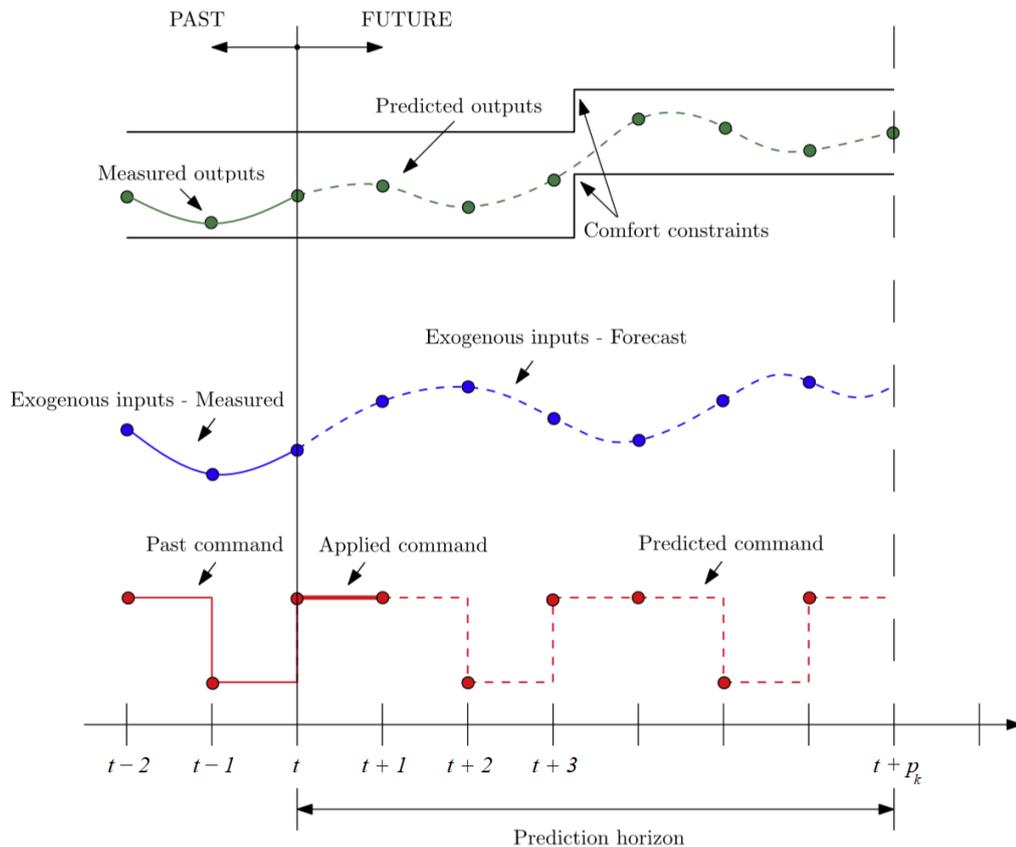


Figure 12: Time evolution in a standard MPC approach [40]

Figure 12 shows how various inputs and outputs, both measured and predicted, determine the commands from the control output. Here, the command is a direct signal from the control. The signal affects the behaviour of the dynamic system, i.e. the output. The predicted effect is shown by the line for the predicted

output in Figure 12. The decisions for the control commands are based on forecasts for exogenous factors and predictions for the behaviour of the system. These predictions for the system output are based on a model of the system. Such control is therefore called model predictive control. The constraints resulting from this model are implemented in an optimisation algorithm to obtain optimal control values for the components in the system. The predicted values considered as inputs and outputs are only considered up to a certain time horizon, called the prediction horizon, p_k . Each time window, $\mathcal{K} = \{t, \dots, t + p_k\} \forall t$, stretches from the present time step t to the prediction horizon p_k .

Figure 13 highlights the receding horizon principle by showing the prediction horizon for multiple time steps. The time steps within the prediction horizon are denoted by $k \in \mathcal{K}$. It is important that only the control command for the first time step within the prediction horizon is applied at that time step; the other control commands are merely predictions. The prediction horizon then shifts one time step forward. It is then important that the newly measured and predicted inputs and outputs are considered by the optimisation algorithm for the next control horizon.

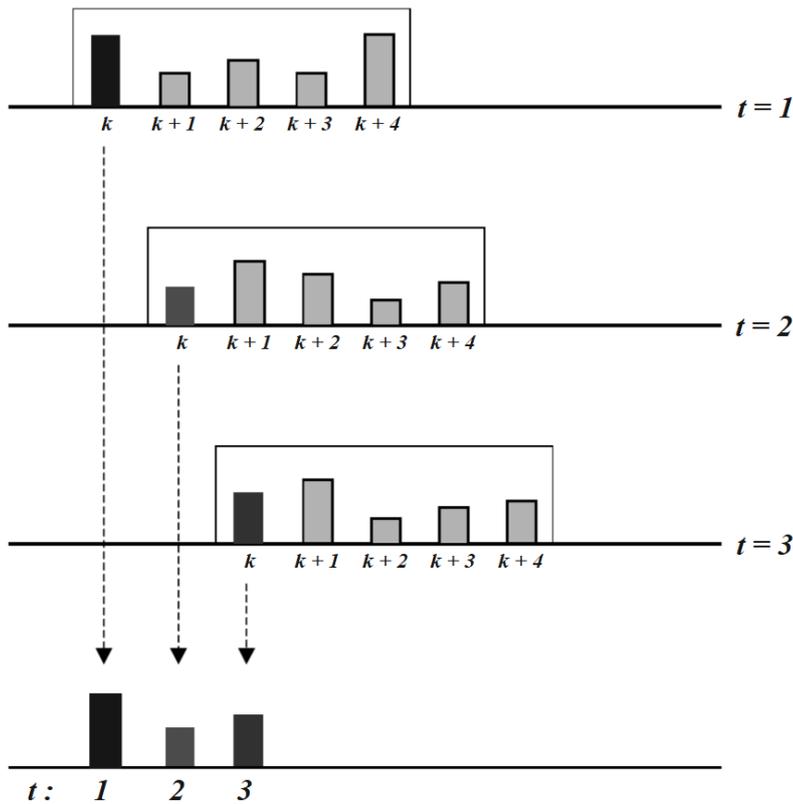


Figure 13: Output of MPC over time, adapted from Kwon and Han [62, p. 88]

Although the optimisation problem for each time window is the same as that for an optimal scheduling problem, the receding horizon principle ensures the model predictive control is up to date with the latest measured and predicted parameters. This allows the behaviour and flexibility of the system in ensuing time steps to be re-evaluated, together with time-based constraints and parameter predictions. This way, the most updated prediction for the exogenous input values can be used, which should lead to a more optimal and reliable control output.

4.2 Wind power simulation model

It is assumed the wind farm operator and the heat network operator have an advanced method of wind power forecasting at their disposal, which uses real-time weather forecasts, historical data, statistical models and computer intelligence. In this report, such advanced forecasting tool is considered as a black box, because wind power forecasting itself falls outside the scope of this research. Instead of modelling the input data and this complex forecasting tool, the output data is modelled by creating wind power production profiles with imitated characteristics, simulated using historical wind power data. For modelling a day-ahead wind power forecast, the persistence model is used. This will be explained in more detail.

Assumption 12. A 'black box', based on statistical models, functions as an advanced forecasting tool and produces wind power forecasts.

4.2.1 Daily wind power production profile

Now, the method to simulate daily wind power production profiles will be explained. These wind power profiles will function as the day-ahead wind power forecast in the experiments, in which the forecast errors are to be mitigated by a flexible operation of the heat network. The model for simulating a wind power production profile attempts to mimic the randomness with regards to the change in wind power from one time step to another, found in location-specific wind power data. The persistence model, as used here, assumes the predicted wind power, \hat{P}_w , equals the measured power at time t , P_w , one hour earlier (35).

$$\underbrace{\hat{P}_w(t+1|t)}_{\text{forecast}} = \underbrace{P_w(t)}_{\text{known}} \quad (35)$$

Therefore, the forecast error resulting from the persistence model, equals the change in wind power from one hour to another (36). One is allowed to interpret this error as the white noise term in a random walk process (30).

$$\hat{\varepsilon}_w(t) = P_w(t) - P_w(t-1) = \Delta P_w(t) \quad (36)$$

Wind power data has been collected from Nord Pool for a price area in Denmark [63]. The data consists of hourly wind power measurements over the year 2018. Figure 14 shows the measured wind power over time at an hourly resolution. The choice for Danish data sets is motivated by the availability of forecasts at a 5-minute resolution, which will be relevant in Section 4.2.2. Furthermore, the Swedish wind power prognoses are not technical forecasts, but rather wind production plans [64].

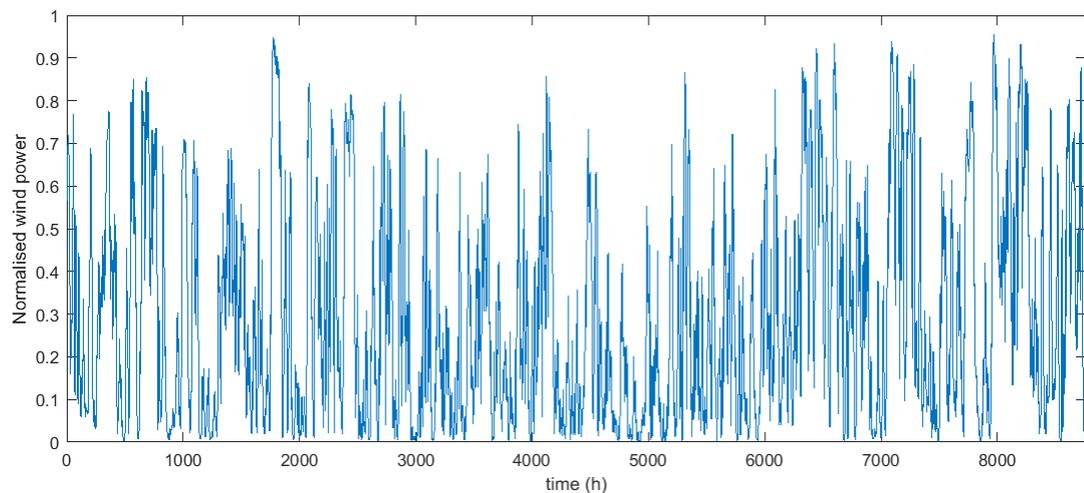


Figure 14: Wind power in the DK2 price area throughout the year 2018

Figure 15 shows a histogram of the hourly changes in wind power production throughout the year, derived from the wind power data shown in Figure 14 and calculated using Equation (36). These hourly

changes are the forecast error if the persistence method would be used each hour to forecast the next hour. Therefore, Figure 15 shows the frequency distribution of the persistence model forecast errors. Now, a continuous probability distribution function can be fitted to the frequency distribution of the errors. Evidently, a normal (Gaussian) probability distribution function is not a proper fit, shown by the red line in Figure 15. Instead, the Cauchy distribution shows an obviously favourable fit, shown by the blue line in Figure 15.

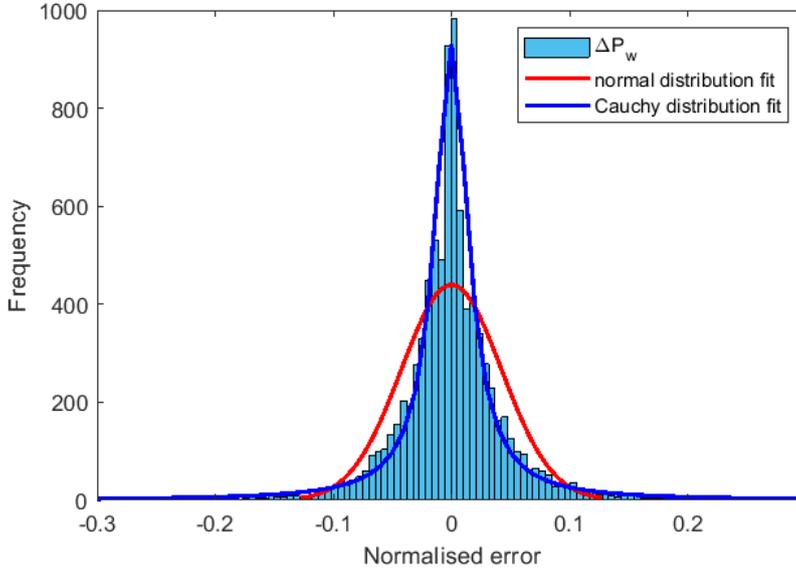


Figure 15: Histogram of hour-to-hour changes in wind power for the DK2 price area in Denmark

It has now been shown that from the persistence model in Equation (35) and a wind power data set for the DK2 price area, a persistence model forecast error distribution — the $\hat{\varepsilon}_w$ probability distribution — can be obtained. As proposed by Hodge and Milligan [65], the Cauchy distribution represents the forecast error distribution for persistence model forecasting. Equation (37) shows its inverse cumulative distribution function which gives a value for the error, $\hat{\varepsilon}_w$, depending on the value from the uniform probability distribution function, \mathcal{U} . The half width at half maximum⁵, γ , is obtained by fitting a Cauchy distribution to the errors in the persistence model for hourly wind power generation data. When time intervals of less than an hour are used and available, their hourly averages are considered for the day-ahead forecast.

$$\hat{\varepsilon}_w = \gamma \cdot \tan\left(\mathcal{U}\left(-\frac{1}{2}, \frac{1}{2}\right) \cdot \pi\right) \quad (37)$$

It follows that a wind power production profile can be modelled according to the relation in Equation (38). Note that this wind power time-series model is similar to the random walk process in Equation (30), but instead of including a Gaussian white noise term, a Cauchy distribution is used to obtain values for the white noise term.

$$\hat{P}_w(t) = \hat{P}_w(t-1) + \hat{\varepsilon}_w(t) \quad \forall t \in \mathcal{T} \quad (38)$$

The constructed wind power production profile functions as the resulting day-ahead ($D-1$) wind power forecast, $\hat{P}_w(t|D-1)$, from the aforementioned black box. The random walk process described by Equations (37) and (38) could be used to simulate how wind power production progresses over time. Naturally, production cannot be negative nor exceed the installed capacity. Script 6 in Appendix C.2 can be implemented in MATLAB together with the appropriate data set to simulate wind power production. An example of the output can be seen in Figure 16, showing a wind power production profile.

Assumption 13. *Wind power production can be modelled as a random walk with Cauchy white noise.*

⁵The full width at half maximum is the width of a probability distribution measured between those probability values which are half the maximum probability. The half width at half maximum is half of the full width at half maximum if the function is symmetric, which is the case in a Cauchy distribution.

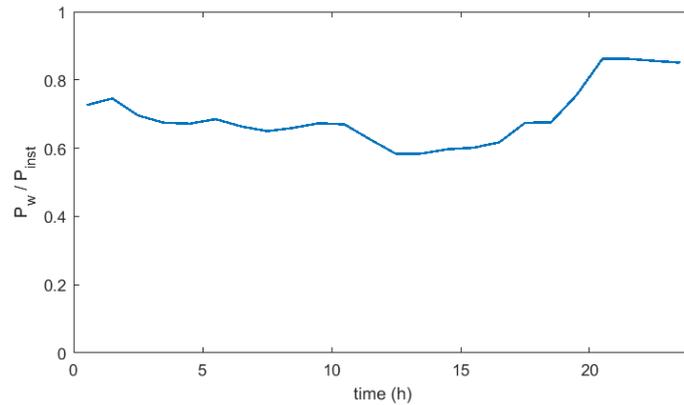


Figure 16: Normalised wind power generation profile modelled for a day

4.2.2 Rolling wind power forecast

During the day, a rolling forecast updates the day-ahead forecast by measuring the actualised wind power and the resulting forecast error, ϕ . This error is assumed to echo — to some extent — through the subsequent forecast errors, which is described here by an autoregressive moving-average model [66], or ARMA(1,1) process (39).

Assumption 14. *Wind power forecast errors for subsequent time steps are statistically related.*

$$\phi(t) = \alpha\phi(t-1) - \beta\gamma(t-1) + \gamma(t) \quad (39)$$

The challenge in modelling the rolling forecast is not to update the forecast with the measured values, but to update the forecast with a new forecast that is on average more accurate. The data set at a 5-minute resolution includes not only day-ahead predictions, but both 5- and 1-hour-ahead forecasts as well (see Figure 17).

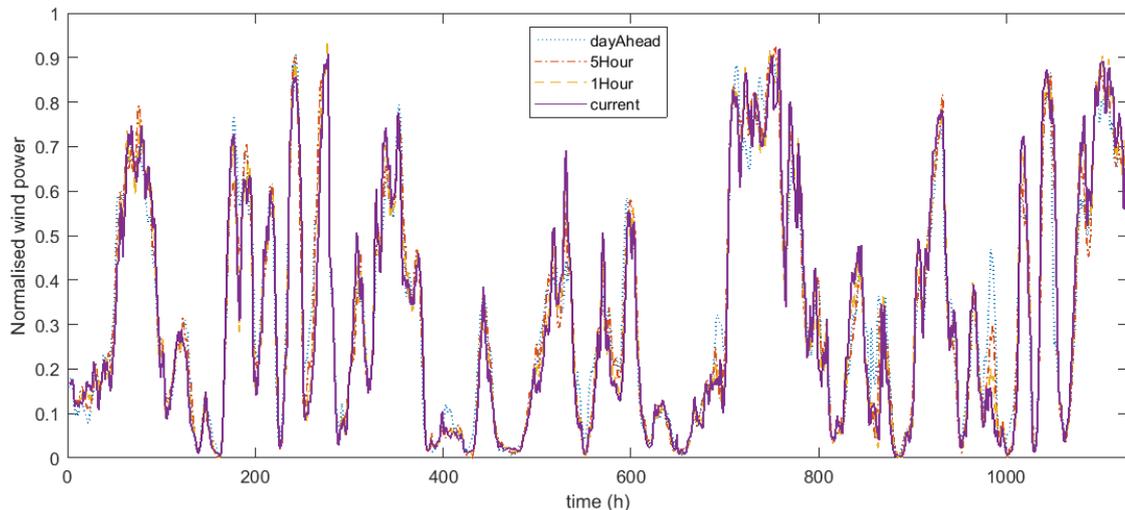


Figure 17: 47 days of forecast data for the DK2 price area [67]

The difference between the day-ahead and the 5-hour-ahead forecasts is assumed to follow an ARMA process. Similarly, the difference between 5- and 1-hour-ahead forecasts and the difference between the 1-hour-ahead forecast and the measured values are all modelled using an ARMA. Therefore, instead of using the resulting forecast errors for the multiple time horizons, the differences in forecasts between the various time horizons are considered for updating the rolling wind power forecast. In short, instead of a single

moment of updating the wind power prediction from the day-ahead forecast, the prediction is updated at the start of the day, 5 hours ahead and 1 hour ahead. The final difference between the day-ahead wind power forecast and the measured wind power is thus a summation of these errors, in the way they relate to the previous forecast update. This summation is shown by Equation (40). Table 4 shows the coefficients for the autoregressive moving-average models used in these forecast updates, including the standard deviations, σ , of the Gaussian white noise.

Assumption 15. *Differences between wind power forecast updates over time can be modelled as autoregressive moving-average processes.*

$$P_w(t) = \hat{P}_w(t) + \phi_{da}(t|t_0) + \phi_{5h}(t|t-5) + \phi_{1h}(t|t-1) \quad (40)$$

Forecast update	day-ahead to 5-hour-ahead (ϕ_{da})	5- to 1-hour-ahead (ϕ_{5h})	1-hour-ahead (ϕ_{1h})
α	.80552	.67273	-.092637
β	.33827	.25759	.40486
σ	.0237	.0257	.0318

Table 4: ARMA(1,1) coefficients

Relevant for the intra-day operation are the hourly short-term forecast updates modelled as ARMA processes. These updates combine the day-ahead forecast and a forecast error consisting of a 5-hour-ahead component and a 1-hour-ahead component (40). Combined with the day-ahead wind power prediction, the forecast can be updated using the most recent forecast error prediction. For every update, a forecast shown as one of the examples in Figure 18 can be modelled. The forecast updates in Figure 18 include an update of the measured hourly wind power production (green) and an update of the wind power forecast for the prediction horizon (blue). Furthermore, the day-ahead wind power forecast is included for comparison.

By running the forecast error simulation for all time steps in a single day a million times, a probabilistic error distribution can be obtained, shown in Figure 19, for the day-ahead wind power forecast in Figure 16. The likely region within which the wind power will progress is clearly visible. The further ahead in time, the larger the uncertainty in the prediction.

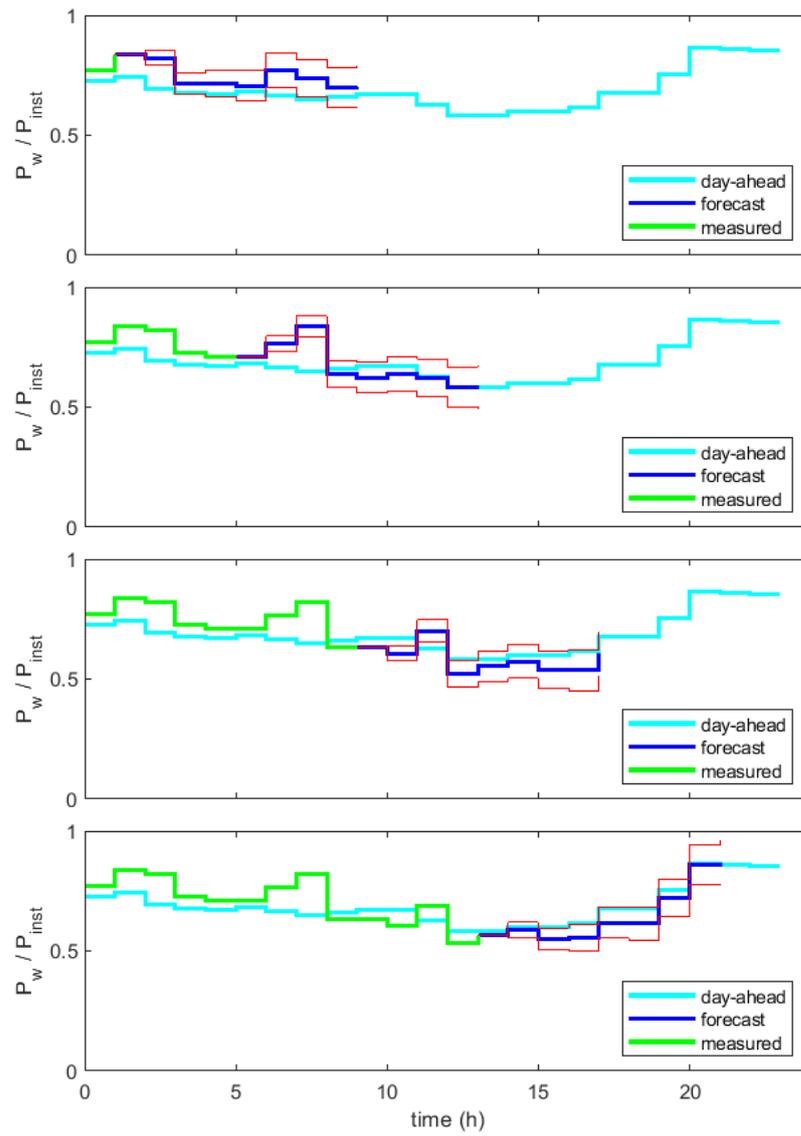


Figure 18: A rolling wind power forecast with a receding 8 hour prediction horizon

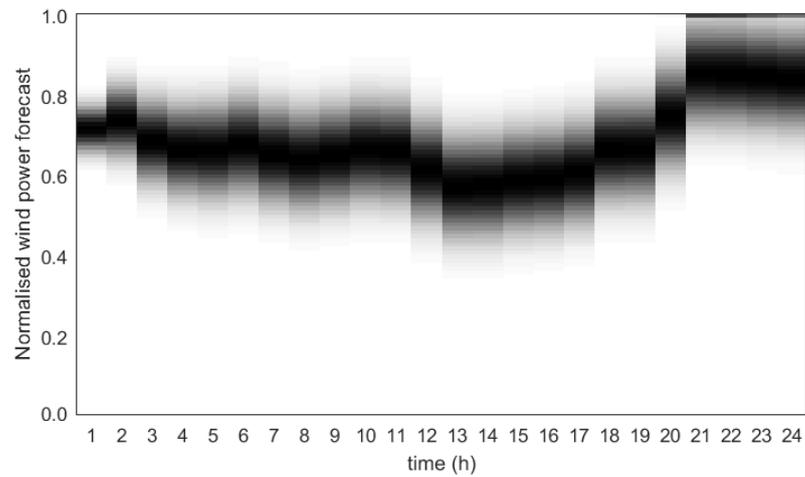


Figure 19: A probabilistic day-ahead wind power forecast

4.3 Problem formulation

Now that a black box model for the wind power forecasts has been developed, it can be linked to the control system of the heat network components, particularly the electric boiler. Figure 20 shows a schematic representation of the intra-day heat network operation with wind power forecasts updated at every time step. Throughout the day, each forecast update is compared to the day-ahead wind power forecast. The difference between the two is called a (day-ahead) forecast error and can be caused by an overproduction or an underproduction of wind power. The forecast error is positive in the case of overproduction and negative in the case of underproduction. The heat network can partially compensate for the forecast error by shifting the electric load of the boiler. This results in a new operation schedule for the heat network. In case of underproduction, the electric boiler turns off if possible, while it turns on in the case of overproduction. The hours of overproduction should balance out the hours of underproduction, while maintaining a reliable supply of heat to the two greenhouses in the heat network. As shown by Figure 20, the boiler shifting is controlled by a model predictive control, introduced in Section 4.1.2. The state of the heat network — temperatures, the day-ahead operation schedule and (physical) constraints — is considered by the model predictive control, together with the wind power forecast error. The output of the model predictive control is a change to the day-ahead schedule for the electric boiler and to the schedule for heat extraction by the two greenhouses — represented by the new operation schedule. After all, heat extraction by greenhouses is dependent on the boiler; if the boiler operation changes, the control of the heat exchangers at the greenhouse substations changes along.

