PROCESSTYRNING 1262

Beslutsstöd för produktionsplanering av fjärrvärme med hjälp av olinjär programmering

Per-Ola Larsson, Stéphane Velut, Håkan Runvik, Sara Modarres Razavi, Andreas Nilsson, Markus Bohlin, Jonas Funkquist

värmeforsk
VÄRMETEKNISK FORSKNING OCH UTVECKLING

Beslutsstöd för produktionsplanering av fjärrvärme med hjälp av olinjär programmering

Decision Support for Short-Term Production Planning of District Heating using Non-linear Programming

Per-Ola Larsson, Modelon AB
Håkan Runvik, Modelon AB
Stéphane Velut, Modelon AB
Sara Modarres Razavi, SICS
Andreas Nilsson, SICS
Markus Bohlin, SICS
Jonas Funkquist, Vattenfall AB.

P 38155

Värmeforsk Serviceaktiebolag
101 53 STOCKHOLM · Tel 08-677 25 80
December 2014
ISSN 1653-1248

Abstract

The short term production planning optimization problem for a district heating system is solved in two steps by integrating physics-based models into the standard approach. The first optimization step solves for the discrete variables of the unit commitment problem (UCP) using mixed integer linear models and standard mixed-integer solvers. The second step, the economic dispatch problem (EDP), considers dynamic optimization using physics-based non-linear models that utilize the unit statuses from the first step. All optimizations aim at maximizing production profit using fuel, electricity and heat prices as well as maintenance and start-up/stop costs. Through the physics-based models in the EDP, it is feasible to optimize over and consider constraints on power flows as well as important physical variables such as supply temperature, supply flow rate, pump speeds and condenser pressures which is not available in today's standard methods.

The modeling has focused on distributed consumption and production. The goal has been to represent the most important units and network distribution of the Uppsala district heating network. The distribution yields that the total heat demand is distributed and the delay times from production to customers are customer individual. The district heating net has been modelled using physics-based pipes, including mass flow dependent delays and temperature dependent (water and outdoor temperature) heat losses. Comparisons between optimizations with and without distribution net models have been performed, showing that careful modeling of the net impacts the production planning in form of reduction of costly production peaks and delay of costly unit start-ups, production compensation for heat losses and time delays as well as usage of the net for heat storage (accumulation). The optimizations also results in production plans where supply temperature and flow is minimized and maximized, respectively, and there is a balance between heat production and heat consumption.

The physics-based modeling and dynamic optimization techniques provides a flexible way to formulate the optimization problem and include constraints of physically important variables such as supply temperature and flow, pressures, heat loads, and start-up/stop trajectories for production units.

Two examples of stochastic optimization have shown that taking probabilities of heat demand predictions into account, the expected profit may be increased.

Sammanfattning

Optimeringsproblemet för den kortsiktiga produktionsplaneringen i ett fjärrvärmenät löses i två steg genom integrering av fysikaliska modeller och standardmetoden. I det första optimeringssteget behandlas start/stop problemet (SSP), där de diskreta variablerna bestäms med hjälp av diskrettidsmodeller och standardlösare för linjärprogrammering. I det andra steget, det optimala lastfördelningsproblemet (LFP), används dynamisk optimering med fysikaliska olinjära modeller där statussignalerna från steg ett ingår. Vid alla optimeringar är målet att maximera den ekonomiska vinsten, genom att el-, värme- och bränslepriser tas i beaktande, liksom kostnader för underhåll, uppstart och nedstängning. Med hjälp av den fysikaliska modelleringen i LFP är det möjligt att optimera effektflöden och inkludera bivillkor för dessa, liksom för andra viktiga variabler såsom framledningstemperatur, massflöden, pumphastigheter och kondensortryck. Detta är inte möjligt med dagens standardmetoder.

Modelleringen har fokuserat på distribuerad konsumtion och produktion. Målet har varit att representera de viktigaste enheterna, samt distributionsnätet i Uppsala. Distributionen leder till att det totala värmebehovet är fördelat och fördröjningstiden från produktion till kunder är kundspecifik. Fjärrvärmenätet har modellerats med hjälp av fysikaliska rör, med massflödesberoende tidsfördröjning och temperaturberoende (vatten- och utomhustemperatur) värmeförluster. Jämförelser mellan optimeringar med och utan den distribuerade nätmodellen har genomförts och resultaten visar att noggrann modellering av nätet påverkar produktionsplaneringen genom att dyra produktionstoppar kan undvikas och dyra uppstarter av produktionsenheter kan fördröjas, produktionen kompenseras för värmeförluster och fördröjningstider och nätet kan användas för värmelagring (ackumulation). Optimeringarna resulterar också i att framledningstemperaturen och massflödet minimeras respektive maximeras och att det är balans mellan värmeproduktion och värmekonsumtion.

Den fysikaliska modelleringen och dynamiska optimeringstekniken ger en flexibel metod för att formulera optimeringsproblemet och inkludera bivillkor på fysikaliskt viktiga variabler såsom framledningstemperatur och -massflöde, tryck, värmelaster och uppstarts-/nedstängningsprofiler för produktionsenheter.

Två exempel med stokastisk programmering har visat att man kan uppnå en större förväntad vinst genom att beakta sannolikheten för olika värmebehovsprognoser.

Executive Summary

The standard approach for production planning of a district heating network today typically involves highly simplified models where most of the physical descriptions are removed. This makes it possible to solve the resulting optimization problem using simple linear optimization methods, but significant physical constraints and degrees of freedom are then excluded from the model.

In this study the production planning problem is divided into two separate optimization problems; the Unit Commitment Problem (UCP) and the Economic Dispatch Problem (EDP).

In the UCP the modeling of the system is simplified to only include piecewise linear dependencies between the variables. The system is furthermore discretized in time, creating a Mixed Integer Linear Programming (MILP) problem. This kind of optimization problem can be solved using a numerical solver.

The EDP model includes much more detailed representations of the producers, customers and distribution net in the district heating network, compared to the UCP model. The modeling is physical, which has benefits in terms of accuracy, interpretation, and possibility to impose constraints corresponding to the limitations of the physical system.

In order to use the EDP model for optimization purposes it is necessary to remove the integer variables from the formulation. The status signals from the UCP optimization results are therefore used as input signals in this optimization. When the EDP optimization formulation is discretized a Non-Linear programming (NLP) problem is created.

The district heating network of Uppsala is modeled in this study. The most important production unit in this network is the cogeneration plant denominated KVV (Kraft- och värmeverk), situated at the production site Boländerna. This plant is represented in detail in the EDP, where a large part of the vapor cycle is modeled. The main components in this model are high and low pressure turbine stages, condensers, and a reheater. For the UCP a polytope in the space of electricity, heat and return temperature is used to describe the KVV. Other production units are not modeled physically in either the UCP or the EDP, but are instead simply adding heat to the district heating water proportionally to their load.

An important unit in the Uppsala network is the accumulator. This is modeled using a finite volume approximation in the EDP, while it works as an integrator for the stored energy in the UCP.

The physical distribution of the customers and producers in the Uppsala district heating network is represented using a one dimensional approach based on the delay time for the customers relative to Boländerna. In order to represent the delay time a simple delay is used in the UCP, while a more complex pipe model is implemented in the EDP. This model consists of a fixed delay combined with a finite volume pipe implementation. In

both the UCP and the EDP a heat loss model based on the outdoor temperature is included in the pipe modeling.

The Pyomo modeling language is used for creating the UCP optimization formulation. Two different solvers are used to solve the problem; Gurobi and GLPK. The Gurobi Optimizer is a commercial solver which can handle several different types of optimization problems and modeling languages. GLPK is an open source software package.

The EDP optimization problem is formulated using the Optimica language. The open source platform JModelica.org is used to solve the problem. This tool translates the problem into an NLP which is solved by the Interior Point Optimizer (IPOPT).

For both the EDP and the UCP the optimization problem is formulated so that the economic profit is maximized. The revenue comes from selling heat and electricity, while fuel and maintenance costs are the main expenses. Only constant prices are considered. Thanks to the physical modeling in the EDP, the optimization formulation in this case includes constraints on critical variables such as mass flows and customer inlet temperatures.

Five optimization cases have been created in order to study the implemented method for solving the production planning problem. The first three consider only deterministic data and are of increasing complexity. In the last two uncertain signals with discrete probability distributions are included in the formulation. The base optimization scenario is a 24 hour period with a demand profile roughly corresponding to the heat demand expected from a residential area on a week day.

In the first optimization case only one producer, the KVV, and one customer is present. A comparison between having no pipes in the model and having supply and return pipes between the production unit and the customer model is conducted.

The results show that the optimal solution in both cases involves maximizing the electricity production. In the EDP this is achieved by maximizing the mass flow through the KVV. Constraints on the pump speed, condenser pressure and customer inlet temperature are limiting. When there are pipes in the model the EDP solution involves using the heat stored in the pipes to satisfy a part of the heat demand. This lowers the customer supply temperature, which again maximizes the electricity production.

In case II there are three customers connected in parallel. The KVV is the production unit in the system and there are pipes between the KVV and the customers, and also between the customers. The effect of adding dissipation of heat in the pipe models is investigated.

The main feature of the results is the lowering of the main production peak due to the distributed customer network. The delay times between the customers distributes the demand peak in time for the producer. Another observation that can be made for this case is that the delay time does not have any influence when heat load changes are handled by only changing the mass flow. The incompressibility of the district heating water makes the heat flow changes instantaneous in this case. The introduction of heat losses has a surprisingly large influence on the optimal solution. The reason for this is that the condenser steam pressure is sensitive to changes in return temperature and mass flow. With heat losses the steam pressure constraint gets active for lower mass flows, resulting in a different production profile.

Case III is a more realistic case with three production units and an accumulator in the system. The optimization interval is four days, making the integration between the UCP and EDP important. A comparison between a distributed and point-wise network is conducted.

The UCP results indicate a clear advantage of considering the distribution of the network in production planning. By doing this it is possible to delay the start-up of additional units, due to the reduced production peaks caused by the distribution. It is also notable that the accumulator is used to handle all load variations and that the production units therefore can be run with constant loads.

The EDP results are very similar to the UCP results without the distributed network. The differences are greater when pipes and delays are added to the system; the reason for this is that is the more detailed pipe modeling in the EDP, which introduces additional possibilities such as using the network as an accumulator.

In case IV a comparison between stochastic programming and robust optimization is conducted. The setup includes three different production units: the KVV, an oil boiler and a solid fuel boiler. The heat demand profile is uncertain for the second half of the optimization time.

A higher expected profit can be achieved with stochastic programming compared to the robust formulation. The reason for this is that the stochastic programming scheme avoids starting additional units unnecessarily. This is highly beneficial when the heat demand turns out to be low. For higher heat demands it becomes necessary to use the more expensive production unit, but due to the probability distribution between the cases the extra cost this introduces is smaller than the gain at lower demands.

In case V the results from stochastic programming are again compared with a robust optimization results. In this model there are two production units and both the heat demand and the electricity price are uncertain. The results show that a higher expected profit can be achieved with stochastic programming.

A scaling test indicates that the commercial solver is clearly superior when optimization problems with many possible scenarios are considered. For less complex problems, such as the UCPs in the deterministic optimization cases, the open-source solver seems sufficient in terms of convergence speed.

Contents

1 Introduction

1.1 Background

1.1.1 Short-Term Production Planning

Running the production of heat and power in a cost efficient manner is desirable both from a producer and consumer perspective. Production planning does not only consider the cost when optimizing unit schedules, but also network heat load demand and operational constraints. The scheduling includes the status (on/off, discrete variable) of each production unit as well as each unit's produced electric power and heat (continuous variables). To the authors' best knowledge, there exists no robust algorithm for solving the resulting optimization problem that involves both the discrete and continuous variables, known as a mixed-integer non-linear programming (MINLP) problem. For tractability reasons, it is therefore necessary to divide the optimization formulation into two separate optimization problems:

1. The Unit Commitment Problem (UCP), where the statuses of the units are optimized and the main difficulty lies in the combinatorial nature of the problem.
2. The Economic Dispatch Problem (EDP), which considers the result from the UCP and optimizes the load level for each active unit. The main difficulty of this optimization problem are the non-linearity of the units and the non-convexity of the optimization problem (local minima may be present).

The major differences between the two optimization problems are the model complexities and resulting optimization problem types. The UCP contains only linear or piecewise linear models and the optimization problem is a mixed integer linear problem (MILP) that can be used to solve both UCP and EDP. The EDP contains non-linear models with higher level of detail than the UCP models and the resulting optimization problem does not contain any discrete variables and is casted as a non-linear program (NLP) to be solved.

1.1.2 Common Approaches

Approaches do not involve solving both the discrete and continuous variables simultaneously due to the difficulty of the MINLP. Instead, a related MILP problem (e.g., Outer Approximation or General Bender Decomposition) or a related NLP problem (eg., Branch and Bound, see [1]), is iterated over.

A few approaches solve both UCP and EDP and most of them are based on Lagrangian relaxation (LR) or on mixed integer linear programs (MILP). The LR approach can handle some non-linearities by using relaxation, but network topology is not considered. The most appealing feature of the LR approach is the decomposition of the global optimization problem into a global master-problem and a small unit-specific problem, which may be beneficial for large networks (over e.g., 100 units) The MILP formulation results in large-scale integer optimization problems, but due to the progress of efficient solvers for such problem, this strategy has increased in popularity.

Typical for today's scheduling of units is that the unit models are highly simplified such that all physical descriptions are almost removed and the resulting optimization problems can be solved by using simple linear optimization methods. In [2], it is mentioned that the following commercial software uses this approach:

- Planner [3]
- Energy Optima 2000 [4]
- OPTIMAX PowerFit [5]

A survey of available approaches for short-term production planning can be found in [6] and [7]. The Värmeforsk-reports [2] and [8] focuses on the effect of an uncertain load prediction on the optimization results and the effect of integrating a model of the distribution network in the MILP formulation, respectively. In [9], the effects of a detailed physics-based model of the EDP are seen together with an integration of UCP and EDP.

1.1.3 Limitations

Current scheduling algorithms performed using MILP formulation of the UCP/EDP problems contain heavily simplified models, often only modelled as algebraic equations with exceptions for storing dynamics (heat and fuel) and delays in the distribution net. The continuous optimization variables are typically heat flows and the effects of the supply temperature and flow as well as return temperature are not directly considered. This is limiting as these temperatures and flows affect many critical parameters such as amount of energy that can be stored in the net, heat losses in the net and efficiencies for electricity production in steam turbines. The simplified modeling approach can be expanded to include supply temperature, see [10], but this strategy introduces several difficulties in the formulation.

The distribution of heat is in many cases point wise, i.e., only one customer is considered. Effects of having a customer delay spread from the production units are not considered. Additionally, distributed production is not considered either.

1.1.4 Proposed Approach

The proposed approach utilizes the advances made in non-linear dynamic optimization. It is based on the decomposition of the discrete problem (UCP) and the continuous problem (EDP) and contains two stages:

1. UCP. The optimization problem is formulated as a MILP and solved using a MILP solver. The result of this optimization problem is the discrete variables, i.e., unit statuses (on/off). These are then passed to the EDP in the second stage.
2. EDP. From the results of the UCP in stage 1, it is known when units should be turned on and off. In the EDP, the desired load is dispatched between the running production units to meet load demands as well as operational and safety constraints.

The second step involves load dispatch considering a non-linear model containing producers, distribution net and customers based on physical laws without any major simplifications. The units are described using mass and energy balances expressed by enthalpy, pressure and mass flow rate based on non-linear steam tables approximated from IAPWS/IF97, see [9]. Dynamics are included to match the real dynamics of the system. The benefits of this model are:

- Reduction of modeling work (no simplification process to comply with solver capabilities)
- Highly accurate models
- Model parameters have physical meaning yielding simpler calibration
- Optimization problem can incorporate constraints on e.g., mass flow rates, temperatures and pressure in units, distribution net and customer.

The higher model complexity and the non-linear dynamic optimization problem require different types of solvers and strategies than the ones used for MILP problems, see [11]. One reliable and efficient method of dynamic optimization, which is used in this work, is to transcribe the dynamic optimization problem into a non-linear programming (NLP) problem by collocation. The transcription parameterizes the trajectories (algebraic variables, states and control variables) in time by a small number of variables, making the optimization problem tractable. The NLP can be solved efficiently using open-source or commercial NLP solvers. In this work, the open-source solver IPOPT (Interior Point Optimizer) was used. The authors have previously used this method in dynamic optimization of a carbon capture plant [12], start-up of a combined cycle power plant [13] and, more significantly, in short-term production planning of district heating in a previous Värmeforsk-project [9].

