NTNU
Norwegian University of
Science and Technology

# Model Predictive Control of District Heating Systems

**Anders Svenheim Rene**

Master of Science in Cybernetics and Robotics
Submission date: June 2016
Supervisor: Bjarne Anton Foss, ITK
Co-supervisor: Brage Rugstad Knudsen, ITK

Norwegian University of Science and Technology
Department of Engineering Cybernetics

# Project description

**NTNU**
**Norwegian University of Science and Technology**

**Faculty of Information Technology,**
**Mathematics and Electrical Engineering**

**Department of Engineering Cybernetics**

## Master assignment

Name of candidate: Anders Svenheim Rene

Subject: Engineering Cybernetics

Title: Model Predictive Control of District Heating Systems
Title (in Norwegian):

1. Describe characteristics and dynamics of typical district heating systems (DHS). Further, include a discussion on operational challenges related to such systems.
2. Discuss relevant control approaches for DHS with focus on model-based techniques.
3. Describe and discuss operations and control of Trondheim DHS, and how control and optimization may improve performance by for instance reducing heat losses.
4. Develop a model of a part of Trondheim DHS and validate this against data. Further, implement a Model Predictive Control (MPC) strategy and evaluate its performance.
5. Make recommendation for future studies of MPC or similar techniques to aid operations of Trondheim DHS.

This master project is done in cooperation with Statkraft.

Starting date: 11.01.2016
End date: 27.06.2016

Co-supervisor: Postdoc Brage Rugstad Knudsen, NTNU

Trondheim, 21.01.2016
Bjarne Foss
Professor/supervisor

# Preface

This thesis is the final work of a five-year study at the Master of Technology program in Engineering Cybernetics at the Norwegian University of Science and Technology (NTNU).

The thesis has been carried out and written over the course of the spring semester 2016, from January 11th to June 27th. The thesis, and the workload behind it has been done in collaboration with Statkraft, and has attempted to lay a foundation for continued exploration of the subject in upcoming work done in a cooperation between Statkraft and NTNU.

I would like to thank my two supervisors, Bjarne Anton Foss, and Brage Rugstad Knudsen, for their contributions. The assistance they have provided over the past year has been detrimental to the work done in this thesis. I would also like to thank the people at Statkraft for helping me with shaping the problem description, providing feedback, and contributing with physical data. In particular, I would like to thank Morten Fossum and Åmund Utne in Statkraft, for exceptionally helpful meetings throughout the semester.

Finally, I would like to thank my colleagues in room G232C for useful discussions and insightful comments.

Trondheim, June 2016

Anders Svenheim Rene

# Abstract

A district heating system is a network of pipelines where heating is delivered to a number of customers from a centralized heating station. For the operators of these systems, the ideal operating scenario is to deliver enough heat to satisfy the heating demand of every customer on the network with the lowest possible costs.
District heating systems are slow processes, and the heated water may travel for several hours from it leaves the heating station until it reaches the most distant customer. Today, many district heating systems in Norway are operated largely based on operational experience, without any available tools for predicting future states or demands.

Model predictive control on district heating systems has been applied in other countries, and has been reported to be very effective. Motivated by its success, this thesis has reviewed the benefits and challenges of applying model predictive control to district heating systems in Norway.

In this thesis, a model predictive controller has been simulated on a small-scale district heating system in Klæbu, a small urban area approximately 19 km south of Trondheim, Norway. This system consists of a single plant which delivers heating on two separated lines. The simulation has been performed on the line that covers the central areas of Klæbu. The line only has five customers, making it a very simple district heating network, well suited for initial testing of predictive control.

The model predictive controller implemented on the district heating system in Klæbu controls the supply temperature of the water delivered from the plant. The controller uses a model that predicts the future heat demand and the transport delay in the network. Based on this model, the simulated controller generates values that meet the predicted heat demand without violating the system constraints. However, the actual performance of the controller has been found challenging to assess. This is due to a low amount of reliable measurements on the network, unclear methods of how to interpet the available measurements, and a lacking definition of economic optimization.

# Sammendrag

Et fjernvarmesystem er et nett med rørledninger som frakter varme til et større antall kunder fra en varmesentral. For fjernvarmesystemets operatører er det ideelle reguleringsforholdet å levere tilstrekkelig varme til samtlige kunder i nettet på en mest mulig økonomisk måte.
Fjernvarmenett er trege prosesser, og vannet kan ha en transporttid på flere timer fra det forlater varmekilden til det ankommer den ytterste kunden. Mange fjernvarmenett i Norge er i stor grad styrt basert på operasjonell erfaring med nettet. Dette innebærer at operatørene ofte ikke har tilgjengelige verktøy for å forutse fremtidige tilstander eller behov.
Modellprediktiv regulering er et reguleringsverktøy som har blitt utnyttet på fjernvarmesystemer i andre land til stor suksess, og dette er bakgrunnen for denne masteroppgaven. I denne oppgaven er det blitt sett på fordeler og utfordringer rundt bruken av modelprediktiv regulering på fjernvarmesystemer.

En modellprediktiv regulator er i denne oppgaven blitt simulert på et lite fjernvarmenett i Klæbu, en kommune som ligger rundt 19 km sør for Trondheim. Dette nettet består av en enkel varmestasjon som leverer varme ut på to separate rørledninger. Simuleringen har vært fokusert på rørnettet som dekker Klæbu sentrum. Det er kun fem kunder påkoblet på denne delen av fjernvarmenettet, noe som gjør at nettet er godt egnet for innledende testing av modellprediktiv regulering.

Den modellprediktive regulatoren er blitt brukt til å bestemme turtemperaturen på vannet som forlater varmekilden. Regulatorens verdier blir generert basert på en modell for fremtidig varmebehov i nettet, samt en evaluering av transporttiden til vannet. Verdiene gitt av den simulerte regulatoren er innenfor systemets begrensninger, og svarer til behovet i nettet. Det har samtidig vært vanskelig å måle ytelsen til den simulerte regulatoren. Dette er grunnet at målinger i nettet gjøres sjeldent, og at det er en viss grad av uklarhet i hvordan de tilgjengelige målingene skal tolkes. Samtidig er det ikke blitt definert hva som konstituerer et ideelt styringsforhold.

# Table of Contents

# List of Tables

# List of Figures

# Abbreviations

| | | |
|---|---|---|
| DH | = | District Heating |
| DHS | = | District Heating System |
| MPC | = | Model Predictive Control |
| RTO | = | Real-Time Optimization |
| MBD | = | Model-Based Design |
| EDP | = | Economic Dispatch Problem |
| UCP | = | Unit Commitment Problem |

Chapter 1

# Introduction

## 1.1 Background and Motivation

District heating is a method of centralizing production of heat, and distributing it for residential and commercial use. In a district heating network, heating is delivered over pipelines, typically on the form of hot water, but also sometimes as steam. The centralized units, or heating stations, control the temperature, pressure, and flow rate of the water, in an attempt to satisfy the heating demands for all consumers connected to the network. Although district heating in Norway is not as widespread as in its neighbor countries, district heating systems are present in the major cities, as well as a number of urban areas.

Compared to individual heating, district heating can offer a greener and more energy efficient way of providing heat. The heat delivered on the network can in some cases be extracted from processes with excess heat that would otherwise be wasted, which poses as a great potential for profit. This is the backbone of a desire of expanding district heating operation in Norway, both in the form of delivering district heating to a larger amount of urban area, as well as improving the current technological level.

District heating systems are slow processes, and changes made to the system often take several hours until they are measured. In spite of this, predictive control is barely present at district heating systems in Norway. When this is the case, the operators will not have access to any model of future behavior, and have to make decisions based on operational experience. Adding a framework for predictive control in district heating systems, that could be used for systems of various size in Norway, would be a large step forward for the Norwegian district heating technological level.

## 1.2 Objective

The aim of this thesis is to explore the benefits and challenges of deploying model predictive control on district heating systems. In order to apply this method of control to district heating systems in Norway, it must be integrated with the current operational regime, and describe properties, challenges, and bottlenecks in a precise manner.

The main objectives of this thesis are as follows:

- Describe the characteristics and dynamics of typical district heating systems (DHS), including a discussion of operational challenges to such systems
- Discuss relevant control approaches for DHS, with focus on model-based techniques
- Describe and discuss operations and control of Trondheim district heating system, and how control and optimization may improve performance by for instance reducing heat losses
- Develop a model of a part of Trondheim DHS and validate this against data. Further, implement a Model Predictive Control (MPC) strategy and evaluate its performance
- Make recommendations for future studies of MPC or similar techniques to aid operations of Trondheim DHS

## 1.3 Previous Work on the Subject

District heating is a well-researched subject, with several topics that are extensively discussed in the literature. Early mentions of model predictive control on district heating systems are found in [1] and [2]. [1] presents an optimization on flow rate and supply temperatures using predictive control and sub-controllers. The sub-controllers optimize parts of the network, and a subsequent full optimization is done at the plant afterwards. [2] presents methods on predicting temperatures on large networks using the so-called node method. In [3] and [4], it is looked into increasing efficiency in district heating systems by reducing return temperatures through identifying and eliminating faults in the system. This is done through evaluation of feedback from consumer measurements. In [5], both linear and nonlinear components for the production, distribution, and consumption parts of a district heating system process are presented. A robust control law is then proposed for the resulting nonlinear model. [6] describes short term minimization of operational costs on a district heating system, including emphasis on the usage of network loading. In [7], the nonlinearities caused by the varying flow rates in the systems are handled by the use of fuzzy regions of local finite impulse response models.

Work regarding modelling of consumer demand is found in [8], where a simple model for prediction of loads in a district heating system, based on the outdoor temperature and social behavior of the consumers is presented. Furthermore [9] describes a dynamic system for load modelling. This is mentioned as especially useful when knowledge of the consumers and their social behaviors is limited.

A report on network aggregation is given in [10]. Here, two aggregation models are presented, that both accomplish to simplify large district heating networks by branch reduction. Reports with mention of specified programs for both modelling and optimization of district heating systems are given in [11] and [12]. [11] contains a case study of modelling temperature dynamics of a district heating system in Naestved, Denmark. The modelling was used by applying a modelling program named TERMIS. [12] presents a program, spHeat designed for performing hydraulic and thermal simulation of meshed pipes in water networks, that optimizes the operating parameters hourly.

Finally, [13] describes an alternative to model-based prediction for predicting temperatures at critical points in a district heating network. This method uses conditional finite impulse response models, for which the coefficients are replaced by coefficient functions of the water flow rate and the time of day.

## 1.4 Thesis Outline

This thesis is outlined as follows: **Chapter 2** presents theory behind district heating systems. Here, different components in a DHS are explained. Furthermore, a description of different aspects of DHS regulation and control is given. In **Chapter 3**, different aspects of modelling and optimization done on large-scale systems are explained. This involves descriptions of model-based techniques, model predictive control, as well as dividing control in additional layers by using real-time optimization. **Chapter 4** is dedicated to a description of the DHS in Trondheim. Here, properties and challenges concerning Trondheim DHS are described on a basic level.
In **Chapter 5**, the DHS in Klæbu is described, and the proposed control structure is explained. **Chapter 6** proposes the results found from simulation of MPC on the Klæbu DHS. These results are further discussed in **Chapter 7**. **Chapter 8** presents a conclusion to the work done in this thesis, and **Chapter 9** describes future work that can be done.

Chapter 2

# District Heating Systems

This chapter will attempt to describe underlying theory surrounding district heating systems. The following sections will deliver an overview over the components found in a district heating system, as well as their purpose.

## 2.1 The Components of a DHS Model

A model of a district heating system with a single heat source can be seen in Figure 2.1. Heat is produced at a centralized heat source, or plant, and delivered on the supply line. A heat pump is used at the plant exit in order to convert the mass flow rate to the level found to be desirable for the network at the given time.

The heating is distributed on the supply line to various heat exchangers that connect the consumers to the network. Hot water from the supply line that passes through a heat exchanger is cooled down, and will need to be reheated at the centralized heat source. The return line transfers all the cooled water back to the plant to be reheated.

The network cycle of the centralized heat source, the supply line, the heat exchangers and the return line is called the primary network. It is important to note that there are many heat exchangers on the primary network, and that these are placed in parallel. This means that for any path the water may take, it will pass through exactly one heat exchanger before it gets returned to the heat source, and that it normally can not go directly from the supply line to the return line without passing through a heat exchanger.

A consumer in a DHS may be a single household, or a large building complex such as a factory or a hospital. Usually smaller buildings are not connected directly on the primary network, but on a network distributing heat extracted from the primary network to several buildings and households. These networks are referred to as secondary networks, and there may be many of them on a single DHS.

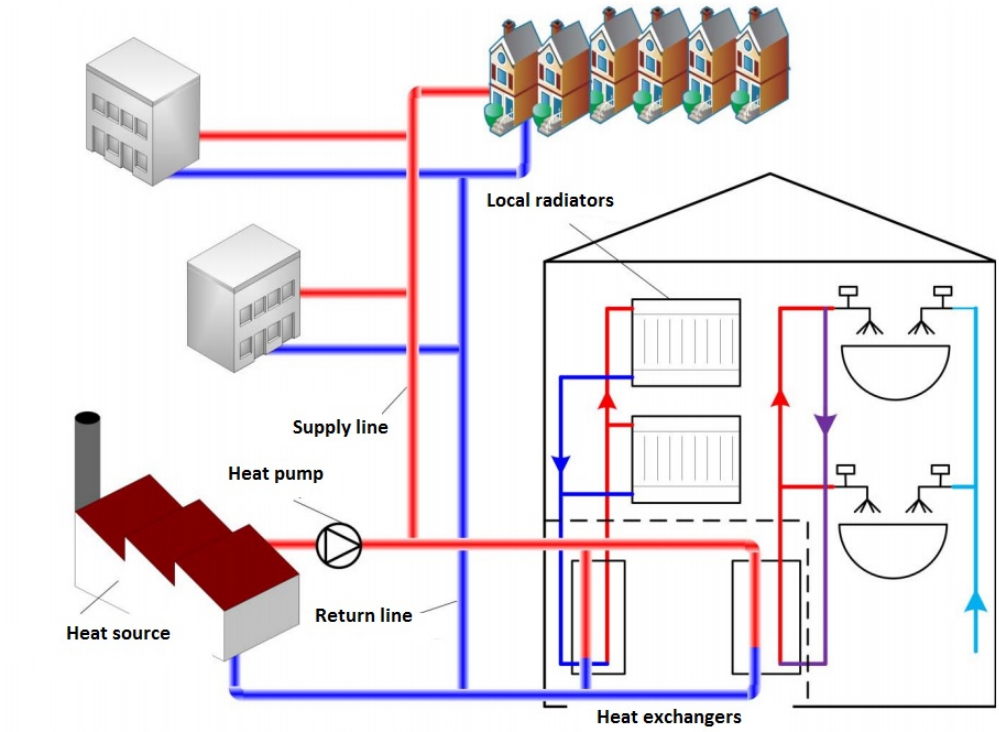


Figure 2.1: Model of a simple DH system with a single heat source

## 2.2 Underlying Physical Equations of a DHS

In this section, the physical aspect and equations of a DHS is assessed. A good model of a DHS will use both the physical equations as well as empirical knowledge in order to model its dynamics. A good understanding of the underlying physics of a DHS model can be very helpful in various ways, such as identifying the main operational challenges, designing the system to best fit its purpose, and identifying errors such as leaks and excessive heat loss.

### 2.2.1 Consumer Heat Requirement Model

As with all consumer driven industry, the challenge for the operators of a DHS is to match production as close as possible to the predicted demand. This means that having a precise prediction of the consumers' heat usage allows for the DHS controllers to operate more economically efficient.

The dynamics of the consumer heat demand is in general a difficult task to precisely model using physical equations. The heating requirements vary largely with the social behaviors of the consumers, creating hourly and seasonal differences, which in turn makes it difficult to fit into a physical model. The modelling of the social component of the heat

requirements is a subject of its own, and it is mentioned later in this chapter, in Section 2.3.

However, having an understanding of the physical aspect may have some usage for evaluating empirical models and observed data. An example of this could be a building that extracts more heat than is realistic based on its physical properties, which could imply there being a leak or a measurement error.

The physical model views the heat requirement as a function of building heat loss, and as a function of hot water required for everyday usage. The physical modelling of heat loss in buildings is an extensively researched topic, with the models having become more and more precise, and also complex. Such models are typically used in projects where the individual properties of a house are of interest, as opposed to placing the requirement of a single house into larger systems as in DHS modelling. For a DHS, a simple model is more practical and realistic. A simple physical model for consumer heat requirements is as follows:

$$Q_h = \dot{q}_{se} A_{tot}(T_i - T_o) \tag{2.1}$$

$$Q_w = \dot{q}_{we} A_{tot} \tag{2.2}$$

with the following parameters:

- $Q_h$ = Total heat required for heating $[W]$
- $Q_w$ = Total heat required for warm water$[W]$
- $\dot{q}_{se}$ = Specific effect $[\frac{W}{Km^2}]$
- $\dot{q}_{we}$ = Specific effect for hot water$[\frac{W}{m^2}]$
- $A_{tot}$= Total residential area $[m^2]$
- $(T_i - T_o)$ = Temperature difference between inside and outside the building $[°C]$

The first equation describes the heating part as a function of the total area and the difference between the inside and the ambient temperatures. The second line describes the total heating required for the warm-water usage of the consumer. This has been modelled as a function of the specific effect, and the total area.