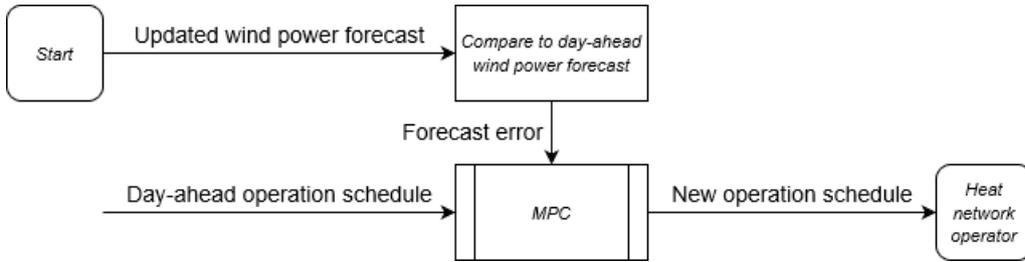


Figure 20: MPC scheme for the boiler in the heat network under changing wind power forecasts

This change to the day-ahead boiler schedule should be such that the wind power forecast errors are mitigated. The model predictive control is therefore set up with the objective of optimising the boiler shifting for each prediction horizon. This will be discussed in more detail.

4.3.1 Objective function and control variables

As stated before, the objective of the control system is to minimise intra-day market participation for the wind farm operator. The load shifting of the boiler can, when properly timed, decrease the amount of power to be traded on the market at a disadvantageous price. The objective to be minimised consists then of the change in boiler use, ΔP , at time step k caused by the boiler shift, subtracted from the day-ahead wind power prediction error (41). This difference is denoted by the residual forecast error, $\hat{\phi}_r$. The prediction error could be negative, thus motivating the operation of the boiler to withhold at that time step. Hence, it is the absolute value of the difference between the prediction error of the day-ahead forecast, $\hat{\phi}$, and ΔP that is to be minimised.

$$\hat{\phi}_r(t) = \hat{\phi}(t) - \Delta P(t)$$

$$\min \sum_{k=t}^{t+p_k} |\hat{\phi}_r(k)| \quad \forall t \in \mathcal{T} \quad (41)$$

The objective function in (41) is nonlinear, because it is a modulus function. When this nonlinear objective function is to be implemented in an optimisation problem formulation, it would turn the problem into an integer nonlinear programming problem. Nonlinear programming is computationally complex. In addition, the runtime of the algorithm is already fairly long (more on that in Section 5.2.1). However, another solution is at hand, in which the optimisation problem remains an integer linear programming problem. The minimisation in (41) can be separated into a linear part for $\hat{\phi}_r > 0$ and a linear part for

$\hat{\phi}_r < 0$, because whether the forecast value for $\hat{\phi}_r$ is positive or negative is determined beforehand [68]. Since $\hat{\phi}_r$ can never be both positive and negative, the linearisation looks like Equation (42).

$$\min \sum_{k=t}^{t+p_k} \begin{cases} \hat{\phi}_r(k) & \text{if } \hat{\phi}_r(k) > 0 \\ -\hat{\phi}_r(k) & \text{if } \hat{\phi}_r(k) < 0 \end{cases} \quad (42)$$

The control variable introduced here, u_k , is a unit commitment variable that determines a change, ΔP , in the already existing boiler schedule, \hat{P} , as shown in Equation (43). When the boiler schedule dictates that the boiler is to be turned on at hour t , the change can then only be to turn it off instead; the change can only be negative, thus the lower bound is minus one, while the upper bound is zero. Conversely, when the boiler is scheduled to be turned off at hour t , the change in boiler power can then only be positive (44).

$$\begin{aligned} \Delta P(t) &= u_k(t) \cdot P_b \\ &= P(t) - \hat{P}(t) \end{aligned} \quad (43)$$

$$u_k(t) \in \begin{cases} \{-1, 0\} & \text{if } \hat{P}(t) = P_b \\ \{0, 1\} & \text{if } \hat{P}(t) = 0 \end{cases} \quad (44)$$

Since market interaction is not to be included in the objective function due to the complexity of modelling intra-day prices, boiler power that is shifted in favour of the wind farm operator is valued equally for every hour of the day. For every hour t , the cost function is expressed according to Equation (45). Upon close examination, it can be deduced that in case of a negative wind prediction error ($\phi < 0$), the cost is minimised when the unit commitment variable is minimised to counter the prediction error. Furthermore, a positive prediction error ($\phi > 0$) drives a maximisation of the boiler power by attaching a negative cost to powering the boiler.

$$c_k(t) = \begin{cases} 1 & \text{if } \phi(t) < 0 \\ -1 & \text{if } \phi(t) > 0 \\ 0 & \text{otherwise} \end{cases} \quad (45)$$

The objective function as an absolute value function in Equation (41) can be linearised due to the formulation of the cost function (45) and the bounds of the control variables. After all, by analysing the wind power forecast error and the corresponding boiler schedule, it can already be determined when the absolute function will negate an error — if $\phi < 0$ — and when it will not. This linearisation is shown by Equation (42). The cost function is formulated with the aim of separating the hours of a wind power surplus and the hours of a wind power deficit. Equation (46) shows the objective function after linearisation.

$$\min_{\mathbf{u}} \vec{c}_k \cdot \vec{u}_k \quad (46)$$

4.3.2 Constraints

In addition to the constraints formulated in the boiler scheduling problem in Section 3.2.2, only a set of equality constraints are to be defined. After all, the equality constraints describing a physical model of the heat network are still valid, together with the inequality constraints that are necessary to satisfy a temperature demand on the greenhouse side of the network. The additional equality constraints necessary for the receding horizon control are there to ensure that the system remains within its limits, even with a shorter prediction horizon in the optimisation problem. This can be explained in the following way. The optimal solution for a certain time window within a period might not be the same as the global optimum for the period as a whole. For example, if the boiler use is optimised for a single hour, it will likely not be turned on, no matter how low the price — assuming prices are positive. The explanation is that the system simply allows the temperature to drop somewhat, thereby avoiding any cost. When the boiler use is optimised for an entire day, it is likely that the boiler will be turned on at that specific hour if the electricity price for that hour is particularly low. By way of explanation, the global optimum differs from the local optimum. To prevent the model predictive controller from diverging from the optimal path, the additional equality constraint is introduced here (47).

$$\sum_{k=t}^{t+p_k} \Delta P(k) = 0 \quad (47)$$

4.3.3 Optimisation problem

The optimisation problem for each sliding time window in the model predictive control algorithm can now be formulated in full (48). The cost-minimising objective function consists of a cost vector, c , and a control vector, \mathbf{u} , as in the optimisation problem (28) in Section 3.2. However, both vectors are adjusted. Because the electricity for the boiler has already been bought at the day-ahead market, no costs are associated with using the electric boiler according to its day-ahead schedule. Instead, costs are coupled to changes to be made to the boiler schedule, ΔP , and thus to the related control variable, u_k , described by Equation (43). This control variable u_k is therefore to be added to the control vector. Matrix A and vector b again describe the system dynamics, defined by Equations (14), (15), (17) and (19). Furthermore, an additional equality constraint is added (47), which has been discussed in Section 4.3.2. The inequality constraints have been discussed in Section 3.2.2. The optimisation problem in (48) is to be solved at every time step t , for the time window denoted by \mathcal{K} . Scripts 7 and 8 in Appendix C.2 can be implemented in MATLAB to simulate the intra-day operation, after having run Scripts 1 to 6.

$$\begin{aligned} & \min_{\mathbf{u}} \quad c^T \mathbf{u} \\ & \text{s.t.} \quad \mathbf{A} \mathbf{u} = b \\ & \quad \sum_{k=t}^{t+p_k} \Delta P(k) = 0 \\ & \quad u_i(k) \cdot T_{min,s,i} \leq T_{s,i}(k) \\ & \quad T_{min,i} \leq T_i(k) \leq T_{max,i} \\ & \quad u_b(k), u_i(k) \in \{0, 1\} \quad \forall i \in \mathcal{S}, \quad \forall k \in \mathcal{K} \\ & \quad u_k(k) \in \begin{cases} \{-1, 0\} & \text{if } \hat{P}(k) = P_b \\ \{0, 1\} & \text{if } \hat{P}(k) = 0 \end{cases} \end{aligned} \quad (48)$$

$$c = [0, \dots, 0, \bar{c}_k]$$

$$\mathbf{u} = [u_b(\mathcal{K}), u_l(\mathcal{K}), \dots, u_N(\mathcal{K}), u_k(\mathcal{K})]$$

Section 2.3.1 explains the equation to calculate the temperature on the greenhouse side, T_i in Equation (19):

$$T_i(t) = T_i(t-1) + \dot{Q}_i(t) \frac{\Delta t}{c_w m_i} - [T_i(t-1) - T_a(t-1)] \frac{c_{q,i} \Delta t}{c_w m_i}$$

Section 2.2.1 explains the equation to calculate the supply temperature at each greenhouse, $T_{s,i}$ in Equation (14):

$$T_{s,i}(t) = (1 - \eta)^\tau \cdot T_{s,i-1}(t - \tau) + \eta \int_{t-\tau}^t (1 - \eta)^{t-\vartheta} \cdot T_a(\vartheta) d\vartheta$$

5 Results

In this section of the report, the results from the optimisation problems formulated in Sections 3.2 and 4.3 are presented. Section 5.1 describes the parameter setting that applies to all model simulations and conducted experiments. Section 5.2 shows the results for the day-ahead schedule and Section 5.3 shows the evaluation of the performance of the model predictive control.

5.1 Standard experimental setup

To analyse how changing a parameter affects the behaviour of the system, each result must be compared with a standard parameter setting, which will be defined here. Table 5 shows the values for the consumer-specific parameters, such as the number of greenhouses and their temperature boundaries. Table 6 shows the values for the network-specific parameters, such as the network time delays and the mass flow rates, \dot{m}_i . If only one value is given, e.g. in the case of the mass flow rates, this value holds for each $i \in \mathcal{S}$. All simulations use the standard setting for each parameter, unless for a parameter it is stated otherwise.

N	m [kt]	$c_{\dot{q}}$ [kW/K]	$T_{min,i}$ [K]	$T_{max,i}$ [K]	$T_{0,i}$ [K]	$\dot{Q}_{hx,i}$ [kW]
2	.2	5	{ 295, 294 }	{ 297, 297 }	{ 296, 295 }	{ 216.5, 206.5 }

Table 5: Consumer parameter setting for the simulations

P_b [kW]	\dot{m}_i [kg/s]	τ_s	$c_{\dot{q},s}$ [W/K]	$T_{min,s}$ [K]	$T_{0,s}$ [K]	τ_r	τ_R	$c_{\dot{q},r}$ [W/K]
280.3	2	1	.01	323	326	1	1	.1

Table 6: Heat network parameter setting

Then there are external and environmental parameter settings. Because the wind data in Section 4.2 is from data sets for a location in Denmark, the temperature used for the ambient temperature throughout the day, T_a , is derived from the temperatures in Denmark shown in Figure 21. The daily average ambient temperature used in the simulations is the average temperature of the months January and February. The winter season is chosen for the standard experimental setup because heat demand is highest then.

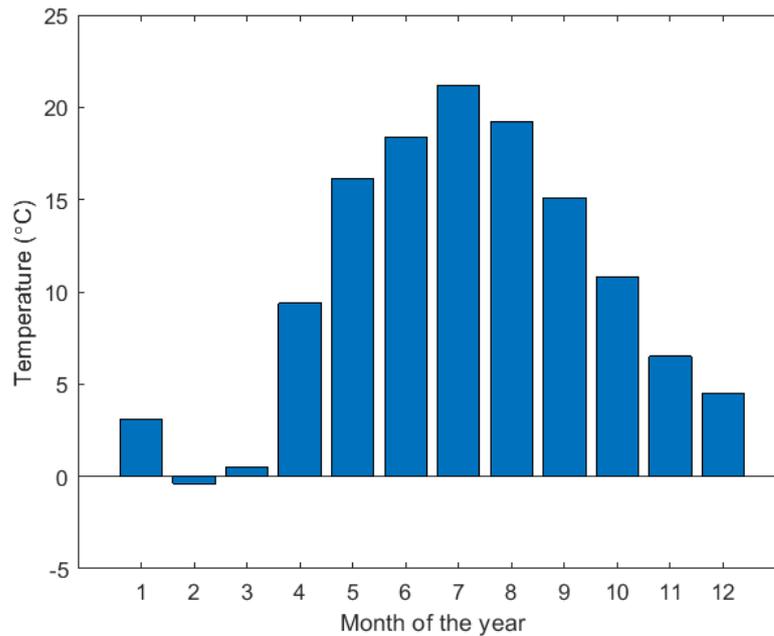


Figure 21: Monthly average outside temperatures in København throughout the year 2018 [69]

For the deterministic electricity price profile in the simulations, an averaged price profile is constructed, using price data for the price area in Denmark [63]. Because seasonal price profiles differ noticeably, different price profiles are constructed for different seasons. Figure 22 shows the averaged price profiles both for the winter months, January and February, and for the summer months, July and August.

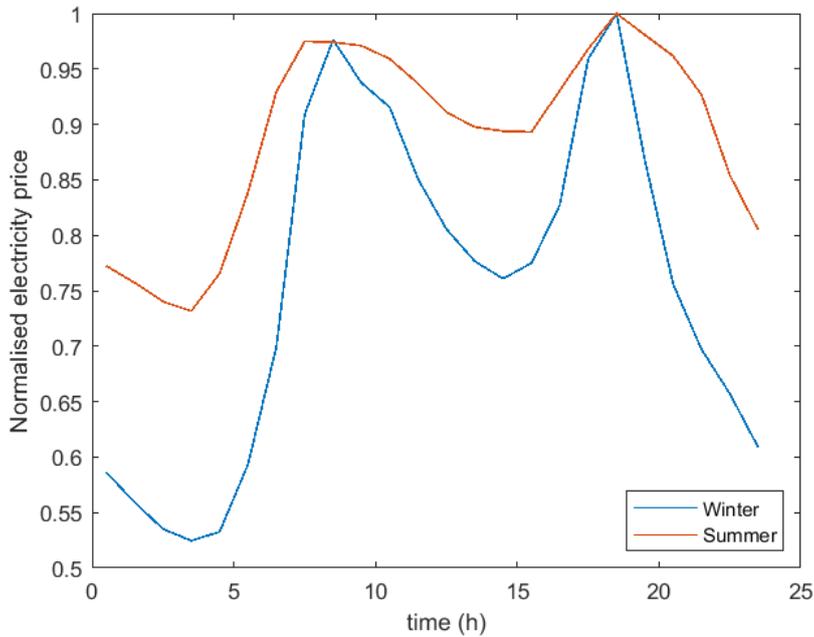


Figure 22: Averaged electricity price profiles for winter and summer [63]

5.2 Day-ahead schedule

The results from the experiments related to research question R1 are now presented and explained. Section 5.2.1 discusses the feasibility of the day-ahead boiler schedule optimisation. In particular, the effect of the initial temperature throughout the heat network and the effects of adjusting the constraints in the optimisation problem formulation are discussed. Section 5.2.2 compares the optimal schedule to the benchmark rule-based operation in terms of cost-efficiency. Section 5.2.4 analyses the effect of longer heat transportation times in the network by doubling the heat network pipeline time delays.

5.2.1 Feasibility and runtime

Before the results of the modelling and the optimisation can be described and discussed, the feasibility of the optimal control must be examined. Namely, it was found that not all simulations of the optimal control led to a feasible solution. For example, if the heat sources are not scaled properly, the heat demand might possibly not be satisfied in any way. The aim here is to investigate how changing parameters affect the ability of the optimisation algorithm to find a feasible solution.

Using the standard experimental setting from Section 5.1, the day-ahead schedules of consecutive days are simulated. The temperature at the start of the first day is set to be exactly 328 K at every node in the heat network. Figure 23 shows a day-ahead boiler schedule, $\hat{P}(t)$, together with the temperature leaving the boiler, $T_{s,0}$, over time. The optimisation algorithm could only find a feasible solution for the first day and therefore stopped prematurely after optimising the first day-ahead schedule. The temperature leaving the boiler shows a downward trend as the day progresses.

The optimisation algorithm has performed simulations with varying initialisation temperatures. For each simulation, the aim was to find seven day-ahead schedules for seven consecutive days. After each simulation, the algorithm runtime for the first day of the week was noted. Figure 24 shows the runtime for

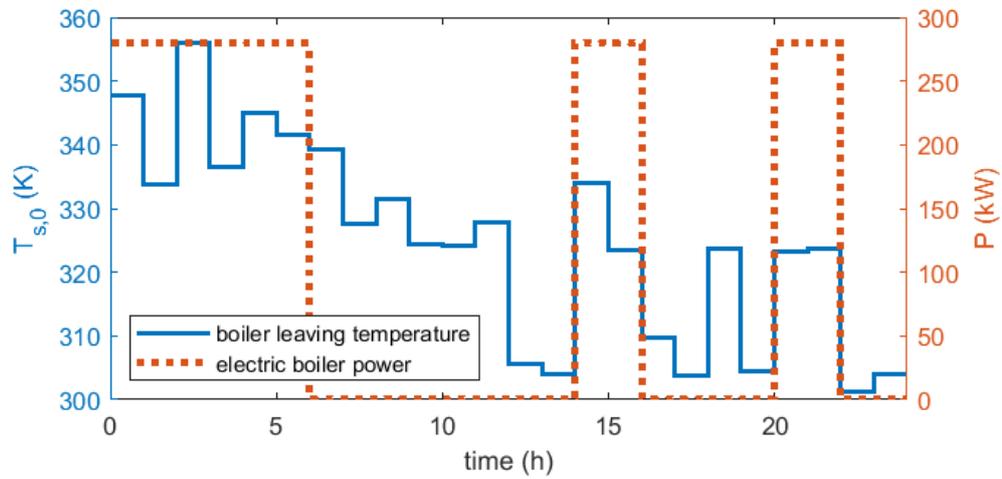


Figure 23: Day-ahead scheduling for a single day with an initial network temperature of 328 K

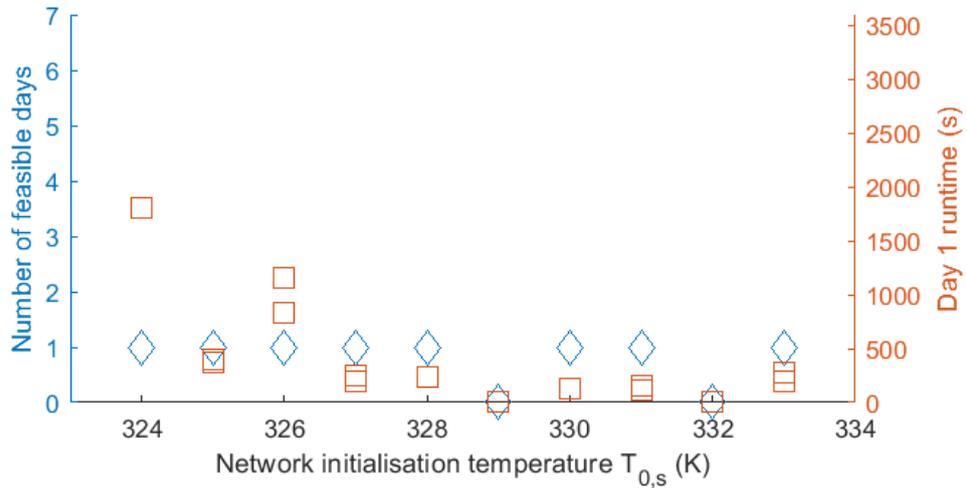


Figure 24: Algorithm runtime in seconds for simulations with varying initial temperatures

different initial temperatures, as well as the number of simulated days until the algorithm stops because no feasible solution can be found.