1.2 Project Goal

The overall project goal is to further develop decision support in district heating production planning based on the results of [9]. From a modelling perspective this includes

- larger and more general district heating networks, including support for distributed production and consumption.
- a district heating network model that supports energy accumulation.
- increased plant model complexities compared to [9], mainly for the unit commitment problem.

On the optimization side, the project goals cover

- maximization of profit for producers by utilizing production and fuel costs as well as heat and electricity prices.
- longer optimization horizons to handle long start-/stops times of units compared to [9].
- robustness against uncertainties in heat load predictions by using stochastic optimization in the unit commitment problem.
- integration between the unit commitment problem and the economic dispatch problem.

1.3 Distribution of Work

The work in the project has been distributed between the three participating parties as follows:

1. Modelon AB
    a. Modelling
        i. Performing all physics based non-linear modelling in Modelica based on [9]
        ii. Deriving net model for UCP and EDP based on results in [14]
    b. Optimization case specifications, including
        i. cases for analyzing the effects of distributed production and consumption
        ii. case with stochastic programming including units with different start-/stop and production costs
        iii. case for scaling test of UCP solvers
    c. Optimizations
        i. Performing non-linear dynamic optimization of the EDP using JModelica.org for all optimization cases
        ii. Performing stochastic optimization of case ii) above
        iii. Analyzing impact of distributed production and consumption as well as heat accumulation in distribution net.
    d. Integration of UCP and EDP optimizations
    e. Major contributions to project report
2. SICS Swedish ICT
    a. Modeling
        i. Modeling of distributed heat production and consumption in UCP.
        ii. Higher fidelity of UCP models compared to [9]
    b. Specification of stochastic optimization case having outcomes with different heat and electricity demands
    c. Optimization
        i. Setup of stochastic optimizations
        ii. Performing UCP optimization of cases where effects of distributed production and consumption should be studied.
        iii. Performing stochastic optimization of case b) above
        iv. Performing optimizations for scalability of different UCP solvers.
    d. Minor contribution to project report
3. Vattenfall AB
    a. Providing measurement data from Uppsala district heating plant and distribution net for validation of physics-based Modelica model
    b. Specifications of optimization cases.
    c. Providing comments and comparisons on optimization results

2 Overview of Uppsala District Heating Network System

The district heat network in Uppsala covers the major part of Uppsala and can handle a heat load range from 50MW to 600MW [14]. One of the main reasons why the Uppsala district heat network was chosen in this study was because it has distributed production of heat. In total, there are three production sites (see Figure 1):

1. Boländerna, containing a peat fired cogeneration plant denominated KVV (kraft- och värmeverk), a waste incineration plant, heat pumps and additional peak load boilers such as electric and oiled fired boilers.
2. Husbyborgsverket, containing oil fired boilers.
3. Värmepumpverket, containing heat pumps.

Figure 1: Production sites in Uppsala: Boländerna, Husbyborgsverket and Värmepumpverket. Figure taken from [15].

Figur 1. Produktionsanläggingar i Uppsala: Boländerna, Husbyborgsverket och Värmepumpverket. Figur tagen från [15].

In the optimization examples performed in this work, only the sites Boländerna and Husbyborgsverket are considered.

At the site Boländerna, also steam and cooling is produced and distributed to customers. However, this will not be part of the current study.

3 Modeling and Optimization Overview

The end goal of the modeling and optimization effort is to deliver an optimal production plan of forthcoming hours for the production units regarding customer heat demands and plant economics. Both modeling and optimization are divided into two separate problems: Unit commitment problem (UCP) and Economic Dispatch Problem (EDP). The two optimization problems are ran in series with UCP first and then EDP, where EDP utilizes the results from UCP.

Modeling efforts and optimizations with focus on plant dynamics, such as return temperature and supply temperature and flow, were made in [9]. Here, the same types of plant dynamics are included and the models and optimizations are extended to contain a more sophisticated distributed net model including pipes and customer models. Thus, focus is shifted towards the net and how it can be modeled, and how it affects the optimization results, while maintaining the dynamic effects of return temperature and supply temperature and flow. The two modeling and optimization problems can be described as follows:

1. Unit commitment problem (UCP): The UCP is constructed using linear, rough models in discrete time only. The optimization formulation provides a mixed-integer linear program to the solver, requiring all constraints to be linear. The solution of the UCP contains production unit statuses (on/off) as well as production unit loads. However, only the statuses are passed on to the economic dispatch problem, where further optimizations of the unit loads are performed.
2. Economic dispatch problem (EDP): The models used in the EDP are high fidelity physics-based models in continuous time. Using non-linear dynamic optimization techniques, the load of each unit is decided. The EDP uses the on/off status results from the UCP as input. With the higher fidelity models compared to the ones used in UCP, the optimization problem can be more tightly related to the physical plant in the form of e.g., maximum flows, temperatures and pressures which the physics-based EDP model contains.

Both types of optimizations need a customer heat demand prediction. In general, this prediction is generated from e.g., weather forecasts and date and time, see e.g. [16]. This type of prediction was used in [9], while in this work predictions manually generated with the same type of characteristics are used.

A more detailed overview of the solution process using UCP and EDP together with measurement data and an external heat demand prediction can be found in [9].

4 Unit and net Models

4.1 Modeling for the Discrete Optimization

The time discrete models in the UCP are very coarse compared to models used in the EDP. The UCP models are linear and describe mainly energy and energy flows. Thus, important variables such as mass flow rate and temperatures are not considered, nor is the influence of actuators such as pumps and valves. The heat production units, the accumulator and the pipe model described in more detail below.

4.1.1 Heat Production Units

4.1.1.1 Kraft- och värmeverk (KVV)

The most advanced model in the UCP is the kraft- och värmeverk (KVV). The heat and electricity production of the KVV is described as a polytope in the space of electricity, heat and return temperature. Thus, for a specific return temperature, the optimization has an electricity-heat plane to place the production in. This is an extension compared to [9], where the return temperature was not considered. The main benefit from including the dependency of the return temperature in the model is that the typical behavior of increased electricity production when the return temperature is decreased, is captured.

The polytope for the KVV is based on the physics-based, non-linear model in the EDP described in Section 4.2.1. By changing the boundary conditions and the control signals of this model, the plausible heat and electricity productions can be found. The load and return mass flow rate were varied according to Table 1, while the by-pass split was fixed to 0% and two different return temperatures, 40ºC and 60ºC, were considered.

Table 1: Load and return mass flows used when generating KVV polytope for UCP model.

Tabell 1. Last och returmassflöde vid generering av KVV-polytopen för UCP modellen.

Boundary condition/Control signal	Minimum	Maximum
Load	50%	100%
Return mass flow	1000kg/s	3000kg/s

The resulting heat and electricity productions for the two return temperatures can be found in Figure 2. In each plane, four extreme points were located and used for finding interpolation planes such that the polytope including the return temperature could be constructed. The resulting and used polyhedron is found in Figure 3. The produced heat of the KVV has constraints in the rate of change; see Table 3.

For calculation of fuel consumption of the KVV, a constant efficiency of $\eta_{KVV} = 88.5\%$ is used as

$$P_{el} + Q_{KVV} = \eta_{KVV} U_{fuel,KVV},$$

where P_{el} and Q_{KVV} are produced electricity and heat, respectively and $U_{fuel,KVV}$ is the fuel energy flow used by the KVV.

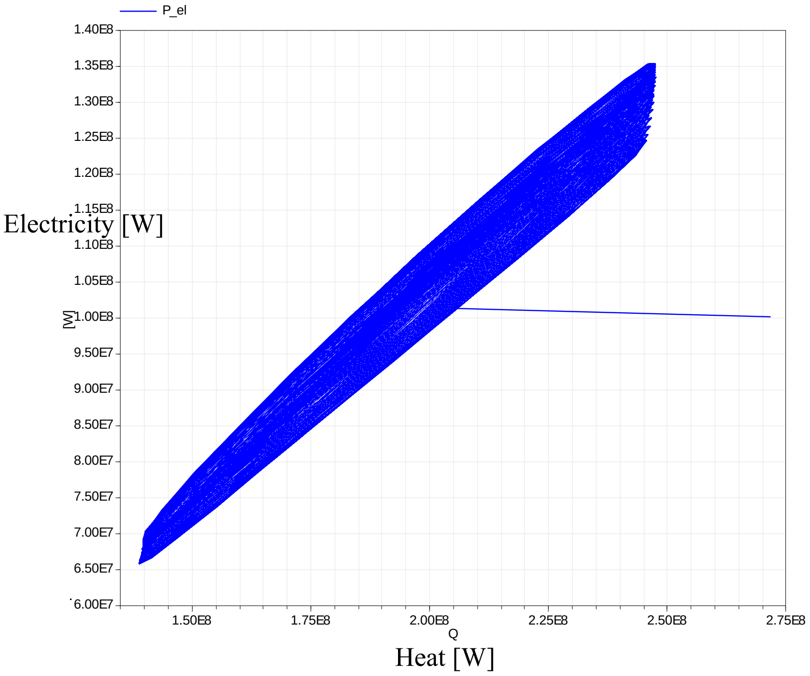

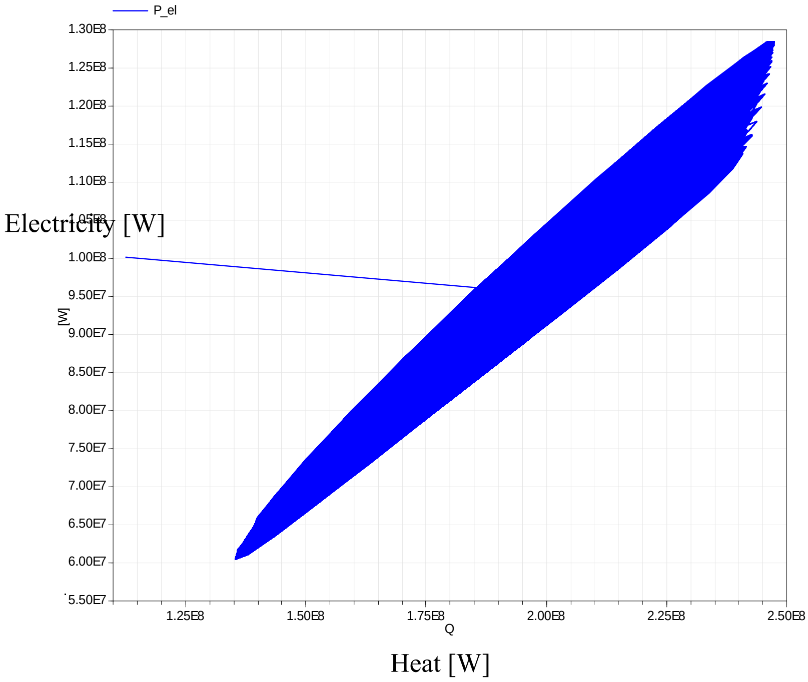

Figure 2: KVV heat and electricity production when varying return flow and load according to Table 1. Top: 40ºC. Bottom: 60ºC.

Figur 2. Värme- och elektricitetsproduktion för KVV när returflödet och lasten varieras enligt Table 1.

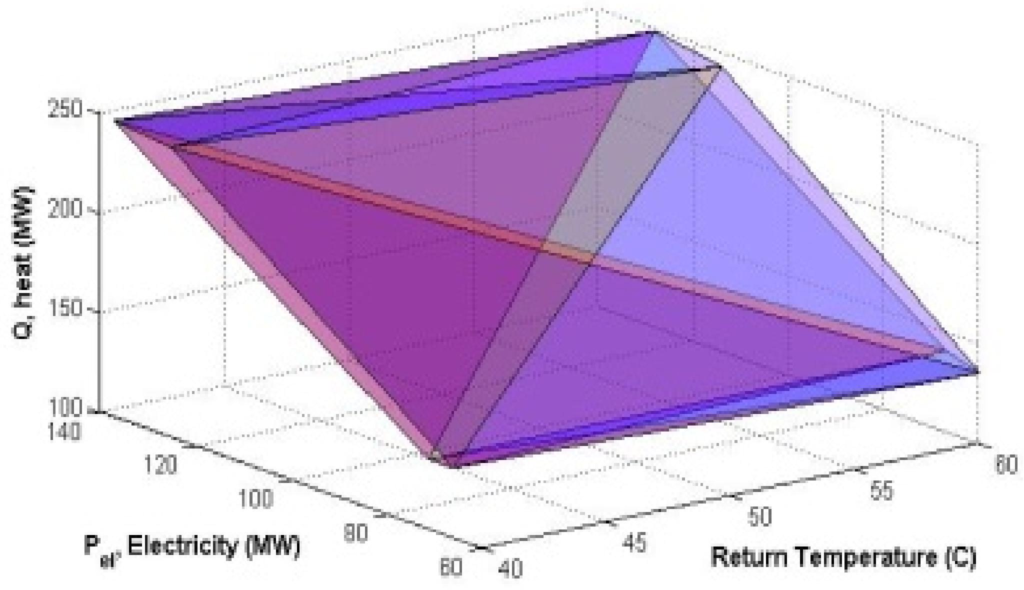

Figure 3: Polyhedron of KVV expressing possible operating regions depending on return temperature.

Figur 3. Polyeder för KVV med möjliga driftområden beroende på returtemperaturen.

4.1.1.2 Other Production Units

None of the production units, except for the KVV, are modeled in more detail than the decision variable, which is the produced heat Q_{unit}. The produced heats have several constraints, see sections 4.1.1.3 and 4.1.1.4, and are related to fuel consumption $U_{fuel,unit}$ through efficiencies η_{unit} as

$$Q_{unit} = \eta_{unit} U_{fuel,unit}.$$

The production units and their efficiencies are found in Table 2.

Table 2: Fuel efficiencies for units in the discrete optimization.

Tabell 2. Bränsleverkningsgrader för enheter i den diskreta optimeringen.

Unit	Type	Efficiency [%]
KVV	Cogeneration plant	88.5
AFA	Waste incineration plant	68.3
Husbyborg	Oil boiler	89.6

4.1.1.3 Production Capacities and Load Changes

Each production unit has minimum and maximum production capacity as well as minimum and maximum load changes. For the considered units in this work, these constraints are found in Table 3.

Table 3: Maximum production capacities and maximum and minimum load changes

Tabell 3. Maximala produktionskapaciteter och maximal och minimala laständringar.

Unit	Min capacity [MW] (when running)	Max capacity [MW]	Min load change [MW/h]	Max load Change [MW/h]
KVV	Polytope	Polytope	-50	50
AFA	89.6	128	-50	50
Husbyborg	15	120	-75	75

4.1.1.4 Start and Stop Trajectories

Different types of heat units have different start and stop times that need to be considered when performing production planning optimization. During the start-up the heat production of a unit typically follows a predefined trajectory. In this study these trajectories are represented in the UCP model using piecewise linear functions for the ramp-up and ramp-down. In Figure 4 the start and stop trajectories for the Husbyborg oil boiler are displayed.

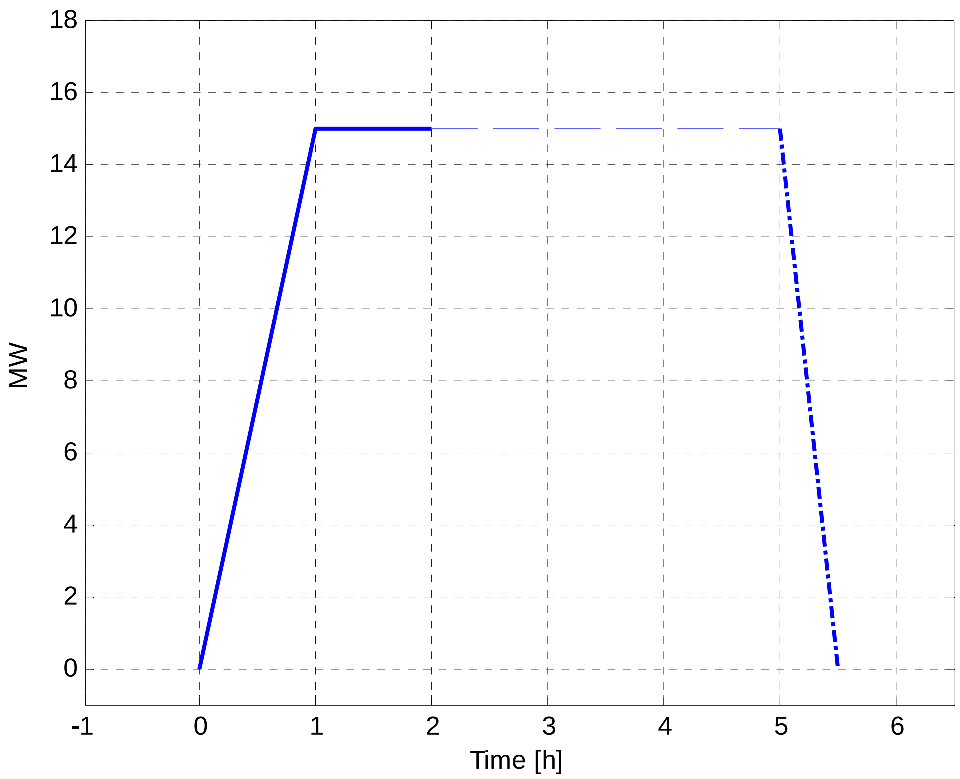

Figure 4: Start-up (solid) and shut-down (dash-dotted) trajectories of Husbyborg. Dashed line shows lower limit on capacity when running.