### 2.2.2 Heat Generation

The heat delivered to the DHS from the heat source at a given time, $t$, depends on the mass flow rate, as well as the temperature difference between the supply water temperature going out from the plant and the return temperature entering the plant. The equation is as follows:

$$Q_{sup}(t) = c_p \dot{m}(T_s(t) - T_r(t)) \tag{2.3}$$

where

- $Q_{sup}(t)$ = Supplied power to the DHS from the plant $[W]$

- $c_p$ = Specific heat capacity of water $[\frac{Ws}{kgK}]$
- $\dot{m}$ = Water mass flow rate $[kg/s]$
- $T_s(t)$ = Temperature of the outgoing water on the supply line from the plant $[°C]$
- $T_r(t)$ = Temperature of the incoming water from the return line to the plant$[°C]$

The values of outgoing power from the pump are measured in mass flow rate, supply and return temperature. However, in DHS theory, it is important to note the difference between the heat delivered from the plant and the heat received by the customers at a certain time, which introduces the concept of heat loads.

### 2.2.3 The Specific Heat Capacity of Water

The value of the specific heat capacity of water, $c_p[\frac{Ws}{kgK}]$ has in this thesis been assumed constant. In reality, this is a parameter that to some extent increases with the temperature of the water. However, the value changes are rather small. The value for $c_p$ at $20°C$ is $4.246$, and the value at $120°C$ is $4.248$. The value for $c_p$ at $100°C$ is $4.219$, which is the constant value that has been used in this thesis.

### 2.2.4 Heat Exchangers and Network Heat Extraction

Heat exchangers are a crucial part of a DHS. With the exception of the plant, they appear whenever heat is transferred from one medium to another. In a DHS this means that all secondary networks are connected to the primary network by a heat exchanger. In addition to this, any customer, whether connected to the primary or the secondary network, will extract its heat through a heat exchanger.

There are several different types of heat exchangers, but the most commonly used heat exchanger in DHS is the counter flow heat exchanger pictured in Figure 2.2. The popularity of the counter flow heat exchanger comes from its high heat transfer efficiency. By having flows with opposite directions, the temperature difference over the heat transfer area is evenly distributed, thus leading to a overall high efficiency.

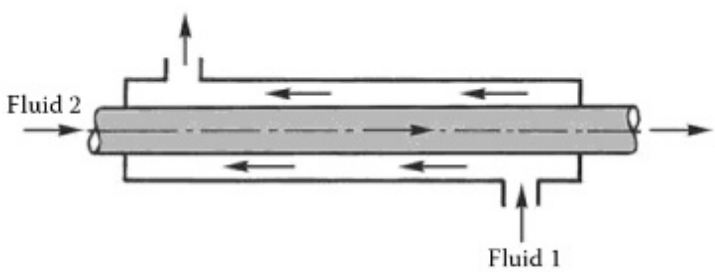


Figure 2.2: Simple schematic of a counter flow heat exchanger

Heat exchangers transfer heat from a hot water flow to a colder water flow, including some minor spill heat to the surroundings. When assuming a zero heat loss, the incoming heat on the cold water flow is the same as the outgoing heat from the hot water flow, and therefore the overall size of the heat transfer can be measured in several ways.

When the assumption of no heat loss during the transfer is made, the thermal power delivered by the hot water can be expressed by:

$$Q = c_p \dot{m}_h (T_{h,in} - T_{h,out}) \tag{2.4}$$

And the power received by the cold water can be expressed by:

$$Q = c_p \dot{m}_c (T_{c,out} - T_{c,in}) \tag{2.5}$$

where

- $Q$ = The exchanged heat $[W]$
- $\dot{m}_h$ = Mass flow rate of the heated water $[\frac{kg}{s}]$
- $\dot{m}_c$ = Mass flow rate of the cold water $[\frac{kg}{s}]$
- $c_p$ = Specific heat capacity of water $[\frac{Ws}{kgK}]$
- $T_{h,in}$ = Temperature of the hot water at the entrance of the heat exchanger $[^\circ C]$
- $T_{h,out}$ =Temperature of the hot water at the exit of the heat exchanger $[^\circ C]$
- $T_{c,in}$ = Temperature of the cold water at the entrance of the heat exchanger $[^\circ C]$
- $T_{c,out}$ =Temperature of the cold water at the exit of the heat exchanger $[^\circ C]$

### 2.2.5 Network Heat Propagation

The dynamics of the power on the network are, as mentioned in previous sections, dependent on the mass flow rate and the supply and return temperatures. These values also describe how the heat propagates through the pipes in the network. Changes in flow rate in a DHS will propagate with approximately the speed of sound in water, $1200\frac{m}{s}$. On the other hand, temperature changes propagate with speeds slightly below the speed of the flowing water, because of the walls in the pipelines absorbing some of the temperature differences.

This difference in response time for the different variables in the network leads to some implications and assumptions having to be made. Firstly, one usually has to assume that changes in flow rate happen instantaneous throughout the network, as the dynamics are too fast to model with a typical sample time used in DHS.

Secondary, establishing a precise relationship between incoming and outgoing power on the network becomes more difficult, and necessitates keeping track of the temperature and mass flow rate history at the plant. To elaborate on this, if one measures the power on a heat exchanger at time $k$, this power is decided by the current mass flow rate $\dot{m}_k$, and the

supply temperature $T_{s,k-\tau}$. $\tau = \tau(\dot{m})$ is the time at which the incoming water has been in the network, a function of the mass flow rates during that time.

In order for having a more intuitive approach to incoming and outgoing heat, as well as the total energy in the network, the concept of heat loads is introduced.

The DHS heat load is defined as the power delivered to the consumers at the given time. This means the total power being extracted at the consumer heat exchangers in the entire network. Furthermore, the terms loading and unloading a DHS refer to the action of generating more heat and less at the plant than the current load respectively, over some time interval. This means that loading the system leads to the system having excess heat which may be used at a later time. Excess heat will usually be on the form of increased network temperatures, both on the supply and return line. Similarly, unloading the system means that one delivers less than is currently being extracted, provided one has excess heat to unload. If one unloads the system long enough, the temperatures on the supply and return lines will reduce until the consumers will not be able to extract their desired level of heat.

An approximation which may sometimes be useful when assessing system characteristics of the DHS is the concept of a point load. This approximation assumes that all the heat exchangers that connect consumers to the network are placed at the same location. The point load is useful for acquiring estimates for the outgoing heat of the system.

## 2.3 Load Models

When applying control to DHS, a model of the heat load is a necessity, and the preciseness of this model is a key factor in the performance of the system controller. When there is a large amount of consumers on a DH network, the estimation of the heat load becomes a complex task.

A good model of the heat demand will allow the controller to make decisions that are closer to reality, thus increasing performance. However, improving the model of the DHS may be costly and time consuming. An improvement on a load model may in many cases involve adding additional factors that contribute to the daily heat demand, such as water consumption patterns, ambient wind profile, and more. As a result of this, there are many different models of heat demand, ranging from very simple to extensive and complex. A load model may have many components, and these can be separated into the stochastic and deterministic categories.

Deterministic components are factors in the model where their impact on the heat demand may be explained by physics. The most important deterministic factor is the ambient temperature. However, other effects such as weather and climate effects also fall under this category.

Stochastic components are factors that depend on the social behavior of the consumers. In order to model these, the use of empirical knowledge and data is necessary. These factors are often strongly dependent on the day of week and day of year. This is solved by having different models for weekdays and weekends or holidays, as well as different models for the different seasons of the year.

### 2.3.1 Social Components

The social component of a heat load covers the dynamics in the heat demand that are caused by social behavior of the consumer, which includes a wide range of factors. Some examples are changes to the heat demand caused by when hot water is used by the consumer, opening hours for commercial and public buildings, and local regulation by thermostats in households and offices where the indoor temperature level is deliberately lowered during night time.

## 2.4 Loading of DH Systems

During peak demands in cold periods, the plant or plants in a DHS may often not be able to supply enough heating to meet the consumer demand over time, thus requiring loading. System loading, or packing, is the concept of supplying more heat than required in anticipation of a demand peak, causing an increase in the total energy in the network. This excess heat will accumulate in the DH network, on the form of increased water temperatures on the supply and return lines. When the peak demand occurs, the DH network will experience unloading, or unpacking. During this phase, the overall temperatures in the pipelines are reduced, as the demand is larger than the maximal heat production from the plants, meaning that the total energy in the system is reduced.

System loading is often necessary in order to deal with peak demands on cold days. However, an increase in the pipeline temperatures leads to larger heat losses. In addition to this, the water temperature increase is not uniform, and the return line temperature often increases more than the supply line temperature. This leads to a reduction in the heat exchanger efficiency in the network.

## 2.5 The Evaluation of Pressure and Pumps in DH-Systems

Pressure plays a key part in district heating systems, and there are several challenges that arise surrounding its regulation.

In order to present these in a orderly manner, it is convenient to emphasize the differences between static and dynamic pressure.

#### The Static Pressure

The static pressure is the pressure the water in the pipelines experiences, and its value is independent of the water speed [14]. The differential pressure is defined as the difference in static pressure between the supply and return lines at separate points. The value of the differential pressure adjusts the water speed between these two.

The static pressure is regulated by the addition of pumps in the network. Pumps are added to the network at chosen locations in order to satisfy the following factors:

- The static pressure generated from the pumping system has to be large enough in order for the water flow rate to cover the entirety of the system. In particular it must be ensured that the pressure differential in the customer that is hydraulically the furthers away from the delivery point is satisfactory

- The static pressure has to be balanced in order to stay within a certain safety margin from the generation of steam pressure.

The decisions on the number of pumps placed in the pipelines, their locations, as well as their respective power levels are made according to these factors. Typically, long stretches of pipelines require several pumps placed with more or less equal distance between them.

### The Dynamic Pressure

The dynamic pressure in a DH network is the velocity dependent pressure. The evaluation of the dynamic pressure in a DH network is a computationally demanding and complex task. It requires and in-depth modelling and calculation of the pipeline layouts and properties, as it is dependent on the friction in the pipes, meaning every single turn or curve affects the dynamic pressure drop.

As a result of this, the evaluation of dynamic pressure in DH-networks is normally done using simulation programs, generating an off-line estimation of the dynamic pressure loss.

An additional component to dynamic pressure is the concept of dissipation, which is the increase in the water temperature as a result of the pipe friction converting the pumping energy into heat energy.

### Bernoulli's Equation

A normal approach to pressure evaluation is calculating the static pressure at a chosen point on the supply line, assuming the pressure at the location of the plant as known. This normally is done using the Bernoulli equation. The Bernoulli equation is an approximate relation between pressure, velocity, and elevation, and is valid in regions of steady incompressible flow where net frictional forces are negligible

Bernoulli's formula for a single circulation pump can be stated as follows:

$$p_x = p_1 + H_p + \rho \cdot g \cdot (z_1 - z_x) - \Delta p_x \tag{2.6}$$

where

- $p_x$ = Static pressure at the observed point $[Pa]$
- $p_1$ = Static pressure before the pump, at the plant. This value is chosen in order to avoid boiling in the system $[Pa]$
- $H_p$ = Pressure rise through the circulation pump $[Pa]$
- $z_1$ = Pressure head at the plant $[m]$
- $z_x$ = Pressure head at the observed point $[m]$
- $\Delta p_x$ = Pressure loss through the pipelines from the plant to the observed point $[Pa]$

### 2.5.1 Water Expansion

Another factor of the water temperature in the pipelines is the water expansion. As the pipes are filled with water, and the system is closed, meaning the only water going in or out of the pipe network is at the plant, water expansion has to be accounted for. This is usually solved by keeping expansion tanks at the plant, and filling or depleting these with water from the pipelines according to the current water temperature.

## 2.6 Heat Loss in the Pipelines

Heat loss in the pipelines is an everlasting problem in DH systems, although modern technology on pipe insulation has greatly reduced its impact. There are several methods of estimating the heat loss in a DHS.

In [5], the heat loss in the pipes is modelled using the following partial differential equation:

$$\frac{\partial T}{\partial t}(x,t) + \frac{\dot{m}(t)}{\pi R_p^2}\frac{\partial T}{\partial x} + \frac{2\mu_p}{c_p \rho R_p}(T(x,t) - T_0) = 0 \tag{2.7}$$

where

- $\dot{m}(t)$ = Mass flow rate $[kg/s]$
- $T(x,t)$ = Pipe water temperature $[K]$
- $\frac{\delta T}{\delta x}$ = Temperature loss in the pipe per unit meter $[K/m]$
- $T_0$ = Temperature of the medium outside the pipes $[K]$
- $\rho$ = Relative density of water $[kg/m^3]$
- $R_p$ = Pipe radius $[m]$
- $\mu_p$ = Thermal loss coefficient $[J/(m^2 sK)]$

This formulation offers precise estimates of the heat load, but requires knowledge of heat loss per meter, and the temperature of the substance surrounding the pipes. A simpler method may be to assume a constant loss percentage on pipelines, based on the quality of the pipe isolation.

In the implementation done in this thesis, the heat loss is not directly modelled. However it is assumed that the heat loss increases with the temperature level. Due to this, emphasis has been placed on reducing the water temperature in the implementation done in this thesis.

## 2.7 Network Structures

Although there are large variations in the size and complexity of DH systems, some specific network structures with different properties can be established. Figure 2.3 shows two

different network structures. In the figure, the black lines denote both the supply line and the return line. Star-shaped and ring-shaped systems are mentioned in[15].

In star-shaped DH networks, the water flows from the plant in a single supply line, and the flow is then distributed into lesser flows. Each sub flow is then passed through a heat exchanger, or further split into more sub flows. There are no loops in the supply line, which means that all water that passes through a heat exchanger in a star-shaped network will have taken the same path from the plant to the heat exchanger. This means that the distance from the plant to each individual heat exchanger or consumer can be calculated.

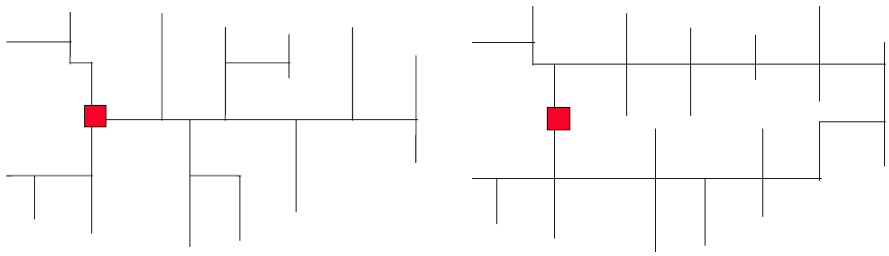

Figure 2.3: Different network structures. *Left*: Star-shaped network. *Right*: Ring-shaped network

In a ring-shaped structure, two supply lines exit the plant, and these are connected at some point in the network, creating a loop. In [15] it is mentioned that this provides more security in terms of stable supply of heating. It is however also mentioned to lead to larger heat losses than the star-shaped structure. As opposed to the star-shaped structure, it is not possible to establish the definite path of the water passing through a heat exchanger.

## 2.8 Different Aspects of DHS Regulation and Control

Regulation and optimization of DH systems is a widespread topic, and there are several aspects of this that can be considered standalone subjects of their own. This section will attempt to describe how the subjects are connected, and which are separable.

### Pumps, Pressures and Network Expansion

The placement of pumps in district heating systems is a static optimization problem. As mentioned in Section 2.5, the network needs to place pumps in the pipelines order to ensure water flow to every customer, but also without violating safety constraints that prevent steam generation in the pipelines. Although this optimization includes flows and temperatures as parameters, these have to be assumed as steady state parameters. This means that this static optimization is separated from the dynamic optimization of satisfying consumer heating demands.

Problems surrounding expansion, maintenance and general changes in the DH network structure are also static optimization problems, and these two problems are connected.

Whenever the pipe layout is changed, a re-evaluation of pumping requirements should be performed.

### Load Models

Having accurate models of the predicted customer demand is essential for achieving good performances when applying control to DH systems. The generation of load models can be done independent of any regulation made in the DHS. The components in these models are mentioned in Section 2.3, and do not include any operation variables.

### Temperature Optimization and Load Distribution

The subjects of distribution of heat production among heat sources in the network, and minimization of supply temperatures are much more intertwined. In early days of model predictive control theory of DH systems, [2] divided the methods into the following three categories:

- Determination of optimal load distribution among different heat producing units.
- Minimization of the supply temperatures without accounting for the possibility of using the DH network as a heat storage
- Full dynamic optimization determining both optimal supply temperatures and optimal load distribution.

The implementation performed in this thesis is applied to a district heating system with only a single source, meaning no evaluation of load distribution is necessary. Distribution of load has therefore not been heavily discussed in this thesis.

Network aggregation is a tool for simplifying complex network structures and finding simple models with equivalent properties and dynamics. Network aggregation is largely used together with load distribution evaluation.

### The Environmental vs. Economic Aspect

The evaluation of the heat production in the district heating system has to factor in both an economic and an environmental aspect. The DHS operators need to have a clear definition of how economic and green each heat production unit is, and how to value these two aspects against each other. When the system contains several heating stations, this problem is combined with the distribution of load, and is used to establish an ideal distribution.