Interestingly, the initialisation temperature — the temperature distribution throughout the network at the start of the day — turns out to affect the time needed for the algorithm to find an optimal solution. In addition, the feasibility of the optimal scheduling algorithm also appears to depend on the initialisation temperature. In each case, the algorithm stops prematurely after simulating one day because no solution can be found for that day or the following day. The explanation is as follows. If the temperature throughout the network lowers at the end of the day, as dictated by the schedule (see temperature trend in Figure 23), the updated temperature initialisation for the next day might not be high enough to satisfy the network requirements. In the optimisation problem formulation (28), there is not yet a constraint included to prevent this. The dashed blue line in Figure 25 shows that the average temperature of the network, calculated with Equation (49), lowers towards the end of the day. Consequently, the optimisation algorithm cannot find a feasible solution for the next day-ahead scheduling problem. This issue leads this analysis to reconsider the constraints.

$$T_{\text{avg}}(t) = \frac{1}{2 \cdot N + 1} \left(T_{s,o}(t) + \sum_{i \in S} [T_{s,i}(t) + T_{r,i}(t)] \right) \quad (49)$$

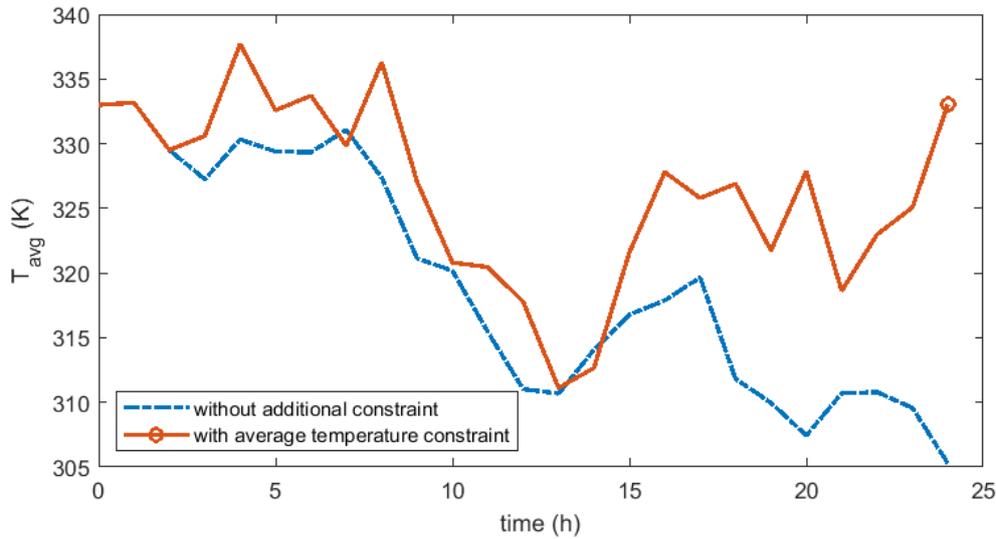


Figure 25: Average temperature throughout the network with an initial temperature of 333 K

The formulation of the optimisation problem in (28) does not include a constraint to prevent the average heat network temperature from lowering at the end of the day. Given that this renders the operation of the heat network unfeasible, a new network constraint must now be introduced to the optimisation problem. The aim of the new constraint is to maintain a sufficiently high temperature distribution for the upcoming daily schedule to be a feasible solution. The constraint should somehow represent the temperature distribution throughout the entire network. To this end, the constraint in Equation (50) is considered. This inequality constraint ensures that at the end of the day, the average network temperature, T_{avg} , is higher than or equal to the initial supply temperature. Henceforth, the minimum average temperature constraint (50) is applied to the optimal scheduling problem. Again in Figure 25, the average network temperature can be seen over time, including a marking for the hour the constraint is active. Due to the constraint, the heat network operates at a higher temperature on average.

$$T_{avg}(d) \geq T_{0,s} \tag{50}$$

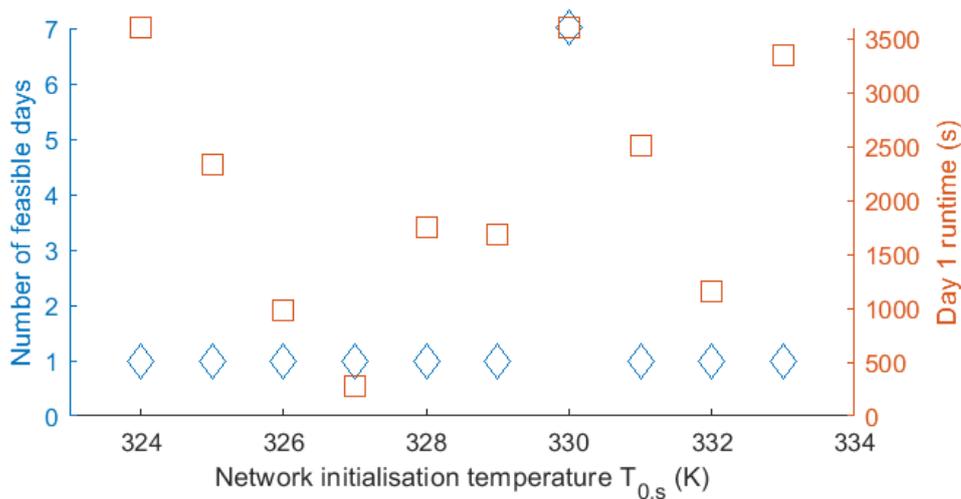


Figure 26: Algorithm runtime for simulations with the average temperature constraint

However, as it turns out, the constraint has not yet resolved the issue concerning the feasibility of the algorithm. The optimisation algorithm including the average temperature constraint could as well only find

a feasible solution for the first day-ahead schedule. Again, the algorithm has performed simulations with varying initialisation temperatures, while this time including the minimum average temperature constraint (50). Figure 26 shows the algorithm runtime for the first day and the number of consecutive days for which a feasible schedule is found. For all initial temperatures, except 330 K, the algorithm stops prematurely after solving for only one schedule, instead of solving for multiple days consecutively. The reason for this is the difference between the temperature in the supply pipe and the temperature in the return pipe. Although the average temperature throughout the network satisfies the constraint, the temperature in the supply pipe might be a lot lower than the temperature in the return pipe. Consequently, for the first few hours of the next day, no heat can be extracted from the network because the water in the supply pipe is too cold. As a result, no feasible schedule for that day can be found.

$$T_{s,i}(d) \geq T_{0,s} \quad \forall i \in \{0, \dots, N\} \quad (51)$$

The issue of an unfeasible operation after implementing the constraint in Equation (50) leads to another reconsideration of the applied constraints, to be discussed here. Instead of applying a constraint for the average temperature throughout the entire network, a similar constraint will be applied for the temperature throughout just the supply pipe (51). The minimum temperature constraint is applied to every node in the supply pipe. This way, enough heat will be available for extraction in those first few hours of the day, during which the boiler cannot raise the temperature of the water that is already in the supply pipe.

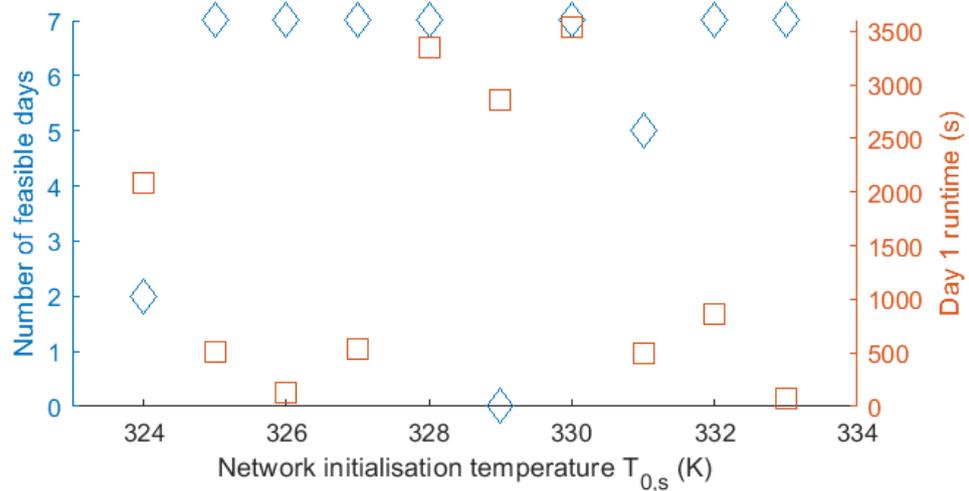


Figure 27: Algorithm runtime for simulations with the supply temperature constraint

Again, the simulation is run for the initial temperatures from 324 K to 333 K. The runtime and the number of feasible consecutive day-ahead schedules for each simulation are presented by Figure 27, which shows that the algorithm can now solve for consecutive day-ahead schedules. For those initial temperatures for which, at some point, no schedule could be found, the algorithm had reached its allowed maximum of total runtime or its allowed maximum number of explored solution nodes.

Figure 28 shows temperatures throughout the heat network over time. Furthermore, the initial average supply pipe temperature for each following day is marked with a circle. Figure 28 includes the average temperature throughout the entire heat network and the average temperature in the network supply pipe. One can interpret the patterns in the fluctuations of the temperatures as a periodicity in the average network temperature. This could be partially an effect of the periodicity of the price profile.

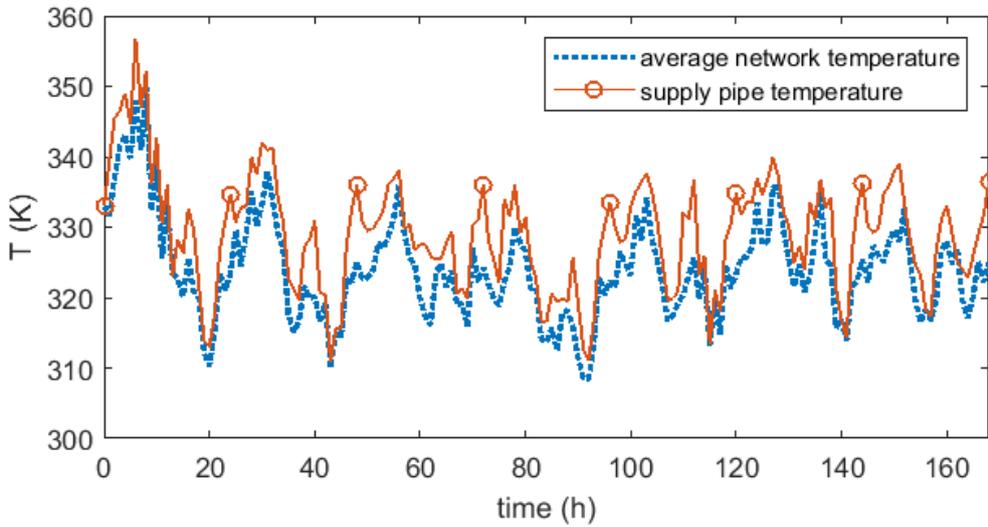


Figure 28: Heat network temperatures with the supply pipe temperature constraint

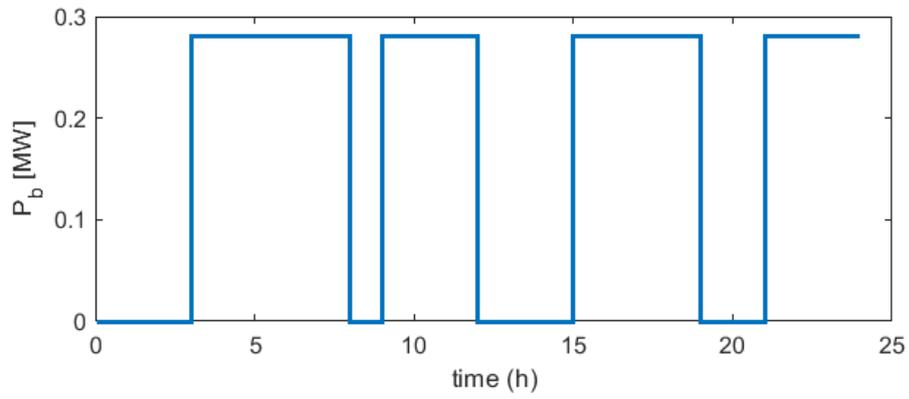
5.2.2 Rule-based control comparison

An important indication of the performance of the optimal control follows from the benchmark comparison. The benchmark was described in Section 1.2.1 as a rule-based control. The rule-based control algorithm is set to turn on the electric boiler if the supply temperature would otherwise not satisfy the minimum supply temperature constraint. Situations in which this constraint is not satisfied, rendering the operation unfeasible, only occur when the system is not scaled properly — e.g. the boiler is too small. In the case of optimal control, the minimum supply temperature constraint is only to be satisfied when heat is in fact extracted by the greenhouse, whereas this constraint must be satisfied all the time in the case of rule-based control. The reason for this is that the rule-based control has no predictive capability. Therefore, the boiler control system receives no signal that a greenhouse downstream will extract heat and must for that reason have a sufficiently high supply temperature at all times. An example of a rule-based boiler schedule is shown by Figure 29a, together with the electricity price profile in Figure 29b. Figure 29c shows the boiler schedule for optimal control. It can be clearly seen, when comparing Figure 29a and Figure 29c, that in the case of optimal control, the resulting schedule depends on the electricity price profile for the day-ahead market. The boiler turns on during the off-peak hours of the price profile. The rule-based control appears to behave completely indifferent towards the price signal. After all, it has not been instructed to be responsive to pricing.

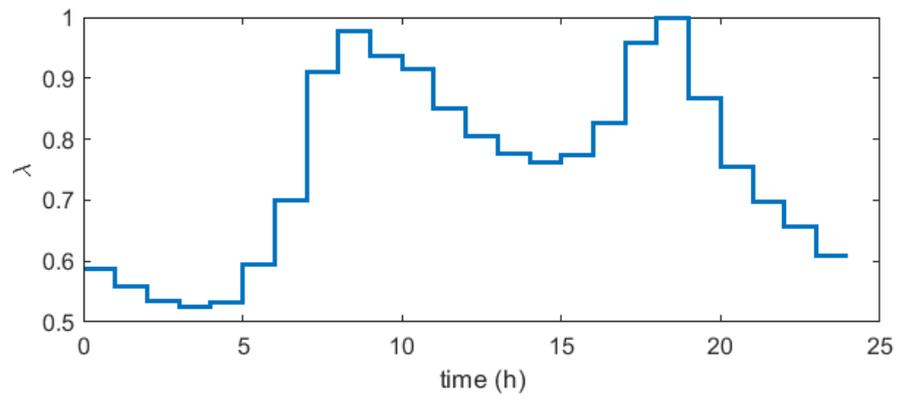
Equation (52) calculates the sum of the hourly electricity costs using the normalised prices, λ , and the maximum hourly electricity price — 48.50 EUR/MWh [63] — in the price profile shown in Figure 22. Figure 30 shows the total normalised electricity costs, $C_{\text{total}}(t)$, from the start of the simulation up to each hour t . It indicates the consistency with which the optimal control leads to a more cost-optimal solution compared to the rule-based control.

$$C_{\text{total}}(t) = 48.50 \cdot \sum_{\vartheta=1}^t P(\vartheta) \cdot \lambda(\vartheta) \quad (52)$$

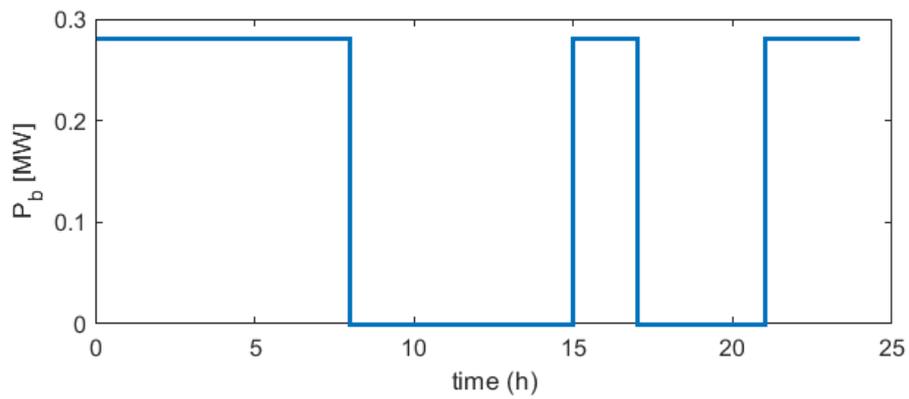
In addition, the total electricity costs for operating the boiler during three consecutive days is calculated for different numbers of greenhouses. Figure 31 shows the total electricity costs per greenhouse over the number of greenhouses throughout the network, for both the rule-based and optimal control. The optimal control is superior to the rule-based control regardless of the number of greenhouses. Naturally, a larger number of greenhouses means a larger total heat demand. Therefore, total operating costs for the boiler increase with the number of greenhouses. Nonetheless, electricity costs per greenhouse show a tendency to slightly decrease with a larger number of greenhouses. However, there is a possibility that with a further increase in the number of greenhouses, electricity costs per greenhouse may rise again, because the distance from the boiler to the furthest greenhouse also increases, thereby increasing heat losses.



(a) Boiler schedule resulting from the rule-based control algorithm



(b) Hourly day-ahead electricity prices



(c) Boiler schedule resulting from the optimal control algorithm

Figure 29: Electric boiler schedules with hourly electricity prices

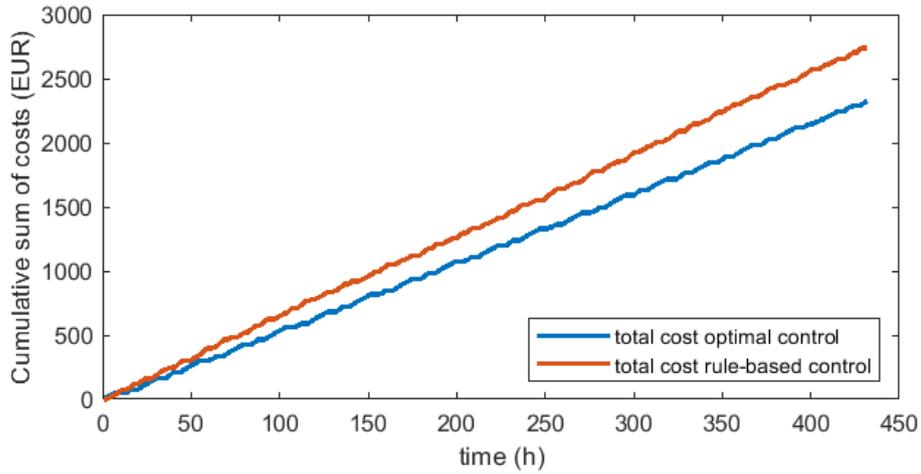


Figure 30: Electricity cost comparison between the optimal control and the rule-based control

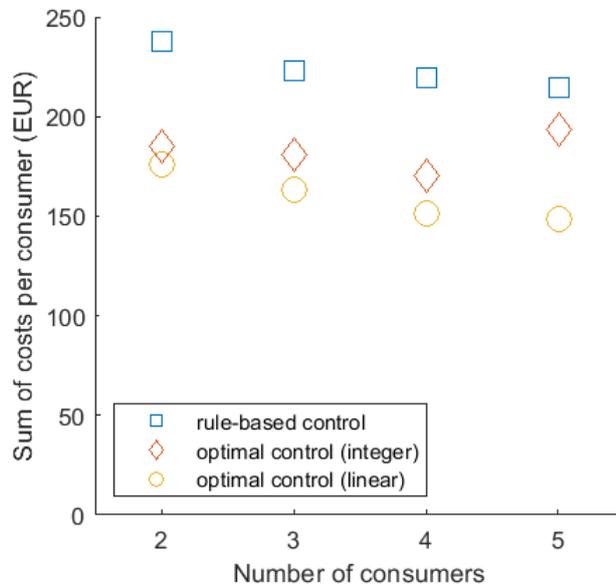


Figure 31: Total electricity costs for different numbers of consumers in the heat network

5.2.3 Integer versus linear programming

Section 3.1.2 already discussed the variety of optimisation problems. The optimal scheduling of the electric boiler in the heat network has been approached as an integer linear programming problem, because the boiler is assumed to be either on or off and each heat exchanger either has a nominal heat flow or no heat flow. The decision variables for the boiler and the heat exchangers are therefore either zero or one (Equation (24) in Section 3.2).

In Section 5.2.1, it was already shown that the algorithm runtime for solving the integer linear programming problem could reach a full hour, after which the algorithm stopped without a feasible solution because the maximum runtime was set at 3600 s. However, the solution for the relaxation of the integer linear programming problem — the first step in the branch-and-bound method — would be found in no time. It is therefore worth to investigate to what extent the integer solution differs from the linear solution. After all, if the solutions differ little, one can consider to use the linear solution and simply round the fractional values to their nearest integer. The gain of a faster runtime must then be weighed against the loss of optimality and reliability. In Figure 31, showing the operating costs for the boiler per greenhouse over the number of greenhouses, the electricity costs resulting from the linear optimal solution are included. As expected, the linear optimal solution leads to a lower total cost than the integer optimal solution, but much of this increase in cost efficiency may be lost when the fractional (linear) solution is rounded to an integer solution.

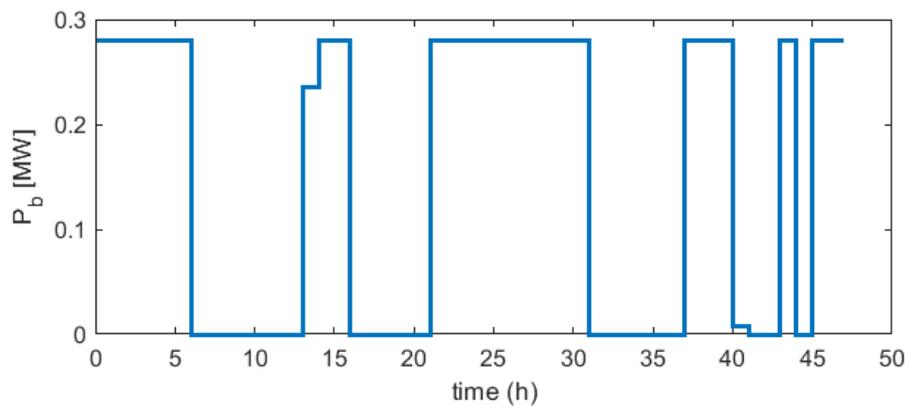


Figure 32: Boiler schedule resulting from optimal linear programming

Figure 32 shows the schedule for the electric boiler use over time, as output from the relaxation of the integer linear programming problem. For most hours, the boiler control signal coincides with an integer; the boiler is either fully on or off. At hours 13 and 40, the boiler is somewhere in between, but clearly close to either on or off. This observation would make the case for using the rounded solution from the linear programming problem. Nonetheless, the linear output for the remaining control variables, those for the heat exchangers, must also be evaluated. In the integer problem formulation, the control variable for each heat exchanger is either zero or one: $u_i \in \{0, 1\}$. Figure 33 shows the heat flow through the exchanger for each greenhouse over time. It can be seen that for many hours, the heat flow is neither zero nor nominal (maximum). Furthermore, the heat flow is often about half the maximum, far from either of the two integer solutions. This means that rounding the linear solution would lead to a completely different operation of the heat exchangers and therefore of the heat network. It is hard to ensure a reliable supply of heat this way, let alone reach a cost-optimal operation of the heat network.

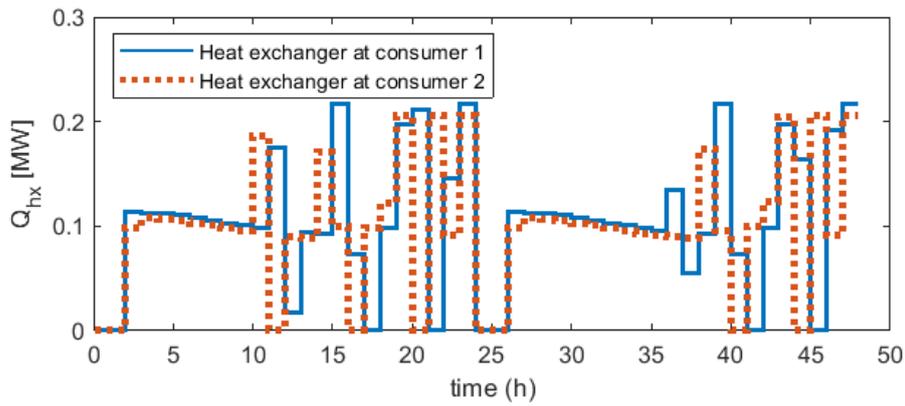
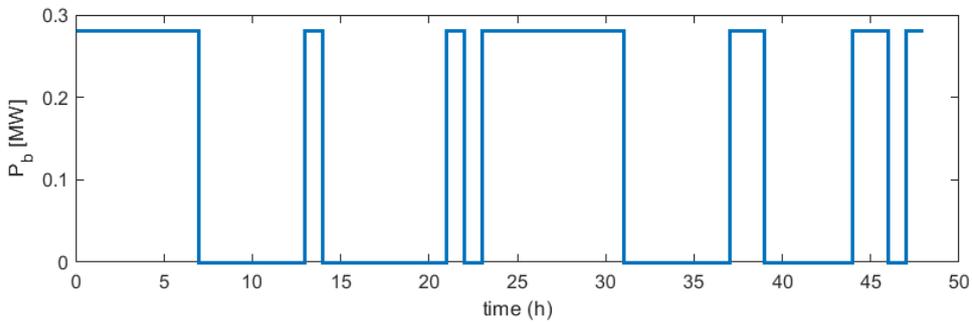


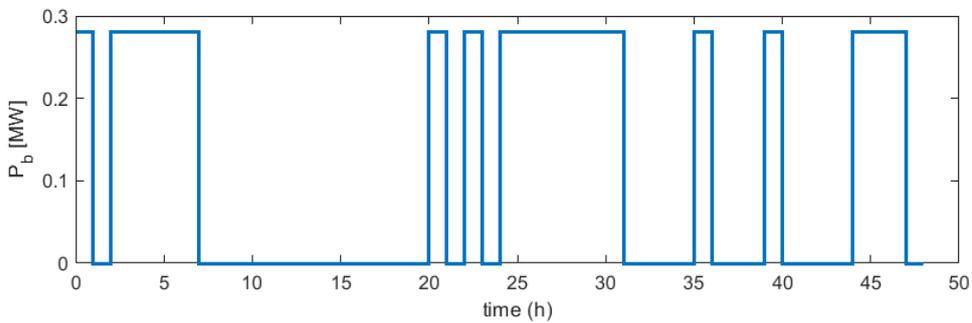
Figure 33: Heat flow through the exchangers as output from linear programming

5.2.4 Doubling heat network time delays

An important aspect of the model in Section 2 are the network time delays for the water to flow from one node to the next (τ in Figures 5 and 11). In the experiments discussed in Sections 5.2.1 and 5.2.2, every time delay equals one time step, $\tau=1$. Next, the boiler schedule will be evaluated for a doubling of the length of every pipeline section, $\tau=2$. Doubling the pipeline section lengths further means that the heat network suffers increased heat losses during transport. This greatly complicates the comparison between single time step time delays and double time step time delays. Therefore, the heat loss coefficients for the pipelines equal zero for this comparison. The heat network with $\tau=1$ over every pipeline section will be called Network 1 and the network with $\tau=2$ will be called Network 2.