Figur 4. Uppstarts (heldragen) och nedstängningsprofiler (streck-prickad) för Husbyborg. Streckad linje visar lägsta kapacitet vid drift.

4.1.2 Pipe

The pipe model used in UCP contains two parts; the delay and the heat loss. The delay is simply a certain number of sample period delays while the heat loss depends on several parameters. The number of samples is calculated by truncating the actual delay to nearest integer number of sample times. The dissipated heat is calculated through

$$\dot{Q} = \frac{2\pi\lambda L(t_0 - t_s)}{\ln\left(\frac{2N}{D} + \sqrt{4\left(\frac{N}{D}\right)^2 - 1}\right)},$$

which describes the heat flow from an underground cylinder with temperature t_0, given the ground temperature t_s [17]. The parameters and values of this equation can be found in Table 4. The resulting value for the heat loss for a pipe of length 3000 meters, as a function of the temperature difference, is presented in Figure 5.

Table 4: Parameters and parameter values of heat dissipation in pipes.

Tabell 4. Parametrar och parametervärden för värmeförlust i rör.

Parameter	Interpretation	Value	Unit
L	Pipe length	Pipe dependent	m
N	Pipe depth	1	m
D	Pipe diameter	0.7	m
λ	Soil heat transfer coefficient	1.2	W/(m^{2}K)

It should be noted that the modeling of the heat loss is highly simplified. The equation above is for instance only valid for subterranean pipes and the fact that the return and supply pipes typically are situated close to each other has been neglected. The fixed parameter values for all pipes are also a clear simplification as all these values typically vary depending on the location of the pipe in the network.

As the UCP model does not contain the district heating water temperature, it is assumed that the supply temperature is 90°C. The return temperature is calculated through a return temperature model that has the outdoor temperature t_s as input, which is an input to the optimization as a prediction. The return temperature model is explained in section 4.3.2.

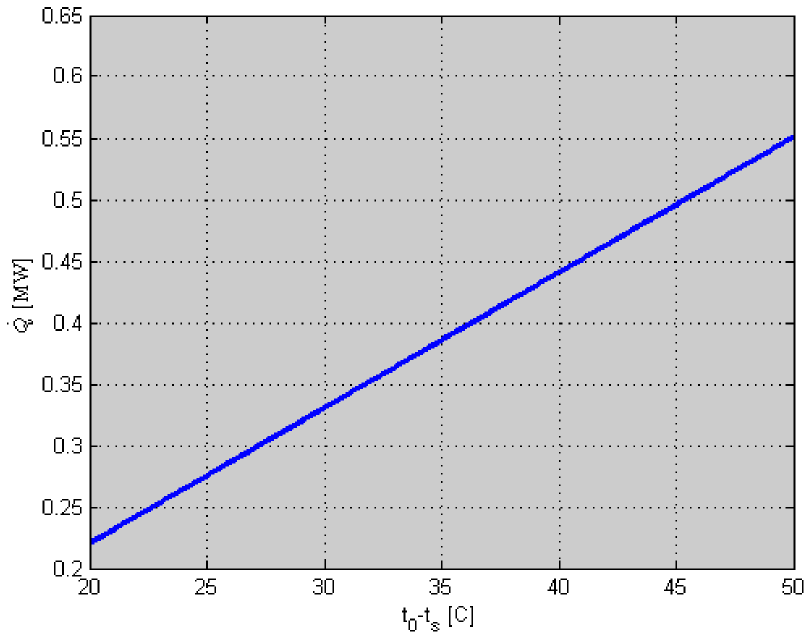

Figure 5: Pipe heat loss depending on temperature difference between water and ground for a 3000 meter pipe.

Figur 5. Värmeförlust i rör beroende på temperaturskillnad mellan vatten och markytan för ett 3000 meter långt rör.

4.1.3 Accumulator

The accumulator is modeled as a simple integrator for the stored energy, i.e.,

$$E_{acc}[t] = E_{acc}[t-1] - hQ_{acc}[t-1]$$

where $E_{acc}[t]$ is the accumulator energy and $Q_{acc}[t-1]$ is the energy flow to or from the accumulator and h is the sampling period. The accumulator has maximum and minimum limits on the stored energy as well as maximum and minimum limits on the heat flows to and from it. The limits are summarized in Table 5.

Table 5: Limits for the accumulator in UCP.

Tabell 5. Begränsningar för ackumulatorn i SSP.

Min storage [MWh]	Max storage [MWh]	Min change [MW]	Max change [MW]
250	1000	-100	100

4.2 Modeling for the continuous optimization

4.2.1 Co-generation plant KVV

The co-generation plant "Kraft- och värmeverket" (KVV) mainly contains

- A boiler providing steam for the turbines
- A reheater raising temperature of the steam after the high pressure turbine to be used by low pressure turbine stages
- Several turbine stages with varying operating points
- Four condensers
    - One condenser mainly used in start-up/dumping mode
    - Two condensers that transfer the major part (~95%) of the heat to the district heating water
    - One condenser transferring a minor part (~5%) of the heat to the district heating water. This condenser uses part of the bleed steam from the high-pressure turbine after the steam has driven the feed water pump.
- A pre-heating systems for pre-heating of water to boiler
- A feedwater tank

The main characteristics needed to be captured by the KVV model for optimization is the influence of plant load and incoming district heating water temperature and mass flow rate on the produced heat and electricity.

The main assumptions when deriving the KVV model have been

- The vapor characteristics from the boiler outlet (pressure and enthalpy) are constant and the vapor mass flow is linearly dependent on the plant load.
- The condensate leaving the condensers is assumed to be at saturation pressure.
- Pre-heating systems are not considered and the bleed streams that normally go from the low-pressure turbines to the pre-heating systems are modeled using a lumped pressure drop and a fixed pressure boundary.
- Bleed steam from high-pressure turbine is fixed.
- The condenser mainly used in start-up/dumping mode is not considered.
- The splits for incoming district heating water between to flue gas cooler and minor condenser (~5% of total transferred heat) are fixed, while the split between main condensers (~95% of transferred heat) and by-pass valve is varied during model validation but fixed during optimization (essentially all flow to main condensers).

The KVV is described by the following units in the Modelica model; see also Figure 6.

- One high-pressure (HP) turbine stage and three low pressure (LP) turbine stages, all with bleed flows.
- Two main condensers, transferring heat from steam to district heating water.
- One reheater for raising temperature of steam from high-pressure turbine to nearly inlet temperature
- Lumped pressure loss for bleed stream of first low-pressure turbine stage.
- Control volumes.
- Flue gas cooler and minor condenser as ideal heat transfers.

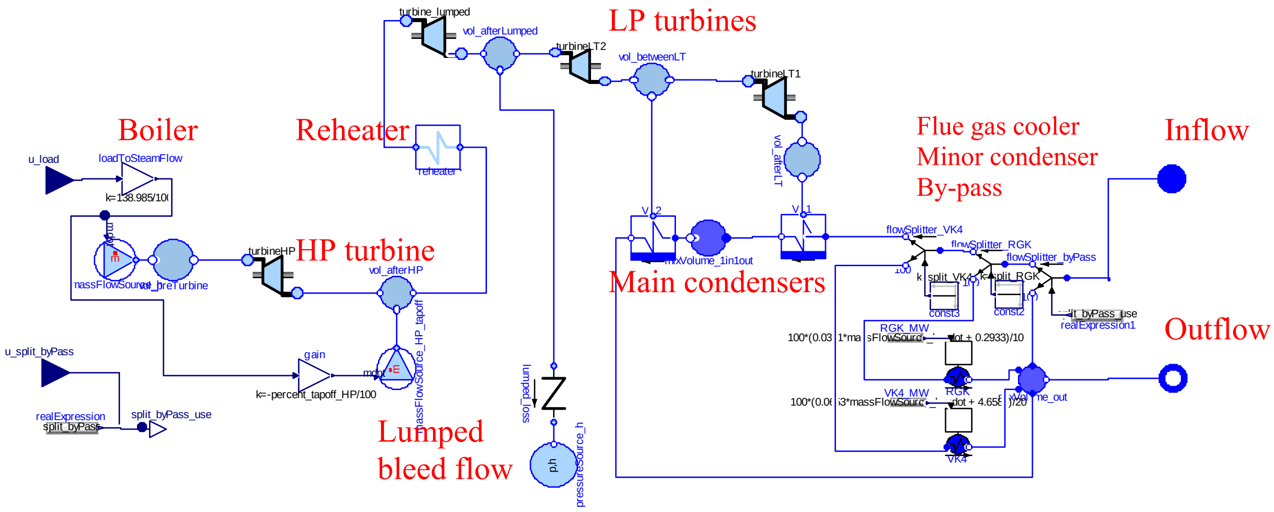

Figure 6: Overview of KVV Modelica model containing boiler, turbines, condensers, reheater and control volumes.

Figur 6. Översikt över Modelica modellen av KVV innehållande panna, turbiner, kondensorer, mellanöverhettare och kontrolvolymer.

4.2.1.1 KVV Units

The following sections will give a more detailed description of the units in the KVV.

4.2.1.1.1 Turbine

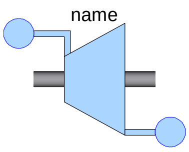

The physics based calculations of the outlet enthalpy and turbine work is defined by an isentropic efficiency, and the mechanical power is calculated using a mechanical efficiency. The pressure drop over the turbine is related to the flow rate using Stodola's law. Generator losses are considered using an electrical efficiency parameter. The efficiency parameters are found in Table 6 with default values.

Table 6: Efficiency parameters with default values for turbine model.

Tabell 6. Verkningsgrader för turbinmodell.

Efficiency type	Value
Isentropic	0.92
Mechanical	0.98
Electrical	0.95

4.2.1.1.2 Condenser

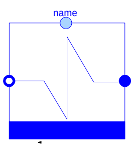

The heat flow rate transferred to the district water is driven by the temperature difference between the incoming water and the saturation temperature in the condenser. This heat flow rate is further used to compute the condensation rate that drives the bleeding flow from the turbine stages.

4.2.1.1.3 Control Volume

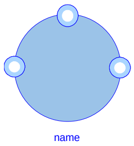

The control volume is a straightforward implementation of dynamic mass and energy balances expressed using pressure and enthalpy as states. Temperature is computed using pressure and enthalpy. The model requires partial derivatives of density with respect to enthalpy and pressure. The control volume is not a component with a physical equivalent in the plant; it is a component for holding physical balances and numerics.

4.2.1.1.4 Pressure Loss

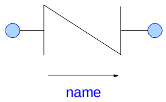

The mass flow through the pressure loss model is calculated using the pressure difference and a quadratic loss function.

4.2.1.1.5 Reheater

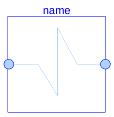

The reheater is assumed ideal in the sense that the temperature of the steam at the outlet is perfectly controlled to a specific set-point given as a parameter.

4.2.1.2 Calibration and Validation

The KVV has been calibrated and validated against both a previously derived plant model and measurement data from the KVV.

The previously derived model of the KVV, provided by Vattenfall AB, was built in Ebsilon. It was derived several years ago and at the time provided good agreement with measurement data. However, as conditions in the plant may have changed, the Ebsilon model was only used as a starting point for the calibration. The Ebsilon model was run in several different load cases, showing that the split factors for the incoming district heating water to the flue gas cooler, the minor condenser and the by-pass were almost constant. The determined splits are found in Table 7. It was also determined that the fixed bleed flow from the high-pressure turbine could be approximated to 12.4% of total steam flow.

Table 7: Split factors in KVV based on Ebsilon simulations.

Tabell 7. Splitfaktorer i KVV baserat på Ebsilon-simuleringar.

Unit(s)	Split factor (%)
Main condensers	93.3
Flue gas cooler	1.0
Minor condenser	4.5
By-pass	1.2

Additionally, from the different load cases, correlations could be derived between the KVV boiler steam flow and heat to district heating water provided by the flue gas cooler and the minor condenser. These could be approximated such that the heat to district heating water depends linearly on the boiler steam flow, see Figure 7 and Figure 8 . The total heat provided to the district heating water is greater than 120 MW at low steam flows, which results in a relative error off less than 1% due to the linear approximation, which was considered acceptable.

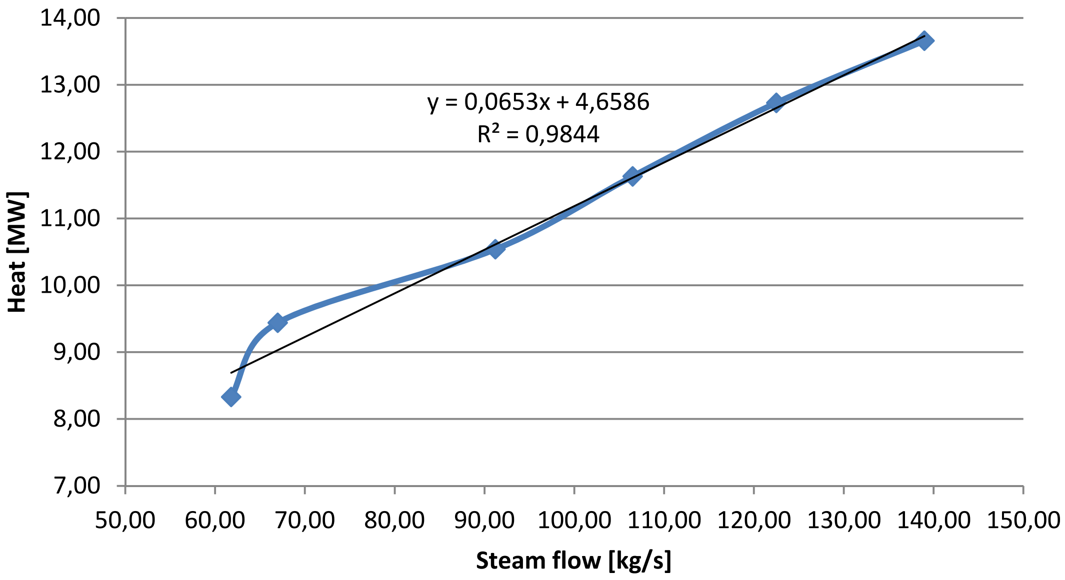

Figure 7: Heat transferred to district heating water from minor condenser as a function of boiler steam flow. Dot-marked line is result from Ebsilon model, while solid line is linear approximation.

Figur 7. Värmeflöde till fjärrvärmevatten från mindre kondensor som funktion av ångflöde. Resultat från Ebsilon-modell markerat med punkter, heldragen linje är linjär approximation.