Chapter 3

# Large-Scale Process Control

When applying control to large-scale processes, as opposed to processes of smaller size, several control challenges arise. This could be because the geographic location of the process spans a large area, or that it involves many lesser processes that require cooperation among them in order to retain operational feasibility. In this regard, having tools to provide transparency and overview of the problem at hand becomes important.

## 3.1 Model-Based Design

Model-based design (MBD) refers to the method of creating mathematical models that imitate the process dynamics one wishes to control. The aim of model-based design is to be able to view each smaller module of a process as an individual component. Every module will have its own purpose and an approach to achieving said purpose. When applying MBD, the module's functionality is modelled using a mathematical formulation. An example of this could be a pump applying power in order to achieve a certain desired pressure level. Using model-based design opens up the possibility of looking at the single module as a standalone component.

There are several benefits of MBD, due to the increased visualization of the process and a better transparency in the system dynamics. Some of the most notable ones are [16]:

- Identifying the domain of each module
- The possibility of validating each module to real data, thus also being able to validate the complete process model
- The possibility of being able to run simulations and case testing on the process or a single module. May help in identifying system properties and constraints
- Hardware-in-the-loop simulations. The possibility of exchanging a simulation model with its physical counterpart. Can be used gradually evolve from controlling a simulation to a real physical system

## 3.2 Model Predictive Control

Since its introduction to the process industry in the 1980's, Model Predictive Control has experienced a steady increase in popularity and usage, and today it is one of the most extensively used control algorithms used in multi-variable control [17]. The concept of model predictive control is to efficiently steer the states towards the optimal values, by generating a prediction of future states. This prediction may be obtained from previous tests on the system, or by applying some system identification technique to the observed plant data.
Normally, MPC operation is done in a discrete time framework. At each sample, a model predictive controller generates a control input sequence, which spans the prediction horizon. This sequence is generated by minimizing some cost function, without violating constraints of the system. The controller then applies a first control move for the system, using the first element in its input sequence, and discarding the rest.

Although most real processes include some amount of nonlinearity, linear model predictive controllers are quite common. This is because nonlinearities often are very small close to stable equilibrium points, so that the linear model is a good approximation of its dynamics. The following description shows the formulation of MPC using a linear quadratic regulation problem for a discrete system:
Consider the discrete state space model

$$x_{k+1} = Ax_k + Bu_k \tag{3.1}$$

where $x_k \in \Re^n$ is the current state, $u_k \in \Re^m$ is the current input, and $A \in \Re^{n\times n}$ and $B \in \Re^{m\times n}$ are constant state and input matrices. $S = S^T \geq 0, \quad S \in \Re^{n\times n}$

The system solves the finite-time optimal control problem:

$$\min_{\bar{u}} J(x,\bar{u}) = x_{k+N}^T S x_{k+N} + \sum_{j=0}^{N-1} x_{k+j}^T Q x_{k+j} + u_{k+j}^{k,T} R u_{k+j}^k \tag{3.2}$$

$$s.t. \qquad x_{k+1+j} = Ax_{k+j} + Bu_{k+j}^k, \qquad j = 0,\ldots,N-1 \tag{3.3}$$

$$u_{\min} \leq u_{k+j}^k \leq u_{\max} \qquad j = 0,\ldots,N-1 \tag{3.4}$$

$$x_{\min} \leq x_{k+j} \leq x_{\max} \qquad j = 0,\ldots,N-1 \tag{3.5}$$

Where $N$ is the prediction horizon, $\bar{u}_k = \left[u_k^{k,T}, u_{k+1}^{k,T}, \ldots, u_{k+N-1}^{k,T}\right]^T$ is the input sequence at time $k$. $S = S^T \geq 0, \quad S \in \Re^{n\times n}$, $Q = Q^T \geq 0, \quad Q \in \Re^{n\times n}$, and $R = R^T > 0, \quad R \in \Re^{m\times m}$ are weighting matrices for the final state of the horizon, remaining states on the horizon, and inputs on the horizon, respectively.

## 3.3 Real-Time Optimization

In large-scale systems, processes will often involve optimization of variables that experience change on very different time scales. Some variables may only change yearly or

seasonally, and although they are to be optimized, their values can usually be considered static.

Real-time optimization (RTO) is a method of optimizing slowly changing variables on a higher level in the control hierarchy, and using the optimized values as setpoints for the faster changing dynamics. The motivation behind this separation is economic gain, by being able to reduce excessive computations on slowly changing dynamics, while still keeping track of fast changing process dynamics. The division of optimization to several layers may also reduce the overall problem complexity, by having two reduced problems instead of one large problem. Another possibility with the usage of RTO is to limit non-linearities and complex problems to only parts of the overall problem, thus allowing simpler dynamics to be solved using quicker solvers that are easier to use and implement.

The time-scale that the different layers in the control hierarchy operate with are not predefined. However, the upper layers need to have significantly slower dynamics than the layers below. When this is not the case, a variable received from an above layer as a static setpoint may in reality have changed, reducing the overall control system performance, and risking infeasible or unstable operation.

In Figure 3.1, a hierarchical control structure for MPC with the addition RTO is shown. In this example, the RTO-layer covers the yearly, seasonally, and daily changing variables, while the model predictive controller dynamics are on the scale of hours to minutes. The bottom layer consists of the basic control. This layer will typically be local regulators, controlling that variables are kept to their desired levels.

## 3.4 The Economic Dispatch Problem and Unit Commitment Problem

The unit commitment problem (UCP) is the problem of meeting a specified demand for a scheduled period, by selecting when different units, or inputs, should be activated. This also covers how long the units should be committed. The units are required to meet the specified demand under operational constraints, and by minimizing the operating cost.

The economic dispatch problem (EDP) deals with the distribution of the demand among the units that are currently units in order to meet it as cheaply as possible For the EDP this means that the input with the lowest cost is actuated until it meets its constraints. Then the input with the second lowest cost is actuated, and this continues until the current demand is met.

An example of unit commitment problem and economic dispatch problem in a DHS framework, could be that the UCP would consist of evaluating which heat sources should be active throughout the day, based on the predicted load demand, and create a scheduling determining when each heating station should be active. The EDP would then at each time step evaluate the active heating stations and the load demand, and provide the distribution of input actuation among the active heat sources.

The unit commitment problem is a problem which typically includes discrete variables. These are reflected on the constraints set by the real plants, concerning issues such as start-up time, shut-down time, associated costs with start-up and shut-down, as well as plants being required to cool down after being active for a while. This makes the UCP a rather

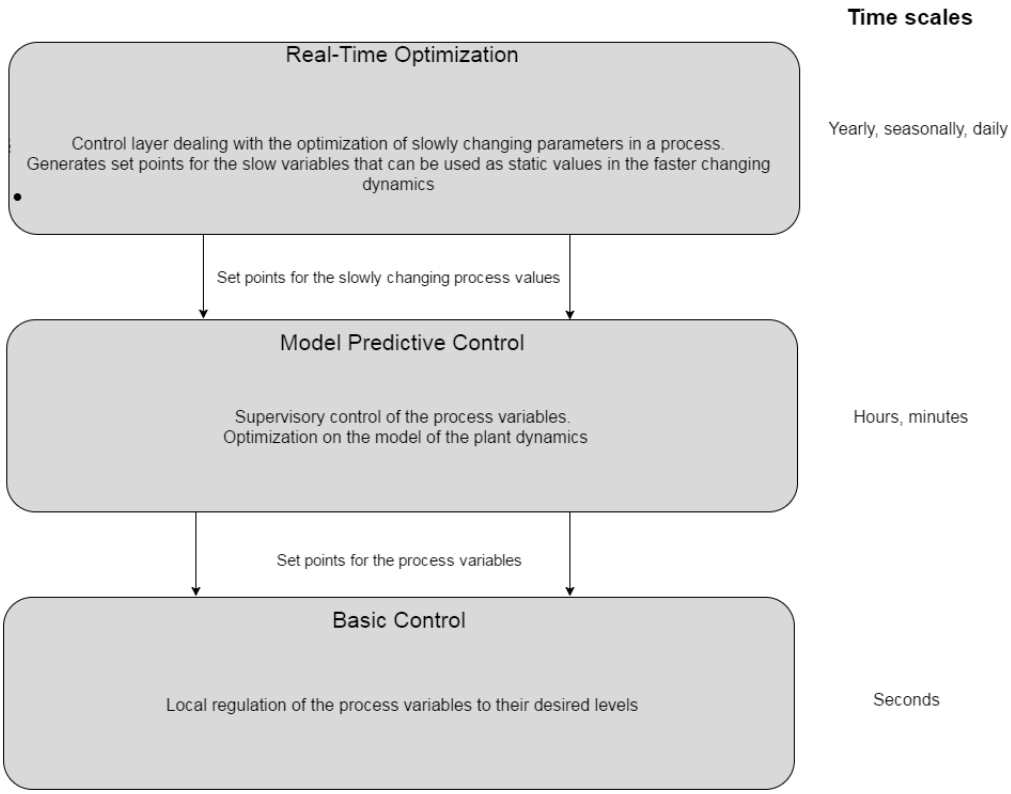


Figure 3.1: A three-layered control hierarchy with a RTO as the top layer, MPC-control as the middle layer, and basic control at the bottom

complex problem to solve, by having to use a mixed-integer solver. A formulation for the thermal mixed-integer linear problem UCP can be found in [18]. However, for the implementation part in this thesis, neither UCP nor EDP were used, due to the simplicity of the DHS subject to simulation.

## 3.5 Robustness

In terms of regulation and control, robustness refers to the ability of handling uncertainties in the system that may be model errors or disturbances. For constrained systems, extra attention is required to constraints close to their boundaries, as disturbances and errors could cause them to be violated.

There are several ways of dealing with robustness of an MPC controller. In [19], a method of guaranteeing robustness through inheritance is mentioned. In this definition of robustness, the nominal closed-loop system, ignoring disturbances is investigated. Given an assumption of bounded disturbances, one is able to guarantee stability of the system

with disturbances within a region of attraction.
A strong formulation of robustness in model predictive control was introduced in [20], through the concept of the tube MPC, having properties such as a large domain of attraction and being linear in complexity. In addition to this, tube MPC can be used when the system is time-varying or subject to parameter uncertainty.

Chapter 4

# Statkraft Varme's District Heating System in Trondheim

This chapter attempts to describe operation of the DHS in Trondheim, and challenges surrounding it. The network is owned and operated by Statkraft. Information of the network and the current operational regime has been attained through visits to the operations center, and through conversations with the operators and employees in Statkraft. Some statements and values may therefore have been subject to human error.

Statkraft AS is an international power company owned completely by the Norwegian government, and the largest producer of renewable energy in Europe [21]. Statkraft Varme is Statkraft's department for district heating, and is today in charge of district heating systems in 13 Norwegian cities and urban areas, delivering a total of over 800GWh. In addition to this, Statkraft Varme delivers 200GWh of district heating in Sweden.

## 4.1 The DH Network in Trondheim

The DH network in Trondheim delivers over 570 GWh, making it Statkraft's largest DH network by far. It is also the second largest DH network in Norway. The Trondheim DH network delivers heating to over 7000 households and more than 500 businesses.

Its structure consists of a primary network and a number of secondary networks. In the primary network there are 10 primary heating stations, and over 150 km of pipelines. The primary and secondary networks are connected through around 40 under-stations. Subsequently, households, and commercial and public buildings are connected on the secondary networks through heat exchangers.

The system generates heat on the heating stations by using several different types of heating supply. More specifically, the system uses biofuel, liquefied petroleum gas (LPG), natural gas (LNG), electric cauldrons, oil cauldrons, heating from waste water treatment, and landfill gas. The largest energy source is however the heat generated by waste inciner-

ation at Heimdal heating station, which is also a renewable energy and the cheapest source of heating of the mentioned.

In total, the system delivers a supply of 500 GWh, which corresponds to around 30 percent of Trondheim's heating requirements. A map of the Statkraft Varme district heating network can be seen in Figure 4.1.

The primary network in the DH network is separated into four parts by pressure separators. The separation points are located at Byåsen, Nedre Leirfoss, Lilleby and Eberg. The reasoning for having separation of the primary network is the large altitude differences in the network. From the highest point in the primary network to the lowest, there is a height difference of about 160 meters, which leads to a pressure difference of almost 16 Bar. The pipe network has been built over a long period of time, with gradual expansions over several decades. The older pipes were not constructed with the capacity of handling the higher pressure flows, and thus the pressure separations were necessary.

The district heating network is organized as follows: The primary network consists of the 10 heating stations, and the heat exchangers previously mentioned. The transported water usually hovers at temperatures between $90-120\ ^{\circ}C$ for the supply line, and between $60-80\ ^{\circ}C$ for the return line. The differential pressure is the pressure differences between the supply line and the return line, and typically varies in value between 0.5 bar and 10 bar.

The secondary networks connect buildings and individual households to the primary network. They are connected to the main network through under-stations. A delivered temperature of $120\ ^{\circ}C$ is usually not desirable in individual households, so the secondary network temperatures are usually chosen to be around $80-90\ ^{\circ}C$. Lastly, buildings and households are connected to the secondary network through their own heat exchangers, with individual regulation.

## 4.2 The Current Operational Regime of Trondheim DHS

Regulation of the DH network in Trondheim is done in the control room at the heating station on Heimdal. In the control room, there are engineers present at all times, that oversee that production is done efficiently, and handle challenges and errors that occur in the system.

The regulation of the heat delivered on the system is done through management of the supply line temperature, the return line water temperature, and the differential pressure between the lines, at different critical points in the network.

The desired values for these variables are generated through regulations for customer specifications. Any customer connected on the DH network has a specified domain for the delivered differential pressure and supply line water temperature measured at the customer's heat exchanger. The DH operator is contractually obliged to meet the values specified by these regulations.

In order to meet the customer demands on all parts of the network, the DH operators have created tables with ideal operational values for the temperatures and differential


Figure 4.1: Schematic representation of the DH network in Trondheim Norway. The heating stations are the buildings with the title "varmesentral". The pipelines are illustrated by the red lines. Several lesser branches are omitted in order for making the figure more readable.

pressures at critical points in the network. These tables set a desired range for the values of these variables based on the outside temperature and the weather. The knowledge of these desired value ranges and the locations of the critical points is based on operational experience.

## 4.3 Challenges surrounding Operation on Trondheim DHS

This section describes some of the special cases and challenges that are present in the Statkraft Varme DH network. Some are constant properties, like delays in the system, while others may be seasonal or sporadic occurrences.

### 4.3.1 Multiple types of Heat Sources and Cost Variation in the System

The large amount of different types of energy sources in Trondheim DHS makes operating economically and environmentally efficient important. This is a large focus of today's operational regime, with the engineers having an interactive feedback of the current operating efficiency of these factors. The engineers are also given an economic incentive for

efficient and green operation, in the form of a bonus payment if certain monthly efficiency levels are met.

There are currently no tools for analyzing the effect each plant has on different parts of the network. When situations arise where an increased demand has to be met on a specific part of the network, the engineers have to make decisions based on experience, both for where to meet the demand, and how to meet the demand.

### 4.3.2 Separation of the Primary Network

As mentioned in the above section, the primary network is separated into four parts by heat exchangers. Though the heat exchangers are placed in designated buildings, and with very low heat losses, they may bring an extra degree of complexity to any operation of the network. The efficiencies of the heat exchangers vary based on the temperatures of the two flows, which both poses as an extra modelling factor and a potential bottleneck in the system.

### 4.3.3 Switching of the Water Flow Direction in the System

Due to the large number of plants in the network, the direction of the flow in the different pipelines is not always be the same. In some cases of the system, especially at the Eberg pressure separator, the flow direction often changes multiple times during a day.

At the Eberg pressure separator, this is due to whether or not the LPG boiler at Dragvoll is running. When it is, it is usually generating more than enough heat for its part of the network, meaning it delivers downwards towards the Lilleby network. When it is turned off, heating is delivered upwards, from the Lilleby network to the Dragvoll network.

An additional property that contributes to this frequent change of flow direction is the varying electricity prices. The butane gas production at Dragvoll is very close in price to electrical heating which, among other places in the network, can be produced at Lilleby heating station. Dragvoll and Lilleby both supply heat to the areas around Tyholt and Lade as can be seen in Figure 4.1. Electricity prices change daily, which means that these two sources of heating often swap places in being the most cost efficient and green energy source.

With regards to operation, the switching flow direction makes it significantly more difficult to establish concrete input output relationships between different points in the network. As a consequence of this, model aggregation of Trondheim DHS could possibly be very challenging.

### 4.3.4 Boiler Fallout

Another major challenge of operating the DH system is the boiler fallouts. The boilers are intricate processes operating under difficult conditions, and faults are unavoidable. The key to minimizing losses associated to boiler fallouts, is to have good procedures for quick and effective response when they occur. The maintenance of boiler units is outside the

scope of this thesis, but it relies on good fault reporting and communication between the operators at different work shifts.

When boiler fallout occurs at important plants, the distribution of delivered heat changes. These changes are especially large if a boiler goes down at Heimdal heating station. Heat generation at Heimdal is normally performed non-stop throughout the year, as Statkraft is paid to incinerate received waste. Therefore the control regime relies on stable heat generation from this plant. However, fallout of one of the three burners at Heimdal is not uncommon. When this happens, it changes the critical points and measurements in the network. This means that the measurements one normally relies on to represent the current operating conditions on different parts of the network may no longer be representative.