(a) Optimal boiler schedule for Network 1



(b) Optimal boiler schedule for Network 2

Figure 34: Electric boiler schedules with hourly electricity prices

Figures 34a and 34b show the day-ahead boiler schedules for two consecutive days for the two differently scaled heat networks. A difference between both schedules that stands out is that Network 2 manages longer periods without turning on the boiler, while Network 1 for example activates the boiler at $t=13$ hours. This observation can be connected to the larger storage capacity of Network 2. Nevertheless, it must be repeated that heat loss coefficients are neglected in the current comparison. Inclusion of these losses would be a limiting effect on the storage capacity of both heat networks, but of Network 2 especially.

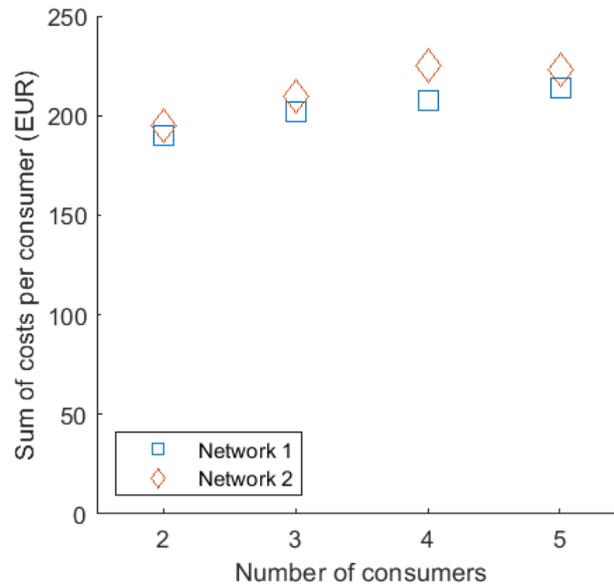


Figure 35: Network time delay cost comparison for different numbers of consumers in the heat network

Furthermore, a comparison can be made regarding the boiler operating costs. Figure 35 shows the electricity costs for both Network 1 and Network 2. The increased storage capacity of Network 2 turns out not to lead to significant reductions in electricity costs related to operating the boiler. Along with the increased storage capacity of the pipelines comes a decrease in the flexibility of the heat demand. In other words, the longer time delays create the need for planning further ahead regarding the boiler use. Once the central boiler is turned off, the temperature of the water leaving the boiler cannot be raised further downstream in the heat network. The increased time delays therefore decrease the flexibility of the system — flexibility necessary to enable the energy demand to respond to electricity prices.

5.3 Boiler load shifting

The results from the experiments related to research questions R2, 2a) and 2b) are now presented and explained. Section 5.3.1 evaluates different sizes for the electric boiler power capacity and how they affect the (boiler) load shifting performance of the heat network. Section 5.3.2 presents a critical analysis of the effectiveness of the boiler shifting and whether the model predictive control is actually productive in mitigating wind power forecast errors. It further proposes changes to the defined constraints for the model predictive control. Section 5.3.3 compares two scenarios in which two different sets of constraints are applied. In addition, the possibility of the heat network operator participating on the intra-day market is discussed.

5.3.1 Electric boiler capacity

Research question 2.b) concerns the performance of the system for different boiler capacities. Equation (53) shows how the boiler power capacity, P_b , has been calculated for each simulation. The boiler capacity depends on the heat demand throughout the network, represented by the nominal heat flow of the various heat exchangers, $\dot{Q}_{hx,i}$, and heat exchanger capacity factors, c_i , which may differ across the set of greenhouses. Furthermore, the boiler capacity depends on the waste heat contribution of the data centre to the network, \dot{Q}_{wh} . The more waste heat there is, the less boiler capacity is needed. Lastly, a virtual — or intended — capacity factor for the boiler, \hat{c}_b , determines the size of the boiler.

$$P_b = \frac{\sum_{i \in \mathcal{S}} c_i \cdot \dot{Q}_{hx,i} - \dot{Q}_{wh}}{\hat{c}_b} \quad (53)$$

Nonetheless, because heat losses throughout the network pipelines are not taken into account in calculating the boiler capacity, the real — or measured — boiler capacity factor, c_b , always turns out higher than the virtual capacity factor. After all, the boiler is scaled too narrowly to maintain a capacity factor low enough while taking into account heat losses during transport. Equation (54) shows how the measured capacity factor is calculated, using the resulting boiler schedule, $P(t)$, and the length of the measured time period in number of time steps, d . Table 7 gives an overview of the boiler capacity factors for the different boiler power capacities.

$$c_b = \frac{\sum_{t \in \mathcal{T}} P(t)}{d \cdot P_b} \quad (54)$$

Table 7: Virtual and measured capacity factors for different boiler power capacities

\hat{c}_b	.15	.20	.25	.30	.40	.50	.60
P_b [kW]	934.3	700.8	560.6	467.2	350.4	280.3	233.6
c_b	.19	.25	.32	.36	.47	.59	.71

After solving for consecutive day-ahead boiler schedules with $\hat{c}_b=.15$ (and thus $P_b=934.3$ kW), the algorithm is set to simulate the continuously changing wind power forecast. The model predictive control described in Section 4.3 is applied to the electric boiler schedule, using the continuously updating wind power forecast as input. The prediction horizon, p_k , is set to 12 hours. The output of the control is a change to the boiler schedule, ΔP . Figure 36 shows the day-ahead boiler schedule, the day-ahead wind power prediction errors, and resulting boiler output after the changes to the existing boiler schedule. It can be seen that during certain hours, the boiler turns on contrary to the day-ahead schedule. At other times, the boiler turns on while the day-ahead schedule dictates otherwise.

Figure 37 shows the day-ahead wind power forecast for each hour and the measured wind power output. The wind power forecast error in Figure 36 is the difference between the two graphs in Figure 37.

With both boiler schedules — the day-ahead, $\hat{P}(t)$, and the actual boiler schedule, $P(t)$ — the total shifted energy shown in Figure 36 can be calculated using Equation (55). Because the power is in megawatts and the time is in hours, the resulting amount of energy, E , is in megawatt hours.

$$E = \sum_{t=1}^{48} |\Delta P(t)| \quad (55)$$

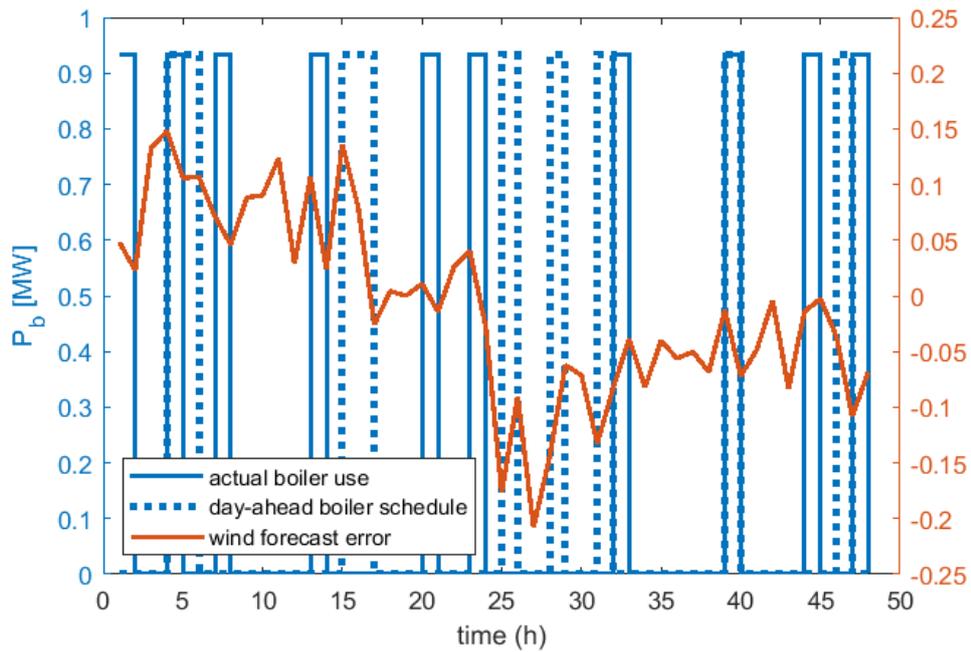


Figure 36: Boiler shifting throughout two days for $\hat{c}_b=.15$ and $p_k=12$

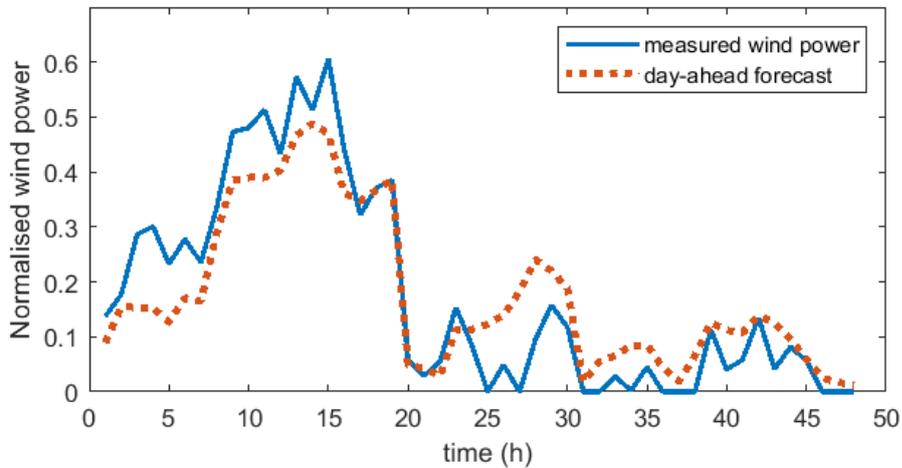


Figure 37: Modelled day-ahead wind power forecasts and the actualisation of the wind power production

Now, the result of the boiler load shifting potential can be presented for different boiler capacities. The algorithm has performed a simulation for every combination of capacity factor, \hat{c}_b from .15 to .6, and prediction horizon, p_k from 6 to 16 hours ahead. The total energy that was shifted, E , is calculated for each simulation. The results for total energy from the simulations with the same combination of \hat{c}_b and p_k are combined into an average of total energy shifted. In Figure 38, the varying capacity factors are positioned on the vertical axis, while the numbers of hours in the prediction horizon for the different experiments are positioned on the horizontal axis. Within the diagram, the averaged results for total energy are presented along the lines of their relevant capacity factor and control horizon from the experiments. A color is assigned to each value in the diagram for the energy shifted, based on the magnitude of the value. Such a color map is meant to help discover trends in the results. Figure 38 shows that the larger the boiler size, the larger the amount of energy is shifted.

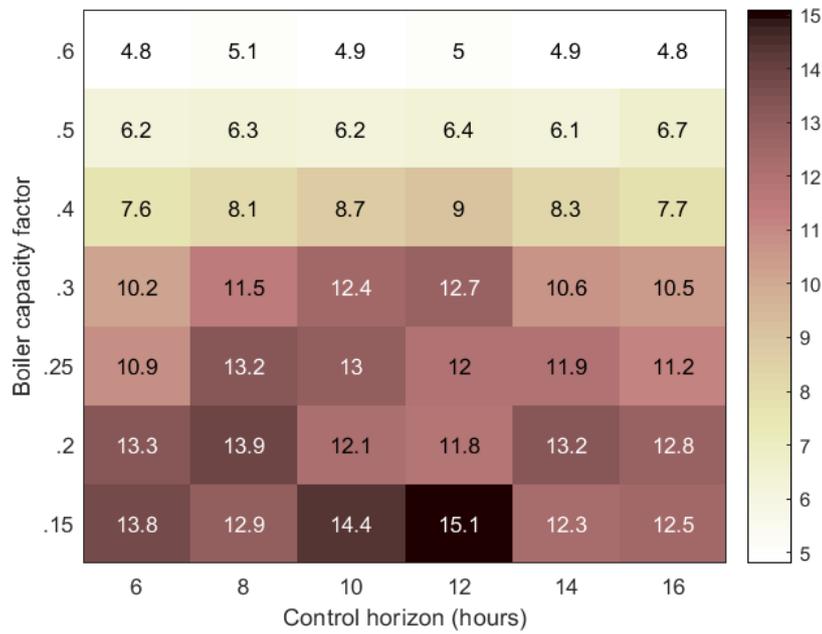


Figure 38: Total shifted energy throughout two days in megawatt hours

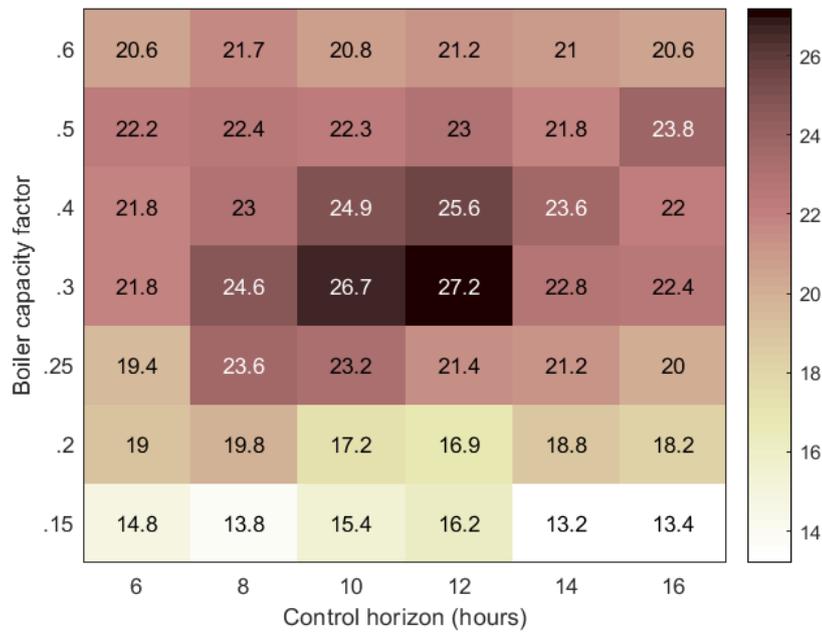


Figure 39: Boiler shifting throughout two days in total shifted load hours

In addition to the total shifted energy, the total number of load hours shifted, N_h , can be calculated for each simulation, using Equation (56). Equation (56) is applied to the values of E in Figure 38 and the boiler power capacities, P_b in Table 7. This results in each value of E in Figure 38 being translated to a value of N_h in Figure 39, which shows the total number of shifted load hours for each simulated combination of boiler capacity and prediction horizon. An optimum appears for a boiler capacity factor of .3 and a prediction horizon of 12 hours.

$$N_h = \frac{E}{P_b} \quad (56)$$

5.3.2 Effectiveness and reliability

Upon close examination of Figure 36, one eventually notices a great failing of the control system. During a number of hours, the model predictive control dictates a change in the boiler schedule to activate the boiler. Problematically, this occurs as well at moments of negative wind power forecast errors. This means that in addition to the wind farm operator having to buy electricity from the intra-day market, the heat network operator also has to buy from the intra-day market. Conversely, the model predictive control sometimes decides for the boiler to turn off, thereby contradicting the day-ahead schedule, while the wind power production exceeds the prediction. Consequently, the heat network operator has to sell the power for that hour on the intra-day market in addition to the wind farm operator having to sell on the intra-day market. In short, the boiler load shifting result from the simulation shown in Figure 36 indicates the model predictive control is, at times, counterproductive because the boiler shifting is ineffective.

An explanation for the problem is that the change in the boiler schedule as output of the model predictive control is directly implemented for the entire prediction horizon. As a result, if the rolling wind power forecast indicates a negative prediction error a number of time steps ahead, the control system does not take the schedule change into account when the prediction error becomes positive in a following forecast update. The model predictive control then does not always change the boiler schedule back to the day-ahead schedule for those particular hours.

In response to this issue, the model predictive control system and the constraints in the optimisation as part of the control are to be re-evaluated. The first proposed adjustment to the control is to implement only the control output for the first time step in the prediction horizon. The model predictive control over the time horizon ahead then only functions to ensure the reliability of the heat network operation — i.e. whether the heat demand is satisfied. The model predictive control is thus no longer used as a continuously updating scheduling tool for the boiler. The predictive aspect of the model predictive control is solely aimed at the response of the heat network to a change in the boiler use during the current time step. This resolves the problem of the changing wind power forecast errors, because whether the wind power forecast error is positive or negative the next hour is highly certain and is not subject to any large changes within that hour.

Another adjustment to the model predictive control is to leave out the constraint in Equation (47). This allows the boiler to extend its load shifting potential across the prediction horizon and, as it turns out, the thermal capacity of the heat network is sufficiently large to handle several hours of either continuously positive wind power forecast errors or continuously negative wind power forecast errors. In short, when load shifting within the prediction horizon is not possible, changes to the boiler schedule are not held back any longer by the constraint. A load can thus be shifted across the prediction horizon.

To be able to compare the effect of these changes, the effectiveness of the shifted loads must be quantified. As explained earlier, the amount of shifted energy can be calculated using Equation (55) and the number of load hours shifted can be calculated using Equation (56). Nonetheless, not all shifted loads effectively reduce the need for the wind farm operator to buy and sell at the intra-day market, as can be seen in Figure 36. A shifted load is effective when a positive change in the boiler schedule ($\Delta P > 0$) coincides with a positive wind power forecast error ($\phi > 0$) or when a negative change in the boiler schedule coincides with a negative wind power forecast error. Effectiveness, ϵ , is thus defined in terms of boiler load shifted for demand response purposes, E_{DR} . The effectively shifted load is calculated by Equation (57) using the cost function (45). Equation (58) then calculates the fraction of effectively shifted load in the total shifted load.

$$E_{DR} = \sum_{t=1}^{48} -c(t) \cdot \Delta P(t) \quad (57)$$

$$\epsilon = \frac{E_{DR}}{E} \tag{58}$$

To summarise the adjustments to the control: the constraint in Equation (47) is left out and only the control output for the first time step is implemented. The performance of the adjusted control system is now evaluated. Figure 40 shows the wind power forecast errors as input and the shifted loads as output of the adjusted model predictive control. The boiler shifting shows a significant improvement compared to the result in Figure 36. The adjusted model predictive control still shifts a comparable amount of energy, E , but has fewer hours of counterproductive boiler shifting. Therefore, more of the energy demand is actually responsive to the wind power forecast errors and therefore the effectiveness, ϵ , is higher. Figures 41 and 42 show the averaged boiler load shifting results of the simulations with different combinations of the boiler capacity factor, \hat{c}_b , and length of the prediction horizon, p_k .

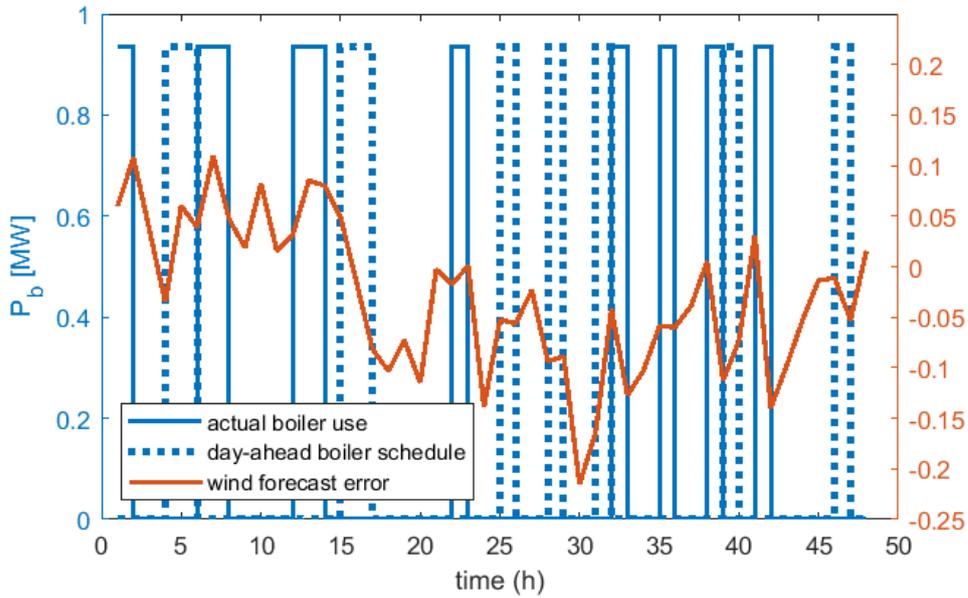


Figure 40: Boiler shifting for $\hat{c}_b=.15$ and $p_k=12$ with adjusted MPC

Even though the boiler only implements the first time step of each control output, intra-day participation by the heat network operator is still observed in the boiler shifting results. The explanation lies in the cost function (45) and the bounds for the control variables (44). The control variable bounds only ensure that the electric boiler does not exceed its power capacity nor uses a negative power. Nonetheless, the bounds are completely indifferent to the wind power forecast error and thus allow the model predictive control to turn on the boiler during negative wind power forecast errors or to turn it off during positive wind power forecast errors. In both cases, the heat network operator would be forced to participate on the intra-day market, as a power exchange with the wind farm is not possible due to the forecast error suffered by the wind farm operator itself.

The cost function (45) further allows it, but does add a cost, or penalty, to the intra-day market participation. To repeat, intra-day prices are not considered in this report. Instead, a cost of -1 or $+1$ is used in the cost function, depending on the wind power forecast error. The total amount of shifted energy in Figure 41 and the total number of shifted load hours in Figure 42 therefore do not present only the boiler loads shifted that prevent intra-day market participation, but also load shifts that force the heat network operator to participate.

The bounds for the control variables can be adjusted to completely prohibit the heat network operator from buying or selling at the intra-day market. A comparison can then be made between the scenario in which the heat network operator is allowed and the scenario in which the operator is not allowed to participate in the market, in terms of the effectively shifted loads — i.e. the loads shifted to prevent intra-day market participation and not the shifted loads that force participation. Section 5.3.3 will discuss the results of that comparison.