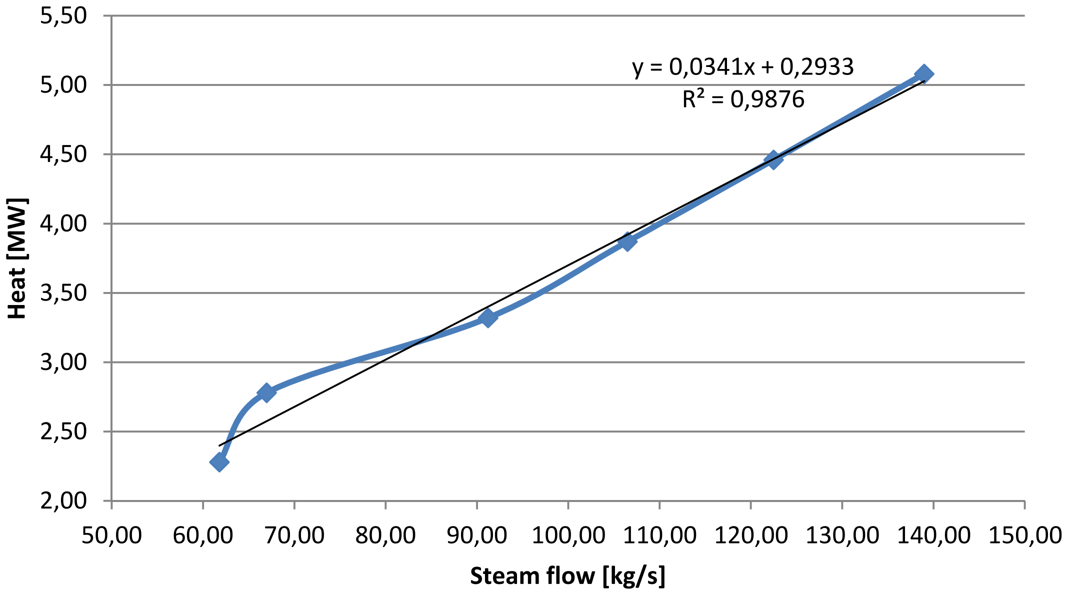

Figure 8: Heat transferred to district heating water from flue gas cooler as a function of boiler steam flow. Dot-marked line is result from Ebsilon model, while solid line is linear approximation.

Figur 8. Värmeflöde till fjärrvärmevatten från rökgaskylare som funktion av ångflöde. Resultat från Ebsilon-modell markerat med punkter, heldragen linje är linjär approximation.

Comparisons between the Ebsilon and Modelica model for total electricity and heat produced are found in Table 8 and Table 9. The relative errors are always less than 5% which is acceptable for the considered application of the Modelica model.

Table 8: Electricity produced in Ebsilon and Modelica model and relative error as a function of boiler steam flow.

Tabell 8. Producerad elektricitet i Ebsilon- och Modelica-modell och relativt fel som funktion av pannans ångflöde.

Steam flow [kg/s]	P el [MW] (Ebsilon)	P el [MW]	Rel. error [%]
61,79	48,05	49,96	3,97
66,97	60,99	62,95	3,21
91,21	74,46	76,32	2,50
106,49	89,37	90,11	0,82
122,49	106,58	104,60	-1,85
138,99	125,12	119,63	-4,38

Table 9: Heat transferred to district heating water in Ebsilon and Modelica model and relative error as a function of boiler steam flow.

Tabell 9. Värmeflöde till fjärrvärmevatten i Ebsilon- och-modell och relativt fel som funktion av pannans ångflöde.

Steam flow [kg/s]	Q [MW] (Ebsilon)	Q [MW]	Rel. error [%]
61,79	125,00	126,30	1,04
66,97	150,00	148,41	-1,06
91,21	175,00	172,85	-1,23
106,49	200,00	198,91	-0,55
122,49	225,00	226,69	0,75
138,99	250,00	256,17	2,47

The Modelica model has been validated against measurement data during the time period 2013-11-20 – 2013-12-10. The placement of sensors in the plant has complicated the validation as

- There is no temperature sensors placed at KVV outlet or inlet. The supply or return temperature is measured after or before the accumulator.
- There is no mass flow sensors placed at KVV inlet or outlet. The supply or return flow is measured after or before the accumulator.
- There is no flow sensor for the by-pass valve in KVV.

The implications of the above sensor placements are that

- Sensor values of the inlet and outlet temperatures and mass flow rates of the KVV can only be used when the accumulator is not in use, which is determined by having a small mass flow rate to or from the accumulator.
- The flow through the by-pass valve is estimated by using total inlet mass flow rate of the KVV when accumulator is not running, and mass flow rates of major and minor condensers and flue gas cooler.

Compared to the validation with the Ebsilon model, the by-pass valve is open more and varies. One reason for this is that it is manually set by operators such that they can guarantee that the pressure in the condensers is high enough for steam to be transported to the feed water tanks. Additionally, the isentropic efficiency of the turbine stages had to be increased compared to the Ebsilon validation.

Figure 9 shows the KVV load, return temperature and flow for the time period used for validation. The load varies within 50-85%, spanning almost the same range as the Ebsilon simulations. The return temperature and flow varies between 45-65ºC and 1400-2600 kg/s, respectively.

The resulting produced electricity and heat to district heating water are shown in Figure 10. The measurement of the heat is slightly noisier than the electricity measurement as it is calculated from several measurement signals in the sensor system of the plant.

Figure 11 shows the relative errors of the electricity and heat. It should be noted that the relative error is set to 0 when the accumulator is running and the sensor values cannot be used. Relative errors are for most of the validation period less than 5%, which is acceptable for the usage of the Modelica model. The validation results show that the Modelica model captures the main characteristics of the KVV in terms of effect of return temperature and flow and load on produced electricity and heat.

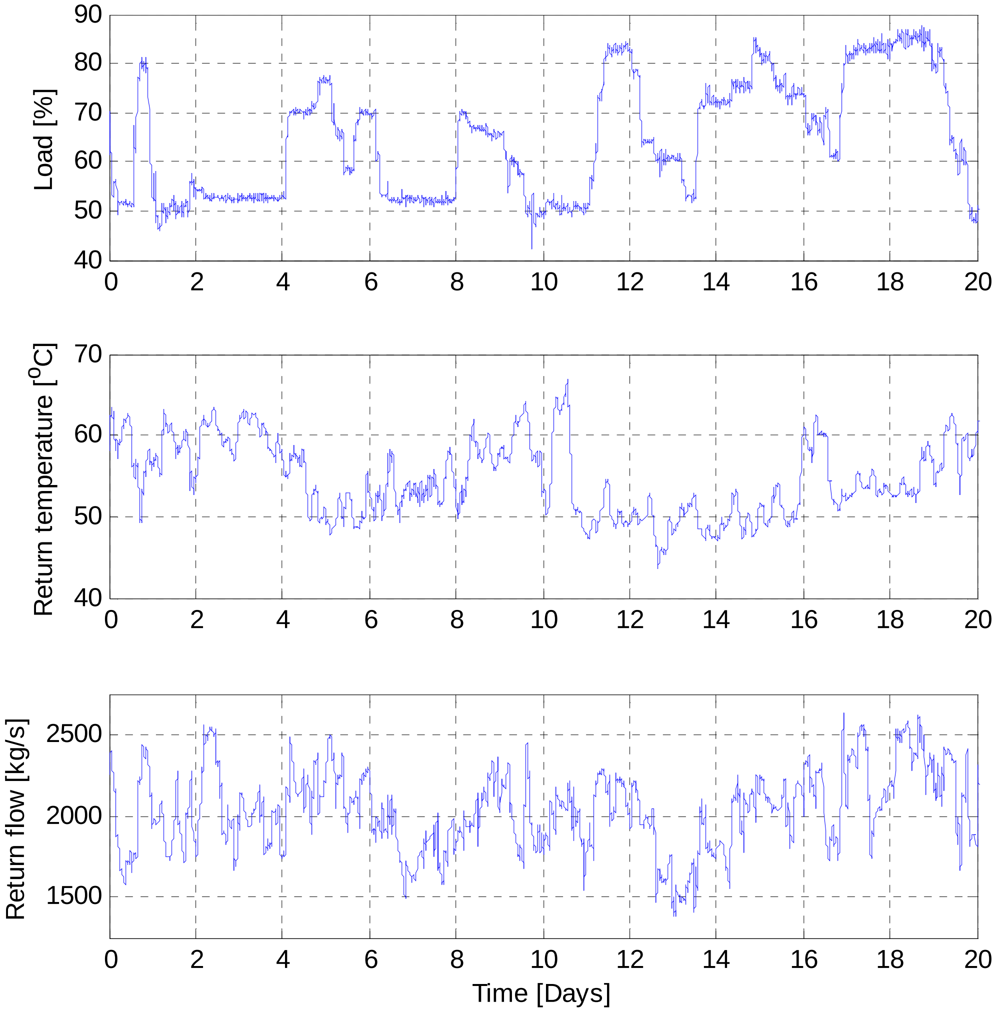

Figure 9: Boundary conditions and load when validating KVV Modelica model with measurement data. Top: KVV load. Middle: Return temperature. Bottom: Return flow.

Figur 9. Randvillkor och last vid validering av Modelica-modellen av KVV med mätdata. Överst: KVV-last. Mitten: Returtemperatur. Nederst: Returflöde

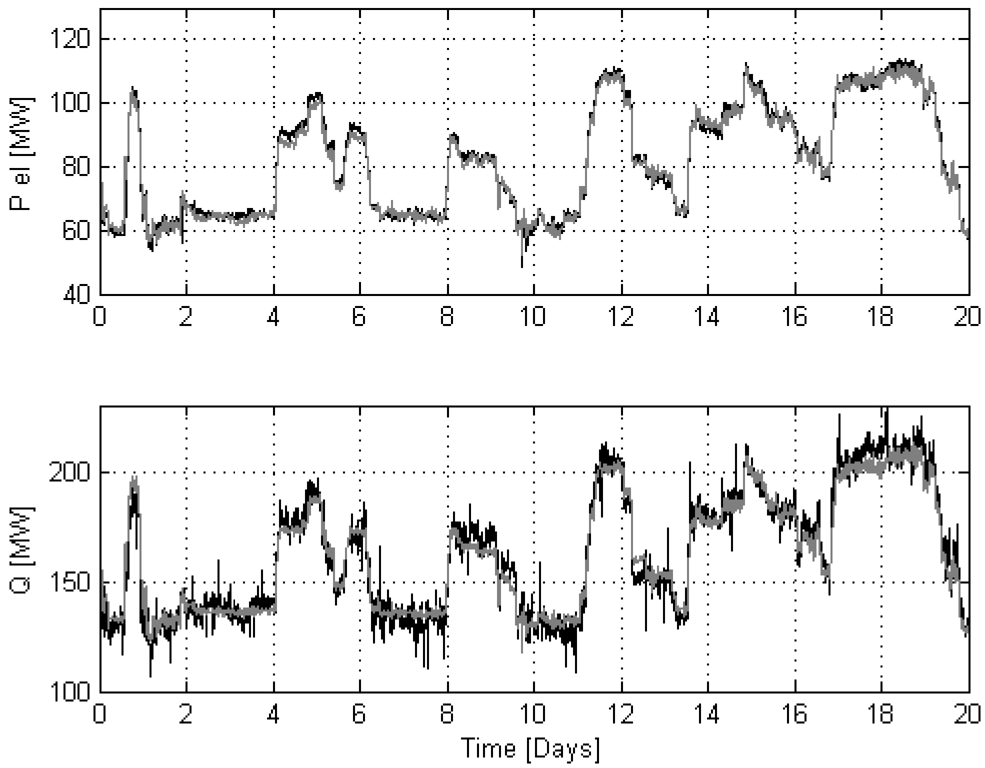

Figure 10: Top: Measured (black) and simulated (grey) electric power from KVV. Bottom: Measured (black) and simulated (blue) heat from KVV.

Figur 10. Överst: Uppmätt (svart) och simulerad (grå) elektricitet från KVV. Nederst: Uppmätt (svart) och simulerad (grå) värme från KVV.

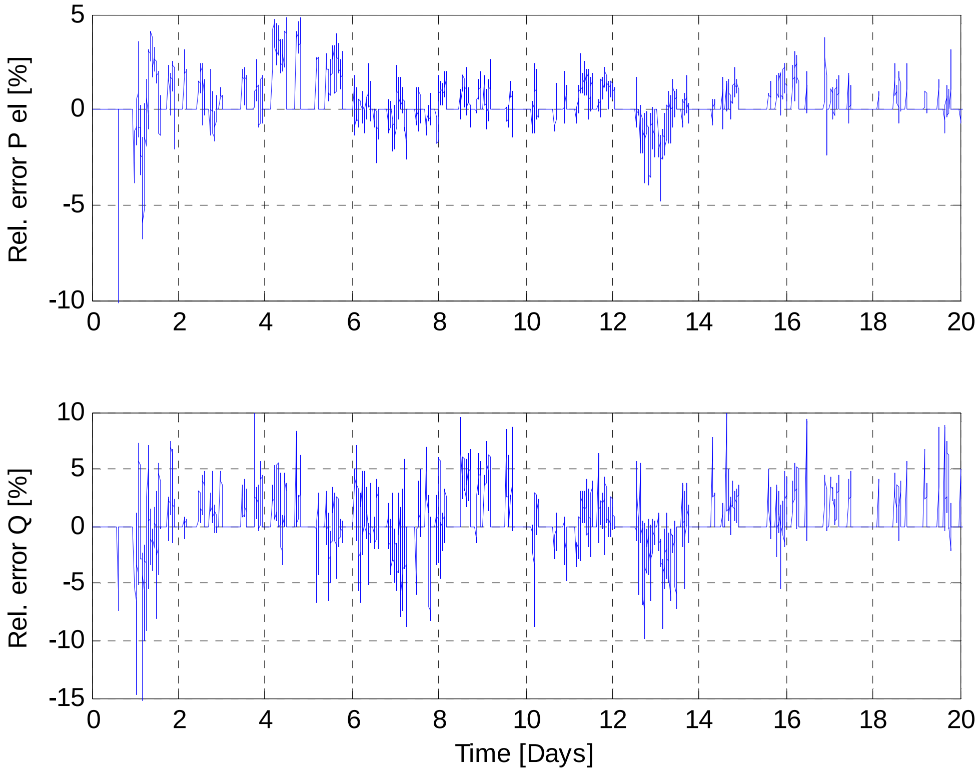

Figure 11: Relative errors between measurement data and model for electricity (top) and heat to district heating water (bottom) when validating KVV Modelica model.

Figur 11. Relativa fel mellan mätdata och modell för elproduktion (överst) och värme till fjärrvärmevatten (nederst) vid validering av Modelica-modellen av KVV.

4.2.2 Pipe

The pipe model in the EDP consists of three different sub-models connected in series: a finite element pipe implementation, a fixed time delay on the water enthalpy, and a heat loss model. The visualization of the pipe model in Dymola is displayed in Figure 12. The heat loss model is the same as for the UCP, explained in Section 4.1.2, but with the actual temperature of the incoming district heating water used instead of a fixed value for the supply pipes. The pipe delay is modeled more accurately, with a flow dependent delay of the water temperature. The details of the modeling are explained below.

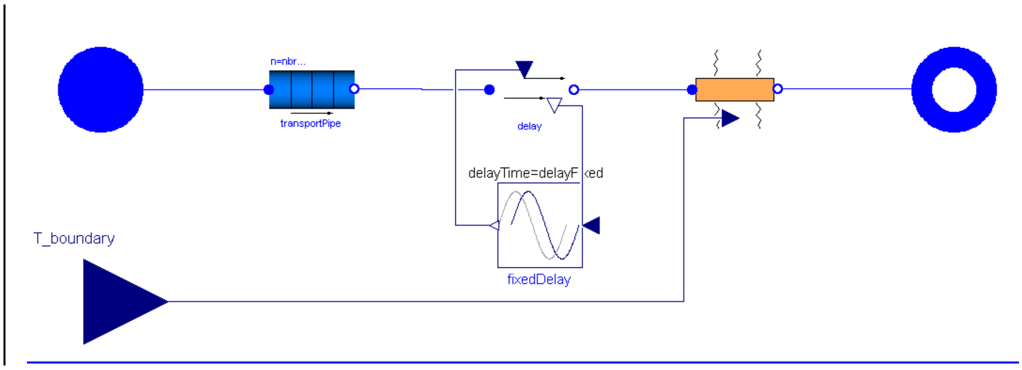

Figure 12: Continuous optimization pipe model in Dymola.

Figur 12. Rörmodell i Dymola för den kontinuerliga optimeringen.

4.2.2.1 Pipe Delay Representation

The pipe and delay models replace the fixed delay in the UCP in order to achieve a more realistic representation of the delay the pipe introduces. The volume of the pipe is divided into the finite volume model and the fixed delay. This division is implemented in order to avoid the downsides of using only a fixed delay or a finite volume pipe model as a representation of the pipe.