### 4.3.5 Meeting Customer Demands on the Complete Network

Due to the large number of customers on Trondheim DHS, measurements on every customer's heat exchanger is not realistic. Instead, values at critical points are measured in order to evaluate whether customers in an area receive satisfactory levels of temperature and pressure. When their levels are too low, customers can report it, and the operators will increase the input to meet the desired levels. The temperature and pressure levels may also be too high for some customers, causing inconvenience in the extraction of heat.

The desired levels for temperatures and differential pressure are adjusted by use of this feedback, thus in many places creating a pragmatic solution to the system operating levels.

### 4.3.6 Integrating Heat Accumulation in the System

A concept which may further increase effectiveness of a DHS is to add an accumulator tank to the system. By adding an accumulator, one can store hot water during periods of low demand or when prices are down, and then use this when the demand rises. At the moment, there are no accumulators in Trondheim DHS. However, it is a proposed project, likely to appear.

This creates the interesting problem of where the ideal placement of the accumulator would be. In order to fill the accumulator tank with cheap heat, the natural choice is the plant with the cheapest source of heating, which for Trondheim DHS is the plant at Heimdal. However, the outgoing pipelines at Heimdal often reach their capacity during the critical periods of cold days. Therefore, an accumulator would not be able to contribute much during these periods. An alternative candidate for the location of an accumulator tank is the plant at Øya. The heating generated at Øya is not as cheap, but the capacity problem here is not as imminent.

### 4.3.7 Heat Exchanger Bypassing

Another proposed improvement of the DHN, which is currently being implemented, is the added option of being able to bypass the Lilleby heat exchanger. By bypassing the heat exchanger, and allowing the water to flow directly through this area instead, the operating system is reduced by one potential bottleneck.

Chapter 5

# Simulating Model Predictive Control on the District Heating System in Klæbu

Klæbu is a small urban area south of Trondheim with a population of approximately 6000 inhabitants. The urban center of Klæbu (hereby mentioned as Klæbu Centre) lies 19 km south of Trondheim and approximately half of the inhabitants in Klæbu live here. The small-scale district heating system that covers Klæbu Centre started operating in 2004.

The Klæbu DHS generates all its heated water from a single heat plant, with a yearly production of about 6 GWh [22]. Biofuel, in the form of incineration of wood pellets, is the main source of heating in this plant. The plant does also have the option of generating heat from an electrical cauldron and oil burners when the biofuel is insufficient in cold periods. The heating capacity of the plant is about 7.1 MW, which is a combined capacity of 1.5 MW biofuel, 2.4 MW from electrical boilers, and 3.2 MW from oil burners. As it is the cheapest and greenest heat resource of the three, the biofuel is the chosen source of heating in the Klæbu DHS, with the addition of the electrical boiler and oil burners in that order, when the heat demand dictates it.

Even though they are fairly close in proximity, the Klæbu network is not connected to the DH network in Trondheim. However, the Klæbu network is almost entirely remotely controlled from Heimdal, where the Trondheim DHS is regulated. The exception to this is regular routine checks at the heat source in Klæbu, to assure that nothing out of the ordinary has happened.

An image of the area covered by Klæbu DHS is given in Figure 5.1. The DHS is divided into two networks that are completely separated at the heat source. These networks are the Centre network, covering the municipality center north-west of the plant, and the Halset-network, covering the area north-east of the plant. The networks are shown in the

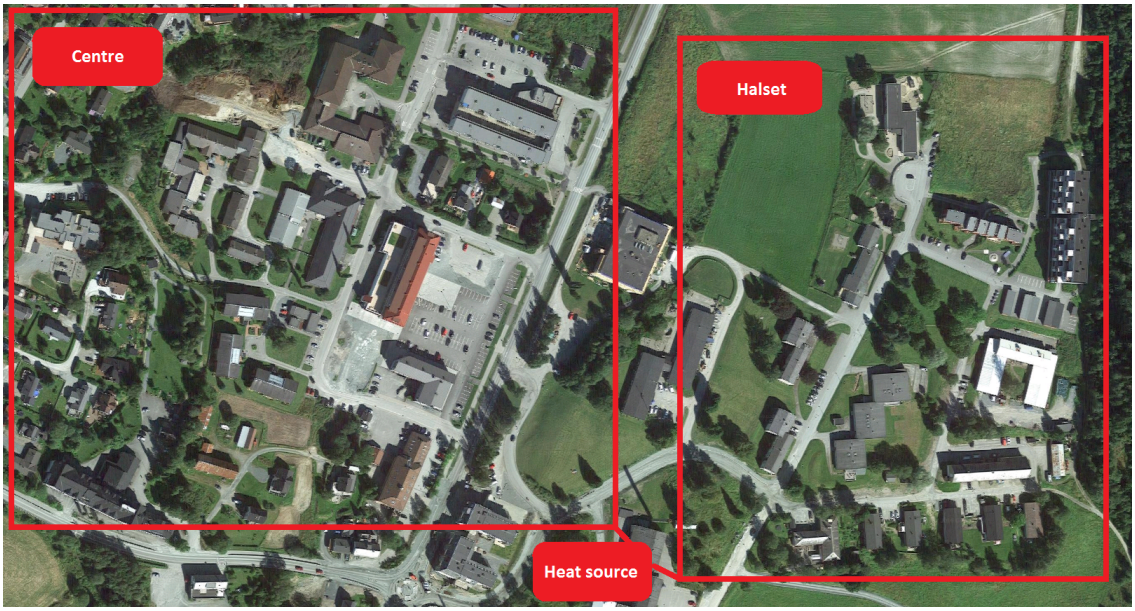


Figure 5.1: The area covered by the district heating system in Klæbu. The system is separated into two networks, divided at the plant. These are the network for the municipality center, indicated by the left box, and the Halset-network indicated by the right box

figure, the left box shows a rough approximation of the area covered by the Centre network, and the right box shows the area covered by the Halset-network. The Halset network is an older part of the DHS in Klæbu. The pipelines in this part of the network were installed without modern isolation technology, which causes heat losses in this part of the network to be much larger than the Centre network.

Due to its simple structure, with a low number of connected customers and single heat source, the western line of Klæbu DHS has been chosen as the network subject to the modelling and optimization done in this thesis.

## 5.1 The Consumer Demand Model

As the objective of operational optimization of DH systems is to meet the heating demand as cost efficiently as possible, a large emphasis has to be put on the modelling of the predicted demands. A natural approach to control of a DHS is to establish a confidence interval on the daily demand, so that the system is known to operate within the boundaries of the confidence interval with a chosen certainty.
By having accurate models, the daily variance between the observed and predicted demand will be limited, thus allowing for a tight bound for the chosen confidence level, and therefore more stable operation.

The daily heating predictions are usually produced by a combination of physical modelling and empirical knowledge of the system. The total demand of a DHS consists of the combined demand from every consumer on the network. However, for large networks, keeping track of every single consumer is too challenging: it is vastly time consuming, as well as requiring individual measurements at each consumer, which the DHS operators normally do not store.

When individual demand prediction is not an option, the alternative approach is to produce generalized demand prediction curves. This is done by separating the consumers in different classes, where each class exhibits similar trends in their daily heating demands, and then generate a demand curve for each class.
This could typically mean that one has a generalized demand curve for heat exchangers delivering to households, and one generalized curve for medium sized commercial and public buildings. Large commercial buildings such as hospitals or large industrial complexes may have their own prediction demand curve.

The Klæbu DHS delivers heat to five building complexes in total, through one heat exchanger for each building complex. Because of the low number of consumers on the network, individual models have been created for each consumer. A schematic figure of the customers in Klæbu can be seen in Figure 5.2. In this figure, the five consumers, mentioned as consumer A-E, are shown along with their location in the network measured in the volume of the supply line pipes. The supply line pipes are assumed to be full of water at all times, so the unit chosen to show the distance of the pipes is given in liters.

For the creation of models for the consumer demands, data has been extracted from the customers' five heat exchangers. This data has been used for making prediction models of consumer demand for weekdays based on time of day, outside temperature, and season, according to theory mentioned in Section 2.3.
The data consists of hourly values of measured power, flow rate, supply temperature and return temperature. The data has been stored, and for the work in this thesis, values ranging back to 1.1.2016 have been used.
In addition to the consumer data, data from the plant has also been used in order to investigate the dynamics of the system. In the plant, a large amount of variables are logged for every ten minute interval. This data is normally used during the regular inspections of the plant, to check if the plant has been operating at the desired levels.

The total demand in Klæbu Centre DHS is the cumulative demand from the five consumers on the network. For the cumulative demand on the network, a point load approximation has been used, which assumes that all loads in the system are at a single location. Comparing the structure of the network to the mentioned structures in Section 2.7, the Klæbu DHS has a star-shaped structure, which implies that individual demands can be added if they are shifted to a common location. The shifting is done in order for the consumer demands to appear equal as viewed from the plant, before and after the point load approximation. Therefore, the shift in the demand curves is chosen by the difference in transport delay between the plant and the original location, and between the plant and the

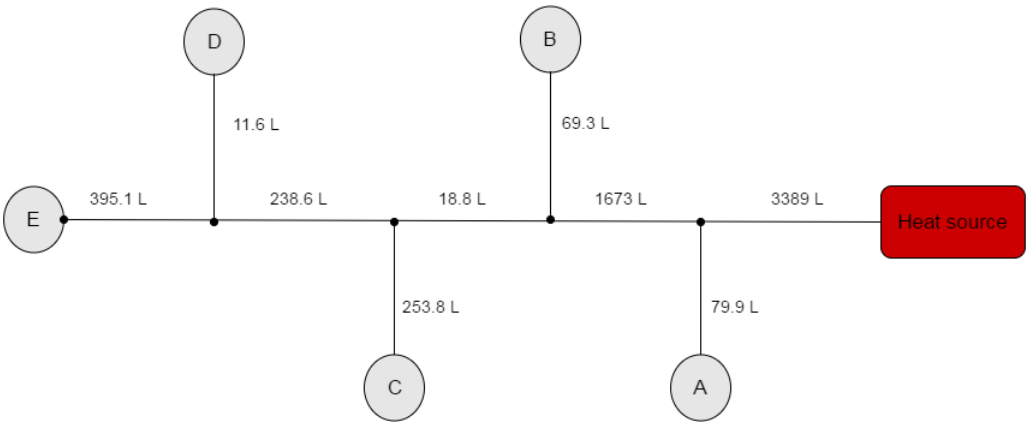


Figure 5.2: Schematic graph that displays the five consumers on the network showing the volume of each pipe segment in the supply line in liters. These are used for calculating the transport time of the water to the different heat exchangers. The grey circles are consumers A to E. The heat source is given by the red rectangle

point load location.

## 5.2 The Proposed Control Hierarchy

Due to the size of the DHS control problem, a multi-layered control hierarchy is necessary, with the addition of a RTO-layer, in a similar fashion as shown in Figure 3.1 in Section 3.3.

For the chosen control structure in this thesis, the bottom layer is be the basic regulation of the process variables such as the plant supply temperatures, flow rate levels and pump pressure levels. The low-level controllers receive their desired levels from the middle level in the hierarchy, the model predictive controller. The MPC captures the dynamics of the water in the pipes. As it must capture the complete dynamics, the MPC must have a prediction horizon that is longer than the slowest dynamic in the pipe network. This is transport delay of the longest pipeline distance in the network, which can be several hours, thus setting the time-scale of the MPC layer from minutes to several hours.

The processes that change slower than on a basis of several hours are placed on the top hierarchical level. This includes seasonal changes and changes to the structure of the pipe system. The latter could be caused by maintenance or expansion of the network to new areas. However, the daily load predictions have also been placed in this layer. The reasoning for this is that even though the weather predictions may change hourly, the horizon for which one must evaluate the predicted temperatures ahead will often be further than the model predictive controller. The predicted demand and pipe dynamics are also decoupled problems, thus it is possible to separate them in the control structure.

With the inclusion of the heat demand predictor in the upper hierarchical layer, this now becomes on the scale of years to hours, which is a very wide formulation. A natural way of reformulating this control hierarchy is to divide the top layer in the hierarchy into two separate layers, one with the yearly and seasonally changing parameters, and one with the heat demand predictor. A figure visualizing the resulting control hierarchy is shown in Figure 5.3.

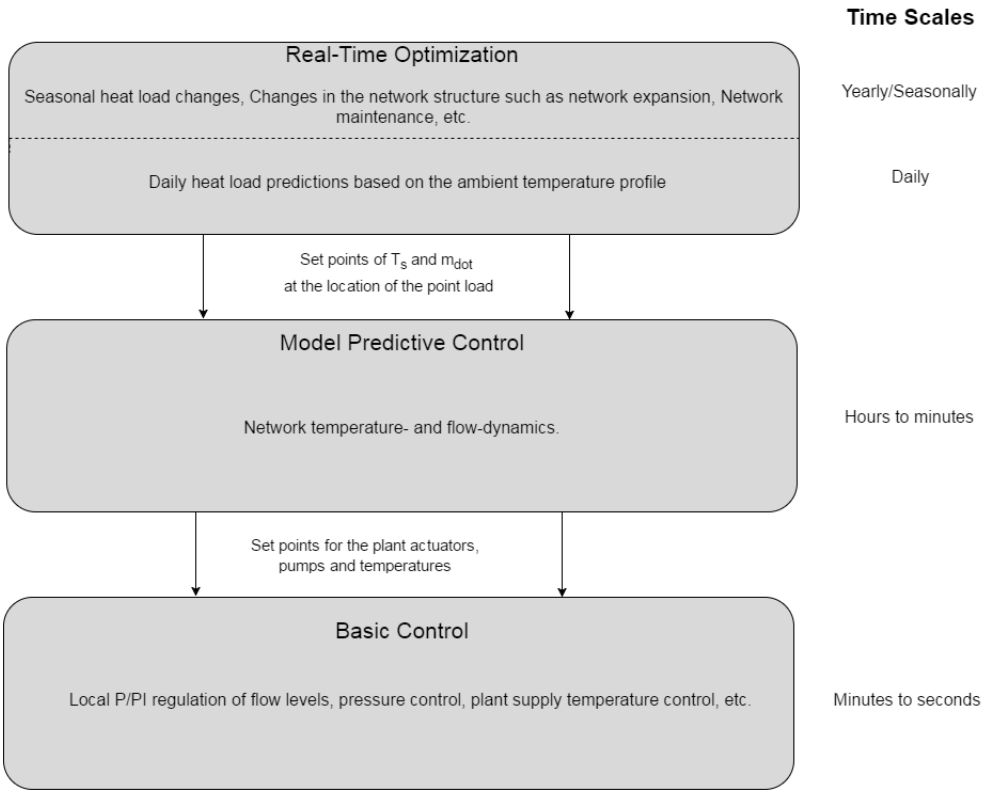


Figure 5.3: The control hierarchy in DHS.

## 5.3 The RTO Control Layer

For the implementation in this thesis, the RTO is separated between the yearly/monthly-changing parameters and the daily changing parameters, as given by the dashed line in Figure 5.3. The upper part of the separated RTO layer consists of the yearly and seasonally changing parameters such as the how the season affects the heat usage, and changes in the network structure. The variables in the upper layer have been assumed to be constant for the duration of implementation done in this thesis.

This means that the daily/hourly-changing variables denote the new top hierarchical control layer, and the slowest changing dynamics in the implementation. This layer is therefore often mentioned as the slow model in this thesis.

The slow model takes in the predicted ambient temperatures as well as the weekday, and estimates the load model. Furthermore, the slow model evaluates the estimated load, and delivers setpoints for supply temperature, $T_s$, and mass flow rate, $\dot{m}$, to the model predictive controller layer.

The setpoints delivered to the MPC from the slow model are found using an optimization with the following formulation:

$$\underset{T_s,\dot{m}}{\text{Minimize}} \quad c_1\dot{m} + c_2|T_s - T_s^*| \tag{5.1}$$

$$\text{subject to} \tag{5.2}$$

$$c_p\dot{m}(T_s - T_r) = \dot{Q}_{\text{pred}} \tag{5.3}$$

$$\dot{m}_{\text{min}} \leq \dot{m} \leq \dot{m}_{\text{max}} \tag{5.4}$$

$$T_{s,\text{min}} \leq T_s \leq T_{s,\text{max}} \tag{5.5}$$

$$T_r = \text{constant} \tag{5.6}$$

$$T_s^* = T_r + 40 = \text{constant} \tag{5.7}$$

Where

- $\dot{m}$ = Mass flow rate on the network. Setpoint for the MPC $[kg/s]$
- $T_s$ = Temperature of the supply line at the location of the load. Setpoint for the MPC $[^\circ C]$
- $c_1$ = The cost for a unit increase in $\dot{m}$ $[\cdot]$
- $c_2$ = The cost for a unit increase in $T_s$ $[\cdot]$
- $c_p$ = The specific heat capacity for water $[\frac{Ws}{kgK}]$
- $T_r$ = The return line temperature at the location of the load. Assumed constant $[^\circ C]$
- $T_s^* = T_r + 40$ = An assumed supply line temperature that maximizes the heat exchanger $[^\circ C]$ efficiency with as little heat loss as possible

Here, $c_2 > c_1$, is used as a simple way to penalize increase in temperature more than increase in mass flow rate. This way, the slow model keeps the supply temperature value at $T_s^*$ whenever it is possible to meet the predicted demand based only on changes of $\dot{m}$ within its feasible area.