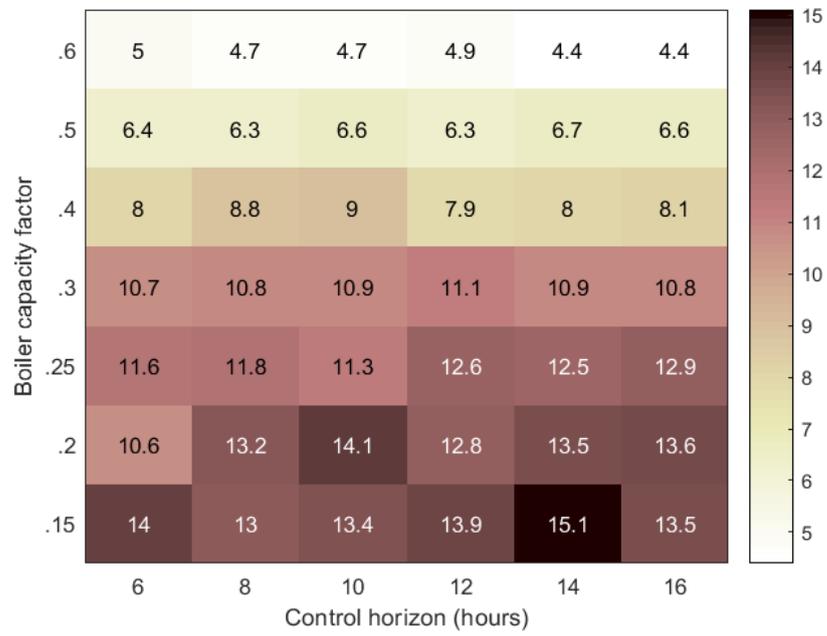


Figure 41: Total shifted energy in megawatt hours with improved MPC

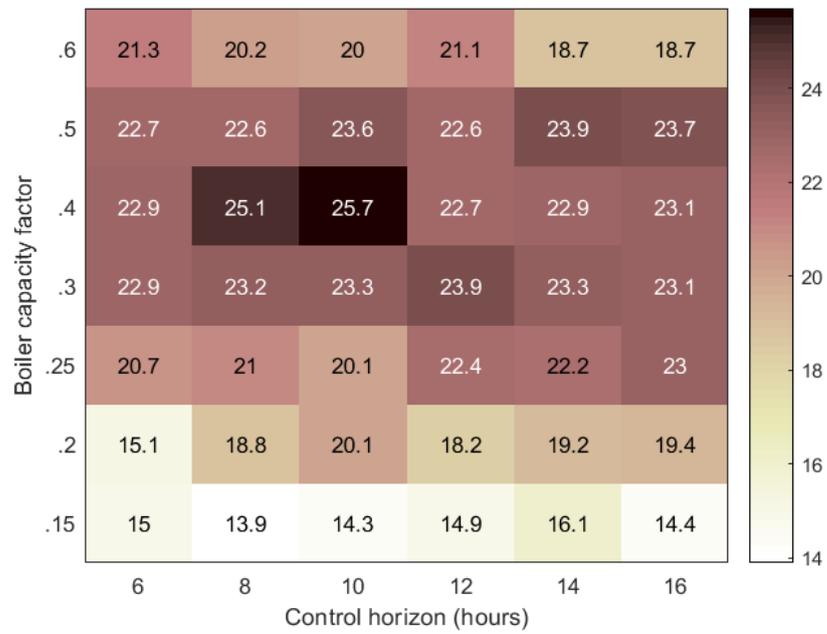


Figure 42: Boiler shifting in total shifted load hours with improved MPC

5.3.3 Heat network operator at the intra-day market: a comparison

Only when the heat network operator is not allowed to buy from the intra-day market, does the effectively shifted load equal the total shifted load. Otherwise, the effectiveness, ϵ , is often less than one. To translate into a constraint whether or not the heat network operator is allowed to exchange power at the intra-day market, a change in the bounds for the control variables is to be made. Table 8 shows the control variable bounds for both scenarios. The additional condition in Scenario 2 ensures each change to the boiler schedule only occurs if each hourly load change effectively decreases the need for intra-day balancing on its own. As described earlier, this is the case when a positive change in the boiler schedule coincides with a positive wind power forecast error or when a negative change in the boiler schedule coincides with a negative wind power forecast error.

Table 8: Control variable bounds for both scenarios of network operator intra-day market participation

Scenario 1: Allowed to participate	Scenario 2: Not allowed to participate
$u_k(t) \in \begin{cases} \{-1, 0\} & \text{if } \hat{P}(t) = P_b \\ \{0, 1\} & \text{if } \hat{P}(t) = 0 \end{cases}$	$u_k(t) \in \begin{cases} \{-1, 0\} & \text{if } \hat{P}(t) = P_b \wedge \phi(t) < 0 \\ \{0, 1\} & \text{if } \hat{P}(t) = 0 \wedge \phi(t) > 0 \end{cases}$

Scenario 1 allows the heat network operator to go around this constraint and buy electricity from the intra-day market when needed. This way, the operator could shift loads more freely, knowing it can buy its way out of an emergency to ensure reliability of the heat network. The comparison will focus on the total effectively shifted loads and load hours.

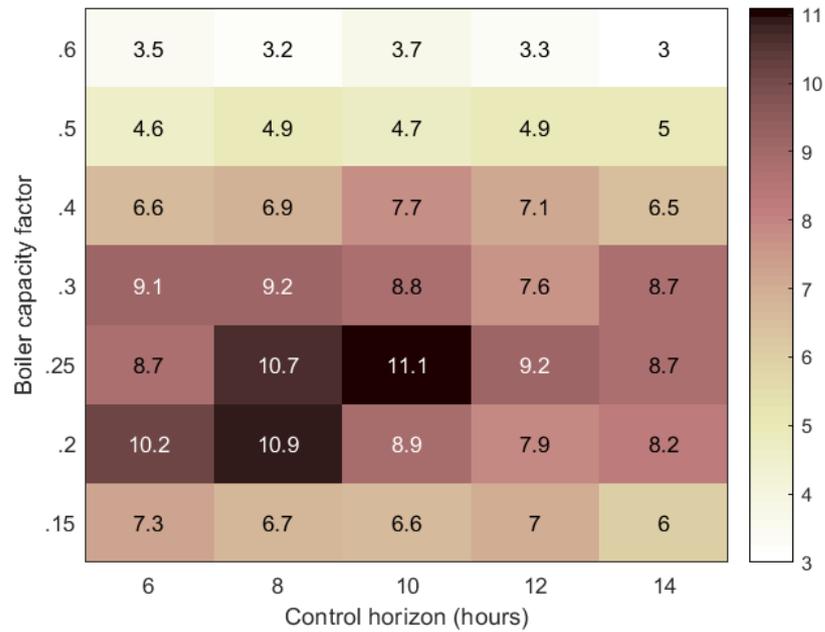
First and foremost, it must be stated that the majority of simulations for Scenario 2 could not simulate more than 48 hours. The simulation would stop because no feasible solution could be found. For the heat network operator not to be allowed to buy energy from the intra-day market is a limiting factor. Seemingly, it is limiting to the point where a reliable operation is not possible. The heat network could potentially run out of heat, while not being able to buy from the intra-day market in such emergency. Still, for the sake of the comparison and the evaluation of the effectiveness of the model predictive control in Scenario 1, the successful simulations were recorded and the boiler load shifting results are shown in Figures 43a, 43b, 44a and 44b. In all results, for both scenarios, some optimum for the shifted energy, E , and the shifted load hours, N_d , could be found for a combination of values for the control horizon, p_k , and the boiler capacity factor, \hat{c}_b . Table 9 gives an overview of the optimal values for E and N_h , together with the parameter values for p_k and \hat{c}_b that lead to these optima.

Optimal	Scenario 1		Scenario 2	
	E	N_h	E	N_h
p_k [h]	10	10	8	8
\hat{c}_b	.25	.4	.25	.4
Optimum	11.1 MWh	21.9 h	10.7 MWh	21.6 h

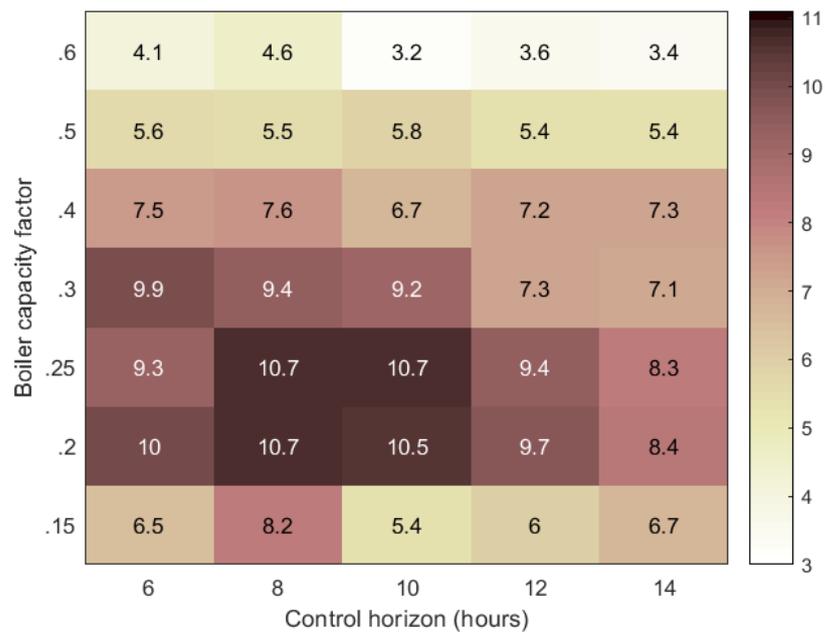
Table 9: Optimal values and their associated parameters, from Figures 43a, 43b, 44a and 44b

Scenario 2 appears to favour a smaller prediction horizon, p_k in comparison with Scenario 1. Namely, p_k equals 8 hours for the boiler load shifting optima in Scenario 2, while p_k equals 10 hours for the load shifting optima in Scenario 1. This makes sense, because a smaller prediction horizon is coupled to smaller uncertainties in the rolling wind power forecasts. In Scenario 2, the heat network operator is more reliant on the forecasts for its own reliability, because no corrections to the boiler schedule can be made using the intra-day market. Remember, most simulations of Scenario 2 failed to reach the end of the period, because no feasible solutions could be found for some time step. The risk of failing to find a feasible solution increases when more load is shifted. More load is shifted on average for a larger prediction horizon, as can be seen in the results for Scenario 1. However, the simulations of Scenario 2 for larger prediction horizons, in which many loads were shifted, may have had a larger chance of failing. Therefore, the results for larger prediction horizons with many shifted loads may not have been recorded.

Table 9 further shows that in both scenarios the most energy was shifted with a lower boiler capacity factor ($\hat{c}_b=.25$), while the most load hours were shifted with a higher boiler capacity factor ($\hat{c}_b=.4$). It is here



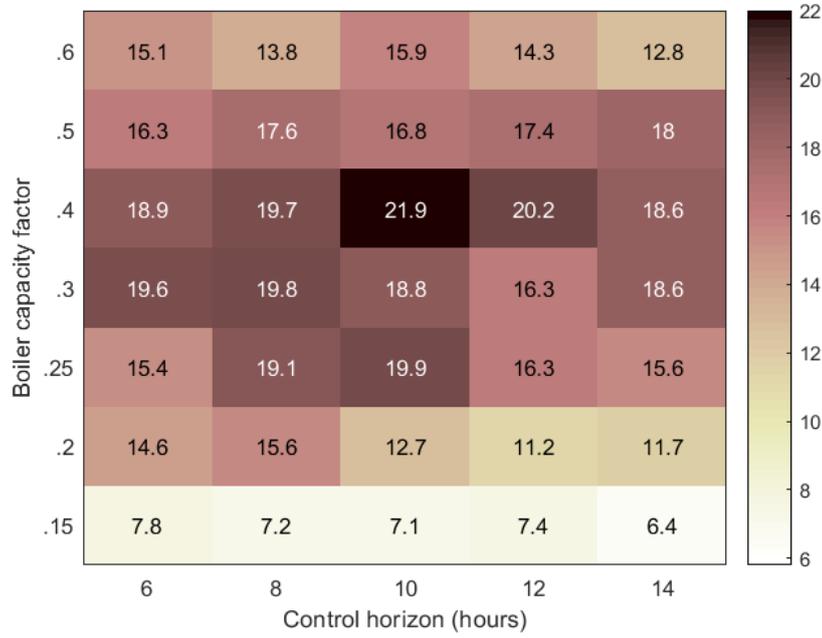
(a) Energy shifted in Scenario 1



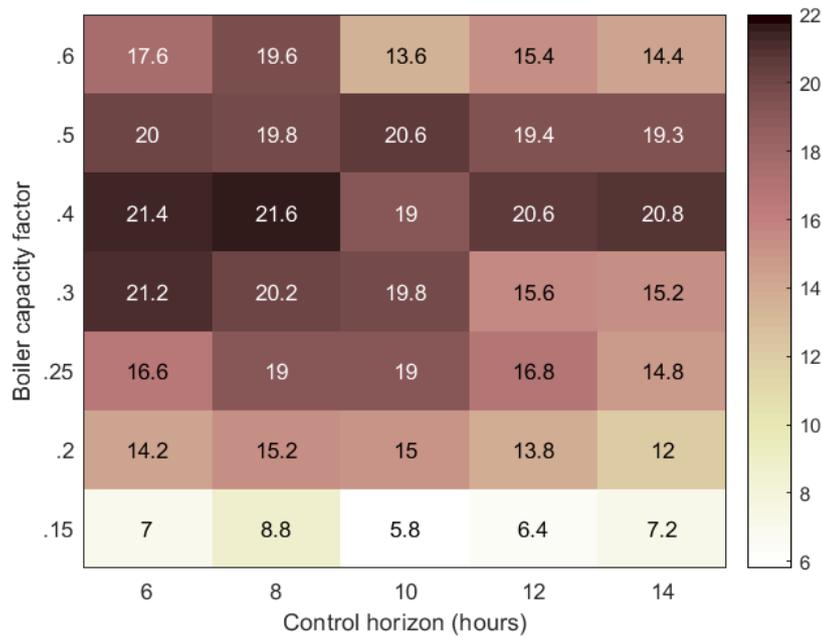
(b) Energy shifted in Scenario 2

Figure 43: Total effectively shifted energy in megawatt hours

that it becomes relevant to put these results in a broader perspective. First, a lower boiler capacity factor is realised by a larger size boiler. Larger electric boilers can be assumed to have higher installation costs, which must be taken into account during the early steps of the construction, expansion or electrification of a heat network. The experimental setting described in Section 5.1 assumes winter season conditions. Naturally, the heat demand is lower during summer. An electric boiler that is scaled to optimally shift megawatt hours of energy during winter, is likely to be considered oversized throughout the rest of the year.



(a) Load hours shifted in Scenario 1



(b) Load hours shifted in Scenario 2

Figure 44: Total effectively shifted load hours

6 Discussion

The results in Section 5 have shown that the optimal control works and can be used for operating power-to-heat in a heat network. In this section of the report, a more detailed discussion of the results is provided, with a focus on the doubling of the time delays and on the shifting of the boiler use during intra-day operation. This section relates to the research questions by highlighting when the method works and how the results are applicable, but also why the carried out research could not always reliably provide a definitive result, along with some limitations of the method presented in this report. Furthermore, a number of recommendations for future work will be provided in Section 6.3. First, however, the implications of the results on the research questions are described.

R1. *How to plan a day ahead the operation of power-to-heat as a centralised source in a heat network?*

- a) *How to minimise electricity costs when scheduling the electric boiler considering heat network time delays?*
- b) *How to schedule the heat demand to maintain consumer temperature levels?*

The results of solving the day-ahead scheduling problem indicate a cost-minimal operation of the electric boiler can be achieved through integer linear programming. For this, the heat demand of the two greenhouses, i.e. when a greenhouse extracts heat from the network, must also be scheduled to follow the boiler activity. Hourly day-ahead electricity prices form the cost function in the objective function used in the optimisation problem. The predicted wind power production is reflected by these prices. Therefore, these results build on existing evidence that power-to-heat can contribute to the integration of wind power in the energy mix. The aspect of the time delays throughout the heat network, related to R1a (and R1b), will be discussed further in Section 6.1.

A limitation in this research lies in the assumption that day-ahead electricity prices are known beforehand. Furthermore, the same averaged electricity price profile has been used in each simulation. Although wind power production and electricity prices are negatively correlated, other factors too play a role, as described in Section 3.1.1. It has therefore been decided not to model the day-ahead price profile using the simulated wind power production forecast, because that would lead to many more assumptions, including ones that cannot be justified within the scope of this report. Other limitations of the results relate to the scalability of the model and the applicability to existing heat networks. The base case described in Section 5.1 assumes a heat network with two greenhouses, unidirectional flow of water in the pipelines and distances between greenhouses represented by a multiple of the time step size. It has been shown that the model is scalable with an increasing number of greenhouses and increasing pipeline lengths. It can therefore be applied to heat networks connecting multiple greenhouses that are kilometres apart. However, some existing heat networks serve thousands of households and have multiple secondary heat networks connected to the primary heat network. The current model does not account for these secondary grids. Neither would the runtime of the computationally complex integer linear programming allow the optimisation for thousands of consumers within the time scope suitable to day-ahead scheduling, as the optimisation problem grows with each consumer. Households in a secondary grid must therefore be grouped as a single consumer at the substation connecting their secondary to the primary grid. This may have consequences for the thermal satisfaction of individual households, as they are not taken into consideration by the approach described here. In short, the model presented in this report is a simplification of a realistic heat network. The method to formulate the model can be used to model heat networks with different components, properties and topologies, but this would require adjustments to the model presented here and a sufficient amount of network-specific data. Furthermore, the operation of a heat network is much more complex in reality. This report focused on operating a centralised electric boiler and the heat flow through exchangers, but did not consider pressure differentials or pump operation. Although the thermodynamics and the hydraulics would likely affect one another when considered simultaneously, the results from the thermal model encourage a further investigation into the potential of power-to-heat in heat networks.

R2. *How to mitigate wind power forecast errors during intra-day operation of the heat source?*

- a) *How does the performance scale with the control horizon?*
- b) *How does the performance scale with the boiler capacity?*

The results for mitigating wind power forecast errors during intra-day operation of the electric boiler indicate a model predictive control with updated wind forecasts can show a significant effect by shifting the electric boiler use. The heat loads must shift to follow the boiler load shifting. Plotting the (averaged) results of the shifted boiler load over the boiler capacity factor on one axis and the prediction horizon on the other — for simulations with different combinations of values for the boiler capacity factor and the prediction horizon — shows a trend towards an optimum at a specific combination of boiler capacity factor and prediction horizon. The colour maps in Figures 38, 39 and 41 to 44 show a clear example of this observation. These results should be taken into account when considering how to scale the boiler capacity and how to control an electric boiler as centralised heat source in a heat network during intra-day operation for the mitigation of wind power forecast errors as a form of demand response. Nonetheless, the capacity factor of the electric boiler will always be higher during colder months than it is during warmer months, because of the correlation between heat demand and outside temperature. Depending on the specific scenario that is to be evaluated, already present waste heat sources may potentially be sufficient to satisfy the heat demand throughout a heat network during summer, resulting then in an infrequent use of the boiler. Therefore, due to a higher capacity factor during winter, the electric boiler might have a higher potential of demand response during cold months for a high wind power penetration into the grid. A noteworthy limitation in the method to evaluate the boiler load shifting potential lies in the wind power simulation model. In the real world, much more advanced forecasting tools exist, relying not only on historical data but also atmospheric measurements. The wind power simulation model described in this report is an attempt to make the best use out of the accessible data for the day-ahead and 5- and 1-hour-ahead wind power forecasts with the actually measured wind power data. Consequentially, the resulting simulation model may not be without limitations. Nonetheless, the simulated wind power production fits the characteristics of the historical wind power data. The model predictive control can therefore still be presented or interpreted to perform well, even with a different wind simulation model. A further discussion of the boiler shifting and the related limitations of the approach presented in this report is provided in Section 6.2.

R. *How to balance uncertain wind power production using a heat network?*

The results in Section 5 indicate an electric boiler as central heat source in a heat network can balance the variability of wind power production. The boiler can adjust its electricity demand to the hourly wind power production. This lowers the necessary installed power generation capacity in the electricity grid. Furthermore, the electric boiler can mitigate the uncertainty of wind power production by responding to both positive and negative wind power forecast errors. This lowers the need for installed balancing power capacity. Overall, demand response through a flexible operation of power-to-heat can be considered a cost-efficient contribution to help solve the intermittency problem. Nonetheless, the method and the results presented in this report come with a number of limitations that cannot be ignored. The model described in Section 2 is formulated using assumptions that might not be valid when applied to an existing heat network, such as that each time delay is a multiple of the time step size. Further simplifications of the model would have to be weighed against the reliability of the results in each case to be studied. Overall, the findings in this report offer valuable insights for operating a heat network with power-to-heat.

6.1 Network time delays

One of the central ideas behind using the heat network as a buffer has been to enable a flexible energy demand through the thermal inertia of the network pipelines. Longer time delays are a result of longer pipelines, which would increase the heat storage capacity of the network. In Figures 34a and 34b, a heat network with doubled time delays ($\tau=2$) and a heat network with $\tau=1$ have been compared. The comparison between the operation schedules for both heat networks shows no significant difference in energy flexibility or electricity costs from operating the boiler. Nonetheless, the one parameter that will always affect energy flexibility in a heat network, is the thermal inertia of the two greenhouses. To be more precise, both the heat capacity of the heat carrier on the greenhouse side and the heat loss factor determine the flexibility in heat demand, because they determine the change in greenhouse temperature over time. Since from early on it was established that the greenhouse demands its temperature to remain within temperature limits, these parameters determining the thermal inertia of the greenhouse are decisive in enabling a shift of heat load over time. After all, an increased thermal inertia enables stretching both the periods without heat extraction from the network and the periods of continuous heat extraction from the network.

A decisive cause of why energy flexibility of the system does not benefit from increased pipeline lengths while the thermal inertia of the greenhouses does affect system flexibility, is the minimum supply temperature constraint (27). This constraint fixates the electricity consumption by the boiler to the supply of heat from the network to the greenhouse. When an active boiler raises the temperature of the supply water for a single hour to a sufficiently high temperature, i.e. to satisfy the constraint, the greenhouse only has a single hour to extract heat from the supply pipe. This, together with the optimal scheduling results, does not prove the time delays in a heat network support energy flexibility of the system.

Nonetheless, there are scenarios to which this conclusion may not apply. Throughout this report, a constant mass flow and variable temperature control strategy was assumed. Variable — and controllable — mass flow rates would enable some flexibility in the network time delays, the time for water to flow from one node to another. This might increase the energy flexibility potential of the storage capacity allowed by the thermal inertia of the heat network pipelines. However, it would also fundamentally alter the model of the heat network used in this research, which at the basis assumes a constant mass flow and variable temperature control strategy. The hydraulics model necessary for variable flow control is nonlinear and would greatly complicate solving the research problems.

6.2 Boiler shifting performance

In Section 1.2.2, it was hypothesised that both boiler size and prediction horizon affect the effectiveness of the load shifting potential of the electric boiler. The results in Figures 38, 41, 43a and 43b show most clearly that the boiler capacity has a clearer effect on the load shifting performance of the boiler than the length of the prediction horizon, for the analysed arrays of values for these parameters, in terms of total energy. Furthermore, when only the total shifted energy is analysed, the optimum is with the largest boiler power capacity — or smallest capacity factor. This contradicts the hypothesis that the optimum would not be at one of the extremes, meaning either the larger boiler the better or the smaller boiler the better. To be more precise, the disadvantage of a large boiler capacity is that the electric load would be too large to shift; any change in the boiler schedule would have a too great of an impact on the ability to satisfy the needs and constraints of the heat network. However, with the evaluated boiler capacities, this theoretical disadvantage has not shown to affect the total shifted electricity demand of the boiler, since the optimum is at the largest boiler capacity. Nonetheless, when evaluating specifically the effectively shifted loads, as in Figure 43a, the optimal boiler capacity factor is .25, which translates to a boiler power capacity of 560.6 kW.

The boiler load shifting performance can be further evaluated based on another factor, the total number of shifted load hours. The results for this factor are shown in Figures 39, 42, 44a and 44b. Interestingly, the performance of the boiler in the simulations in terms of shifted load hours does seem to clearly indicate a dependency on the choice for both the boiler size and the prediction horizon. Furthermore, for the total shifted load hours in Figure 42 and the effectively shifted load hours in Figures 44a and 44b, the results from the simulations show an optimum for both the boiler size ($\hat{c}_b=4$) and the prediction horizon ($p_k=10$ hours). This confirms the hypothesis formulated in Section 1.2.2 stating that there is an optimal boiler size and prediction horizon regarding the load shifting performance of the boiler. After all, if the prediction horizon is too small, the number of time steps available to shift loads might be limited, while a too large horizon might suffer from the increasing uncertainty in wind power forecasts. Similarly, if the boiler is too large,

the heat network might not be able to handle the shifting of such a load. It must be noted, however, that the difference in the results for total shifted energy and total shifted load hours had not been anticipated. The results of the simulations have certain implications for the use of the heat network as a demand response system in support of wind power producers. These implications will be discussed now.

Consider the wind farm integrated into a future energy system with a large share of renewable sources in the energy mix. It is likely there is more demand response capacity, consisting of flexible loads, in the system to provide balancing power for the wind farm than just the heat network. In this case, it may not be the most important that the heat network can shift the most of its energy demand, because the other demand response systems — other than the heat network — can give additional balancing support. Therefore, it might be more desirable to be able to shift a larger number of hours, to increase the chance that the load flexibility of the heat network can be used at moments when the other flexible loads cannot be shifted. Nonetheless, if the heat network is the only flexible energy load supporting the wind farm, the scaling of the electric boiler, from the perspective of maximising the amount of energy that can be shifted, should be weighed against the higher installation costs of a larger electric boiler.