A fixed delay of the water enthalpy has a clear disadvantage whenever the mass flow through the pipe is varying, as this implies that the delay time of the real pipe is also varying, according to

$$T_{delay} = \frac{m}{\dot{m}}$$

where m is the total mass of the water in the pipe. In the Uppsala district heating network the mass flow through the KVV typically varies between 1400 and 2600 kg/s as seen in section 4.2.1, which makes a fixed delay model alone inaccurate. However, whenever the mass flow is fixed, the delay provides an exact solution to the translation of the temperature profile of the water passing through the pipe.

A finite volume pipe is always representing the delay time of a pipe correctly, even for varying mass flows. The weakness of this kind of model is that the temperature profile of the water passing the pipe gets distorted due to numerical dissipation. This effect gets more pronounced as the volume of each element in the model increases. The phenomenon can therefore be reduced by increasing the number of elements in the model, but this is undesirable in a model that should be used for optimization, as this would increase the model complexity.

By combining the two ways of modeling the pipe delay a compromise between the accuracy in delay time and temperature profile is achieved. The exact delay time of the fixed delay and the volume of the finite volume model are decided using the range of mass flows that will be present in each specific pipe. After deciding the level of delay time error that is acceptable on the boundaries of this range the corresponding volume and fixed delay is calculated. A higher accuracy in delay time means that volume of the finite volume model will be increased, introducing more numerical dissipation.

4.2.2.2 Heat Transfer Effects

In the EDP only the enthalpy of the water passing through the component is delayed in the delay and finite volume components, while the mass flow is unaffected. This introduces an important difference compared to the UCP model where the entire heat flow signal gets delayed. In the EDP instantaneous changes in heat flow through the pipe can be achieved by changing the mass flow rate while keeping the temperature constant, something that is not possible in the UCP.

The separate modeling of the water temperature and mass flow also introduces the possibility to use the pipes as accumulators. By feeding the pipes with hotter or colder water than the water that is in the pipe at a specific time, the total amount of energy in the pipe gets manipulated. These changes in energy can for example be used to handle customer load variations.

4.2.2.3 Optimization Model Implementation

In order to implement this kind of function in optimization models where the time is discretized some modifications are necessary. Similarly to the UCP case the amount of fixed delay is rounded to the nearest integer number of time steps. As a part of the delay is handled by the finite volume model the fixed delay times in the EDP are shorter than the corresponding delays in the UCP.

4.2.2.4 Simulation Validation

To investigate the behavior of the pipe implementation it is compared to a more accurate pipe model which is described in [18]. The parameter values in the experiment can be seen in Table 10.

Table 10: Pipe validation parameters.

Tabell 10. Rörvalideringsparametrar.

Parameter	Value
Pipe volume	1000 m^3
Mass flow range	1200 – 1600 kg/s
Accepted delay error	10 %
Number of finite volumes	2

In the optimization pipe model the parameter values result in a finite volume pipe volume of 500m^3 and a fixed delay time of 331 seconds. In the experiment the inlet temperature is ramped between 70^{o}C and 90^{o}C for mass flows of 1200 and 1800 kg/s. A comparison between the output temperatures and the delay times can be seen in Figure 13.

4.2.2.5 Energy balance violation

The fixed delay component introduces a possibility to violate the balance between the inlet and outlet energy flow for the whole pipe. By manipulating the mass flow and inlet temperature it is possible to consistently extract a different amount of energy from the pipe than what is added. The following example is a simple illustration.

Assume that the mass flow and enthalpy are oscillating out of phase with a period two times larger than the delay time. This can be achieved by having $\dot{m}(t) = \dot{m}_0 + A_m \sin(t)$, $h(t) = h_0 - A_h \sin(t)$ with a time unit chosen so that the delay time becomes $T_{delay} = \pi$. The difference between inlet and outlet energy then becomes

$$Q_{out} - Q_{in} = \dot{m}(t)\left(h\left(t - T_{delay}\right) - h(t)\right) =$$
$$(\dot{m}_0 + A_m \sin(t))\left(h_0 - A_h \sin\left(t - T_{delay}\right) - h_0 + A_h \sin(t)\right) =$$
$$(\dot{m}_0 + A_m \sin(t)) A_h (\sin(\pi - t) + \sin(t)) =$$
$$2(\dot{m}_0 + A_m \sin(t)) A_h \sin(t) = 2A_h(\dot{m}_0 \sin(t) + A_m \sin^2(t))$$

The quadratic term makes the expression greater than zero on average, meaning that more energy is leaving the pipe than what is entering. This is illustrated in Figure 14 and Figure 15.

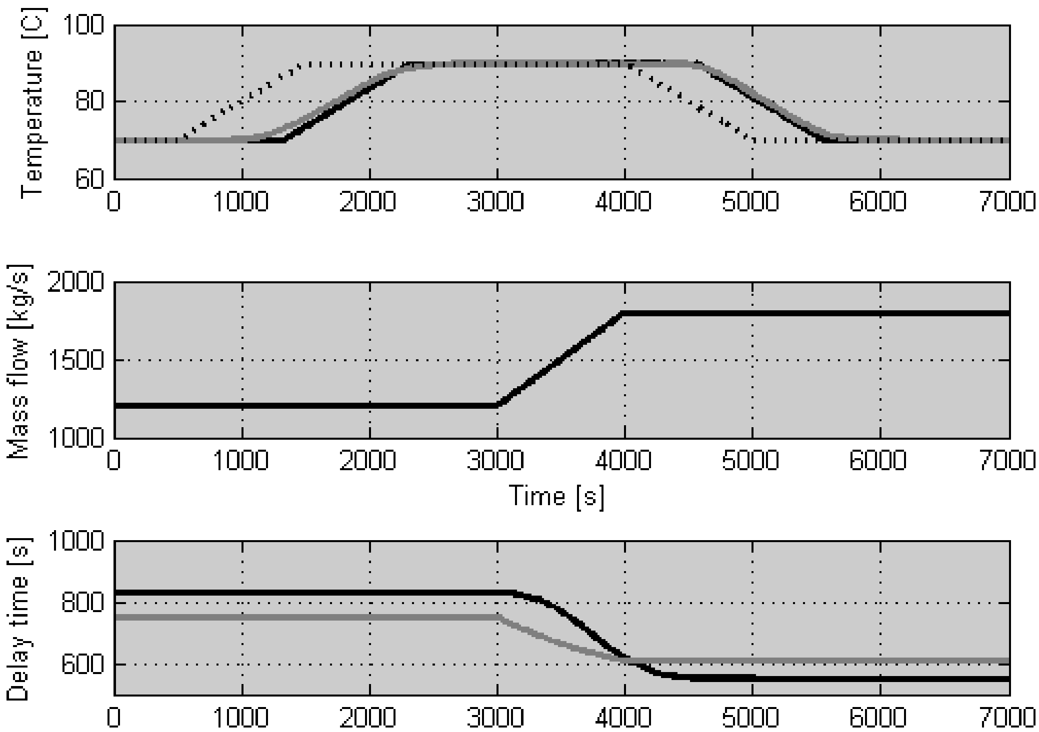

Figure 13: Comparison between detailed pipe implementation (black) and optimization model pipe (grey), when the input water temperature (dotted) and mass flow are varied.

Figur 13. Jämförelse mellan detaljerad rörimplementering (svart) och rörmodell för optimering (grå) när det inkommande vattnets temperatur (prickad) och massflöde varieras.

This phenomenon is not obvious during simulation as the effect will be cancelled out for most input signals, but the optimization software could exploit it as the pipes then become a source for free energy.

The following measures can be taken in order to limit this unwanted effect:

- Introduce pipe energy constraints
- Adjust ratio between fixed and variable delay
- Reduce degrees of freedom in optimization

All of these methods have some drawbacks. A pipe energy constraint is typically implemented as a point constraint for the final time of the optimization. This prevents unbalanced energy flows, but it can be hard to decide the final energy value and often results in transient behavior at the end of the optimization interval.

Reducing the fixed delay and increasing the variable delay correspondingly reduces the modeling error in terms of time delay, but increases the model complexity for two reasons. It typically requires an increased number of finite volume elements to represent the temperature profiles correctly and it requires a finer discretization in order to represent shorter fixed delay times correctly.

By limiting the degrees of freedom in the optimization formulation the possibility to manipulate mass flow and temperature simultaneously can be restricted or even eliminated, but it could also result in a less optimal solution.

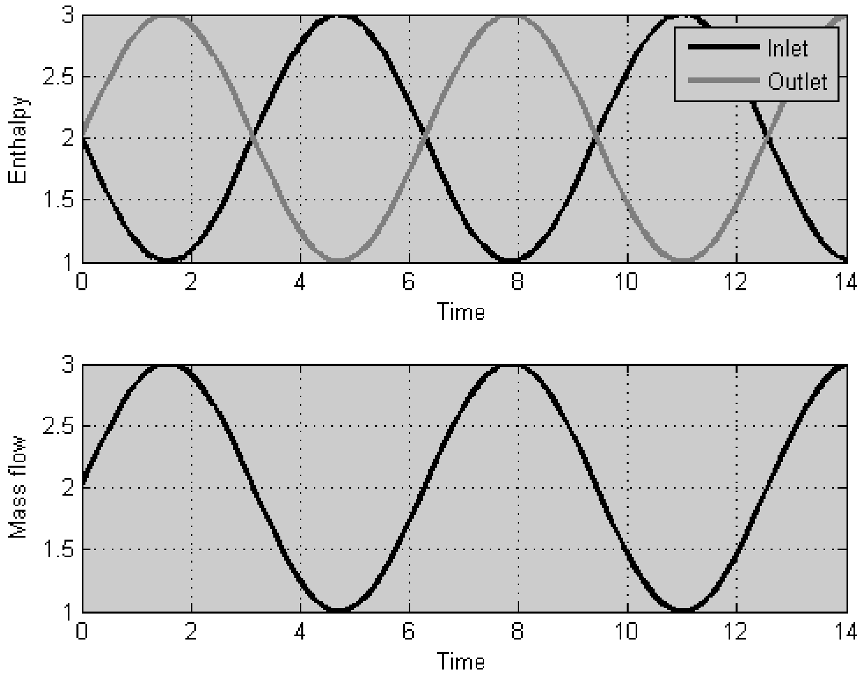

Figure 14: Fluid inlet enthalpy varying with a period of $2 * T_{delay}$, out of phase with the mass flow creating a violation against the energy balance.

Figur 14. Inkommande vätskans entalpi varierar med perioden $2 * T_{delay}$ ur fas jämfört med massflödet, vilket leder till brott mot energibalansen.

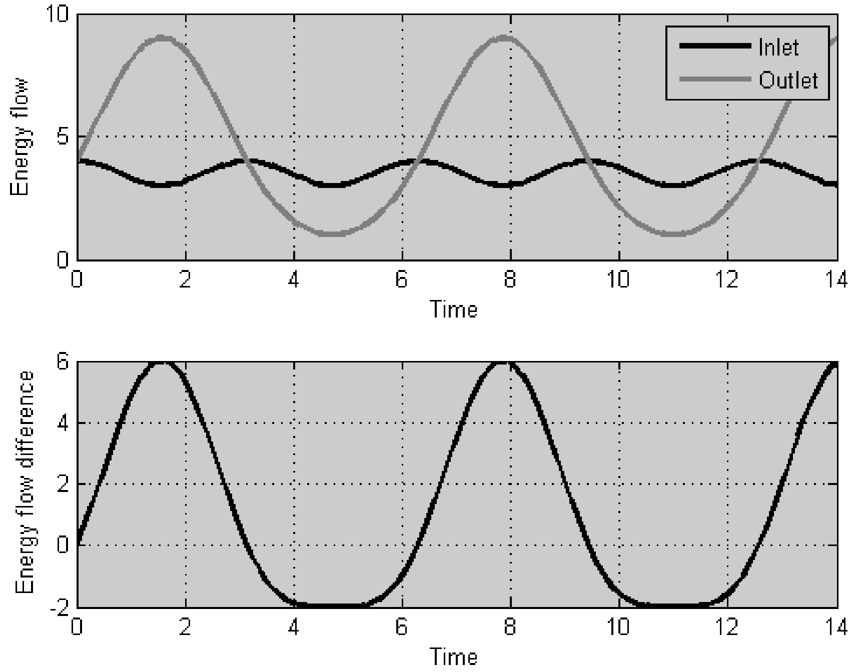

Figure 15: Energy flow at pipe inlet and outlet and the energy flow difference, which is greater than zero on average.

Figur 15. Energiflöde vid rörets in- och utlopp och skillnaden i energiflöde, som i genomsnitt är större än noll.

4.2.2.6 Optimization Validation

In order to determine how much the pipe modeling is affecting the optimization results two experiments are presented in this section. These tests are based on optimization case Ib, which is presented in detail in chapter 7. Firstly, the optimization case is solved with different pipe implementations. Two results are examined; the pipe energy difference between the end points of the optimization interval and the load profile for the production unit. Secondly, the optimization results with the most and least physical pipe representations are compared in simulations where both pipes are present. The energy levels in the pipes are compared.

For the first test, models with different pipe implementations are created, using the same setup as in case Ib. The difference in the pipe modeling consists of different ratios between the fixed and variable delay. A longer fixed delay indicates a more unphysical pipe model. The corresponding optimization problem for each model is solved and for each case the difference in return pipe energy between hour 0 and 24 is calculated. This difference is a measure of how much the energy balance is violated in each case. The results are presented in Table 11.

In this table the fixed delay and pipe volume used in the actual optimization case in chapter 7 correspond to the row marked with base. The settings marked with case III ratio represent the delay ratio used in the pipes in case II and III. A pipe implementation without any fixed delay and 10 finite volume elements is handled as the exact solution.

In a more complex model with more pipes this implementation would be impossible to optimize due to the increased number of states in the model introduced by the extra volumes.

Table 11: Pipe energy comparison for different pipe implementations.

Tabell 11. Jämförelse av rörenergi för olika rörimplementationer.

Finite volume elements	Finite volume volume [m^3]	Fixed delay [min]	Energy difference [MWh]	Comment
10	14462	0	0.86	”Exact”
5	12051	20	0.39	
4	9642	40	-0.65	
2	7231	60	-2.1	Base
2	4821	80	-4.4	Case III ratio

When the energy differences in the different optimization results are compared one can note that the most unphysical pipe implementation results in an energy difference that is approximately 5 MWh lower than the most accurate model. This indicates that the energy balance in the least favorable case is violated to a similar extent over the 24 hour optimization interval. In Table 12 this value is compared to the total heat demand over the total optimization time in the considered system, the total capacity of the accumulator, and the energy variations in the supply pipe during one day. From this comparison the energy gained from the unphysical pipe modeling is considered small enough to be acceptable.

Table 12: Pipe energy difference for case Ib compared to other data.

Tabell 12. Energiskillnad för rör i fall 1b jämfört med annan data.

Energy difference [MWh]	Total energy demand [MWh]	Accumulator capacity [MWh]	Pipe energy variations [MWh]
5	4300	750	74

The impact of the pipe modeling on the optimization result must also be investigated. A comparison between the KVV load profiles for three different pipe implementations are presented in Figure 16. The “exact” pipe implementation above is compared with the implemented settings for case Ib and the less favorable fixed-variable delay ratio used in for example case III. There are some differences, but the general characteristics are similar. As expected, the greater fixed delay does result in a solution further away from the “exact pipe” solution.

For the second test, simulation models containing the most and least physical pipe implementations were created. The inlet mass flow and temperature for the return pipes in the optimization results were used as input. By examining the energy level in the most accurate pipe model when the optimization results for the more unphysical pipe is used as input, and vice versa, the behavior of the optimization pipe model can be investigated.

The results from this test can be seen in Figure 17. When the optimal trajectories for the more unphysical pipe are used as input the difference in final pipe energy again is approximately 5 MWh, indicating that the optimization results involves a violation of the energy balance to a similar degree.

However, when the optimal solution to the optimization problem without a fixed delay is used as input there still is a difference in final pipe energy of approximately 2 MWh, when the final energy level of the two pipe models are compared. This implies that the usage of a pipe energy constraint would result in a too conservative solution, as the optimal solution without a fixed delay results in a lower energy level in the fixed delay model.

The conclusion from these tests is that the return pipe errors introduced by the pipe model implementation are small enough to be tolerated without the adding extra constraints for the considered pipes and that the implementation of energy constraints on these pipes would result in sub-optimal solutions.