As the optimization variables in the slow model are used as setpoints for the model predictive controller in the layer below, setpoints for the full control horizon of the MPC-controller need to be delivered before the MPC can perform its optimization. Therefore the slow model needs run its optimization and generate setpoints for each ten minute time step of the MPC-controller, several hours ahead in time. It is important to note that these optimizations are independent of each other, and that the slow model has no memory. This means that there are no specified constraints on the rate of change for $\dot{m}$ and $T_s$, these have

to be handled by the MPC.

For the implementation done in this thesis, the slow model delivers setpoints for the full 24 hour simulation duration to the model predictive controller prior to the MPC simulation. However, these setpoints are dependent on the outside temperature predictions that are updated regularly.

If a simulator run in real time were to be implemented, a natural scenario could be for the slow model to deliver setpoints 12 hours into the future every 4 hours. The updated values would then replace the older ones where there would be an overlap, as they would be more reliable. As long as the control horizon of the MPC does not extend 8 hours, the model predictive controller would at all times have the setpoints necessary for operation in this scenario.

### 5.3.1 The Constraints in the Slow Model

The constraints in the slow model are the maximal and minimal values for the mass flow rate and supply temperature delivered from the plant. These values are set as realistic as possible, by using the data from the plant and the five users' heat exchangers during the period January 1st 2016 to May 4th 2016. The values can be seen in table 5.1. The maximum and minimum values for the supply temperatures are chosen as the observed maximum and minimum temperatures delivered from the plant during that period.

However, the plant does not have separate measurements for the flow rate going out to the Klæbu Centre network. The maximal value for flow rate are chosen using measurement data from the heat exchangers.

This value is chosen somewhat roughly, as the sum of the five maximal flow rates passing through the heat exchangers, with a value found to be $2.481[kg/s]$, and this has been chosen to be the absolute maximal flow rate value in the implementation.

Another factor that plays a key part in choosing the mass flow rate limit, is to reserve some capacity so one can account for sudden variations in prediction demand. When the daily demand changes rapidly from the predicted demand, extra heat often has to be delivered at quick notice, which can not be done by increasing the supply temperature, as its effect is too slow.

Therefore, a nominal maximal value of the flow rate is chosen to be roughly two thirds of the measured maximal value, so that the final third is reserved for compensating for daily variations. The nominal maximal value for the flow rate was chosen to be $1.667[kg/s]$, and is value is used as the upper mass flow rate constraint in the slow model.

## 5.4 The Model Predictive Controller

The task of the model predictive controller in the control hierarchy is to evaluate the dynamics in the pipelines and meet the desired temperatures at the location of the point load. This is achieved by evaluating the current and future delays from the plant to the load, and apply the reference temperature levels for the load at the correct time.

The MPC receives setpoints for the values of $T_s$ and $\dot{m}$ at the heat exchanger's position for all time steps of the prediction horizon.

Table 5.1: Maximum, minimum and average values for mass flow rate and supply temperature

| **Constraint** | **Value** | **Unit** |
|---|---|---|
| Maximum $T_s$ | 111 | $°C$ |
| Minimum $T_s$ | 72.700 | $°C$ |
| Absolute maximum $\dot{m}$ | 2.481 | $kg/s$ |
| Nominal maximum $\dot{m}$ | 1.667 | $kg/s$ |
| Minimum $\dot{m}$ | 0 | $kg/s$ |
| Average $\dot{m}$ | 0.547 | $kg/s$ |

The MPC-controller is formulated as follows:

$$\text{Minimize} \sum_{k=1}^{N} c_y|y_k - y_{\text{ref},k}| + c_u u_k + c_{\Delta u}(u_k - u_{k-1}) \tag{5.8}$$

subject to (5.9)

$$u_{min} \leq u_k \leq u_{max}, \qquad k = 1, \ldots, N \tag{5.10}$$

$$x_{min} \leq x_k \leq x_{max}, \qquad k = 1, \ldots, N \tag{5.11}$$

$$\Delta u_{min} \leq u_k - u_{k-1} \leq \Delta u_{max}, \qquad k = 1, \ldots, N \tag{5.12}$$

$$x_{k+1} = Ax_k + Bu_k \qquad k = 1, \ldots, N \tag{5.13}$$

$$y_k = C_j x_k \qquad k = 1, \ldots, N \tag{5.14}$$

Where:

- $k$ = Prediction horizon variable. Denotes which step on the prediction horizon is currently being evaluated
- $N$ = The prediction horizon
- $y_k$ = Temperature at the heat exchanger at time step $k$
- $y_{\text{ref},k}$ = Reference temperature at the heat exchanger at step $k$, given by the slow model
- $u_k$ = Temperature of the water delivered from the plant at step $k$. The MPC-controller's actuator variable
- $c_y$ = The cost of deviating $y_k$ from its setpoint
- $c_u$ = The input cost
- $c_{\Delta u}$ = The delta penalty. The cost of change in the input
- $A$ = State matrix. Constant for the duration of the simulation
- $B$ = Input matrix. Constant for the duration of the simulation

- $C_j$ = Output matrix. Chosen from a set of static C-matrices according to the current delay from the plant to the heat exchanger

The MPC-formulation drives the plant to deliver the desired temperatures at the plant with as little input and delta input as possible. The state space formulation describes how the delivered water propagates in the network, and is defined as follows:

$$\begin{bmatrix} x_{1,k+1} \\ \vdots \\ x_{14,k+1} \end{bmatrix} = \begin{bmatrix} 0 & 0 & \dots & 0 & 0 \\ 1 & 0 & \dots & 0 & 0 \\ 0 & 1 & & \vdots & \vdots \\ \vdots & & \ddots & 0 & 0 \\ 0 & \dots & 0 & 1 & 0 \end{bmatrix} \begin{bmatrix} x_{1,k} \\ \vdots \\ x_{14,k} \end{bmatrix} + \begin{bmatrix} 1 \\ 0 \\ \vdots \\ 0 \end{bmatrix} u_k \tag{5.15}$$

$$\tag{5.16}$$

$$y_k = \begin{bmatrix} c_1 & c_2 & \dots & c_{13} & c_{14} \end{bmatrix} x_k \tag{5.17}$$

This state space formulation causes the dynamics to work in the following way: the input is applied to $x_{1,k}$ at each step. For each time step, the value of the states are shifted downwards on the form $x_{i+1,k} = x_{i,k+1}$. This means that $x_{2,k} = x_{1,k+1}$, and $x_{14,k} = x_{1,k+13}$. In other words, $x_{14,k}$ is the temperature level from the plant 13 samples ago.

One challenge with the MPC-controller in the DHS is the varying delay time from the plant to the point load. In this thesis, this is solved by using a piecewise affine system, a linear time invariant system with different constant C-matrices for defined operating areas, specified by the current value of $\dot{m}$. The different C-matrices and their corresponding operating areas are given in Table 5.2. The first column denotes which one of the four matrices is currently considered. The second column is the mass flow rate value that corresponds to the exact delay of the matrix, while the third column is the interval of the mass flow rate for which the matrix is in use in the MPC simulations. The fourth column is the transport delay for the chosen C-matrix, and the fifth column is the value of the C-matrix. This column indicates which value in the C-matrix is nonzero. All unmentioned elements in the C-matrix have the value 0.

It should be noted that as $\dot{m}$ is a continuous variable, the delay from the plant to the point load is a continuous function. This means that the C-matrices will not provide an accurate description of the delays, but rather an approximation. This approximation is based on the middle value of the operating area, specified by $\dot{m}_k$. This means that when one moves further away from the center of the current operating area, the current approximated C-matrix provides less accuracy. This model error is further discussed in Chapter 7.

The constraints in the model are limits for applied temperatures and the applied rate of the temperatures changes. The applied temperature constraints are the same as in the slow model. The constraint on the supply temperature rate in the MPC are set to be $2°C$ per time step or $0.2[°C/min]$.

Table 5.2: Table showing the 4 different output matrices, the value of $\dot{m}$ they correspond to, and the mass flow rate interval for which they are used in the system. The table also shows the corresponding delay from input to output in the system with the respective matrix. The C-value column shows which value in the output matrix is non-zero. For each matrix, all elements in the matrix except the mentioned value are zero.

| C-Matrix | $\dot{m}$-value $[kg/s]$ | $\dot{m}$-operating area $[kg/s]$ | Transport delay | C-value |
|---|---|---|---|---|
| 1 | 1.667 | 2.481 - 1.371 | 5 samples | $c_5 = 1$ |
| 2 | 1.075 | 1.370 - 0.940 | 8 samples | $c_8 = 1$ |
| 3 | 0.805 | 0.939 - 0.719 | 11 samples | $c_{11} = 1$ |
| 4 | 0.632 | 0.718 - 0 | 14 samples | $c_{14} = 1$ |

### 5.4.1 Feedback in the Model Predictive Controller

A fundamental part in achieving good results when applying model predictive control, as well as model-based control in general, is to update the system using measurement feedback in order to compensate for the model error.

The application of measurement feedback in DHS-control is however a challenging task. This is partly due to a lack of frequent measurements, and partly because identifying the source of measurements deviating from their predicted values is difficult.

For the case of Klæbu Centre DHS, the following measurements are available for feedback.

- Delivered supply line temperature from the plant $[°C]$. Momentary value measured at the plant every 10 minutes
- Measured supply line temperature at the heat exchangers $[°C]$. Momentary value measured every hour
- Measured return line temperature at the heat exchangers $[°C]$. Momentary value measured every hour
- Water volume passed through the heat exchanger for the previous hour $[m^3/hr]$. Measured hourly
- Energy transferred over the heat exchanger for the previous hour $[kWh]$. Measured hourly

In the model predictive controller simulated in this thesis, the only feedback which has been applied is the actual demand for the different consumers. It is impossible to perfectly predict the daily demand, thus this measurement is vital for operating the DHS.

The chosen procedure for handling the daily demand variations is by evaluating the system's input and output energy, and attempting to achieve an overall zero energy difference in the system over the course of the simulation.

This is done by evaluating the hourly measurements of the true consumer demands, and comparing these to the delivered power the previous hour. The difference between them is the output energy of the system. The controller then changes the mass flow rate

of the system to balance the system's energy as fast as possible in the subsequent control steps. Following a measurement of the true demand, the controller adjusts the mass flow rate by the following factor:

$$\dot{m}_{k,new} = \dot{m}_{k,old} + \frac{\Delta E_{measured}}{c_p(T_{s,k} - T_r)} \tag{5.18}$$

However, still subject to the following constraints:

$$\dot{m}_{min} \leq \dot{m}_{k,new} \leq \dot{m}_{abs,max} \tag{5.19}$$

$$\Delta\dot{m}_{min} \leq \dot{m}_{k,new} - \dot{m}_{k-1,new} \leq \Delta\dot{m}_{max} \tag{5.20}$$

Where

- $\dot{m}_{k,new}$ = The new mass flow rate for the system to apply $[kg/s]$
- $\dot{m}_{k,old}$ = The setpoint of the mass flow rate delivered from the slow model prior to the measurement feedback $[kg/s]$
- $\Delta E_{measured}$ = The total energy difference between the predicted power demand and the measured power demand for the previous hour. $[Wh]$
- $T_{s,k}$ = The current applied supply temperature $[^\circ C]$
- $\dot{m}_{min}$ = The minimum value of the mass flow rate, given in Table 5.1 $[kg/s]$
- $\dot{m}_{abs,max}$ = The absolute maximum value of the mass flow rate, given in Table 5.1. $[kg/s]$
- $\Delta\dot{m}_{max}$ = The upper bound on the allowed in mass flow rate between two samples $[kg/s^2]$
- $\Delta\dot{m}_{min}$ = The lower bound on the allowed difference in mass flow rate between two samples $[kg/s^2]$

# Chapter 6

# Results

For the results in this thesis, the proposed control structure and hierarchy described in chapter 5 is applied to the system. The controller is then simulated, using data from March 31st as the real physical system to which the control is applied. March 31st was one of the coldest days in the period of which the data for the models was used, which reflects on its prediction curves.

This chapter is structured as follows. The first section covers the demand curves generated the Klæbu Centre network. Following this, the results from the slow model are presented, and finally the simulation of the MPC-controller is shown, using the setpoints from the slow model, and data from March 31st.

## 6.1 The Consumer Demand Models

In this section the five consumer demand models are presented. A unique demand curve is created for each of the five consumers, as they display major differences in their daily trends. Furthermore, the curve for the total demand, the estimated system point load, is presented.

In order to generate demand prediction functions that best fit the observed demand curves, regression using data from the measurement period of weekdays in March and April 2016, from March 29th to April 22nd is used. The resulting demand functions are thought to be representative for normal weekdays during spring time.

The measured demand curves for each individual consumer are evaluated, and a nonlinear function describing the trends as accurately as possible is chosen. The parameters for the function are then chosen using the function fitnlm in MATLAB. Fitnlm is a nonlinear regression function for model fitting that fits the parameters of a nonlinear function to a set of observed data, by using a least-squares approach.

For the nonlinear regression, the measured heat demand curves for each individual consumer are used as response variables, and the ambient temperatures along with the corresponding observation times are used as the predictor variables, as shown in table 6.1.

Table 6.1: Observation variables and predictor variables used for the nonlinear regression model

| **Observation variables** | **Predictor variables** | **Predictor variables** |
|---|---|---|
| Vector of observed demand for the given day $[kW]$ | Vector of the 24 observation times for each day $[hours]$ | Vector of the ambient temperatures for the given day $[°C]$ |
| $Y_1$ (24x1)<br>$Y_2$ (24x1)<br>$\vdots$<br>$Y_{20}$ (24x1) | $t$ (24x1)<br>$t$ (24x1)<br>$\vdots$<br>$t$ (24x1) | $T_1$ (24x1)<br>$T_2$ (24x1)<br>$\vdots$<br>$T_{20}$ (24x1) |

The resulting nonlinear regression functions are mathematical expressions of the predicted demand based on time of day and ambient temperature.

For the consumer models presented in the upcoming sections, the graphs show the resulting demand model both based on the mean temperatures for the measurement period, as well as the demand based on the temperatures from March 31st. This day was one of the coldest days in the measurement period, shown in Figure 6.1, and has been used as the simulation day for the MPC-controller.

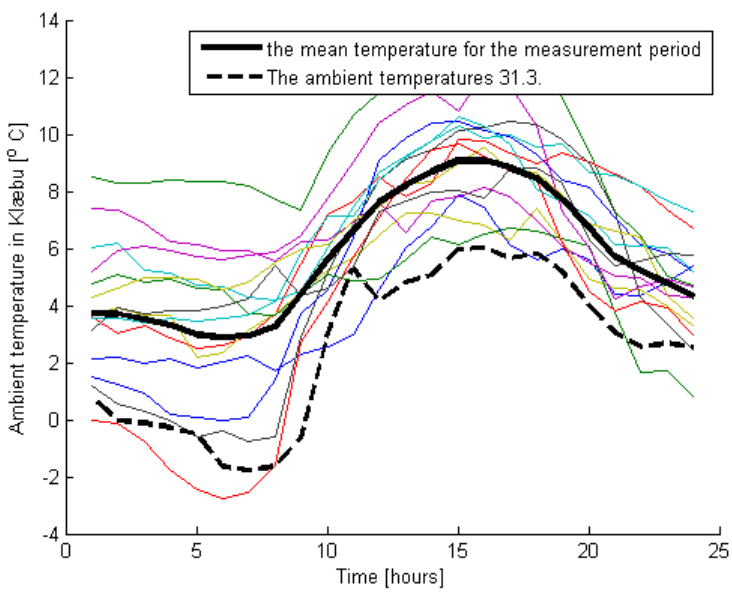


Figure 6.1: The mean temperature along with its measurement data for weekdays in the period 29.3 - 22.4

### 6.1.1 Consumer A

The model fit for consumer A is shown in figure 6.2. The daily variations are rather large in consumer A, but a trend towards an early peak is visible. The early peak is located around approximately 07:30 in the morning, and except for this, the demand curve for consumer A appears rather flat.