Based on the boiler load shifting performance using a signal from the wind, it can be concluded it is worth to investigate load shifting using a price signal from the intra-day market. This would add another dimension along which to optimally control the heat network. The cost function used here is uniform — every hour of intra-day trading has the same cost — whereas when using intra-day prices, the cost function changes from hour to hour. A different cost function likely leads to a different control output. This is especially relevant when the wind power producer has a number of demand response systems available for load shifting. The first energy demand to respond to the changing wind power availability is the load with the lowest cost to shift it. This would result in a more diversified objective function in the control system that optimises the load shifting.

Now, another aspect of the cost function will be discussed. The conditions regarding ϕ in the cost function (45) and in Table 8 compare the wind power forecast error, ϕ , to the number zero. This determines whether the forecast error is positive or negative and, in turn, whether the error is positive or negative affects whether the boiler schedule should change or not. However, if the wind power forecast error is so close to zero that the total residual load is less than half of the boiler power capacity, P_b , a change in the boiler schedule may in fact result in a larger residual load. A negative residual load would become a larger positive residual load and a positive residual load would become a larger negative residual load. In short, the absolute value of the residual load would increase due to the output of the model predictive control. A solution to this would be to change the conditions regarding ϕ and compare the wind power forecast error not to zero but to the boiler power capacity. Because the installed power capacity of the wind farm, P_{inst} , was assumed to be that much larger than the boiler capacity, the number zero was used in the conditions — if $P_{inst} \gg P_b$ and $P_{inst} = 1$, then $P_b \approx 0$. Furthermore, instead of a wind power forecast error in the unit for power, watt, a normalised wind power forecast error was used, relative to the power capacity of the wind farm. With a given size of the wind farm in question, it would be recommended to translate the normalised wind power forecast error, ϕ , to a power value in watt and to compare that value with the boiler power capacity in the cost function.

6.3 Future work

A set of recommendations regarding different aspects of this research are discussed here for future work. These alternative approaches include a variable flow control of the water in the heat network pipelines, the decentralisation of electric boilers or heat pumps, and a fine-grained control for operating the heat flow through the exchangers at the consumer substations. A more general recommendation here relates to scaling the system dimensions; it is recommended to use the rule-based operation to scale the various components and parameters throughout the heat network. If a component is not scaled properly, the optimal control would simply indicate there is no feasible operation possible after running for a long time, whereas the rule-based operation would still show results, making it quicker and easier to spot the abnormalities and correct the relevant parameter. In addition, the output from the wind power simulation model might not be a realistic representation of a possible wind power forecast. A lack of accessible data has played a limiting role here. In future work, it is recommended to use an advanced wind power forecasting tool with access to more data.

Variable flow control

As discussed before, a variable flow control of the heat network would surely complicate optimising the operation of the network due to the necessary inclusion of the nonlinear hydraulics model. Future research on the matter of operating power-to-heat as a source for heat networks could build on the linear thermal model by adding the nonlinear hydraulics model to enable the variable mass flow and variable temperature control strategy. For example, it would be interesting to find out the effect of variable flow control on changing heat network pipeline lengths and time delays.

Decentralised power-to-heat

Another solution to enable heat storage and energy flexibility through the heat network pipelines is to install multiple decentralised power-to-heat systems, instead of a single centralised electric boiler as used in this report. This way, the location where electric energy is converted to heat can be taken into account in the optimal control of the boilers. The electricity consumption by the boilers is therefore not fixed to a single time delay until the added heat reaches the consumer. A flexible power demand by the boilers is thereby enabled while ensuring the consumer can extract heat from the network to maintain its temperature levels.

Fine-grained control

In Section 5.2.3, it has been shown that the relaxation of the integer linear programming problem is already a good approximation of the electric boiler control variables, since the output from the linear programming problem largely coincides with an integer solution. Nonetheless, the linear programming problem cannot function as a proper substitute, because the linear output variables for the heat exchangers are far from a close approximation of an integer solution. A recommendation resulting from these observations is to implement a fine-grained control for the heat exchangers. A variable flow pump on the consumer side of the heat exchanger can control the water flow which affects the temperature difference in and thus the heat flow through the exchanger. This is feasible with a proper control of the pump which would allow a variable heat flow — other than the nominal heat flow — and would enable the use of linear programming to calculate a close approximation of the optimal solution. Even if the hydraulics model is not included, the variable heat flow can be modelled by substituting the discrete bounds of the heat exchanger control variables by continuous bounds, without having to model the heat exchanger itself.

Another recommendation in this light is to have a fine-grained controlled electric boiler in the load shifting experiments. Although the linear programming problem did not lead to overwhelmingly non-integer solutions for the day-ahead operation schedules, a boiler with a fine-grained control might show the benefits of linear programming especially in the load shifting experiments. The argument for why especially the load shifting might benefit from a fine-grained controlled boiler is because not necessarily the entire load needs to be shifted. Already fraction of the boiler load might contribute to the load shifting potential, without greatly affecting the operation and reliability of the heat network. Note that a fine-grained control of the central boiler would be the same as having a multitude of smaller boilers as the central heat source. Multiple smaller boilers could trigger the investigation into the effect of decentralising the electric boilers throughout the heat network.

7 Conclusion

One of the challenges central to the energy transition is the intermittency of renewable energy sources for electricity. Fluctuations in wind power production can be balanced by power-to-heat technologies such as an electric boiler. In this report, it has been shown that an electric boiler as central heat source in a thermal grid could alleviate uncertainty in power generation by a wind farm. The modelled heat network has a sufficient thermal inertia to allow a flexible operation of the centralised electric boiler. Furthermore, shifting of the boiler use can respond to updating wind power forecasts, thereby relying on the most recent and therefore most accurate prediction for the wind.

Planning a day ahead the operation of an electric boiler as a centralised source in a heat network can be done by formulating this objective as an optimisation problem. Because the day-ahead wind power forecast is already to some extent reflected by the day-ahead electricity price profile, these electricity prices comprise the objective function in the optimisation problem. The solution to the optimisation problem includes a day-ahead schedule for when to turn the boiler on and when to turn it off. This report shows that the day-ahead boiler schedule is price-responsive and therefore responsive to wind power availability, since electricity prices have a negative correlation with wind power production. The following can be concluded further in relation to the day-ahead boiler scheduling:

- The linear relaxation of the integer linear programming problem does not approximate an integer solution and can therefore not substitute integer linear programming, unless a fine-grained control can be applied to the extraction of heat through the exchangers at the greenhouse substations.
- In case of a heat network with a constant flow and variable temperature control, the lengths of the heat network pipeline sections and resulting time delays do not contribute to increased energy flexibility and lower electricity costs related to the electric boiler use.

Mitigating wind power forecast errors during intra-day operation can be done using model predictive control. A continuously updating wind power forecast is compared to the day-ahead wind power prediction and the difference indicates the need for balancing power. Instead of intra-day electricity prices, the signal from the wind comprises the objective function for the model predictive control. Changes can be made by the model predictive control to the day-ahead boiler schedule to respond to this balancing need. This report has focused on the effects of various constraints and of changing certain system parameters, in particular the boiler capacity and the time horizon considered by the model predictive control. The following can be concluded further in relation to the intra-day boiler control:

- In terms of boiler load shifting, there is an optimal combination of parameter values for the power capacity of the boiler and the length of the prediction horizon for the model predictive control. The optima in terms of shifted energy and in terms of shifted load hours do not coincide.
- Allowing the heat network operator to participate on the intra-day market greatly improves the reliability of the heat network but shows no significant improvement of the load shifting performance of the boiler, compared to the scenario in which the heat network operator is not allowed to participate on the intra-day market.

Within the context of heat and renewable power integration, this report is a contribution of insights that could enable heat networks to help mitigate imbalances in the power grid caused by the intermittency of wind power production. In a broader context, it can be concluded that operators of power-to-heat (in heat networks) could participate in intra-day trading and provide balancing power through the intra-day market to help mitigate imbalances caused by any time-varying supply and demand of power.

There is a great uncertainty as to how the future will unfold. It is unknown how the electric power production, electric power consumption, the role of energy storage systems and sector-specific heat demand will exactly develop. Nonetheless, it is important to have insights into the potentialities of each development, so that any consideration can be justified accordingly, when investment decisions are to be made. If there is anything certain, it is that the near-future energy mix is diverse with an increasing share of intermittent sources. The potential of demand side management and demand response will be further exploited, as these are profitable - not only with the current development of injecting renewable power into the grids - considering electricity prices have fluctuated with controllable power sources too.

References

- [1] IEA, "Tracking buildings," International Energy Agency, Paris, Tech. Rep., 2019.
- [2] M. Papapetrou, G. Kosmadakis, A. Cipollina, U. L. Commare, and G. Micale, "Industrial waste heat: Estimation of the technically available resource in the EU per industrial sector, temperature level and country," *Applied Thermal Engineering*, vol. 138, pp. 207–216, 2018.
- [3] L. Lundström and F. Wallin, "Heat demand profiles of energy conservation measures in buildings and their impact on a district heating system," *Applied Energy*, vol. 161, pp. 290–299, 2016.
- [4] P. Y. Liew, W. L. Theo, S. R. W. Alwi, J. S. Lim, Z. A. Manan, J. J. Klemeš, and P. S. Varbanov, "Total site heat integration planning and design for industrial, urban and renewable systems," *Renewable and Sustainable Energy Reviews*, vol. 68, pp. 964–985, 2017.
- [5] X. Liu, J. Wu, N. Jenkins, and A. Bagdanavicius, "Combined analysis of electricity and heat networks," *Applied Energy*, vol. 162, pp. 1238–1250, 2016.
- [6] S. J. Self, B. V. Reddy, and M. A. Rosen, "Geothermal heat pump systems: Status review and comparison with other heating options," *Applied Energy*, vol. 101, pp. 341–348, 2013, Sustainable Development of Energy, Water and Environment Systems.
- [7] L. Brand, A. Calvén, J. Englund, H. Landersjö, and P. Lauenburg, "Smart district heating networks - a simulation study of prosumers' impact on technical parameters in distribution networks," *Applied Energy*, vol. 129, pp. 39–48, 2014.
- [8] H. Lund, S. Werner, R. Wiltshire, S. Svendsen, J. E. Thorsen, F. Hvelplund, and B. V. Mathiesen, "4th Generation District Heating (4GDH): Integrating smart thermal grids into future sustainable energy systems," *Energy*, vol. 68, pp. 1–11, 2014.
- [9] R. Abdurafikov, E. Grahn, L. Kannari, J. Ypyä, S. Kaukonen, I. Heimonen, and S. Paiho, "An analysis of heating energy scenarios of a Finnish case district," *Sustainable Cities and Society*, vol. 32, pp. 56–66, 2017.
- [10] U. Persson, B. Möller, and S. Werner, "Heat Roadmap Europe: Identifying strategic heat synergy regions," *Energy Policy*, vol. 74, pp. 663–681, 2014.
- [11] Helge Averbalk and Sven Werner, "Essential improvements in future district heating systems," *Energy Procedia*, vol. 116, pp. 217–225, 2017.
- [12] T. Naegler, S. Simon, M. Klein, and H. C. Gils, "Quantification of the European industrial heat demand by branch and temperature level," *International Journal of Energy Research*, vol. 39, pp. 2019–2030, 2015.
- [13] W. van Heddeghem, S. Lambert, B. Lannoo, D. Colle, M. Pickavet, and P. Demeester, "Trends in worldwide ICT electricity consumption from 2007 to 2012," *Computer Communications*, vol. 50, pp. 64–76, Sep. 2014.
- [14] M. Dayarathna, Y. Wen and R. Fan, "Data center energy consumption modeling: a survey," *IEEE Communications Surveys & Tutorials*, vol. 18, no. 1, pp. 732–794, 2016.
- [15] M. G. Nielsen, J. M. Morales, M. Zugno, T. E. Pederson, and H. Madsen, "Utilizing data center waste heat in district heating - impacts on energy efficiency and prospects for low-temperature district heating," *Energy*, vol. 140, pp. 1228–1238, 2017.
- [16] S. Werner, "International review of district heating and cooling," *Energy*, vol. 137, pp. 617–631, 2017.
- [17] W. Su, J. Wang, and J. Roh, "Stochastic energy scheduling in microgrids with intermittent renewable energy sources," *IEEE Transactions on Smart Grid*, vol. 5, pp. 1876–1883, 4 2014.
- [18] H. Ü. Yilmaz, D. Keles, A. Chiodi, R. Hartel, and M. Mikulić, "Analysis of the power-to-heat potential in the European energy system," *Energy Strategy Reviews*, vol. 20, pp. 6–19, 2018.
- [19] B. Baeten, F. Rogiers, and L. Helsen, "Reduction of heat pump induced peak electricity use and required generation capacity through thermal energy storage and demand response," *Applied Energy*, vol. 195, pp. 184–195, 2017.

- [20] D. Böttger, M. Götz, M. Theofilidi, and T. Bruckner, “Control power provision with power-to-heat plants in systems with high shares of renewable energy sources - an illustrative analysis for Germany based on the use of electric boilers in district heating grids,” *Energy*, vol. 82, pp. 157–167, 2015.
- [21] D. Böttger, M. Götz, N. Lehr, H. Kondziella, and T. Bruckner, “Potential of the power-to-heat technology in district heating grids in Germany,” *Energy Procedia*, vol. 46, pp. 246–253, 2014.
- [22] Mohammed Yekini Suberu, Mohd Wazir Mustafa and Nouruddeen Bashir, “Energy storage systems for renewable energy power sector integration and mitigation of intermittency,” *Renewable and Sustainable Energy Reviews*, vol. 35, pp. 499–514, 2014.
- [23] O. Schmidt, A. Gambhir, I. Staffell, A. Hawkes, J. Nelson, and S. Few, “Future cost and performance of water electrolysis: An expert elicitation study,” *International Journal of Hydrogen Energy*, vol. 42, no. 52, pp. 30 470–30 492, 2017.
- [24] O. Schmidt, A. Hawkes, A. Gambhir, and I. Staffell, “The future cost of electrical energy storage based on experience rates,” *Nature Energy*, vol. 2, no. 17110, pp. 1–8, 2017.
- [25] N. O’Connell, P. Pinson, H. Madsen, and M. O’Malley, “Benefits and challenges of electrical demand response: A critical review,” *Renewable and Sustainable Energy Reviews*, vol. 39, pp. 686–699, 2014.
- [26] W. Gu, J. Wang, S. Lu, Z. Luo, and C. Wu, “Optimal operation for integrated energy system considering thermal inertia of district heating network and buildings,” *Applied Energy*, vol. 199, pp. 234–246, 2017.
- [27] J. Zheng, Z. Zhou, J. Zhao, and J. Wang, “Integrated heat and power dispatch truly utilizing thermal inertia of district heating network for wind power integration,” *Applied Energy*, vol. 211, pp. 865–874, 2018.
- [28] —, “Effects of the operation regulation modes of district heating system on an integrated heat and power dispatch system for wind power integration,” *Applied Energy*, vol. 230, pp. 1126–1139, 2018.
- [29] Peter D. Lund, “Clean energy systems as mainstream energy options,” *International Journal of Energy Research*, vol. 40, pp. 4–12, 2016.
- [30] S. Hers, M. Afman, S. Cherif, and F. Rooijers, “Potential for power-to-heat in the Netherlands,” CE Delft, Tech. Rep., Aug. 2015.
- [31] G. Schweiger, J. Rantzer, K. Ericsson, and P. Lauenburg, “The potential of power-to-heat in Swedish district heating systems,” *Energy*, vol. 137, pp. 661–669, 2017.
- [32] A. Bloess, W.-P. Schill, and A. Zerrahn, “Power-to-heat for renewable energy integration: A review of technologies, modeling approaches, and flexibility potentials,” *Applied Energy*, vol. 212, pp. 1611–1626, 2018.
- [33] J. Huang, Z. Li, and Q. Wu, “Coordinated dispatch of electric power and district heating networks: A decentralized solution using optimality condition decomposition,” *Applied Energy*, vol. 206, pp. 1508–1522, 2017.
- [34] P. A. Østergaard and A. N. Andersen, “Booster heat pumps and central heat pumps in district heating,” *Applied Energy*, vol. 184, pp. 1374–1388, 2016.
- [35] L. Schibuola, M. Scarpa, and C. Tambani, “Demand response management by means of heat pumps controlled via real time pricing,” *Energy and Buildings*, vol. 90, pp. 15–28, 2015.
- [36] X. Chen, C. Kang, M. O’Malley, Q. Xia, J. Bai, C. Liu, R. Sun, W. Wang, and H. Li, “Increasing the flexibility of combined heat and power for wind power integration in China: Modeling and implications,” *IEEE Transactions on Power Systems*, vol. 30, no. 4, pp. 1848–1857, 2015.
- [37] D. Fehrenbach, E. Merkel, R. McKenna, U. Karl, and W. Fichtner, “On the economic potential for electric load management in the German residential heating sector - an optimising energy system model approach,” *Energy*, vol. 71, pp. 263–276, 2014.
- [38] Jyri Salpakari, Jani Mikkola and Peter D. Lund, “Improved flexibility with large-scale variable renewable power in cities through optimal demand side management and power-to-heat conversion,” *Energy Conversion and Management*, vol. 126, pp. 649–661, 2016.

- [39] A. Nilsson and N. Brandt, "Proposing an hourly dynamic wind signal as an environmental incentive for demand response," in *Advances and New Trends in Environmental Informatics*, V. Wohlgemuth, F. Fuchs-Kittowski, and J. Wittmann, Eds., Springer International Publishing, 2017, pp. 153–164.
- [40] G. Bianchini, M. Casini, A. Vicino, and D. Zarrilli, "Demand-response in building heating systems: A model predictive control approach," *Applied Energy*, vol. 168, pp. 159–170, 2016.
- [41] S. Chan, K. Tsui, H. Wu, Y. Hou, Y. Wu, and F. Wu, "Load/price forecasting and managing demand response for smart grids: Methodologies and challenges," *IEEE Signal Processing Magazine*, vol. Special Issue on Signal Processing Techniques for Smart Grid, pp. 68–85, Sep. 2012.
- [42] M. Sameti and F. Haghghat, "Optimization approaches in district heating and cooling thermal network," *Energy and Buildings*, vol. 140, pp. 121–130, 2017.
- [43] E. Heydarian-Forushani, M. P. Moghaddam, M. K. Sheikh-El-Eslami, M. Shafie-khah, and J. P. Catalão, "Risk-constrained offering strategy of wind power producers considering intraday demand response exchange," *IEEE Transactions on Sustainable Energy*, vol. 5, pp. 1036–1047, 4 2014.
- [44] Richard Scharff and Mikael Amelin, "Trading behaviour on the continuous intraday market Elbas," *Energy Policy*, vol. 88, pp. 544–557, 2016.
- [45] C. Weber, "Adequate intraday market design to enable the integration of wind energy into the European power systems," *Energy Policy*, vol. 38, pp. 3155–3163, 2010.
- [46] Anders Skajaa, Kristian Edlund and Juan M. Morales, "Intraday trading of wind energy," *IEEE Transactions on Power Systems*, vol. 30, pp. 3181–3189, 6 2015.
- [47] Z. Li, W. Wu, J. Wang, B. Zhang, and T. Zheng, "Transmission-constrained unit commitment considering combined electricity and district heating networks," *IEEE Transactions on Sustainable Energy*, vol. 7, no. 2, pp. 480–492, 2015.
- [48] A. Benonysson, "Dynamic modelling and operational optimization of district heating systems," PhD thesis, Technical University of Denmark, DK 2800 Lyngby, Sep. 1991.
- [49] Theodore L. Bergman, Adrienne S. Lavine, Frank P. Incropera and David P. Dewitt, *Fundamentals of Heat and Mass Transfer*, 7th ed. John Wiley & Sons, Apr. 2011.
- [50] Chenyu Wu, Wei Gu, Ping Jiang, Zhenyuan Li, Hongyi Cai and Baoju Li, "Combined economic dispatch considering the time-delay of district heating network and multi-regional indoor temperature control," *IEEE Transactions on Sustainable Energy*, vol. 9, no. 1, pp. 118–127, Jan. 2018.
- [51] H. Fang, J. Xia, K. Zhu, Y. Su, and Y. Jiang, "Industrial waste heat utilization for low temperature district heating," *Energy Policy*, vol. 62, pp. 236–246, 2013.
- [52] A. Capozzoli and G. Primiceri, "Cooling systems in data centers: State of art and emerging technologies," *Energy Procedia*, vol. 83, pp. 484–493, 2015, 7th International Conference on Sustainability in Energy Buildings.
- [53] M. Wahlroos, M. Pärssinen, S. Rinne, S. Syri, and J. Manner, "Future views on waste heat utilization - case of data centers in Northern Europe," *Renewable and Sustainable Energy Reviews*, vol. 82, pp. 1749–1764, 2018.
- [54] R. Weron, "Electricity price forecasting: A review of the state-of-the-art with a look into the future," *International Journal of Forecasting*, vol. 30, pp. 1030–1081, 2014.
- [55] Nord Pool Group, *Day-ahead market*, 2019. [Online]. Available: www.nordpoolgroup.com/the-power-market/Day-ahead-market/.
- [56] T. Hong, "Probabilistic energy forecasting: Global energy forecasting competition 2014 and beyond," *International Journal of Forecasting*, vol. 32, no. 3, pp. 896–913, 2016.
- [57] —, "Probabilistic electric load forecasting: A tutorial review," *International Journal of Forecasting*, vol. 32, no. 3, pp. 914–938, 2016.
- [58] A. M. Foley, P. G. Leahy, A. Marvuglia, and E. J. McKeogh, "Current methods and advances in forecasting of wind power generation," *Renewable Energy*, vol. 37, pp. 1–8, 2012.
- [59] S. Soman, H. Zareipour, O. Malik, and P. Mandal, "A review of wind power and wind speed forecasting methods with different time horizons," in *North American Power Symposium 2010*, Oct. 2010, pp. 1–8.

- [60] M. Lei, L. Shiyan, J. Chuanwen, L. Hongling, and Z. Yan, "A review on the forecasting of wind speed and generated power," *Renewable and Sustainable Energy Reviews*, vol. 13, pp. 915–920, 2009.
- [61] Z. Yu, G. Huang, F. Haghghat, H. Li, and G. Zhang, "Control strategies for integration of thermal energy storage into buildings: State-of-the-art review," *Energy and Buildings*, vol. 106, pp. 203–215, 2015.
- [62] W. Kwon and S. Han, *Receding Horizon Control: Model Predictive Control for State Models*, ser. Advanced Textbooks in Control and Signal Processing. Springer, 2005.
- [63] Nord Pool Group, *Historical market data*, 2018. [Online]. Available: www.nordpoolgroup.com/historical-market-data/.
- [64] L. Herre, T. Matusевичius, J. Olauson, and L. Soder, "Exploring wind power prognosis data on Nord Pool: The case of Sweden and Denmark," *IET Renewable Power Generation*, pp. 690–702, Jan. 2019. DOI: 10.1049/iet-rpg.2018.5086.
- [65] B.-M. Hodge and M. Milligan, "Wind power forecasting error distributions over multiple timescales," in *Proceedings of the IEEE Power and Energy Society General Meeting*, Jul. 2011, pp. 1–8.
- [66] B. C. Ummels, M. Gibescu, E. Pelgrum, W. L. Kling, and A. J. Brand, "Impacts of wind power on thermal generation unit commitment and dispatch," *IEEE Transactions On Energy Conversion*, vol. 22, no. 1, pp. 44–51, 2007.
- [67] Energinet, *Forecasts 5min*, 2019. [Online]. Available: https://www.energidataservice.dk/dataset/forecasts%5C_5min.
- [68] B. A. McCarl and T. H. Spreen, *Applied Mathematical Programming Using Algebraic Systems*. Texas A&M, 1997. [Online]. Available: <http://agecon2.tamu.edu/people/faculty/mccarl-bruce/mccspr/new09.pdf>.
- [69] D. M. Institut, *Vejrarkiv*, 2018. [Online]. Available: <https://www.dmi.dk/da/vejrarkiv/>.
- [70] G. P. Panayiotou, G. Bianchi, G. Georgiou, L. Aresti, M. Argyrou, R. Agathokleous, K. M. Tsamos, S. A. Tassou, G. Florides, S. Kalogirou, and P. Christodoulides, "Preliminary assessment of waste heat potential in major European industries," *Energy Procedia*, vol. 123, pp. 335–345, 2017.
- [71] N. Maruoka, T. Mizuochi, H. Purwanto, and T. Akiyama, "Feasibility study for recovering waste heat in the steelmaking industry using a chemical recuperator," *ISIJ International*, vol. 44, no. 2, pp. 257–262, 2004.