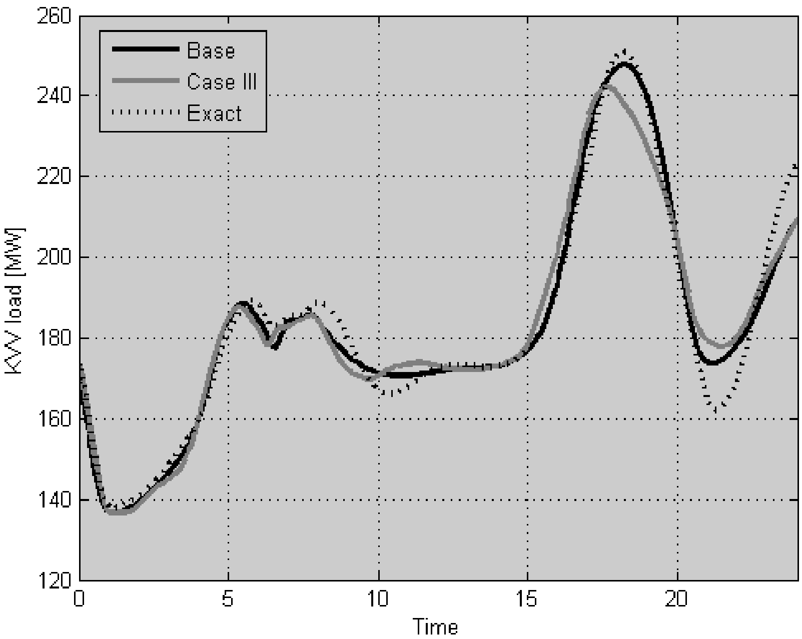

Figure 16: KVV load profile for different pipe implementations.

Figur 16. Lastprofil för KVV vid olika rörimplementationer.

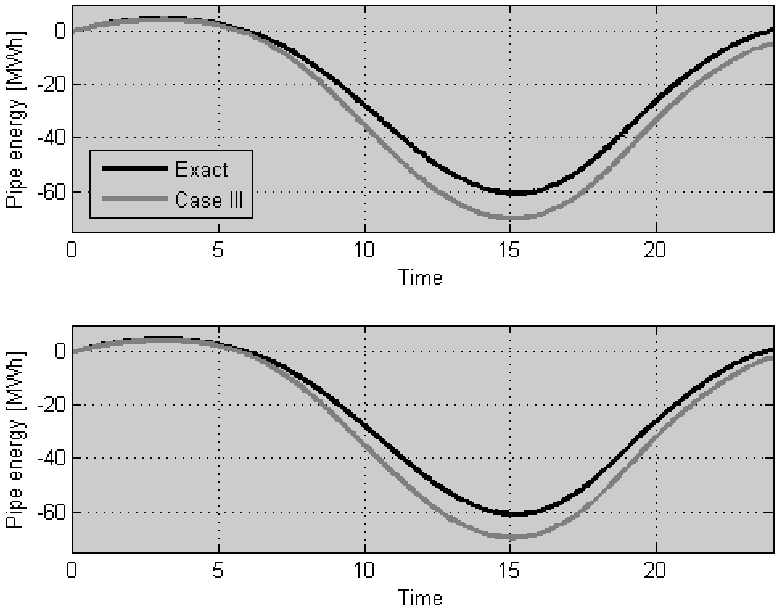

Figure 17: Pipe energy profiles in simulations with optimization input from model with optimization pipe (top) and exact pipe (bottom).

Figur 17. Rörenergiprofiler vid simuleringar med insignaler från optimering med optimeringsrör (överst) och exakt rör (nederst).

4.2.3 Pump

The pump is modelled in an ideal fashion, where the delivered mass flow depends linearly on the control input.

4.2.4 Heat production unit

A simple model that transfers heat to the water based on a control input corresponding to the load and a parameter for maximum heat transfer. Used for modeling various production units in the network.

4.2.5 Accumulator

A finite volume approximation is used to model the accumulator. Buoyance effects are neglected, which means that no mixing is assumed when the accumulator is not charging or discharging. The accumulator is connected to the supply water on the top and return water on the bottom and these ports are used for charging and discharging the tank.

4.2.6 Dynamic volume

For models without pipes a volume is added to the network. The extra state this adds to the system improves the behavior of the model significantly. It makes the system much less sensitive to load changes and therefore easier to control. The delay the volume introduces to the network is negligible compared to the time constants of the rest of the system.

4.3 Customer

4.3.1 Customer Model

Each customer demands a predefined amount of heat from the network, based on the heat load prediction. The district heating water mass flow through each customer is decided so that this demand is fulfilled based on the supply temperature from the network and the specific return temperature of the customer.

4.3.2 Return Temperature

The return temperature is calculated based on the outdoor temperature. The model described in [16] is used for this purpose. The correlation between outdoor and return temperature, shown in Figure 18, is implemented in the EDP model using a table.

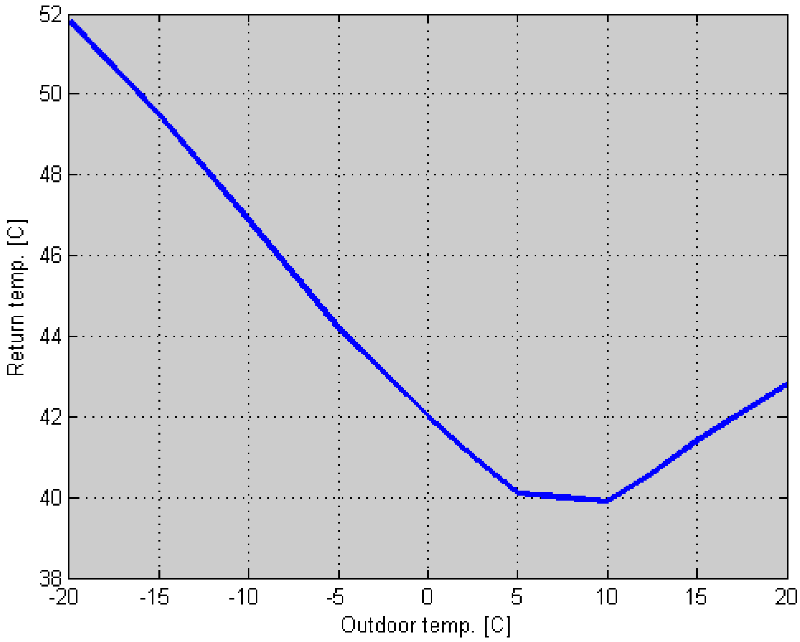

Figure 18: Customer return temperature depending on the outdoor temperature.

Figur 18. Kundernas returtemperatur beroende på utomhustemperaturen.

4.4 District Heating Network Model

The mathematical model of a district heating network should describe the mutual influence of the production sites, the distribution network and the consumers for different load and weather conditions. An exact mathematical description of the network shown Figure 1 requires a two-dimensional representation containing the geographical position of every customer and production site. Such a model with all customers, pipes and production units is only valuable for detailed simulations aiming at designing or extending a network but it would not be tractable in the context of dynamic optimization for production planning. A simpler network representation has therefore been used in the project to capture the properties that are relevant for production planning, namely:

- The transport delay between the produced heat and consumed heat
- The geographical distribution in the heat load
- The influence of the outdoor temperature on the supply water temperature at the customer stations
- The influence of the flow rate (or total load) on the transport delay
- The distributed nature of the production

A variant of the one-dimensional district heating model presented in [14] turns out to be sufficient to capture the previously listed features and to substantially improve the quality of the production plans compared to the standard approach without any distribution model at all. Figure 19 shows a schematic representation of the network structure that is used for the UCP and the EDP sub-problems and Figure 20 shows the distribution of heat as function of the distance from KVV.

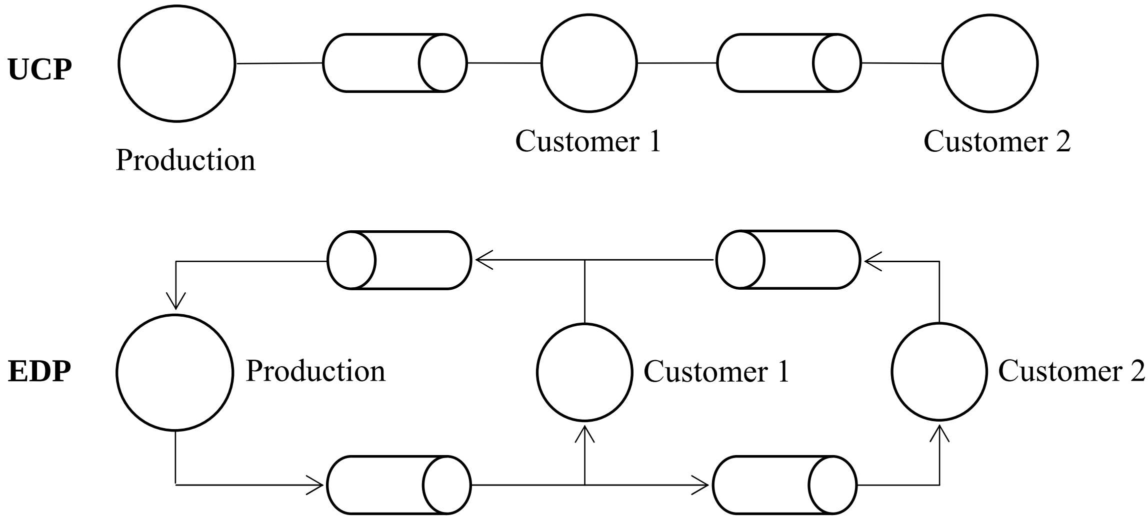

Figure 19: Schematic representation of network structure used in the discrete (upper structure) and the continuous optimization (lower structure).

Figur 19. Schematisk representation av den använda distributionsnätstrukturen i den diskreta (överst) respektive den kontinuerliga optimeringen (nederst).

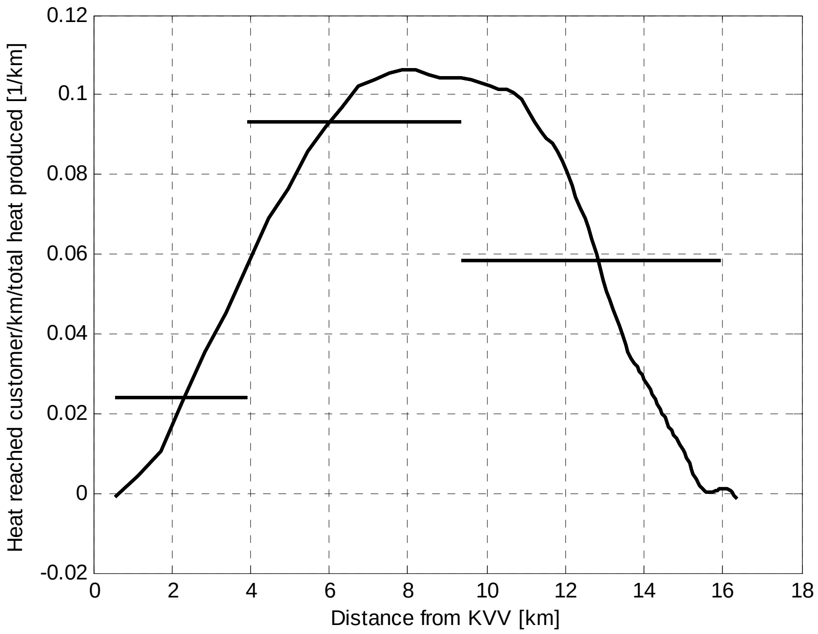

Figure 20: Heat reached customer/km/total produced heat as a function of distance from KVV, taken from [14]. Discretized values shows the three customers used to describe the Uppsala district heat network.

Figur 20. Andel av total värmeproduktion till kund per km som funktion av avstånd från KVV, hämtad från [14]. De diskretiserade värdena visar de tre kunderna som används för att beskriva Uppsalas fjärvärmenät.

In addition to its simplicity and flexibility, this network model can easily be combined with a global load prediction model, which is more common than a distributed load prediction. The total predicted customer load $Q_{tot}(t)$ is indeed distributed between the modeled customers using distribution factors α_i that are calibrated for a given network, i.e.,

$$Q_{cust,i}(t) = \alpha_i Q_{tot}(t).$$

The Uppsala network has previously been described and calibrated in [14], using 13 customers and delays. For the optimization, the model has been further simplified to contain three customers and delays. Two different distribution setups are used for optimization. The delays and distribution factors can be found in Table 17 in section 7.1.1.1.

One limitation of the chosen network model is the pre-defined direction of the flow in all pipes.

4.4.1 Discrete Optimization District Heat Network Model

The district heat network model in the UCP optimization model uses approximations that are not physical. They are:

- The delays are constant instead of flow dependent. This is limiting in some cases as an increased customer load met by an increased flow rate does not result in decreased transport delay. The transport delay should only be large when the plant follows the load by changing its supply temperature.
- The heat loss is not dependent on actual supply temperature; a fixed supply temperature is used to compute the heat loss in supply pipes.
- There is no return pipe. Heat loses in return pipes are however calculated using an outdoor temperature model and a return temperature model.

4.4.2 Continuous Optimization District Heat Network Model

The network model used in the continuous optimization is physics-based, meaning that

- The water in the pipes is characterized by a flow and a temperature instead of a transported heat flow rate.
- The heat loss is varying with outdoor temperature and water temperature, see Section 4.1.2.
- The transport delay is time-varying and dependent on the flow and the pipe volume, see Section 4.2.2.1.
- The return pipes are modeled and captures heat loss from the customers to the plant.

5 Optimization Tools

The two types of optimizations performed require two distinctly different optimization frameworks. The UCP requires a framework supporting formulation of and solving mixed-integer optimization problems while the EDP requires a framework for solving non-linear dynamic optimization problems.

5.1 Discrete Optimization

The aim of this section is to provide a short introduction to the programming language and solvers used for the UCP.

5.1.1 Pyomo

The UCP in this project is implemented in Python using the Pyomo modeling language [19]. To use Pyomo, three sets of files should be prepared for an optimization process:

- The **model file** in which all the sets, parameters, constraints (rules) and the objective function are defined.
- The **data file** in which all the input data for a specific scenario is set. To be more efficient in implementing these data files for different case studies, a MATLAB code was constructed to generate them.
- The **run file** in which the model and the data files together with the solver are called to solve the optimization problem. Handling the results in terms of saving or plotting is also defined in this file.

5.1.1.1 Gurobi

The UCP has been solved by the Gurobi commercial solver, see [20]. The Gurobi Optimizer is a commercial optimization solver for several types of optimization problems, including mixed-integer linear programming (MILP) problems. The UCP is of this type. The Gurobi Optimizer supports a variety of programming and modeling languages including: object-oriented interfaces for C++, Java, .NET, and Python and also matrix-oriented interfaces for C, MATLAB, and R. Gurobi provides free licenses for academic purposes.

5.1.1.2 GLPK

The UCP has also been solved by GNU Linear Programming Kit (GLPK), see [21], which is an open-source software package intended for solving large-scale linear programming (LP), mixed integer programming (MIP), and other related problems. GLPK uses the revised simplex method and the primal-dual interior point method for non-integer problems and the branch-and-bound algorithm together with Gomory's mixed integer cuts for (mixed) integer problems.

5.2 Continuous optimization

Modelica models can be simulated in many different environments and optimization platforms are emerging as well. JModelica.org, which is an open-source platform for both simulation and optimization of Modelica models, has been used in this work for the EDP simulation and optimization. Next section will give an overview of it, see also [22] and [23].

5.2.1 JModelica.org

The main focus of the Modelica language has been on simulation. However, as tools for this purpose have matured, focus has also been directed towards using the models for optimization. Integration tools that can interface several design tools (analysis, simulation and optimization tools) have emerged as well as simulation tools with limited optimization packages as add-ons. Optimization often requires high-level model details, including derivatives and sparsity structures, as well as a skilled user as the original optimization problem often needs to be transcribed in some way.

To circumvent some of these problems it is beneficial if the optimization problem is stated in a high-level fashion, which can directly be associated with the mathematical formulation and the physics is represented in the model. One software package that supports this is JModelica.org, which is an open-source platform for both simulation and optimization of Modelica models. JModelica.org extends the Modelica language with Optimica constructs that allows for straightforward formulation of optimization problems. The software package contains a compiler and supporting functions to compile a Modelica model and translate it to a dynamic optimization problem that is transcribed into a non-linear optimization program using the method of collocation. JModelica.org uses the automatic differentiation tool CasADi [24], giving the possibility to provide symbolic Jacobian and Hessian matrices. The resulting non-linear program is solved by the Interior Point Optimizer (IPOPT), see [25]. Thus, JModelica.org bridges the gap between simulation models and state-of-the-art optimization algorithms.