The nonlinear regression function which is used is as follows:

$$Y = \beta_1 + \beta_2 \cdot T(t) + \beta_3 \cdot e^{-1/\beta_4 (t-t_1)^2} \tag{6.1}$$

Where

- $Y$ = Heat demand $[kW]$
- $t$ = Hour of day, starting at midnight, 00:00 $[hour]$
- $T(t)$ = Ambient temperature for the given time of day $[°C]$
- $t_1 = 7.5$: The time of the observed morning demand peak at 07:30 $[hours]$
- $\beta_1$ = Parameter for the constant factor $[\cdot]$
- $\beta_2$ = Parameter for the temperature dependence $[\cdot]$
- $\beta_3$ = Parameter for the magnitude of the morning peak $[\cdot]$
- $\beta_4$ = Parameter for the steepness of the morning peak$[\cdot]$

The parameters $\beta_1$ to $\beta_4$ are fitted in for the demand curve to match the observed data as well as possible. The resulting values are given in table 6.2:

Table 6.2: Estimated parameters for the regression function of consumer A

| Parameter | value |
|---|---|
| $\beta_1$ | 21.0328 |
| $\beta_2$ | -0.0672 |
| $\beta_3$ | 5.0451 |
| $\beta_4$ | 1.2100 |

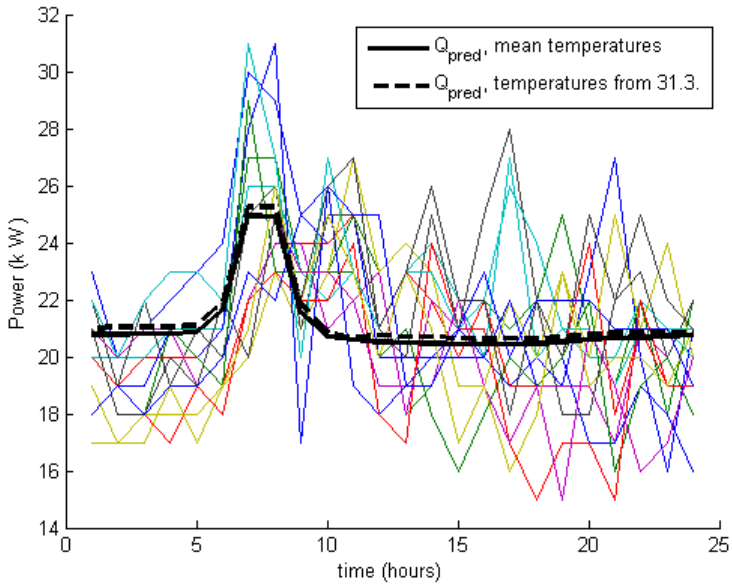


Figure 6.2: The model fit for consumer A along with its measurement data.

### 6.1.2 Consumer B

The model fit for consumer B is shown in figure 6.3. Similar to consumer A, consumer B exhibits a morning peak, then a very low demand during the midday, with the demand slowly increasing over the night. The demand peak for consumer B is located around approximately 09:00 in the morning.

The nonlinear regression function is the same as for consumer A, but with a different peak time. The fitted parameters are also different:

$$Y = \beta_1 + \beta_2 \cdot T(t) + \beta_3 \cdot e^{-1/\beta_4 (t-t_1)^2} \tag{6.2}$$

Where

- $Y$ = Heat demand $[kW]$
- $t$ = Hour of day, starting at midnight, 00:00 $[hour]$
- $T(t)$ = Ambient temperature for the given time of day $[°C]$
- $t_1 = 9$: The time of the observed morning demand peak at 09:00 $[hours]$
- $\beta_1$ = Parameter for the constant factor $[\cdot]$
- $\beta_2$ = Parameter for the temperature dependence $[\cdot]$

- $\beta_3$ = Parameter for the magnitude of the morning peak [·]
- $\beta_4$ = Parameter for the steepness of the morning peak[·]

The resulting estimated model parameters for consumer B are given in 6.3:

Table 6.3: Estimated parameters for the regression function of consumer B

| Parameter | value |
|---|---|
| $\beta_1$ | 4.5896 |
| $\beta_2$ | -0.3945 |
| $\beta_3$ | 3.3584 |
| $\beta_4$ | 1.1825 |

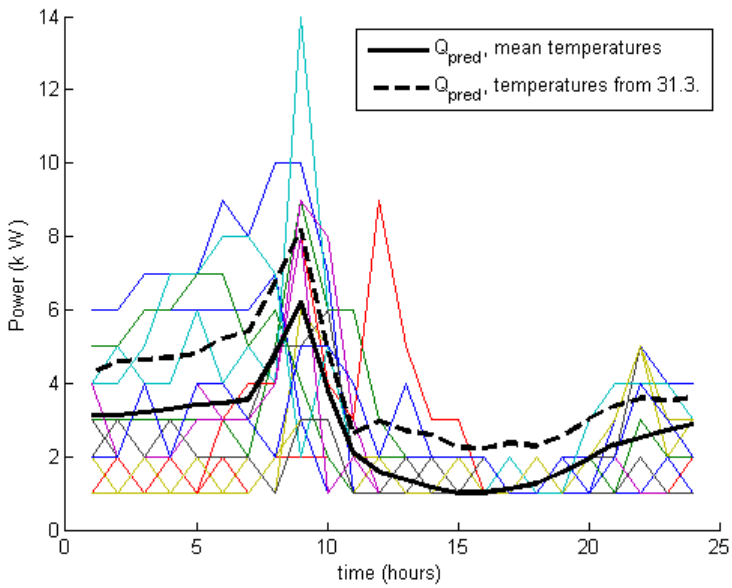


Figure 6.3: The model fit for consumer B along with its measurement data.

### 6.1.3 Consumer C

Because of the data forming a structure which is difficult to mathematically describe, as well as not displaying too much variance in demand from day to day, the function for customer C is generated by a set of points. More specifically, the function is defined by a specific power demand for each hour, that best match the observed data. The resulting function is not dependent on ambient temperatures. This is justified by the very low daily variations from the measurement data. The model fit for consumer C is given in figure

6.4. The black bold curve shows the function made to best fit the model, along with the measurement data it is decided from.

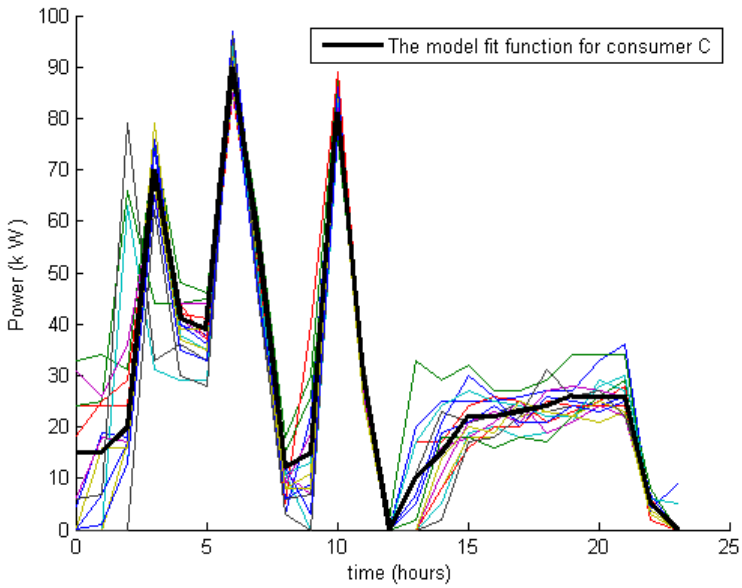


Figure 6.4: The model fit for consumer C along with its measurement data.

### 6.1.4 Consumer D

The model fit for consumer D is shown in figure 6.5. Similar to consumer A and B, the curve for consumer D is generated using the exponential function. Consumer D exhibits a clear peak in heat demand at 06:00 in the morning, as well as a drop in demand during the midday. From around 12:00 to 16:00, the consumer extracts very low amounts of heat from the DHS. The demand rises to some extent in the evening, with a small peak at 19:00.

The regression function used for consumer D is an exponential function with two peaks, structured as follows:

$$Y = \beta_1 + \beta_2 \cdot T(t) + \beta_3 \cdot e^{-1/\beta_5 (t-t_1)^2} + \beta_4 \cdot e^{-1/\beta_6 (t-t_2)^2} \tag{6.3}$$

Where

- $Y$ = Heat demand $[kW]$
- $t$ = Hour of day, starting at midnight (00:00) $[hour]$
- $T(t)$ = Ambient temperature for the given time of day $[°C]$
- $t_1 = 6$: The time of the morning demand peak at 06:00 $[hours]$

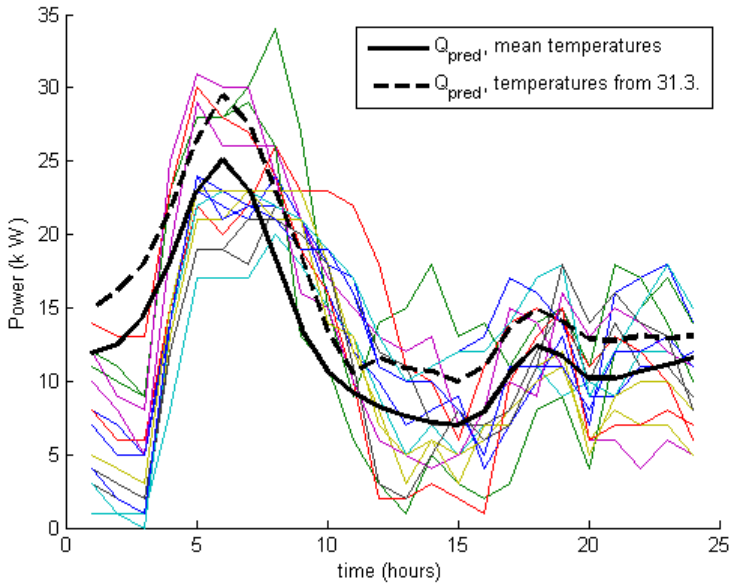


Figure 6.5: The model fit for consumer D along with its measurement data.

- $t_2 = 19$: The time of the evening demand peak at 19:00 [$hours$]
- $\beta_1$ = Parameter for the constant factor [·]
- $\beta_2$ = Parameter for the temperature dependence [·]
- $\beta_3$ = Parameter for the magnitude of the morning peak [·]
- $\beta_4$ = Parameter for the magnitude of the evening peak[·]
- $\beta_5$ = Parameter for the steepness of the morning peak [·]
- $\beta_6$ = Parameter for the steepness of the evening peak [·]

The parameters $\beta_1$ to $\beta_6$ are set in a least-squares fashion to best fit the observed data. The resulting values are as given in 6.4:

### 6.1.5 Consumer E

The model fit for consumer E is shown in figure 6.6. Consumer E has a ramp-like shape, with the demand being high from around 06:00 in the morning until 17:00 in the afternoon. In order to achieve the ramp shape, the hyperbolic tangent function, or tanh function, has been used. The tanh function works to some extent similar to a switch, meaning using one tanh function to start the ramp at the desired time, and another tanh function to end the

Table 6.4: Estimated parameters for the regression function of consumer D

| Parameter | value |
|---|---|
| $\beta_1$ | 15.5094 |
| $\beta_2$ | -0.9636 |
| $\beta_3$ | 12.4413 |
| $\beta_4$ | 5.0155 |
| $\beta_5$ | 5.4008 |
| $\beta_6$ | 2.8504 |

ramp, gives a curve with close to the desired properties. The regression function which is used is formulated as follows:

$$Y = \beta_1 + \beta_2 \cdot T(t) + \beta_3 \cdot \tanh(\beta_4(t - t_1)) - \beta_3 \cdot \tanh(\beta_4(t - t_2)) \tag{6.4}$$

Where

- $Y$ = Heat demand $[kW]$
- $t$ = Hour of day, starting at midnight (00:00) $[hour]$
- $T(t)$ = Ambient temperature for the given time of day $[°C]$
- $t_1 = 6$: The starting time of the ramp-like duration of high demand at 06:00 $[hours]$
- $t_2 = 17$: The end time of the ramp-like duration of high demand at 17:00 $[hours]$
- $\beta_1$ = Parameter for the constant factor $[\cdot]$
- $\beta_2$ = Parameter for the temperature dependence $[\cdot]$
- $\beta_3$ = Parameter for the magnitude of the ramp $[\cdot]$
- $\beta_4$ = Parameter for the steepness of the ramp start and end $[\cdot]$

The resulting values are given in table 6.5:

Table 6.5: Estimated parameters for the regression function of consumer E

| Parameter | value |
|---|---|
| $\beta_1$ | 6.8586 |
| $\beta_2$ | -0.5757 |
| $\beta_3$ | 9.5418 |
| $\beta_4$ | 0.3214 |

### 6.1.6 The Total Heating Demand

As all the individual consumer heat demand models are achieved, it is now possible to obtain a model of the total heating demand. However, as the consumers are all placed at different locations in the pipeline network, the times of their heating demands as seen from the plant are shifted differently for each consumer.

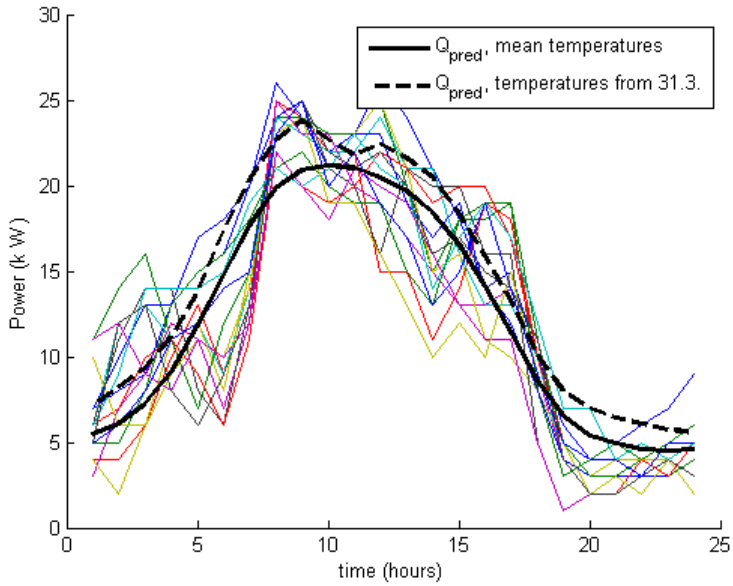


Figure 6.6: The model fit for consumer E along with its measurement data.

The point load approximation is therefore used. In [9] the total load estimated as a point load was placed in the network's volumetric center. However, for the total demand of the Klæbu network, the placement of the point load is chosen to be at the location of consumer C. The reasoning for this is that the dynamics of consumer C dominate the dynamics of the total heating demand, as this dynamic has a larger magnitude and its hourly changes are more drastic. By placing the total load at the location of consumer C, the time delays of consumer C become precisely calculated, at the cost of some loss in the accuracy of the delay approximation for the other consumers.

In order to generate the total demand as viewed from the location of the heat exchanger of consumer C, the curves for the individual demands from consumer A,B,D and E are shifted.

This shifting is equal to the difference in transport delay between the plant and the original location and between the plant and the location of consumer C.

In order to calculate these delays, the pipe segments in the network are enumerated, so different paths can be defined. The individual pipe segments are enumerated as in Figure 6.7.

The calculation of the demand curve shift is done as follows:

1. Calculate the average delay in each pipe segment, based on pipe segment volume and the average flow rate in the pipe segment

2. Calculate the average delay of each consumer with respect to the location of consumer C. The delay is calculated by subtracting pipe segments the water no longer

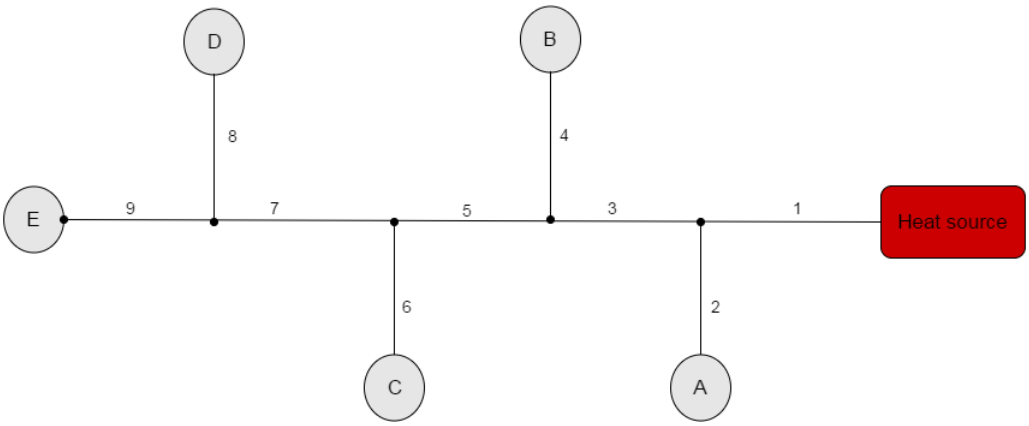


Figure 6.7: The pipe segments in the network

takes, and adding the pipe segments from the plant to consumer C that were not in the original consumer path

3. Shift the demand curve of each consumer forward or backwards according to the calculated delay. Consumer C remains unshifted.