A Background on waste heat

Industrial waste heat is the excess heat produced during industrial processes. The calculated potential of waste heat recovery depends not only on the temperature at which it is recovered, but also on the temperature at which the waste heat can be used after recovery, naturally. Various industries and processes produce heat at a wide range of temperatures. For example, the steel industry operates at high temperatures, leading to high temperature waste heat. Waste heat temperatures are categorised into low (< 373 K), medium, and high (> 573 K) [12], [51], [70]. Others categorise low (< 473 K), medium, and high (> 773 K) temperature levels [2].

Total industrial heat consumption amounts to over 1800 TWh per year in the EU, of which over 500 TWh is attributed to the iron and steel industry [2]. Energy losses take the form of waste heat, called the theoretical or physical potential of waste heat [70]. As expected, not all waste heat can be recovered for use. Most research in waste heat recovery has therefore focused on the technical potential of waste heat. The theoretical technical potential is determined using a process-related analysis and Carnot's potential, while the applicable technical potential depends on plant specific parameters. Lastly, the feasible waste heat potential considers economic criteria on top of the determinative technical parameters. The total technical waste heat potential in the EU is estimated to be 300 TWh yearly, which equals 9.5 percent of total industrial energy demand [2].

The activity of recovering waste heat faces numerous challenging aspects such as the waste heat temperatures [10]. At higher temperatures, thermal energy can be used more effectively in terms of exergy. High temperatures of waste gases from e.g. the steel industry should therefore be maintained and directly supplied to processes that demand those temperatures and should never be used to heat water [71]. Similarly, the exergetic potential of natural gas is largely wasted when used to heat water. Although a lot of heat is involved in high temperature industrial processes, the waste heat potential after recovery is relatively low because consequent energy losses are much higher. Additionally, the potential for low temperature waste heat is considered to be relatively low, due to a small temperature difference compared to ambient temperatures and a relatively small thermal energy consumption at low temperatures, while heat losses are consequently smaller at lower temperatures. Nonetheless, the amount of heat at low temperature levels in absolute terms of energy demand is still considerable, making its neglect unwise.

A measure of heat recovery efficiency is the outlet water temperature after the process of collecting industrial waste heat. In a cascade waste heat collection process, water cools lower temperature industrial processes first, after which it cools gradually higher temperature processes until the outlet water temperature is the highest at the highest temperature process that needs cooling [51]. Similarly, in a heat network of consumers connected to a pipeline, high temperature consumers are connected upstream, i.e. closer to the heat source.

B Derivations

B.1 Consumer temperature

A discretisation of Equation (18) leads to Equation (62), through Equations (59) to (61).

$$\frac{T_i(t) - T_i(t-1)}{\Delta t} = \dot{Q}_i(t) \frac{1}{c_w m_i} - [T_i(t-1) - T_a(t-1)] \frac{c_{q,i}}{c_w m_i} \quad (59)$$

$$T_i(t) = T_i(t-1) + \dot{Q}_i(t) \frac{\Delta t}{c_w m_i} - [T_i(t-1) - T_a(t-1)] \frac{c_{q,i} \Delta t}{c_w m_i} \quad (60)$$

$$\eta_i = \frac{c_{q,i} \Delta t}{c_w m_i} \quad (61)$$

$$T_i(t) = \underbrace{(1 - \eta_i) \cdot T_i(t-1) + \eta_i \cdot T_a(t-1)}_{\text{temperature after heat loss}} + \underbrace{\dot{Q}_i(t)}_{\text{heat gain}} \frac{\Delta t}{c_w m_i} \quad (62)$$

The recursive dependence of the temperature can be seen in Equation (62). As $T_i(t-1)$ is known beforehand solely for the calculation of $T_i(1)$ — namely, $T_i(t-1)$ then equals the assumedly predetermined initial temperature $T_{0,i}$ — a derivation for $T_i(t)$ is necessary for every case in which $t \neq 1$. To further clarify, merely the formulation of the constraints, that are essential for solving the optimisation problem, has been derived at this point. That means that the control variables, e.g. whether or not to extract heat from the network through the heat exchanger, that affect the constraints, e.g. the temperature $T_i(t-1)$ for $t > 1$, are yet to be determined by solving the optimisation problem. To resolve this chicken and egg situation, the constraint formulated in Equations (25) and (62) needs further derivation, through the process of substitution.

The derivation starts with the aforementioned calculable case of $T_i(1)$ in Equation (63), of which the result is substituted into the case of $t = 2$, in which $T_i(t-1)$ is now known (64).

$$T_i(1) = (1 - \eta_i) \cdot T_{0,i} + \eta_i \cdot T_a(0) \quad (63)$$

$$\begin{aligned} T_i(2) &= (1 - \eta_i) \cdot \underbrace{[(1 - \eta_i) \cdot T_{0,i} + \eta_i \cdot T_a(0)]}_{T_i(1)} + \eta_i \cdot T_a(1) + \dot{Q}_i(1) \frac{\Delta t}{c_w m_i} \\ &= (1 - \eta_i)^2 \cdot T_{0,i} + (1 - \eta_i) \cdot \eta_i \cdot T_a(0) + \eta_i \cdot T_a(1) + \dot{Q}_i(1) \frac{\Delta t}{c_w m_i} \end{aligned} \quad (64)$$

As can be seen in Equation (65), as t grows, so too does the expression for $T_i(t)$. Therefore, a more concise general expression for $T_i(t) \forall t \in \mathcal{T}$ is found, presented in Equation (66), based on the derivation in Equations (63) to (65). In short, no matter what decisions for the control variables are made, the constraints remain valid, as they are intrinsically dependent on the control variables. This means that, even though the temperature $T_i(t)$ depends on $\{T_i(t-1), T_i(t-2), \dots\}$ and $\{\dot{Q}_i(t-1), \dot{Q}_i(t-2), \dots\}$, the only parameters in the constraints that are yet to be determined are the control variables.

$$\begin{aligned}
T_i(3) &= (1 - \eta_i) \left[\underbrace{(1 - \eta_i)^2 \cdot T_{0,i} + (1 - \eta_i) \cdot \eta_i \cdot T_a(0) + \eta_i \cdot T_a(1) + \dot{Q}_i(1) \frac{\Delta t}{c_w m_i}}_{T_i(2)} \right] \dots \\
&\quad + \eta_i \cdot T_a(2) + \dot{Q}_i(2) \frac{\Delta t}{c_w m_i} \\
&= (1 - \eta_i)^3 \cdot T_{0,i} + (1 - \eta_i)^2 \cdot \eta_i \cdot T_a(0) + (1 - \eta_i) \cdot \eta_i \cdot T_a(1) + \eta_i \cdot T_a(2) \dots \\
&\quad + (1 - \eta_i) \cdot \dot{Q}_i(1) \frac{\Delta t}{c_w m_i} + \dot{Q}_i(2) \frac{\Delta t}{c_w m_i} \tag{65}
\end{aligned}$$

$$\begin{aligned}
&= \sum_{\vartheta=1}^3 \left[\left(\dot{Q}_i(\vartheta - 1) \cdot \frac{\Delta t}{c_w m_i} + \eta_i T_a(\vartheta - 1) \right) (1 - \eta_i)^{3-\vartheta} \right] + (1 - \eta_i)^3 \cdot T_{0,i} \\
T_i(t) &= \sum_{\vartheta=1}^t \left[\left(\dot{Q}_i(\vartheta - 1) \cdot \frac{\Delta t}{c_w m_i} + \eta_i T_a(\vartheta - 1) \right) (1 - \eta_i)^{t-\vartheta} \right] + (1 - \eta_i)^t \cdot T_{0,i} \tag{66}
\end{aligned}$$

B.2 Supply temperature

Essential for deriving the model of a heat network is to consider the heat losses in the pipeline. To illustrate, Figure 45 shows a single consumer heat network, with the sole purpose of exploring a method to formulate the pipeline heat losses. It is assumed the temperature leaving the boiler T_s is measurable and known. To calculate the supply temperature at the substation $T_{s,l}$, the temperature drop due to heat losses to the ambient is to be deducted from the temperature leaving the boiler.

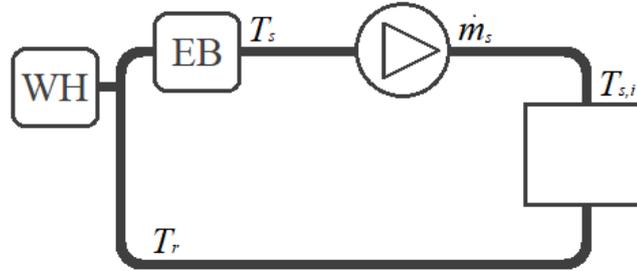


Figure 45: Heat sources with a consumer in a heat network

Equation (67) shows the temperature $T_{s,l}$ if substation 1 was exactly one time step away from the electric boiler ($\tau_{s,l}=1$). Notice that the same temperature dynamics are used as in the energy balance in Equation (18), leaving out any heat gain.

$$\begin{aligned}
T_{s,l}(t) &= T_s(t - 1) - \eta_s \cdot [T_s(t - 1) - T_a(t - 1)] \\
&= (1 - \eta_s) \cdot T_s(t - 1) + \eta_s \cdot T_a(t - 1) \tag{67}
\end{aligned}$$

If the substation would be moved an extra time step away from the electric boiler ($\tau_{s,l} = 2$), the temperature drop of the water over the first time step would be the same as in Equation (67). The temperature after the extra time step is then calculated in the same way (68).

$$\begin{aligned}
T_{s,l}(t) &= (1 - \eta_s) \cdot T_{s \rightarrow 1}(t - 1) + \eta_s \cdot T_a(t - 1) \\
&= (1 - \eta_s) \cdot [(1 - \eta_s) \cdot T_s(t - 2) + \eta_s \cdot T_a(t - 2)] + \eta_s \cdot T_a(t - 1) \\
&= (1 - \eta_s)^2 \cdot T_s(t - 2) + \eta_s \sum_{\vartheta=t-1}^t (1 - \eta_s)^{t-\vartheta} \cdot T_a(\vartheta - 1) \tag{68}
\end{aligned}$$

A formulation of calculating the water temperature after $\tau_{s,l}$ time steps, corresponding with a certain pipeline length, is shown in Equation (69).

$$T_{s,l}(t) = (1 - \eta_s)^{\tau_{s,l}} \cdot \underbrace{T_s(t - \tau_{s,l})}_{\text{temperature leaving boiler}} + \eta_s \int_{t-\tau_{s,l}}^t (1 - \eta_s)^{t-\vartheta} \cdot T_a(\vartheta) d\vartheta \tag{69}$$

$$\eta_s = \frac{c_{q,s} \Delta t}{c_w \dot{m}} \tag{70}$$

The division of the mass flow over a multitude of consumers has an implication for the calculation of the supply temperatures $T_{s,i} \forall i \in N$ in Equation (69). After all, Equation (69) takes a single value for the temperature leaving the electric boiler, whereas in a parallel network configuration with separating mass flow rates, this upstream temperature changes with every consumer the supply pipe passes. Therefore, a general expression for $T_{s,i}$, valid for every consumer $i \in \mathcal{S}$, is formulated in Equation (71). Equation (72) shows the heat loss factor for the different parts of the supply pipeline.

$$T_{s,i}(t) = (1 - h_{s,i})^{\tau_{s,i}} \cdot T_{s,i-1}(t - \tau_{s,i}) + h_{s,i} \int_{t-\tau_{s,i}}^t (1 - h_{s,i})^{t-\vartheta} \cdot T_a(\vartheta) d\vartheta \tag{71}$$

$$h_{s,i} = \frac{c_{q,s} \Delta t}{c_w \dot{m}_{s,i}} \tag{72}$$

The implication of the mass flow rate distribution among consumers on the supply water temperature at consumer i can be seen in Equations (71) and (72). After the supply pipe passes consumer 1, the temperature used to calculate $T_{s,i}$ is not any longer the temperature leaving the electric boiler T_s , but the temperature at the previous network node from $\tau_{s,i}$ time steps ago, $T_{s,i-1}(t - \tau_{s,i})$. Furthermore, the heat loss factor $\eta_{s,i}$ (72) now depends on the pipeline section that corresponds with the time delay $\tau_{s,i}$ or, to be more precise, the mass flow rate. The mass flow rates at the various supply pipe sections are calculated by taking the main network mass flow rate \dot{m} and subtracting all mass flows \dot{m}_i that have already found their way to a consumer (73).

$$\dot{m}_{s,i} = \sum_{j=i}^N \dot{m}_j = \dot{m} - \sum_{j=1}^{i-1} \dot{m}_j \tag{73}$$

From Equation (71) it follows that $T_{s,l}$ depends on $T_{s,0}$ or, in words, the temperature leaving the boiler. Seen from the electric boiler, there is no upstream heat supply water temperature $T_{s,i-1}$. Instead, the known upstream temperature — measured or calculated — useful for finding an expression for $T_{s,0}$, is the temperature in the return pipe at consumer 1, $T_{r,1}$. Furthermore, the temperature increase, due to the heat gain from the waste heat source \dot{Q}_{wh} and the electric boiler P , is included. Equation (75) shows the heat loss factor for the main part of the return pipeline, i.e. from consumer 1 to the waste heat source (see Figure 5). The difference with Equation (72) is the heat transfer factor c_q which depends on the level of insulation. It is assumed that, since insulation material is of better use in the supply pipeline, $c_{q,s}$ is lower than $c_{q,r}$.

$$T_{s,0}(t) = \frac{P(t) + \dot{Q}_{wh}(t)}{c_w \dot{m}} + (1 - \eta_r)^{\tau_r} \cdot T_{r,1}(t - \tau_r) + \eta_r \int_{t-\tau_r}^t (1 - \eta_r)^{t-\vartheta} \cdot T_a(\vartheta) d\vartheta \tag{74}$$

$$\eta_r = \frac{c_{q,r} \Delta t}{c_w \dot{m}} \tag{75}$$

B.3 Return temperature

Throughout the return pipeline, different temperatures are measured at different return pipe connection nodes. These temperatures are a result of the mixing of two water flows, one from the consumer substation i and the other from the upstream return pipe ($i + 1$), as shown in Equation (76). The water flow from the upstream return pipe suffers a drop in temperature due to heat losses before reaching the point of mixing, while, from the water flow through the consumer substation, heat is removed at the substation, lowering the substation inlet temperature $T_{s,i}$.

$$\begin{aligned}
 T_{r,i}(t) &= \frac{\dot{m}_x T_x + \dot{m}_y T_y}{\dot{m}_x + \dot{m}_y} = \frac{\dot{m}_i \cdot \left(T_{s,i}(t) - \frac{\dot{Q}_i(t)}{c_w \dot{m}_i} \right) + \dot{m}_{r,i} \cdot (T_{r,i+1}(t - \tau_{r,i}) - \Delta T_{\dot{q}\text{-loss}})}{\dot{m}_i + \dot{m}_{r,i}} \\
 &= \frac{\dot{m}_i T_{s,i}(t) - \frac{\dot{Q}_i(t)}{c_w} + \left[(1 - h_{r,i})^{\tau_{r,i}} \cdot T_{r,i+1}(t - \tau_{r,i}) + h_{r,i} \int_{t-\tau_{r,i}}^t (1 - h_{r,i})^{t-\vartheta} \cdot T_a(\vartheta) d\vartheta \right] \dot{m}_{r,i}}{\dot{m}_i + \dot{m}_{r,i}} \quad (76)
 \end{aligned}$$

Equation (77) shows the heat loss factor for the different parts of the return pipeline. The difference with Equation (75) is the mass flow rate $\dot{m}_{r,i}$. Similar to the supply pipeline, different parts of the return pipeline have different mass flow rates, because a water mass flow is added at every return pipe connection node shown in Figure 5. The mass flow rate $\dot{m}_{r,i}$ corresponding to the part of the return pipe associated with a time delay $\tau_{r,i}$ and a heat loss factor $\eta_{r,i}$, is calculated in Equation (16).

$$h_{r,i} = \frac{c_{q,r} \Delta t}{c_w \dot{m}_{r,i}} \quad \forall i \in \{1, \dots, N - 1\} \quad (77)$$

C Optimisation problem implementation

C.1 Temperature control variables

Again, in Equations (71), (74) and (76), the temperature is recursively dependent on a temperature at another time t , while all the different temperatures at various time steps are influenced by control decisions made along the way for the boiler and heat exchanger unit commitment variables \mathbf{u} . This recursion in the formula is an issue because, as stated earlier, for a properly functioning constraint, the only parameters that are yet to be determined must be the control variables. Therefore, the constraint must be formulated such that it consists of control variables and known parameters — in this case, *known* means independent of any control variable. In Equations (66), (71), (74) and (76), the respective temperatures $T_i(t-1)$, $T_{s,i-1}(t-\tau_{s,i})$, $T_{r,l}(t-\tau_r)$, $T_{s,i}(t)$ and $T_{r,i+1}(t-\tau_{r,i})$ are neither known nor control variables. One way to solve it is to rewrite Equations (71), (74) and (76) such that they depend just on their known starting temperatures and the control variables. This approach approximates the derivation for T_i in Equation (66) in Appendix B.1, but would be of a higher degree of complexity for the case of a heat network with parallel mass flow rates that includes different time delays. Nonetheless, another approach can be used, with the aim of ultimately solving the optimisation problem in MATLAB. This approach encompasses the introduction of control variables for the network temperatures (78), which cannot freely become any value, but are bound to the system equations of the linear heat network model. In light of this, an addition to the set of control variables is introduced. From now on, $T_{s,0}(t)$, $T_{s,i}(t)$, $T_{r,i}(t)$ and $T_i(t)$ are called control variables (78), as are u_b and u_i . These network temperature control variables u_T cannot just take any value; they must abide the laws of the physical model. That means Equations (14), (15), (17) and (19) become equality constraints in the optimisation problem, next to the already defined inequality constraints (25) (26) (27). The physical model further favours the water to be in the liquid phase for use as an energy carrier in the heat network. Therefore, the lower and upper bounds for the temperature control variables are set to the freezing and boiling points of water, respectively.

$$\vec{u}_T = [T_{s,0}(\mathcal{T}), T_{s,1}(\mathcal{T}), \dots, T_{s,N}(\mathcal{T}), T_{r,1}(\mathcal{T}), \dots, T_{r,N-1}(\mathcal{T})] \quad (78)$$

The optimisation problem that includes the temperature control variables is formulated in (79). A given $T_{0,s}$ is a substitute for $T_{s,i}(t < 1)$, $T_s(t < 1)$ and $T_r(t < 1)$, as it is assumed that from early on the control decisions for boiler use and heat extraction are of much greater influence on the temperature at any point in the heat network, rather than the temperatures of $t < 1$. Due to this temperature initialisation, past boiler and heat exchanger activity is neglected.

$$\begin{aligned} \min_{\mathbf{u}} \quad & c^T \mathbf{u} \\ \text{s.t.} \quad & \mathbf{A}\mathbf{u} = \mathbf{b} \\ & u_i(t) \cdot T_{min,s,i} \leq T_{s,i}(t) \\ & T_{0,i} \leq T_i(d) \\ & T_{min,i} \leq T_i(t) \leq T_{max,i} \\ & u_b(t), u_i(t) \in \{0, 1\} \quad \forall i \in \mathcal{S}, \quad \forall t \in \mathcal{T} \\ & 0 \leq \vec{u}_T \leq 100 \end{aligned} \quad (79)$$

$$c = [P_b \lambda(\mathcal{T}), 0, \dots, 0]$$

$$\mathbf{u} = [u_b(\mathcal{T}), u_1(\mathcal{T}), \dots, u_N(\mathcal{T}), \vec{u}_T]$$

C.2 MATLAB scripts

MATLAB script 1: parametersettings.m

```
1 %% system parameters
2 c_w = .004181; % specific heat capacity water 'MJ/kgK'
3 dt = 3600; % sample time 's'
4 N_d = 2; % number of days to schedule consecutively
5 n = 8; % prediction horizon
6 season = 'winter'; % determine temperature and price profiles
7
8
9 %% consumer parameters
10 N = 2; % number of consumers
11 m = 2*10^5.*[1 1]; % consumer water mass 'kg'
12 cq = 5*10^-3.*[1 1]; % heat transfer coefficient 'MW/K'
13 Tmin = [22 21]; % minimum temperature consumer water volume
14 Tmax = [24 24]; % maximum temperature consumer water volume
15 T0 = [23 22]; % initial temperature consumer water volume
16 cf = .5.*[1 1]; % consumer heat exchanger capacity factor
17
18
19 %% network parameters
20 cf_b = .5; % electric boiler capacity factor
21 dm = 2.*[1 1]; % flow rate through different substations 'kg/s'
22
23 % supply
24 tau_s = [1 1]; % supply pipe time delays in # time steps
25 cq_s = 10^-8; % heat transfer coefficient of supply pipes
26 Tmin_s = [50 50]; % minimum supply temperature for consumer i
27 T0_s = 53; % initial heat network supply temperature
28
29 % return
30 tau_R = 1; % time delay from consumer 1 to data centre
31 tau_r = [1 1]; % return pipe time delays in # time steps
32 cq_r = 10^-7; % heat transfer coefficient of return pipes
33
34
35 %% data centre parameters
36 cq_dc = 10^-3; % heat transfer coefficient of data centre
37 T_dc = 10; % data centre temperature
38 P_dc = 4*10^-2; % data centre power use (per consumer) 'MW'
```

MATLAB script 2: parametercalculations.m

```

1 %% system calculations
2 dataT = csvread('kopenhagen-temperature-2018.csv');
3 dataEP = xlsread('elspot-prices_2018_hourly_eur');
4 EP = dataEP(isfinite(dataEP(:,8)),8);
5 d = round(24*3600/dt); % number of time steps in a day
6 for t = 1:d % price profiles
7 pWinter(t) = mean(EP(t:24:60*24));
8 pSummer(t) = mean(EP(max(0:24:length(EP)/2)+(t:24:60*24)));
9 end
10 if season == 'winter' % temperature profiles
11 T_a = mean(dataT(1:2))-2*sin((1:d)*2*pi/d+pi/4);
12 lambda = pWinter;
13 else
14 T_a = mean(dataT(7:8))-5*sin((1:d)*2*pi/d+pi/4);
15 lambda = pSummer;
16 end
17 for t = 2:d
18 P_dc(t) = P_dc(t-1); % waste heat production profile
19 end
20 dt = 24*3600/d; % time step size
21
22 %% consumer calculations
23 S = 1:N;
24 e = dt./m(S)/c_w;
25 h = cq(S).*e; % heat loss time constant at consumer i
26 Qhx = cq(S).*(T0(S)-mean(dataT(1:2)))./cf(S);
27 for i = S
28 lbT_k(:,i) = [Tmin(i).*ones(d-1,1); T0(i)];
29 ubT_k(:,i) = Tmax(i).*ones(d,1);
30 end
31 lbT = reshape(lbT_k,1,N*d);
32 ubT = reshape(ubT_k,1,N*d);
33
34 %% data centre calculations
35 Qdc_a = cq_dc*(T_a-T_dc);
36 Qwh = N*P_dc + Qdc_a;
37
38 %% network calculations
39 P_b = (sum(cf(S).*Qhx)-mean(Qwh))/cf_b;
40 C_w = c_w*sum(dm(1:N));
41 dm_s = zeros(1,N);
42 for i = S
43 dm_s(i) = sum(dm(i:N));
44 end
45 C_s = dm_s*c_w;
46 h_s = cq_s*dt./C_s;
47 h_r = cq_r*dt./C_w;
48 dm_r = zeros(1,N-1);
49 for i = 1:N-1
50 dm_r(i) = sum(dm(i+1:N));
51 end
52 C_r = dm_r*c_w;
53 h_r = cq_r*dt./C_r;