Using models for optimization puts certain restrictions and demands on the models. An important demand is that all model equations should be C^2-continuous (twice continuously differentiable). It is also often very helpful if the model variables are scaled such that they all have the same magnitude and that they have minimum and maximum values declared. The Modelica language contains attributes for all variables for providing these properties. The models used in the EDP optimization all obey the above and a more detailed explanation can be found in [9].

A more detailed description of JModelica.org and its capabilities in simulation and optimization can be found in [22] and [23].

6 Optimization Formulations

6.1 Plant Economics

The goal of the both the UCP and the EDP optimization is to satisfy the customer heat demand with as low economic cost as possible. The plant economics is therefore an important part of the problem formulation in both of the optimization problems.

In this project constant fuel and electricity prices have been assumed. Start-up costs are also considered, as well as variable and fixed maintenance costs for the production units. Pumping costs have been neglected.

Prices and costs normalized with the electricity price are presented in Table 13 and Table 14. The fuel cost for the AFA is negative because the unit incinerates waste.

Table 13: Fuel and sell prices.

Tabell 13. Bränsle- och försäljningspriser.

Unit	Fuel price
KVV	0.52
Husbyborg	1.31
AFA	-1.21

Type	Sell price
Electricity	1
Heat	0.61

Table 14: Maintenance costs.

Tabell 14. Underhållskostnader.

Unit	Maintenance cost fixed	Maintenance cost variable
KVV	10,6	0,021
Husbyborg	2,02	0,015

6.2 Discrete Optimization

In this section, the UCP optimization formulation is presented. The problem results in a mixed-integer linear program (MILP). The status variables of the solution to this problem are used as inputs to the EDP. The formulation follows the lines of [26], which can be used for a more detailed explanation. The formulation was also used in [9].

6.2.1 Degrees of Freedom and Inputs

The UCP optimization problem contains two types of variables; continuous and binary. The continuous variables can take any value between its minimum and maximum constraint values, while the binary may only take the values 0 or 1. The two types of variables have in common that they are both discontinuous in time due to the sampled nature of the UCP.

For a heat and electricity producing unit i, the following variables are available as decision variables:

- $Q_i[t]$ – Heat produced by the unit.
- $P_{el}[t]$ – Electricity produced by the unit (only available for KVV)
- $status_i[t]$ – Status of the production unit (on/off). 1 when on, 0 otherwise.

The produced heat and electricity are continuous variables, while the status variable is a binary variable.

The units are constrained to follow certain start-up and stop sequences, modeled as piecewise linear sequences. Two additional decision variables are needed for these sequences:

- $start_i[t]$– variable indicating the time point for when a unit is beginning its starting sequence by being equal to 1 at that sample, 0 otherwise.
- $stop_i[t]$– variable indicating the time point for when a unit is beginning its stopping sequence by being equal to 1 at that sample, 0 otherwise.

These variables are communicated to the EDP optimization problem at the integration stage of the production planning optimization.

The accumulator, as modeled in Section 4.1.3, has the heat flow $Q_{acc}[t]$ to and from it as decision variable.

6.2.2 Optimization Problem

6.2.2.1 Cost Function

The cost function expresses the profit of the heat and electricity production, taking into account heat and electricity sell prices as well as fuel prices, unit start costs and other production costs. The instantaneous profit $R[t]$ can be approximated as follows:

$$R[t] = P_{el}[t]p_{el} + Q[t]p_Q - \sum_{\substack{i=all\ production \\ units}} (U_{fuel,i}[t]p_i + U_{fuel,i}[t]p_{varcost,i} + status_i[t]p_{fixcost,i} + start_i[t]s_i) - \sum_{\substack{i=all\ production \\ units}} \gamma_i|\Delta Q_i[t]|$$

The first two terms contain the electricity price $p_{el}$and heat price p_Q and the total production of electricity P_{el} and by customers received heat Q, and thus provide the revenue.

The first sum contains the production costs, of which there are four. The first is the fuel cost where p_i is the fuel price, the second is a variable maintenance cost depending on the amount of fuel used, the third is a fixed maintenance cost for running the unit and the fourth and last term is the start-up cost where s_i is the price for one start-up.

The second sum is a penalty on change in heat production as $\Delta Q_i[t] = Q_i[t] - Q_i[t - h]$. This sum removes undesirable rapid changes of unit loads and constitutes the only part that is not directly coupled with plant economics. However, this term is very small part of the total cost function. The absolute value function is implemented using linear functions.

The total approximated profit to be maximized at optimization is thus

$$J = h \sum_{t=t_0}^{t_f} R[t],$$

where t_0 and t_f are start and end time of the optimization interval, respectively.

6.2.2.2 Constraints

The major constraints of the UCP are presented here while a more detailed and mathematical description can be found in [26].

The production units have minimum and maximum production capacities. Note that the minimum capacity limit is greater than 0, since there is always a base load on the boiler. For all production units, except for the KVV, the constraints are

$$Q_{i,min} \leq Q_i[t] \leq Q_{i,max},$$

while for the KVV, the heat and electricity production is confined inside the polytope defined in Section 4.1.1.1.

Each unit has a minimum and maximum heat production change, which is formulated as

$$\Delta Q_{i,min} \leq \Delta Q_i[t] \leq \Delta Q_{i,max}.$$

The change $\Delta Q_i[t]$ is the same change used in the cost function to remove undesirable rapid load changes.

Heat units may take a long time to start and stop and must follow specific start-up and stop trajectories. Thus, these constraints are important to include in the optimization of the production plan. For this, each unit has a predefined start-up trajectory $Q_{i,start}[t]$ and predefined stop trajectory $Q_{i,stop}[t]$ which are used in constraints as

$$\begin{aligned} Q_i[t] &= Q_{i,start}[t], \quad t \in [t_{start}, t_{start} + t_{startdelay}] \\ Q_i[t] &= Q_{i,stop}[t], \quad t \in \left[t_{stop}, t_{stop} + t_{stopdelay}\right], \end{aligned}$$

where $t_{startdelay}$ and $t_{stopdelay}$ are the durations of the constraints. The formulations of the start-up and stop constraints involves the $start_i[t]$ and $stop_i[t]$ variables and the details can be found in [26].

The accumulator has minimum and maximum limits on how much energy it can hold and also the rate of which it may change. These are formulated as

$$E_{acc,min} \leq E_{acc}[t] \leq E_{acc,max}$$
$$Q_{acc,min} \leq Q_{acc}[t] \leq Q_{acc,max}.$$

Additionally, to prevent the accumulator to be emptied at final time, a constraint is set as

$$E_{acc}(t_f) \geq E_{acc}(t_0).$$

It is reasonable not to empty the accumulator at final time as the accumulator is the heat buffer of the system that should be used at unforeseen events.

The constraint on fulfilling customer heat demand depends on how the production and consumption is distributed. The main characteristics of this constraint are that the accumulator heat and total produced heat should cover the customer demands and possible heat losses in pipes. This can be formulated as

$$Q_{acc}[t] + \sum_{\substack{i=all\,prod.\\units}} Q_i[t] \geq \sum_{\substack{i=all\\cust.}} C_i[t+t_i] + \sum_{\substack{i=all\,supply\\pipes}} PL_i^s[t+t_i^s] \sum_{\substack{i=all\,return\\pipes}} PL_i^r[t+t_i^r]$$

where $C_i[t+t_i]$, $PL_i^s[t+t_i^s]$ and $PL_i^r[t+t_i^r]$ is the customer demand, supply pipe heat loss and return pipe heat loss, respectively, time shifted for correction for pipe delays.

6.2.3 Stochastic Methodology

In a real world problem such as the production planning for a district heating network there are always uncertainties present. Data such as the customer heat demand or the electricity price cannot be known exactly in advance and predictions with some degree of uncertainty must therefore be used. This makes it valuable to include stochastic aspects in the optimization formulation.

It is natural to include the stochastic formulation to the UCP, as it is here the long term decisions concerning the statuses of units are made. These plans needs to be made in a way so that the customer demand can be fulfilled in all future scenarios. Having included the distribution of future scenarios into the formulation, the resulting stochastic UCP can be solved using the same technique as the deterministic problem.

In this study the uncertainty is included in the optimization formulation using a two-stage stochastic programming formulation. The optimization time is then split into two stages. The outcome in the first stage is known, while one or more signals are uncertain during the second stage. The probability distributions of the unknown variables are discretized into a number of possible outcomes and the discrete probability distribution is assumed to be known. For signals such as the heat demand this is a reasonable approach as the accuracy of a prediction concerning the nearest future often is quite high, but the certainty decreases for longer horizons.

For the stochastic cases, the following three different methods are implemented:

- **Perfect information approach:** Separate plans are derived for each outcome, given that the outcomes are known in advance. In other words, it is assumed that at the start of stage 1, the planner already knows the outcome in stage 2 and the profit is maximized for each scenario separately.
- **Stochastic programming:** Plans maximizing the expected profit are derived, under the assumption that the outcome of stage 2 is unknown during stage 1. This means that the same stage 1 plan must be used for all stage 2 scenarios. This method is well-known in the mathematical optimization community. For an application to the unit commitment problem, see [27].
- **Wait-and-see**: Implemented through the following steps:
    a. Stages 1 and 2 are optimized for the expected outcome in stage 2, computed using the maximum demand and expected price in each time step.
    b. The solution to stage 1 is fixed.
    c. Stage 2 is re-optimized separately for each scenario, simulating a robust approach where the maximum demand and the expected price are used to plan, and re-planning is done when the outcome in stage 2 becomes known.

6.3 Continuous Optimization

In this section the continuous time optimization formulation used for solving the EDP is presented.

6.3.1 Degrees of Freedom

Depending on the specific case considered, a number of the following model inputs are present in the EDP model:

- U_{KVV}, KVV fuel load
- U_{HVC}, HVC fuel load
- $U_{Husbyborg}$, Husbyborg fuel load
- U_{Acc}, Accumulator pump speed

These inputs will however not be used as decision variables in the optimization formulations. Instead their derivatives will be used, which is achieved by introducing equations of the form

$$U_j(t) = \int_t \dot{U}_j(t)dt$$

into the model. This construction allows the formulation of constraints on input signal derivatives in a simple way.

6.3.2 Optimization Problem

6.3.2.1 Cost Function

Similarly to the UCP formulation, the cost function is based on the economic costs and incomes from running the district heating network. However, as the model is running in continuous time, this is implemented as an integral, rather than a sum. The definition of the instantaneous profit $R(t)$ is very similar to the EDP formulation in section 6.2.2.1:

$$R(t) = P_{el}(t)p_{el} + Q(t)p_Q - \sum_{\substack{i=all\ production \\ units}} \left(U_{fuel,i}(t)p_i + U_{fuel,i}(t)p_{varcost,i} + status_i(t)p_{fixcost,i}\right),$$

with the same denominations as in the UCP function, but for continuous time. The main difference is that the start-up cost has been removed as the start-up of a unit is pre-defined in the UCP formulation. The load change penalty is furthermore handled differently, as explained next.

In order to achieve a well-behaved optimization problem, it necessary to introduce additional terms to the cost function, penalizing the input derivatives. This is typically done by adding weighted quadratic terms for all input derivatives, according to

$$\dot{W}(t) = \sum_{\substack{i=all\ model \\ inputs}} q_{\dot{U}_i} \dot{U}_i^{\,2}$$

The weights $q_{\dot{U}_i}$ are chosen so that each derivative only contributes to a small degree to the total cost function. The need for this type of added cost is explained in more detail in [9].

The cost function to be minimized is now formulated in the following way:

$$J = \int_{t_0}^{t_f} \left(\dot{W}(t) - R(t)\right) dt,$$

which is the continuous counterpart to the cost function in the UCP, except for the start-up costs and the penalties for load variations and input signal changes.

6.3.2.2 Constraints

A main benefit of the detailed physics-based EDP model is that constraints can be set on variables present in reality. These constraints can be based on the limits of physical components in the network. One such example is the maximum flow in the network set by the maximum flow rate by the pumps. Relevant constraints are summarized in Table 15.

All production units have constraints on their load change rates as explained in the UCP section. This can be seen as the inertia of the boiler or operational constraints.

The accumulator has two types of constraints. The first one is limits on its pump capacity, limiting the rate of which the accumulator may be charged or discharged. Secondly, a constraint is set on the contained energy at optimization end time. This is needed as the optimization results would otherwise be an emptied accumulator at end time. There is furthermore a constraint on the change rate of the accumulator pump. This is a decision variable in the optimization formulation and the main purpose of the constraint is to obtain a more well-behaved optimization problem and consequently an improved optimization convergence.

The KVV, which is the most detailed physics-based model, has constraints mainly on the condensers VK1 and VK2. The steam pressure shall not be lower than 0.1bar and the inlet water temperature should be above 39°C.

It is necessary to provide water with a temperature of at least 55^{o}C in order to avoid growth of bacteria [16]. In order to have a safety margin a minimum constraint of 57^{o}C is therefore used on the district water inlet temperature for each customer.
Constraints on pipe energy of the form $E_{pipe}(t_f) \geq E_{pipe}(0)$ have been introduced into the formulation. The main reason for this is to keep a balance between production and consumption combined with heat losses, for each optimization case.

Table 15: Continuous optimization constraints.

Tabell 15. Optimeringsbivillkor i den kontinuerliga optimeringen.

Variable	Min	Max	Comment
Accumulator end energy	From UCP optimization	-	EDP accumulator strategy based on UCP planning
Accumulator mass flow	-600 kg/s	600 kg/s	Pump limitation
Customer inlet temperature	57 C	110 C	Min value is to limit bacterial growth
VK 1/2 steam pressure	0.1 bar		Plant operating conditions
VK 1/2 inlet water temp.	39°C		Plant operating conditions
District heating water mass flow	1000 kg/s	3000 kg/s	Pump limitation
Pipe energy	Pipe energy initial value	-	Balance between production and consumption
$\dot{U}_{Acc}$	-0.02	0.02	Improves optimization convergence
$\dot{U}_{Husbyborg}$	-0.024	0.024	Plant load change limitation
$\dot{U}_{KVV}$	-0.00575	0.00575	Plant load change limitation

For units that change status during the EDP optimization interval the start-up and/or shut-down times and the corresponding trajectories constitute additional constraints. These are described in the integration between UCP and EDP.

6.3.2.3 Pipe Energy Balancing Measures

In this section the specific parts of the EDP optimization formulation that are used to prevent pipe energy unbalances are presented.

For the supply pipes, the energy constraints guarantee that the optimization results do not contain any unbalanced energy flows. In the return pipes the predefined temperature profiles make it harder to impose corresponding constraints whenever the temperatures of start and final time in the optimization do not match, which is the case in case IIIb. This case is presented in detail in the next chapter. However, the fixed temperature profile also limits the possibility to exploit the unphysical modelling, since the mass flow mainly must be determined by the main optimization objectives. The Husbyborg pump mass flow not being a decision variable in case IIIb was decided specifically with this in mind.

6.3.3 Problem Complexity

In this section, data describing the complexity of the EDP is presented in relation to the computational power of the system used to solve the problem.

For the EDP, the most complex optimization case considered is case IIIb, which is presented in detail in section 7.4. The continuous formulation contains 307 variables and 33 states. This system is discretized with an element length of 20 minutes and a 20 hour optimization horizon, creating a nonlinear programming problem containing approximately 69000 variables.

All EDP solving during the project was conducted on a laptop with 8 GB RAM and four 2.6 GHz CPUs. On this system the convergence time for the most complex optimization problem is approximately 5 to 10 minutes. The corresponding number of iterations for the IPOPT solver is between 60 and 150.

At the moment the limiting factor for the complexity of the nonlinear programming problems that can be solved is the amount of memory needed to perform the optimization. It is mainly JModelica and not IPOPT that is using up the memory capacity.

6.4 Integration of Discrete and Continuous Optimization

The optimization horizon for the EDP is limited to a maximum of 20-24 h, depending on the complexity of the considered case. The reason for this is the large amount of memory required to perform the optimization, as explained above. The typical production planning optimization horizon is however several days, which cannot currently be handled through EDP optimization for the entire scenario.

With the separation of the production planning problem into the UCP and the EDP, it is however not needed, or even relevant, to use the EDP for optimization intervals longer than one day. In the EDP model the focus lies on faster dynamics, while the UCP is used to determine the long term plans. Having a longer time scale in the UCP is no problem, due to the significantly simpler modeling.

In order to achieve a useful result both the faster dynamics and the long term planning must be considered. This makes the integration between the two optimization problems very important.