The calculation of the average transport time in each pipe segment is done using the average flow rate through each segment, as well as the water volume in the segment. Table 6.6 shows the pipe segments and their transport times.

| pipe segment | volume $[L]$ | flows passing through | avg. flow rate $[L/s]$ | avg. transport time $[s]$ |
|---|---|---|---|---|
| 1 | 3389 | A,B,C,D,E | 0.5470 | 6196 |
| 2 | 79.9 | A | 0.1352 | 591 |
| 3 | 1673 | B,C,D,E | 0.4119 | 4062 |
| 4 | 69.3 | B | 0.0321 | 2159 |
| 5 | 18.8 | C,D,E | 0.3798 | 49 |
| 6 | 253.8 | C | 0.2280 | 1113 |
| 7 | 238.6 | D,E | 0.1518 | 1572 |
| 8 | 11.6 | D | 0.0853 | 136 |
| 9 | 395.1 | E | 0.0665 | 5941 |

Table 6.6: Table showing the average transport time of the water for each pipe segment in the network. The first column shows the pipe segment. The second column indicates the supply pipeline water volume. The third column indicates which flows pass through the respective pipe segment. The average flow rates are given in the fourth column, and the average transport times are given in the fifth column

| Consumer | Old path | New path | Difference in pipe segments | Total delay shift $[s]$ | samples shifted, rounded to nearest 10 min sample |
|---|---|---|---|---|---|
| A | 1,2 | 1,3,5,6 | -2 + 3,5,6 | 4633 | 8 samples |
| B | 1,3,4 | 1,3,5,6 | -4 + 5,6 | -997 | -2 samples |
| C | 1,3,5,6 | 1,3,5,6 | unchanged | 0 | 0 samples |
| D | 1,3,5,7,8 | 1,3,5,6 | -7,8 + 6 | -595 | -1 sample |
| E | 1,3,5,7,9 | 1,3,5,6 | -7,9 + 6 | -4964 | -8 samples |

Table 6.7: Table showing the shift times of the five consumers. The first column is the current consumer. The second and third columns are the old and new paths of the consumer, given in pipe segments. The fourth column shows the difference between the two. The fifth column is the shifted delay in seconds, and the sixth column is the shift measured in 10 minute samples

Figure 6.8 shows the five consumers' shifted predicted demand curves. The colored curves are the five individual demand curves using temperatures from March 31st. The dashed line shows the total demand for March 31st, generated by adding the shifted consumer models. The black line shows the total shifted demand when the mean temperatures for the measurement duration is used.

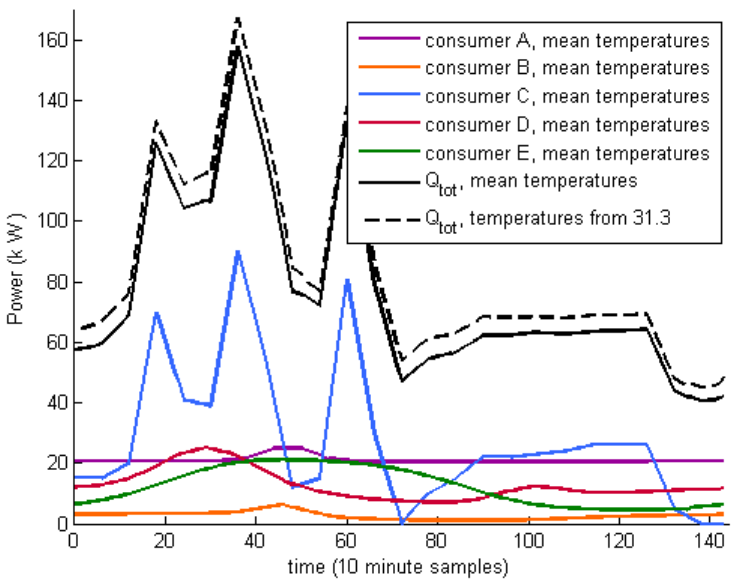


Figure 6.8: The total demand, before shifting the consumers based on their average transport delay

## 6.2 The Results from the Slow Model

After the completion of a total heat demand model, the set points for the supply temperature and mass flow rate can be evaluated. For the implementation done in this thesis, the slow model generates setpoints for the full 24 hour simulation. These are based on the total demand curve from March 31st, the dashed line in Figure 6.8.

However, upon generating the setpoints, it was discovered that for the given demand model and temperatures, the setpoints generated from the slow model provided the MPC with a rather uninteresting control problem. For the given total demand model and the applied ambient temperatures, the slow model was almost solely able to meet the predicted demand by applying an increased $\dot{m}$. This meant that the slow model would evaluate the supply temperature setpoints to a nearly constant value, resulting in the MPC-controller problem being close to keeping the temperature at a steady state value.

In order to achieve a more interesting control problem for the MPC-controller, the demand used in the simulation of the model predictive controller is scaled by a factor of 1.2, or a 20% increase in the predicted demand for the duration of the simulation. Figure 6.9 shows the original and rescaled demands, as well and their respective setpoints.

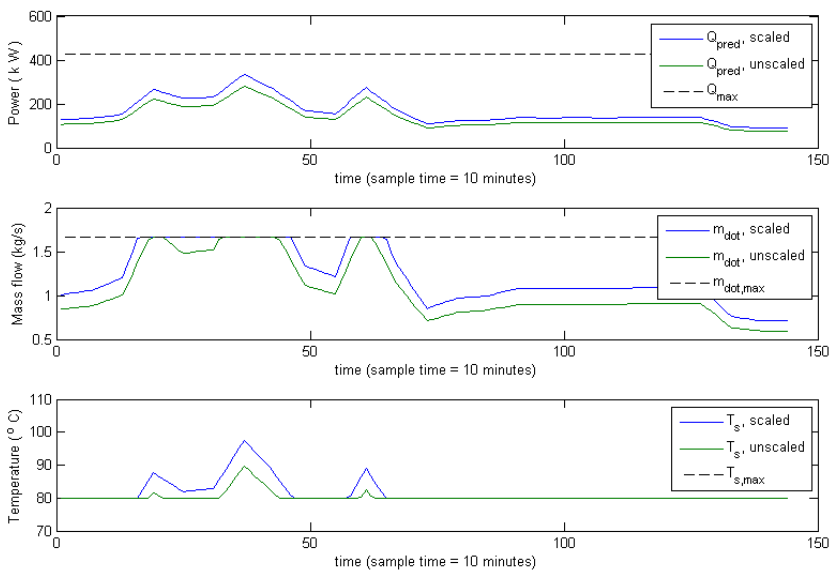


Figure 6.9: The evaluated setpoints for $\dot{m}$ and $T_s$ for both the unscaled demand from 31.3., and the demand from 31.3. scaled up by 20% of its original value

Figure 6.10 shows a zoomed in view at the temperature peaks for the scaled and unscaled case. In the unscaled case, the predicted demand, $Q_{pred}$, can be met completely by increasing the mass flow rate except for the time steps $t = 18 - 21$, $t = 32 - 44$, and

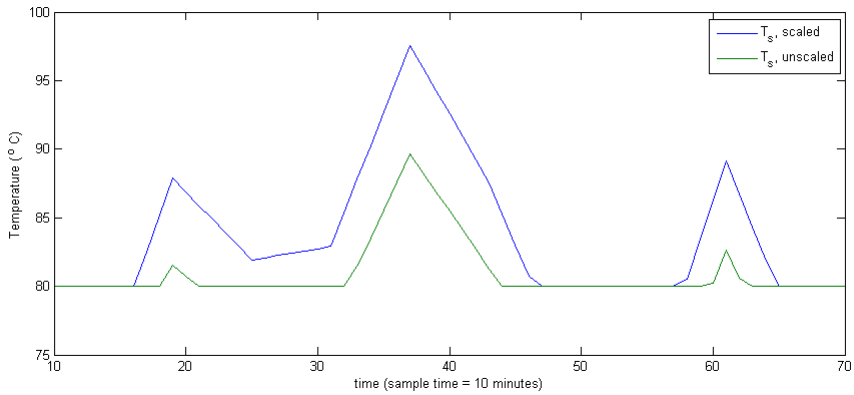


Figure 6.10: A zoomed look at the peaks in the supply temperature setpoints for the scaled and unscaled cases

$t = 59 - 63$, where the mass flow rate reaches saturation, and the temperature has to be increased. For these three demand peaks, the temperature has to be increased by $1.6°C$, $9.7°C$, and $2.7°C$, respectively, in order for the plant to meet the demand.
For the up-scaled case, $\dot{m}$ reaches saturation from $t = 17 - 46$ and $t = 58 - 64$. During this period, the demand peaks force a temperature increase of $7.8°C$, $17.6°C$, and $9.2°C$, respectively for the three peaks, with the desired supply temperature rising from $81.9°C$ to $82.9°C$ between the first and the second peak.

## 6.3 The Results from the Model Predictive Controller

In this section, the model predictive controller is implemented using the formulation from Section 5.4, and with the setpoints for $\dot{m}$ and $T_s$ from the scaled slow model. The difference between the predicted and true demand has been fed back to the controller, and mass flow rate has been applied accordingly to compensate for unmodelled errors whenever necessary.

Figure 6.11 shows the predicted and true demands for March 31st. The true demand is affected by stochastic occurrences during the day, and deviates both upwards and downwards sporadically.

During the day, the curves switch several times between which showing the largest demand. The true demand is slightly above the predicted from 00:00 to 02:00, while the predicted demand is higher than the true demand from 04:00 to 06:00. From 08:00 to 12:00 the largest difference during the day occurs, when the true demand is much higher than the predicted demand. The true demand is also higher than the predicted demand around 16:00 before it drifts down in the evening.

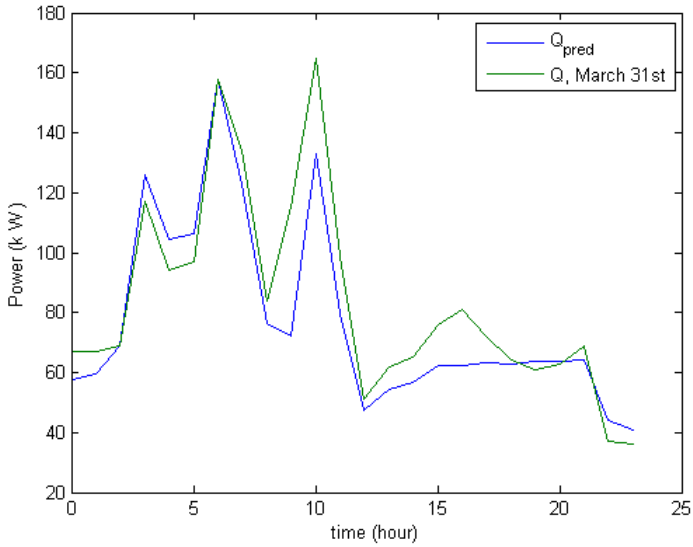


Figure 6.11: The true and predicted demand for March 31st

Figure 6.12 shows how the MPC-controller attempts to meet the supply temperature reference values given by the scaled slow model. The blue steps are the reference values for the point load at each time step. The red steps are the applied supplied temperatures from the MPC-controller. The green steps are the MPC-controller's modelled temperatures at the point load. The output follows the reference trajectory close, but with slight deviations. As the system is an open loop simulation, these deviations are caused by the difference in the evaluated delay and the current applied delay decided by the current choice of output matrix, C.

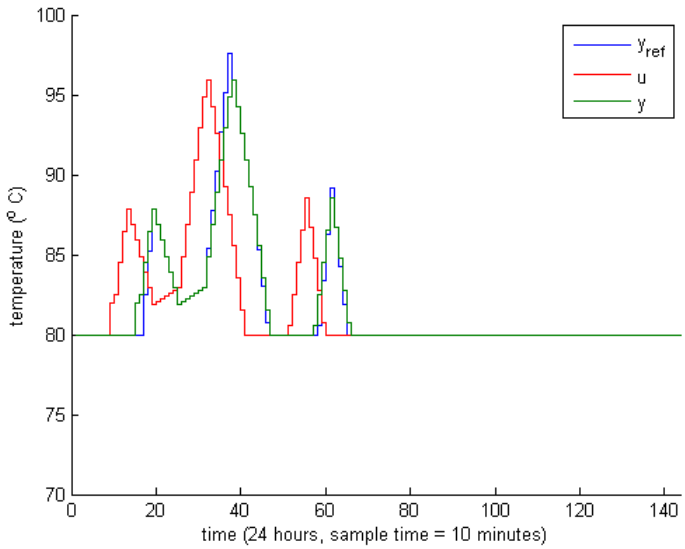


Figure 6.12: The simulated input and output values of the MPC-controller, along with the reference value

Figure 6.13 shows the evaluated differential energy in the network, and Figure 6.14 shows the applied $\dot{m}$ along with its original trajectory given by the slow model. The energy differential is a function with spikes for every hour of simulation. Each of these spikes corresponds to a measured difference between the predicted and measured energies $[kWh]$. The energy differential shows mostly positive spikes for the first 48 samples, denoting the first 8 hours of the simulation. During this part of the simulation duration, the true demand is mostly lower than the predicted demand. At the 60th sample, the controller receives a large negative value, corresponding to peak in the true demand at 10:00 in Figure 6.11. Following this drop, the energy function oscillates between a positive and negative value for a series of time steps. These oscillations are due to the mass flow rate meeting the saturation for its delta constraint, caused by the demand curves having very large changes in references in a short amount of time during this part of the simulation. This means that the changes in mass flow rate are both affected by changes in the demand curve and measured energy differences.

In Figure 6.14 one can observe that the applied $\dot{m}$ is increased and decreased with spikes wherever the energy function has spikes, but in the opposite direction. These are the compensatory actions, attempting to force the system back to the original energy level. At around $t = 60$, a large spike can be seen, which exceeds the nominal maximum for the mass flow rate. However, the absolute max flow rate constraint is not activated, meaning that the consumer demand is still satisfied. This would not be the case if the absolute maximal constraint for the mass flow rate was activated for a long time, causing a gradual decline in the system energy, and the temperature on the lines.

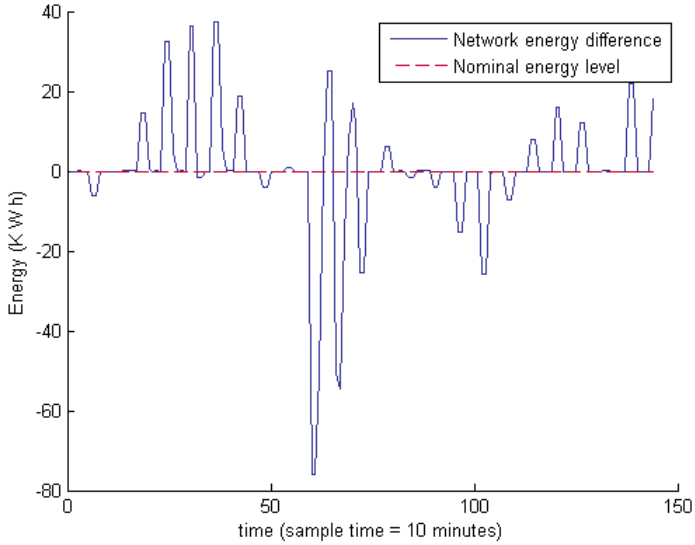


Figure 6.13: The measured energy-differential, caused by differences between the predicted and true demand

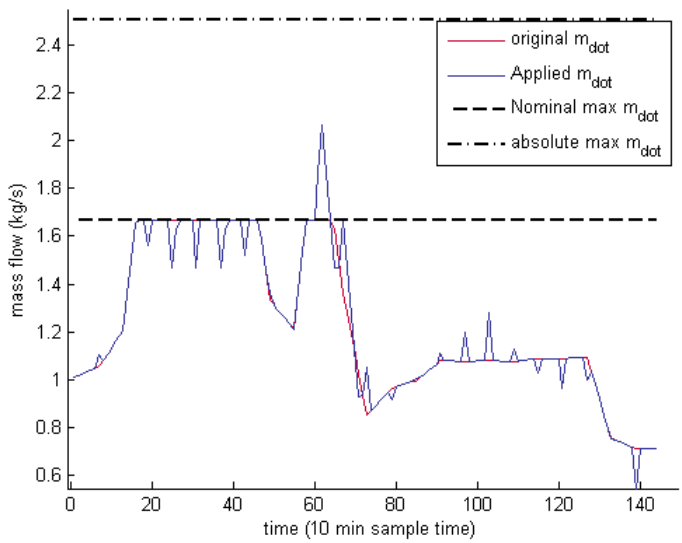


Figure 6.14: The estimated mass flow rate along with the applied mass flow rate of the controller, after the energy feedback

## 6.4 The Economic Profit of applying MPC to the Klæbu DHS

A formulation for the potential profit by reducing supply temperatures has not been made in this thesis. The potential economic gain is based on several factors such as pump efficiency at the plant, heat exchanger efficiency in the network, and the reduced heat loss due to reduced temperatures. Despite a lack of this formulation, the implementation done in this thesis is done with the goal of reducing supply temperatures, which it accomplishes. Figure 6.15 shows the logged supply temperatures at the plant from the period of simulation, from March 29th to April 22nd. The figure shows that the temperatures were kept at a steady high, either with values hovering around approximately $102°C$ or approximately $92°C$. Two days during this period exhibit large deviations from these values, possible caused by a fault or a boiler fallout at the plant. The reasoning for the two different temperature levels delivered from the plant is not known, and an explanation for it is not attempted in this thesis. However, for the duration of the simulation period, the mean delivered supply temperature was found to be $100.4°C$ in the true system, whereas the average supply temperature value for the MPC simulation on the day March 31st, is $81.7°C$.