```

MATLAB script 3: rulebasedcontrol.m

```

1 Q_rb = zeros(d,N);
2 P_rb = zeros(d,1);
3 T_rb = Q_rb;
4 T_s = Q_rb;
5 T_r = Q_rb;
6 T_s0 = P_rb;
7
8 for t = 1:d*N_d
9
10 for i = S
11
12 if t > 1
13
14 if Tmax(i) < (1-h(i))*T_rb(t-1,i)+e(i)*Qhx(i)+h(i)*T_a(rem(t-1,d)+1)
15 Q_rb(t,i) = 0;
16 elseif Tmin(i) > (1-h(i))*T_rb(t-1,i)+h(i)*T_a(rem(t-1,d)+1)
17 Q_rb(t,i) = Qhx(i);
18 else
19 Q_rb(t,i) = Q_rb(t-1,i);
20 end
21
22 T_rb(t,i) = (1-h(i))*T_rb(t-1,i)+e(i)*Q_rb(t,i)+h(i)*T_a(rem(t-1,d)+1);
23 else
24 T_rb(t,i) = (1-h(i))*T0(i)+h(i)*T_a(rem(t-1,d)+1);
25 end
26
27 B = zeros(1,d*N_d);
28 if t <= tau_s(i)
29 for j = 1:t
30 B(j) = (1-h_s(i))^(t-j)*T_a(rem(j-1,d)+1);
31 end
32
33 T_s(t,i) = (1-h_s(i))^t*T0_s+h_s(i)*sum(B);
34 elseif i > 1
35 for j = t-(1:tau_s(i))
36 B(j) = (1-h_s(i))^(t-j)*T_a(rem(j-1,d)+1);
37 end
38
39 T_s(t,i) = h_s(i)*sum(B)+(1-h_s(i))^tau_s(i)*T_s(t-tau_s(i),i-1);
40 else
41 for j = t-(1:tau_s(i))
42 B(j) = (1-h_s(i))^(t-j)*T_a(rem(j-1,d)+1);
43 end
44
45 T_s(t,i) = h_s(i)*sum(B)+(1-h_s(i))^tau_s(i)*T_s0(t-tau_s(i));
46 end
47
48 B = zeros(1,d*N_d);
49 if i == N
50 T_r(t,i) = T_s(t,i)-Q_rb(t,i)/c_w/dm(i);
51 elseif t > tau_r(i)
52 for j = t-(1:tau_r(i))
53 B(j) = (1-h_s(i))^(t-j)*T_a(rem(j-1,d)+1);
54 end
55

```

```

56 T_r(t,i) = (dm(i)*T_s(t,i)-Q_rb(t,i)/c_w+((1-h_r(i))^tau_r(i)*...
57 T_r(t-tau_r(i),i+1)+h_r(i)*sum(B))*dm_r(i))/(dm(i)+dm_r(i));
58 else
59 for j = 1:t
60 B(j) = (1-h_s(i))^(t-j)*T_a(rem(j-1,d)+1);
61 end
62 T_r(t,i) = (dm(i)*T_s(t,i)-Q_rb(t,i)/c_w+((1-h_r(i))^t*T0_s+...
63 h_r(i)*sum(B))*dm_r(i))/(dm(i)+dm_r(i));
64 end
65 end
66
67 B = zeros(1,d*N_d);
68 if t > tau_R
69 for j = t-(1:tau_R)
70 B(j) = (1-h_R)^(t-j)*T_a(rem(j-1,d)+1);
71 end
72
73 if max(Tmin_s(S)) > (1-h_R)^tau_R*T_r(t-tau_R,1)+...
74 h_R*sum(B)+Qwh(rem(t-1,d)+1)/C_w
75 P_rb(t) = P_b;
76 else
77 P_rb(t) = 0;
78 end
79
80 T_s0(t) = (1-h_R)^tau_R*T_r(t-tau_R,1)+h_R*sum(B)+...
81 (P_rb(t)+Qwh(rem(t-1,d)+1))/C_w;
82 else
83 for j = 1:t
84 B(j) = (1-h_R)^(t-j)*T_a(rem(j-1,d)+1);
85 end
86
87 if max(Tmin_s(S)) > (1-h_R)^t*T0_s+h_R*sum(B)+Qwh(rem(t-1,d)+1)/C_w
88 P_rb(t) = P_b;
89 else
90 P_rb(t) = 0;
91 end
92
93 T_s0(t) = (1-h_R)^t*T0_s+h_R*sum(B)+(P_rb(t)+Qwh(rem(t-1,d)+1))/C_w;
94 end
95 end
96
97 T_rb(t,i) = (1-h(i))*T_rb(t-1,i)+e(i)*Q_rb(t,i)+h(i)*T_a(rem(t-1,d)+1);

```

MATLAB script 4: optimalschedule.m

```

1  %% objective function
2  J = [P_b*lambda ... cost of boiler power
3  zeros(1,N*d) ... "cost" of heat extraction
4  zeros(1,3*N*d+d)];
5
6  %% constraints
7  A = zeros(1,(4*N+2)*d);
8  Aeq = A;
9  b = zeros(N*d,1);
10 beq = zeros(3*N*d+d,1);
11 for t = 1:d
12 for i = S
13 Aeq((2*N+i)*d+t,i*d+t) = -e(i)*Qhx(i);
14 Aeq((2*N+i)*d+t,(3*N+i+1)*d+t) = 1;
15 if t > 1
16 Aeq((2*N+i)*d+t,(3*N+i+1)*d+t-1) = h(i)-1;
17 beq((2*N+i)*d+t) = h(i)*T_a(t);
18 else
19 beq((2*N+i)*d+t) = h(i)*T_a(t)+(1-h(i))*T0(i);
20 end
21 A(i*d-d+t,i*d+t) = Tmin_s(i);
22 A(i*d-d+t,(N+1+i)*d+t) = -1;
23 B = zeros(1,d);
24 Aeq(i*d+t,(N+1+i)*d+t) = 1;
25 if t > tau_s(i)
26 Aeq(i*d+t,(N+i)*d+t-tau_s(i))=-(1-h_s(i))^tau_s(i);
27 for j = t-(1:tau_s(i))
28 B(j) = (1-h_s(i))^(t-j)*T_a(j);
29 end
30 beq(i*d+t) = h_s(i)*sum(B);
31 else
32 for j = 1:t
33 B(j) = (1-h_s(i))^(t-j)*T_a(j);
34 end
35 beq(i*d+t) = h_s(i)*sum(B)+(1-h_s(i))^t*T0_s;
36 end
37 Aeq((N+i)*d+t,(2*N+1+i)*d+t) = 1;
38 if i < N
39 B = zeros(1,d);
40 Aeq((N+i)*d+t,(N+1+i)*d+t)=-dm(i)/(dm(i)+dm_r(i));
41 Aeq((N+i)*d+t,i*d+t) = Qhx(i)/c_w/(dm(i)+dm_r(i));
42 if t > tau_r(i)
43 Aeq((N+i)*d+t,(2*N+2+i)*d+t-tau_r(i))=...
44 -(1-h_r(i))^tau_r(i)*dm_r(i)/(dm(i)+dm_r(i));
45 for j = t-(1:tau_r(i))
46 B(j) = (1-h_r(i))^(t-j)*T_a(j);
47 end
48 beq((N+i)*d+t) = h_r(i)*sum(B)*dm_r(i)/...
49 (dm(i)+dm_r(i));
50 else
51 for j = 1:t
52 B(j) = (1-h_r(i))^(t-j)*T_a(j);
53 end
54 beq((N+i)*d+t) = (h_r(i)*sum(B)+...
55 (1-h_r(i))^t*T0_s)*dm_r(i)/(dm(i)+dm_r(i));

```

```

56 end
57 else
58 Aeq((N+i)*d+t, (N+1+i)*d+t) = -1;
59 Aeq((N+i)*d+t, i*d+t) = Qhx(i)/c_w/dm(i);
60 end
61 end
62 B = zeros(1,d);
63 Aeq(t, (N+1)*d+t) = 1;
64 Aeq(t,t) = -P_b/C_w;
65 if t > tau_R
66 Aeq(t, 2*(N+1)*d+t-tau_R) = -(1-h_R)^tau_R;
67 for j = t-(1:tau_R)
68 B(j) = (1-h_R)^(t-j)*T_a(j);
69 end
70 beq(t) = Qwh(t)/C_w+h_R*sum(B);
71 else
72 for j = 1:t
73 B(j) = (1-h_R)^(t-j)*T_a(j);
74 end
75 beq(t) = mean(Qwh(1:t))/C_w+...
76 h_R*sum(B)+(1-h_R)^t*T0_s;
77 end
78 end
79 for i = 1:N+1
80 A(N*d+i, d*(1+N+i)) = -1;
81 b(N*d+i) = -T0_s;
82 end
83
84 lb = [zeros(1, (3*N+2)*d), lbT];
85 ub = [ones(1, d+N*d), 100*ones(1, 2*N*d+d), ubT];
86
87 %% MILP to solve optimal scheduling problem
88 u = intlinprog(J, 1:d+N*d, A, b, Aeq, beq, lb, ub);
89 P_da = u(1:d)*P_b;
90 Q_da = zeros(d,N);
91 T_s = Q_da;
92 T_r = Q_da;
93 T_s0 = u(N*d+d+(1:d));
94 for i = S
95 Q_da(:,i) = u(d*i+(1:d)).*Qhx(i);
96 T_s(:,i) = u((N+1+i)*d+(1:d));
97 T_r(:,i) = u((2*N+1+i)*d+(1:d));
98 T_da(:,i) = u((3*N+1+i)*d+(1:d));
99 end

```

MATLAB script 5: optimalcontrol.m

```

1  for day = 2:N_d
2
3  %% constraints
4  beq = zeros(3*N*d+d,1);
5  for t = 1:d
6  for i = S
7  if t > 1
8  beq((2*N+i)*d+t) = h(i)*T_a(t);
9  else
10 beq((2*N+i)*d+t) = h(i)*T_a(t)+(1-h(i))*T_da(day*d-d,i);
11 end
12
13 B = zeros(1,d);
14 if t > tau_s(i)
15 for j = t-(1:tau_s(i))
16 B(j) = (1-h_s(i))^(t-j)*T_a(j);
17 end
18 beq(i*d+t) = h_s(i)*sum(B);
19 else
20 for j = 1:t
21 B(j) = (1-h_s(i))^(t-j)*T_a(j);
22 end
23 beq(i*d+t) = h_s(i)*sum(B)+(1-h_s(i))^t*T_s0(day*d-d+t-tau_s(i));
24 end
25 if i < N
26 B = zeros(1,d);
27 if t > tau_r(i)
28 for j = t-(1:tau_r(i))
29 B(j) = (1-h_r(i))^(t-j)*T_a(j);
30 end
31 beq((N+i)*d+t) = h_r(i)*sum(B)*dm_r(i)/(dm(i)+dm_r(i));
32 else
33 for j = 1:t
34 B(j) = (1-h_r(i))^(t-j)*T_a(j);
35 end
36 beq((N+i)*d+t) = (h_r(i)*sum(B)+...
37 (1-h_r(i))^t*T_r(day*d-d+t-tau_r(i),i+1))*dm_r(i)/(dm(i)+dm_r(i));
38 end
39 end
40 end
41
42 B = zeros(1,d);
43 if t > tau_R
44 for j = t-(1:tau_R)
45 B(j) = (1-h_R)^(t-j)*T_a(j);
46 end
47 beq(t) = Qwh(t)/C_w+h_R*sum(B);
48 else
49 for j = 1:t
50 B(j) = (1-h_R)^(t-j)*T_a(j);
51 end
52 beq(t)=mean(Qwh(1:t))/C_w+h_R*sum(B)+(1-h_R)^t*T_r(day*d-d+t-tau_R,1);
53 end
54 end
55

```

```
56 |  
57 | %% MILP to solve optimal scheduling problem  
58 | u = linprog(J,1:d+N*d,A,b,Aeq,beq,lb,ub);  
59 | P_da(day*d-d+(1:d)) = u(1:d)*P_b;  
60 | T_s0(day*d-d+(1:d)) = u(N*d+d+(1:d));  
61 | for i = S  
62 | Q_da(day*d-d+(1:d),i) = u(d*i+(1:d)).*Qhx(i);  
63 | T_s(day*d-d+(1:d),i) = u((N+1+i)*d+(1:d));  
64 | T_r(day*d-d+(1:d),i) = u((2*N+1+i)*d+(1:d));  
65 | T_da(day*d-d+(1:d),i) = u((3*N+1+i)*d+(1:d));  
66 | end  
67 |  
68 | end
```

MATLAB script 6: windpowerforecast.m

```

1 %% Nord Pool hourly wind power data
2 dataFcst = xlsread('wind-power-dk-prognosis_2018_hourly');
3 dataReal = xlsread('wind-power-dk_2018_hourly');
4 dataFcst(isnan(dataFcst))=(dataFcst(find(isnan(dataFcst))+1)+...
5 dataFcst(find(isnan(dataFcst))-1))/2;
6 Pwind = dataReal(:,2)./max(dataFcst(:,2));
7 Pwind(isnan(Pwind))=(Pwind(find(isnan(Pwind))+1)+...
8 Pwind(find(isnan(Pwind))-1))/2;
9 wbl = wblfit(Pwind);
10 fcstErr.da = Pwind-dataFcst(:,2)./max(dataFcst(:,2));
11 %% Energinet 5-minute interval data
12 data=xlsread('forecasts_5min-11b36cc9bad24bf6b5259c8b6b47ecf4');
13 P5min = (data(1:2:length(data),:)+data(2:2:length(data),:));
14 P5min = P5min./max(dataFcst(:,2));
15
16 %% forecast errors
17 P1h = zeros(floor(length(P5min(:,1))/12),4);
18 for t = 1:length(P1h)
19 P1h(t,:) = mean(P5min((t-1)*12+(1:12),:));
20 end
21 fcstErr.da5h = P1h(:,1)-P1h(:,2);
22 fcstErr.h5h1 = P1h(:,2)-P1h(:,3);
23 fcstErr.h1rt = P1h(:,3)-P1h(:,4);
24 P5min = P5min(:,4);
25 %% forecast errors in the persistence model
26 pstErr = P5min(2:length(P5min))-P5min(1:length(P5min)-1);
27 gamm = cauchyfit(pstErr);
28 gamm = gamm(2);
29
30 %% day-ahead wind power forecast
31 Pw_da = min(max(Pwind),max(min(Pwind),wblrnd(wbl(1),wbl(2)))));
32 for t = 2:d*12*N_d
33 dPw_da(t) = gamm*(tan(pi*(rand-.5))+randn)/2;
34 if dPw_da(t) > 0
35 Pw_da(t) = Pw_da(t-1)+ min(max(pstErr),dPw_da(t));
36 else
37 Pw_da(t) = Pw_da(t-1)+ max(min(pstErr),dPw_da(t));
38 end
39 if Pw_da(t) <= min(Pwind)
40 Pw_da(t) = max(pstErr)*min(1,abs(randn)/4);
41 elseif Pw_da(t) > max(Pwind)
42 Pw_da(t) = 1+min(pstErr)*min(1,abs(randn)/4);
43 end
44 end
45 P_h = zeros(floor(length(Pw_da)/12),1);
46 for t = 1:d*N_d
47 P_h(t) = mean(Pw_da((t-1)*12+(1:12)));
48 end
49 Pw_da = P_h; clear P_h

```

MATLAB script 7: recedinghorizoncontrol.m

```

1 p = N_d*d-d-n+1;
2 Pb = P_da';
3 T_rt = T_da;
4
5 %% wind power forecast update
6 armada = estimate(arima(1,0,1),fcstErr.da5h);
7 arma5h = estimate(arima(1,0,1),fcstErr.h5h1);
8 armalh = estimate(arima(1,0,1),fcstErr.h1rt);
9 sigma.h5 = sqrt(arma5h.Variance);
10 sigma.h1 = sqrt(armalh.Variance);
11 gamma5h = [sigma.h5*randn(1,1+5) zeros(1,N_d*d-1-5)];
12 gammalh = [sigma.h1*randn zeros(1,N_d*d-1)];
13 phi.da = simulate(armada,N_d*d)';
14 phi.h5 = gamma5h;
15 for t = 2:5
16 phi.h5(t) = (arma5h.AR{1,1})*phi.h5(t-1)+...
17 (arma5h.MA{1,1})*gamma5h(t-1) + gamma5h(t);
18 end
19 phi.h1 = gammalh;
20 phi.rt = phi.da+phi.h5+phi.h1;
21 Pw = Pw_da'+phi.rt;
22 t=1;
23
24 %% constraints
25 A_k = zeros(1,(4*N+3)*n);
26 Aeq_k = A_k;
27 b_k = zeros(N*n,1);
28 beq_k = zeros(3*N*n+1,1);
29 k = 1:n;
30 for i = 0:size(A,1)/d-1
31 for j = 0:size(A,2)/d-1
32 A_k(i*n+k,j*n+k) = A(i*d+k,j*d+k);
33 end
34 b_k(i*n+k) = b(i*d+k);
35 end
36 for i = 0:size(Aeq,1)/d-1
37 for j = 0:size(Aeq,2)/d-1
38 Aeq_k(i*n+k,j*n+k) = Aeq(i*d+k,j*d+k);
39 end
40 beq_k(i*n+k) = beq(i*d+k);
41 end
42 for i = k
43 Aeq_k(size(Aeq,1)*n/d+i,size(Aeq,2)*n/d+i)=-1;
44 Aeq_k(size(Aeq,1)*n/d+i,i) = 1;
45 end
46 beq_k(size(Aeq,1)*n/d+k) = Pb(k)./P_b;
47 Aeq_k(3*N*n+1,size(Aeq,2)*n/d+k) = 1;
48
49 lb_k = [zeros(1,(3*N+2)*n), reshape(lbT_k(k,:),1,N*n), -(Pb(k)>0)];
50 ub_k = [ones(1,N*n+n), 100*ones(1,2*N*n+n), ...
51 reshape(ubT_k(k,:),1,N*n), (Pb(k)==0)];
52
53 %% objective function
54 Jmpc = [zeros(1,size(A,2)*n/d) (phi.rt(t,k)<0)-(phi.rt(t,k)>0)];
55

```

```

56
57 %% MILP to solve for first prediction horizon
58 u_k = intlinprog(Jmpc, [1:N*n+n (4*N+2)*n+k], ...
59 A_k, b_k, Aeq_k, beq_k, lb_k, ub_k);
60 Pb(k) = u_k(k)*P_b;
61 T_s0(k) = u_k(N*n+n+k);
62 for i = S
63 Q_rt(k,i) = u_k(n*i+k).*Qhx(i);
64 T_s(k,i) = u_k((N+1+i)*n+k);
65 T_r(k,i) = u_k((2*N+1+i)*n+k);
66 T_rt(k,i) = u_k((3*N+1+i)*n+k);
67 end
68
69 %% receding horizon control
70 for t = 2:p
71 k = t:t+n-1;
72
73 % wind power forecast update
74 gamma5h(t+5) = sigma.h5*randn;
75 gamma1h(t) = sigma.h5*randn;
76 phi.h5(t+5) = (arma5h.AR{1,1})*phi.h5(t+4)+...
77 (arma5h.MA{1,1})*gamma5h(t+4)+gamma5h(t+5);
78 for kk = t+5+1:d
79 phi.h5(kk) = (arma5h.AR{1,1})*phi.h5(kk-1)+...
80 (arma5h.MA{1,1})*gamma5h(kk-1)+gamma5h(kk);
81 end
82 for kk = t+(1:5-1)
83 phi.h1(kk) = (arma1h.AR{1,1})*phi.h1(kk-1)+...
84 (arma1h.MA{1,1})*gamma1h(kk-1)+gamma5h(kk);
85 end
86 phi.rt(t,:) = phi.da+phi.h5+phi.h1;
87 Pw(t,:) = min(1,max(0, Pw_da'+phi.rt(t,:)));
88 for kk = 1:t
89 Pw(t,kk) = Pw(kk,kk);
90 end
91
92 % objective function
93 Jmpc = [zeros(1,size(A,2)*n/d)...
94 (phi.rt(t,k)<0)-(phi.rt(t,k)>0)];
95
96 % constraints
97 recedinghorizoncontrolConstraints
98
99 %% MILP to solve for prediction horizon
100 u_k(:,t) = intlinprog(Jmpc, [1:N*n+n (4*N+2)*n+(1:n)], ...
101 A_k, b_k, Aeq_k, beq_k, lb_k, ub_k);
102 Pb(k) = u_k(1:n,t)*P_b;
103 T_s0(k) = u_k(N*n+n+(1:n),t);
104 for i = S
105 Q_rt(k,i) = u_k(n*i+(1:n),t).*Qhx(i);
106 T_s(k,i) = u_k((N+1+i)*n+(1:n),t);
107 T_r(k,i) = u_k((2*N+1+i)*n+(1:n),t);
108 T_rt(k,i) = u_k((3*N+1+i)*n+(1:n),t);
109 end
110 end

```

MATLAB script 8: recedinghorizoncontrolConstraints.m

```

1 beq_k = zeros(3*N*n, 1);
2 for x = 1:n
3
4     for i = S
5         if x > 1
6             beq_k((2*N+i)*n+x) = h(i)*T_a(rem(k(x)-1,d)+1);
7         else
8             beq_k((2*N+i)*n+1) = h(i)*T_a(rem(t-1,d)+1) + (1-h(i))*T_rt(t-1,i);
9         end
10
11        B = zeros(1,n);
12        if x > tau_s(i)
13            for j = k(x)-(1:tau_s(i))
14                B(j) = (1-h_s(i))^(k(x)-j)*T_a(rem(j-1,d)+1);
15            end
16
17            beq_k(i*n+x) = h_s(i)*sum(B);
18        else
19            for j = t:k(x)
20                B(j) = (1-h_s(i))^(k(x)-j)*T_a(rem(j-1,d)+1);
21            end
22
23            if i > 1
24                beq_k(i*n+x) = h_s(i)*sum(B) + (1-h_s(i))^x*T_s(k(x)-tau_s(i),i-1);
25            else
26                beq_k(i*n+x) = h_s(i)*sum(B) + (1-h_s(i))^x*T_s0(k(x)-tau_s(i));
27            end
28        end
29
30        if i < N
31            B = zeros(1,n);
32            if x > tau_r(i)
33                for j = k(x)-(1:tau_r(i))
34                    B(j) = (1-h_r(i))^(k(x)-j)*T_a(rem(j-1,d)+1);
35                end
36
37                beq_k((N+i)*n+x) = h_r(i)*sum(B)*dm_r(i)/(dm(i)+dm_r(i));
38            else
39                for j = t:k(x)
40                    B(j) = (1-h_r(i))^(k(x)-j)*T_a(rem(j-1,d)+1);
41                end
42
43                beq_k((N+i)*n+x) = (h_r(i)*sum(B) + (1-h_r(i))^x*...
44                T_r(k(x)-tau_r(i),i+1))*dm_r(i)/(dm(i)+dm_r(i));
45            end
46        end
47    end
48
49    B = zeros(1,n);
50    if x > tau_R
51        for j = k(x)-(1:tau_R)
52            B(j) = (1-h_R)^(k(x)-j)*T_a(rem(j-1,d)+1);
53        end
54        beq_k(x) = Qwh(rem(k(x)-1,d)+1)/C_w+h_R*sum(B);
55    elseif k(x) > tau_R

```

```

56 for j = t:k(x)
57 B(j) = (1-h_R)^(k(x)-j)*T_a(rem(j-1,d)+1);
58 end
59 beq_k(x)=mean(Qwh(rem((t:k(x))-1,d)+1))/C_w+...
60 h_R*sum(B)+(1-h_R)^x*T_r(k(x)-tau_R,1);
61 else
62 for j = t:k(x)
63 B(j) = (1-h_R)^(k(x)-j)*T_a(j);
64 end
65
66 beq_k(x) = mean(Qwh(t:k(x)))/C_w+h_R*sum(B) + (1-h_R)^x*T0_s;
67 end
68 end
69
70 beq_k(size(Aeq,1)*n/d+(1:n)) = Pb(k)./P_b;
71
72 lb_k=[zeros(1,(3*N+2)*n), ...
73 reshape(lbT_k(rem(k-1,d)+1,:),1,N*n),-(Pb(k)>0)];
74 ub_k = [ones(1,N*n+n), 100*ones(1,2*N*n+n), ...
75 reshape(ubT_k(rem(k-1,d)+1,:),1,N*n), (Pb(k)==0)];

```