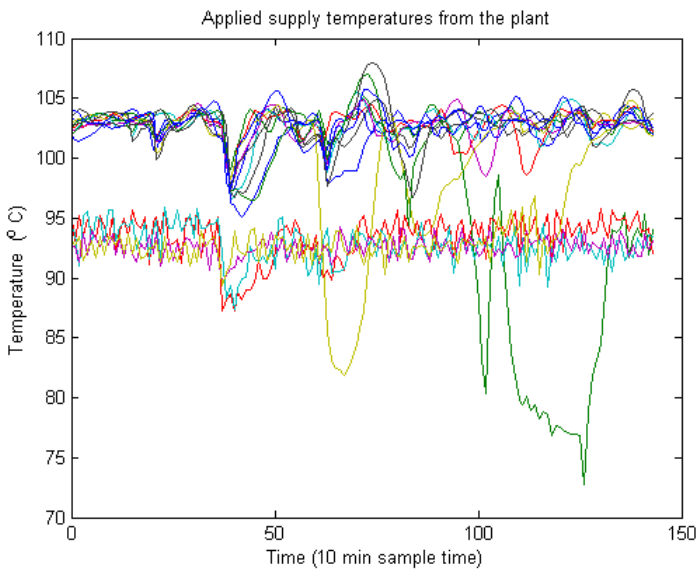


Figure 6.15: Measured temperatures on the supply line delivered from the plant for the simulation period March 29th to April 22nd

Chapter 7

# Discussion

## 7.1 On the Total Load from the Five Consumers

The method which is used to estimate the total load on the network, performed in Section 6.1, is based on an approximation that causes it to be somewhat inaccurate.

The method for shifting the models, described in Section 6.1.6, shifts the curve for each model according to the difference in the transport delay of the water between the plant and the original location, and between the plant and the new location. These transport delays are here assumed to be constant, and their values are approximated as a mean flow rate for each pipe segment. Furthermore, the mean flows rate are found by taking the average flow rate from the heat exchangers, and evaluating which pipelines each heat exchanger's flow passes through. To illustrate this: if a pipeline is located near the far end of the network, and the only water passing through it is the water going to consumer D and E, then the average flow rate for this pipe segment is assumed equal to the sum of the average flow rates for consumer D and consumer E.
These assumptions cause the shifting for every consumer in the total demand model to be slightly inaccurate, with the exception of consumer C.

There are two ways this error could potentially be reduced. The most apparent way would be to reduce the deviations in flow rate. The varying flow rate is not only a cause of inaccuracy in the creation of the total load, but also on other parts of the modelling and control.

A second way of reducing these inaccuracies would be to estimate the delay for every single optimization in the RTO. However, this would cause a major increase in computational complexity, as a mutual dependence between the total load and the setpoints for the mass flow rate would most likely require an iterative procedure in the RTO optimization. The RTO would evaluate an initial estimate of the total load and generate an initial sequence of flow rates for the future samples. These would then be used to update the total model with new delay shifts, before a new sequence of flow rates could be generated.

## 7.2 The Modelling of Return Temperatures

For the implementation done in this thesis, no model for the temperature on the return line is included. Despite the return line temperature being measured at each heat exchanger, the dynamics were found difficult to model. A reason for this is that the return line receives a number of incoming flows, with variable flow rates and temperatures, at different points on the line, causing its modelling to be rather complex. The measurements at the heat exchangers are also only measured every hour, meaning that the return line has to be estimated for the samples where no measurement is available.
A possible simple method of estimating the return temperature is to assume it uniform and equal to its previously measured value at the plant. This was stated in [23] to be the estimation giving the best results.

### 7.2.1 Loading the System

Due to the lack of a model for the return line temperature, the implementation done in this thesis does not allow for any loading of the system to be made. In order for a controller to be able to incorporate loading possibilities, it would require modelling of either the overall energy level or the overall temperature level on the network.

Although a network energy difference function is implemented and used for feedback, this only covers the difference between the modelled and measured consumer load, without saying much about the overall level in the network.

A possible way of implementing loading possibilities could be by adding a relation between the temperatures of the recently arrived return line flow at the plant, and the maximal supply temperature applied to the system. This would reflect reality, as hotter return water will require lesser amounts of reheating.

## 7.3 On Feedback in the Model Predictive Controller

The usage of feedback in the model predictive controller for the implementation done this thesis is rather lacking. The only data which is used in the control loop, is the measured load for each separate consumer. This measurement is compared to the predicted demand, and the difference is used in order to evaluate the network energy difference. There are however more measurements in the network that are available for use in the controller. The following items are the available measurements in the network, mentioned previously in Section 5.4.1, as well as proposed methods of using them:

- Delivered supply line temperature from the plant $[^{\circ}C]$. Momentary value measured at the plant every 10 minutes. Can be fed back to a plant module in order to get a more precise relationship between the plant burner module and the temperature of the outgoing water
- Measured supply line temperature at the heat exchangers $[^{\circ}C]$. Momentary value measured every hour. For consumer C, this can be compared to the modelled temperature at the load for the given time. However, this measurement is a momentary

value, and may be subject to fluctuations in the water temperature. The measurement may be a bad indicative of the temperature passing through, without there currently being any way of identifying this

- Measured return line temperature at the heat exchangers $[°C]$. Momentary value measured every hour. Subject to the same uncertainties as the supply line temperature measurement
- Water volume passed through the heat exchanger for the previous hour $[m^3/hr]$. Measured hourly. Can be used as feedback to the plant's model of pump power to validate the relationship between power and delivered mass flow rate. May also be used to validate flow distribution models in an aggregated network
- Energy transferred over the heat exchanger for the previous hour $[kWh]$. Measured hourly. These values have to be fed back in order to assure that the requested effect level for the consumers is being met by the DHS operator. These should also be used to calibrate the demand models

**Additional Measurements in the Network**

In order for feedback to be integrated into the controller, an increased measurement frequency in the network would be advisable. The most important source for feedback would be the supply temperature, and as mentioned above, the current temperature measurement at the consumers is a momentary measurement every hour.

If measurements for the supply temperature for the consumers, especially consumer C, were to be made more frequent, it would potentially be possible to eliminate bad measurements. Bad temperature measurements could be measurements based on the fluctuations in the water temperature, not indicative of the actual temperature level at the time. If measurements were made frequent enough, these values could potentially be filtered away.

Another difficulty with temperature measurements is to establish if an unexpected high or low value is a result of an actual temperature increase or decrease, or that the estimated transport delay is incorrect. If the latter is the case, the temperature could indicate wrong due to peak temperature timings coming at unexpected times. The measured difference between the modelled temperatures and the measured temperatures will then be a result of the model comparing for two different parts of the peak curve.
If the measurements of temperature are frequent enough, it could potentially be possible to see when the actual peak hits, thus being able to identify the source of any measured errors.

### 7.3.1 On the Accuracy of the Load Models

The creation of load models is work that requires time and effort. Typically, the DHS operators have to evaluate a trade-off between the level of precision that is desired for the system, and the time it takes to improve the models. Small improvements in the models may be creating separate models for consumers that differ from the generalized curves,

adding additional factors such as daily weather to the models, or improving the mathematical functions to which the curves are approximated.

The assumption of the ambient temperature having a linear relationship to the heat demand may also be inaccurate, especially when the outside temperatures are very high or very low. For very hot days, the heating is most likely kept minimal, though the temperature may differ slightly. Similarly, for extremely cold days, the heating is kept on a close to constant very high level.
In [8], the temperature dependent part is modelled as a piecewise linear function, causes the temperature dependency to be weaker for very high and very low temperatures. This could potentially be a factor that would increase the modelling accuracy, however for the simulation period during March and April, it would most likely not have had a large impact.

**Consumer C and its Demand Prediction Curve**

As mentioned in the results in Chapter 6, the observed demand curve for consumer C is hard to mathematically explain, and shows extremely low daily variance. The regulation of heat extraction is done on the consumer side, and therefore not a part of this thesis. However, for this particular case it could be of interest for the DHS operator to research the reasoning for the unnatural shape in demand. The reason for the unusual heat extraction could potentially be as simple as an outdated thermostat with faulty logic, or the consumer having some economic profit in extracting in this particular fashion.

The exclusion of temperature dependence in the chosen prediction demand function for consumer C is justified by the low daily variance. However, the curves do show an increased variance from midday and during the night, from 13:00 to approximately 02:00. Here, one would most likely benefit from separating the prediction demand function into two functions, where one would be the static function chosen in this thesis, and the other would be a temperature dependent curve for the time period mentioned above.

## 7.4 On how to deal with the Varying Time Delay in the MPC-Controller

The decision to assume a set of linear time invariant systems for equality constraints of the model predictive controller is as mentioned earlier a source of error. The constant C-matrices are computed for a specific $\dot{m}$ and its corresponding delay, and when the current $\dot{m}$-value deviates from the value that is chosen in the C-matrix approximation, the modelled time delay from input to output becomes inaccurate.

An alternative method to the piecewise linear system would be a non-linear system in which the time delay is calculated for each step in the horizon at each sample in the model predictive controller. This would naturally have caused a significant increase in the complexity of the controller, which is why it was not done.

However, for a large part of the simulation the mass flow rate is equal to its nominal maximum value. whenever the temperature setpoint is increased in the slow model, the mass flow rate is in saturation, and whenever the mass flow rate is not in saturation, the

desired temperature at the point load is static. As one of the chosen C-matrices is chosen for the exact value of the mass flow rate saturation, the delay will be correct whenever the mass flow rate equals its nominal maximum. Unfortunately the daily heat demand variations, and the subsequent compensations in mass flow rate will require the mass flow rate to adjust the energy level, thus moving away from its nominal maximum.

A different alternative could be to make changes to the slow model in order for the system to operate with a static flow rate, before any compensations for the energy feedback. The mass flow rate in the system could be set to the nominal maximum, which would mean that the temperature would have to be decreased in places where the mass flow rate was not previously in saturation. The benefits of this would be that it would make an overall reduction of the temperatures in the network, thus decreasing heat losses. It would also simplify the MPC-controller, and reduce the delay errors as the transport delay more stable. However, the mass flow rate would have to be used for compensating for the daily load variations, meaning some transport delay modelling error would still occur.

The downside of changing the mass flow rate to a static value would be that the heat exchanger efficiency in the system might be reduced.

Chapter 8

# Conclusion

The aim of this thesis was to explore the benefits and challenges of applying model predictive control to district heating systems, and the possibilities of applying this technique to DH systems in Norway. This was approached by simulating MPC on a small-scale DHS in Klæbu, as well as observing the current operational regime of the large-scale DHS in Trondheim, and identifying challenges involved with applying predictive control to this system.

For the MPC simulated on Klæbu DHS, only the western-going distribution line was looked at. This distribution line only has five customers, which allowed for a very specific consumer model. The creation of accurate models for the future heat demands of the network customers is essential for effective control. The demands for the five consumers were all very different from the others, and because of this, an individual consumer demand model was made for each of the five consumers. These consumer demand models were based on social behavior and outdoor temperatures, the factors mentioned in theory as the most influential. Through usage of measured data of each consumer's actual demand, the generated models were found to be very accurate for describing the daily demand, although some variance was unavoidable.

Model aggregation is an important tool in district heating control. It describes simplifications done to the network in an attempt to describe complex structures by simple models with equivalent properties. The models for the predicted heat demand were all generated as models covering 24 hours, representing a weekday. For the model aggregation, these curves were shifted forwards or backwards in order to represent them all at a common location. This was done in order for them to be added for creating a model of the total demand in the network. The varying transport delays in the network were discussed to be a cause of errors in the network aggregation. The time shift of the curves was performed using a mean of the flow rate, which causes the aggregated demand to be inaccurate when the flow rate deviates from this value. For the implemented system, this error was reduced to some extent by placing the location of the aggregated model to the consumer with the largest impact on the total demand curve.

One of the largest challenges when applying model predictive control to district heating systems is to precisely model the varying transport delays of the water in the pipelines. For the implemented controller, this was attempted by using a piecewise linear system, with different regions corresponding to different levels of flow rate. This resulted in a better accuracy of the modelled delay, but did not completely solve the problem. The addition of measurement feedback to the MPC was found to be very challenging for the given control system. The only feedback that was applied to the controller was the measured difference between the predicted load and the actual load of each consumer. Although this was accounted for in order to balance the input and output energy of the network, it could not be used to adjust the system model, as this difference is mainly caused by daily variations in consumer heat consumption.

Chapter 9

# Future Work

The implemented controller in this thesis is only applied to a very simple small-scale DHS. The long-term goal of this work and future continuations of it is to extend this to a control framework for model predictive control that can be applied to DH systems of different size, structure and complexity.

Even though the control scheme presented in this thesis covers a large part of the control problem, from creation of consumer models to the evaluation of supply temperatures, there still is much to improve. The performance of the controller is not yet tested, and focus should therefore lie with further improvement of the scheme before ascending in operational complexity. The main focus points for improvement should be better integration of feedback, as well as a revision of the slow model optimization.

In Section 7.4 in the discussion, an alternative evaluation approach of how to meet the predicted demand was mentioned. In this alternative method, the mass flow rate is set to the nominal maximum for the duration of the simulation, which ultimately decreases the supply temperatures, but forces it to be constantly changing.

This thesis has not made an attempt to evaluate which of the two methods is the most economically beneficial. This is dependent on multiple factors such as heat exchanger efficiency, heat loss in the pipelines, and knowledge of the heat generation in the plant. In continuations of this work, a clear definition of which of these two methods is desirable should be achieved.

Another point that should be heavily focused on in future work on the subject, is the addition of feedback to the controller. As mentioned in Section 7.3 in the discussion chapter, adding more frequent measurements at the consumers could improve the reliability of the measurements, as well as making it easier to identify the source of modelling errors.

Due to the simple structure in Klæbu and the availability of measurements at each consumer, it could be beneficial to use the DHS in Klæbu as a basis for calibration and

testing of a model predictive control scheme. Calibration of these models is work that there is little experience on in Norway, and this could uncover new challenges or properties of the overall control problem. Furthermore, calibration of the modelled heat dynamics in the system would increase the transparency, and could identify leaks or faulty measurements, even without actually implementing control.

# Bibliography


[1] H. Madsen and K. Sejling. On flow and supply temperature control in district heating systems. *Heat Recovery Systems and CHP*, 14:613–620, 1994.

[2] A. Benonysson, B. Bøhm, and H. Ravn. Operational optimization in a district heating system. *Energy Conversion and Management*, 36:297–314, 1995.

[3] H. Gadd and S. Werner. Achieving low return temperatures from district heating substations. *Applied Energy*, 136:59–67, 2014.

[4] H. Gadd and S. Werner. Daily heat load variations in swedish district heating systems. *Applied Energy*, 106:47–55, 2013.

[5] G. Sandou, S. Font, and S. Tebbani et al. Predictive control of a complex district heating network. *Proceedings of the 44th IEEE Conference on Decision and Control*, pages 7372–7377, 2005.

[6] E. Ikonen, I. Selek, and J. Kovacs et al. Short term optimization of district heating network supply temperatures. *Energycon 2014 - IEEE International Energy Conference*, pages 996–1003, 2014.

[7] S. Grosswindhager, A. Voigt, and M. Kozik et al. Predictive control of district heating network using fuzzy dmc. *Proceedings of 2012 International Conference on Modelling, Identification and Control, Wuhan, China, June 24-26, 2012*, 2012.

[8] E. Dotzauer. Simple model for prediction of loads in district-heating systems. *Applied Energy*, 73:277–284, 2002.

[9] A. Heller. Heat-load modelling for large systems. *Applied Energy*, 72:371–387, 2002.

[10] H. Larsen, B. Bøhm, and M. Wigbels. A comparison of aggregated models for simulation and operational optimisation of district heating networks. *Energy Conversion and Management*, 45:1119–1139, 2004.

[11] I. Gabrielaitiene, B. Bøhm, and B. Sunden. Modelling temperature dynamics of a district heating system in naestved, denmark - a case study. *Energy Conservation and Management*, 48:78–86, 2006.

[12] I. Hassine and U. Eicker. Simulation and optimization on the district heating network in scharnhauser park. *2nd European Conference on Polygeneration*, 49:1–18, 2011.

[13] P. Pinson, T. Nielsen, and H. Nielsen et al. Temperature prediction at critical points in district heating systems. *Applied Energy*, 194:163–176, 2009.

[14] Yunus A. Cengel and John M. Cimbala. *Fluid Mechanics - Fundamentals and Applications*. McGraw-Hill Companies, 2010.

[15] Rolf Utseth. *Vannbåren varme - systemer og egenskaper*. Institutt for energi- og prosessteknikk, NTNU, 2012.

[16] Fabio Paterno. *Model-based design and evaluation of interactive applications*. Springer Science & Business Media, 2012.

[17] S. Qin and T. Badgwell. A survey of industrial model predictive control technology. *Control Engineering practice*, 11:733–764, 2003.

[18] M. Carrión and J. Arroyo. A computationally efficient mixed-integer linear formulation for the thermal unit commitment problem. *IEEE Transactions on Power Systems*, 21:1371–1378, 2006.

[19] D. Mayne, J. Rawlings, C. Rao, and P. Scokaert. Constrained model predictive control: Stability and optimality. *Automatica*, 36:789–814, 2000.

[20] W. Langson, I. Chryssochoos, S. Raković, and D. Mayne. Robust model predictive control using tubes. *Automatica*, 40:125–133, 2004.

[21] Statkraft. About Statkraft. `http://www.statkraft.no/om-statkraft/`, 2016. [Online; accessed 22-June-2016].

[22] Statkraft. Statkraft's operation in Klæbu. `http://www.statkraft.no/Energikilder/vaare-kraftverk/norge/Klabu/`, 2016. [Online; accessed 22-June-2016].

[23] Linn Saarinen. Modelling and control of a district heating system. Master's thesis, Uppsala Universitet, Sweden, 2008.