The unit statuses from the UCP are used to construct constraints in the EDP formulation, determining when the production units are started and stopped. This is explained in detail below. The UCP also determines the constraint on the accumulator energy at the end of each EDP optimization. This is an improvement compared to the formulation used in [9], where the accumulator energy constraint always was based on the initial energy level.

6.4.1 Start-up Trajectories

In scenarios where production units change status, the start-up and/or shut-down times from the UCP results are included as parameter values into the EDP formulation.

In order to make a unit start and follow a start-up trajectory at a specific time, constraints on the unit load are introduced in the EDP. An example of this implementation can be seen in Figure 21, where the upper plot shows the result from the UCP optimization and the lower plot the result from the EDP optimization.

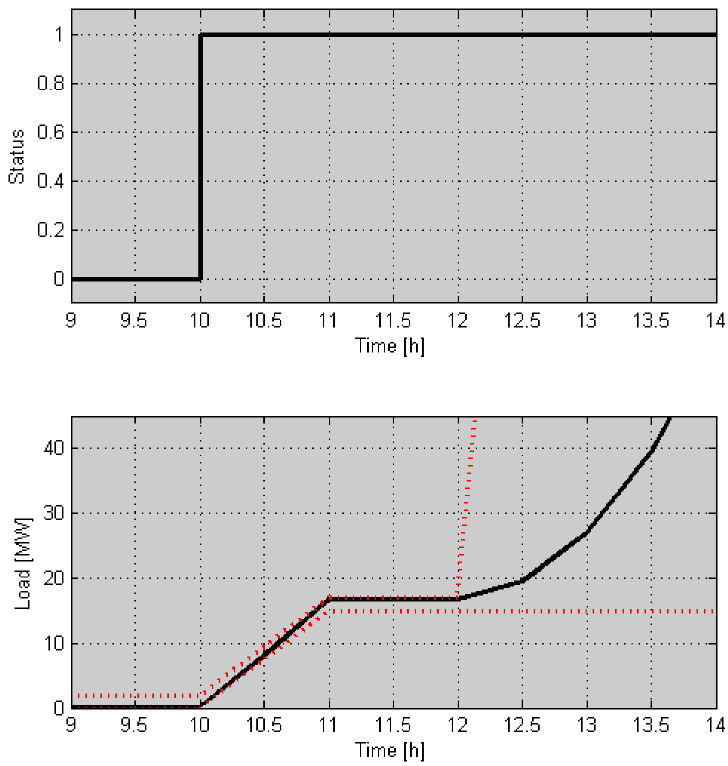

Figure 21: Husbyborg status (upper) and the corresponding load constraints and resulting load (lower).

Figur 21. Status (överst) och motsvarande bivillkor för last tillsammans med lastprofil (nederst).

When a unit is running, the load constraints are defined by the normal minimum and maximum load of that unit. When a unit is turned off, the constraints force the load to be between 0% and 2% of maximal load. For the start-ups and shut-downs, piecewise linear trajectories corresponding to the specific start-up and shut-down time of each unit, have been constructed. The load constraints are formulated so that the load of each unit follows these trajectories within 1 percent of full load.

A tabular implementation is used to construct these constraints in the Modelica optimization model. In order to keep the model C^2-continuous the table is constructed using C^2-continuous approximations of min and max functions.

6.4.2 Modeling Differences and Feasibility

Due to the difference in model complexity there are inevitable differences between the UCP and EDP models. The two main differences are:

- **KVV modeling:** The polyhedral representation of the KVV in the UCP is an approximation of the UCP model of the plant. The polyhedral representation of the feasible volume in the electricity production-heat production-return temperature space will contain some errors as the polyhedron must either cut away some extreme point or include some unfeasible points. In order to avoid unfeasibility problems a conservative approach was decided upon where the UCP polyhedral in general is slightly smaller than the actual KVV operating region.
- **Pipe modeling:** Due to the difference in pipe modeling between the UCP and EDP, explained in section 4.2.2, the production profile for a certain load profile is different in the two models. In the UCP the production profile is completely determined by the load, as both the delay time and the heat loss are fixed. This is not the case in the EDP where the separate mass flows and temperature modeling introduces more freedom.

The conservative approach when determining the polyhedron and the extra freedom in production introduced by the more detailed pipe description should in most cases guarantee that the input from the UCP results in a feasible EDP. The only situation when a feasibility problem could arise would be when the fixed delay in the UCP spreads and therefore lowers a load peak more than what is possible in the EDP, due to e.g. a higher mass flow or temperature constraints. Such a scenario has not yet been found in this project.

7 Optimization examples

7.1 Overview of Cases

The following section contains five optimization cases. The first two are used for analyzing important characteristics of the implementation of the district heating network. Case III is a more realistic scenario where the integration between the UCP and EDP is included and the optimization runs over several days. These cases are based on deterministic data, whereas the last two include stochastic programming. In these cases the load profile is uncertain and in case V also the electricity price is nondeterministic.

The first three cases are divided into two sub-cases, to clarify the effects of increasing the complexity of the models. The characteristics of each case are summarized in Table 16.

Table 16: Production planning case characteristics.

Tabell 16. Karakteristik för produktionsplaneringsfall.

Case	Number of customers	Delay	Heat losses	Number of prod. units	Accumulator	Optimization time
Ia	1	No	No	1	No	24 h
Ib	1	Yes	Yes	1	No	24 h
IIa	3	Yes	No	1	No	24 h
IIb	3	Yes	Yes	1	No	24 h
IIIa	1	No	No	3	Yes	96 h
IIIb	3	Yes	No	3	Yes	96 h
IV	1	Yes	No	3	No	24 h
V	1	Yes	No	2	No	

7.1.1 Common Settings

For the cases I, II and IV the optimization time is 24 hours, starting at midnight. In cases I and II the same time horizon is used for the EDP and the UCP and there is no integration between the optimizations. For the stochastic cases only an UCP optimization is considered. In case III the optimization time horizon is four days for the UCP. This time is split into 18 and 15 hour sections for the respective subcase in the EDP. The starting point for each optimization is then the final point of the previous one. In all cases the element length is 30 minutes in the UCP and 20 minutes in the EDP.

In the first four cases the customer heat demand is defined by a base load with two peaks added at 07:00 and 19:00, as seen in Figure 22. This is roughly corresponding to the heat demand expected from a residential area on a week day. In case I and II the base load is constant, while a linear increment is added in case III, making the load increase from one day to the next. In the stochastic cases the load profile is uncertain and therefore split into several independent scenarios.

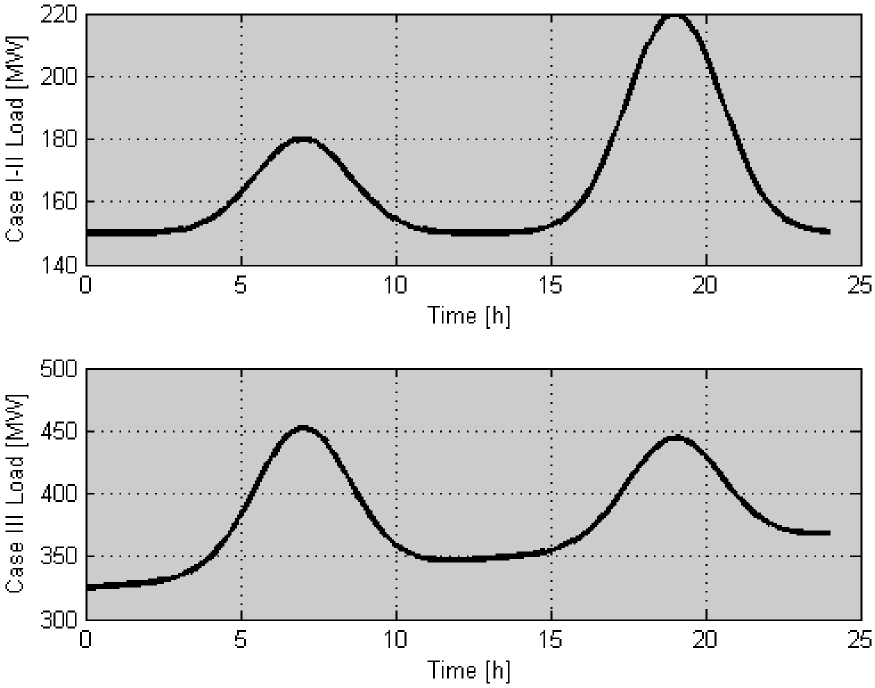

Figure 22: Load profiles for the deterministic optimization cases.

Figur 22. Lastprofil för de deterministiska optimeringsfallen.

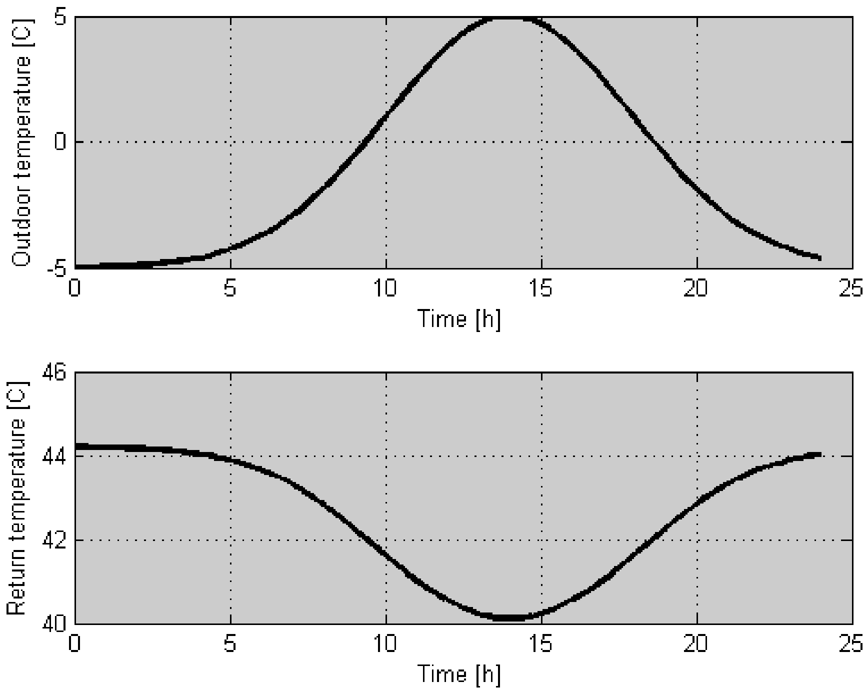

Figure 23: Outdoor temperature profile and corresponding customer return temperature.

Figur 23. Utomhustemperatur och motsvarande returtemperatur från kunderna.

In the deterministic cases an outdoor temperature model is used with a base temperature of -5 ^{o} C and a peak temperature of 5 ^{o} C at 14.00. This temperature is used for determining the heat loss in pipes and the temperature of the water leaving the customers. The temperature profiles are displayed in Figure 23.

7.1.1.1 Topology

The network topology for each case is illustrated in Figure 24. The method for deriving the customer distribution from data, used in case II and III, is explained in detail in Chapter 4.3. The distribution of the heat demand between the customers and the pipe delays for these cases are described in Table 17. Case II and III have slightly different customer distributions due to the addition of the Husbyborg oil boiler in case III.

Table 17: Customer distribution data. The exact pipe delay cannot be used in the optimization models due to the discretization.

Tabell 17. Kundfördelningsdata. Den exakta rörfördröjningen kan inte användas i optimeringsmodellerna på grund av diskretiseringen.

Case	Customer	Exact pipe delay [h]	Implemented Pipe delay [h]	Customer demand [%]
Ib	1		2	100
II	1	0.4455	0.5	10
II	2	0.9616	1	70
II	3	2.3036	2.5	20
IIIb	1	0.5815	0.5	10
IIIb	2	0.9429	1	57
IIIb	3	2.1427	2	33
IV	1		0.5	100

7.1.2 Result presentation

The most important optimization results are presented in plots for each deterministic sub-case. In a summary of what these plots contain is presented. The profit in each case is normalized with the profit of case Ia. For each individual case additional plots are added in order to highlight the specific features of that case.

Table 18: Summary of plot content for the different optimization cases.

Tabell 18. Sammanfattning av innehåll i plottar för de olika optimeringsfallen.

Plot description	UCP	EDP
Heat demand and production	Yes	Yes
KVV electricity production	Yes	Yes
Normalized profit	Yes	Yes
District heating water mass flows and temperatures	No	Yes

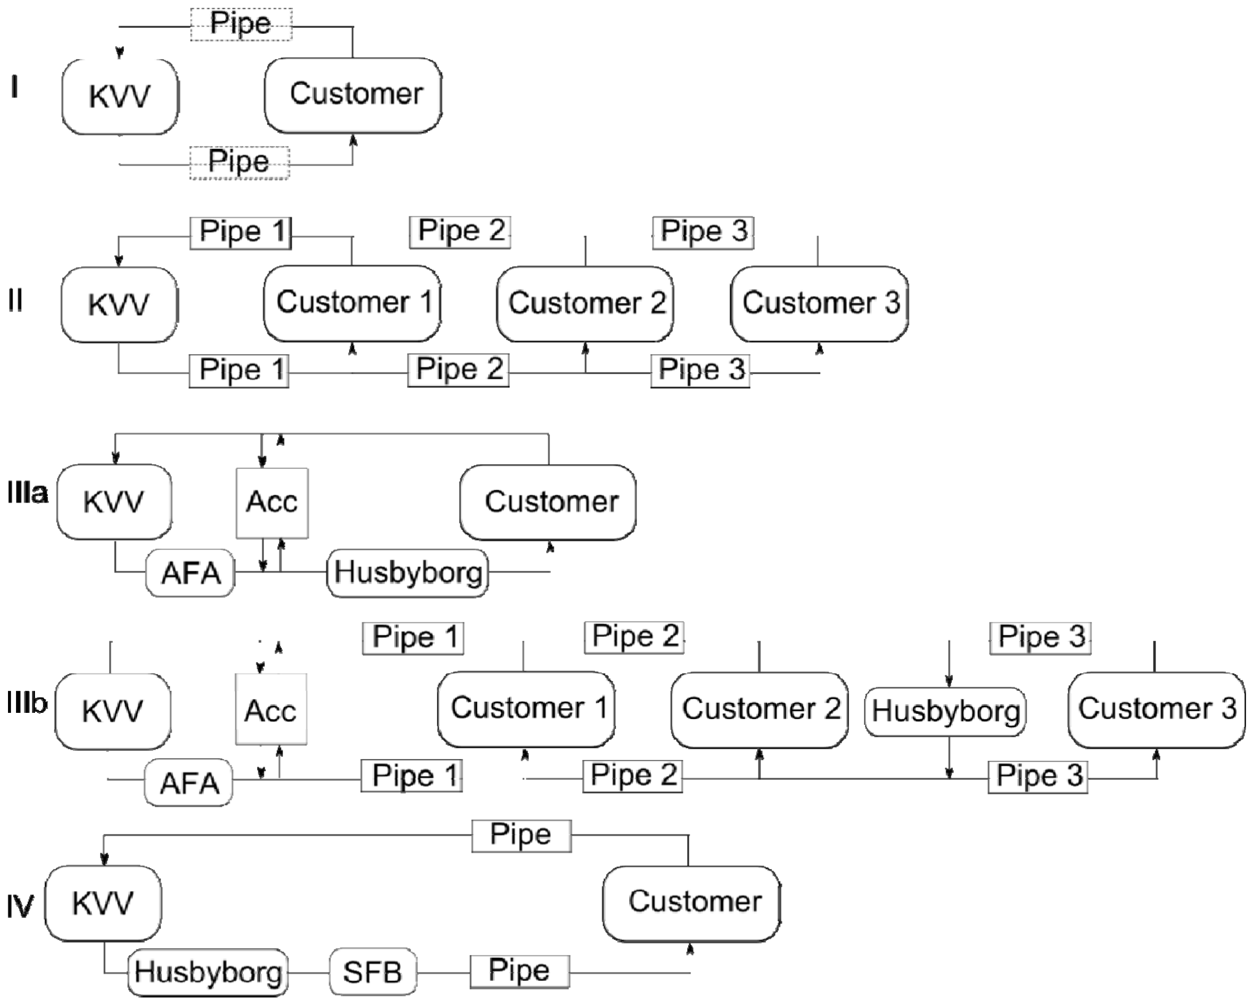

Figure 24: Network topology for the optimization cases.

Figur 24. Topologi för distributionsnätet i optimeringsfallen.

7.2 Case I: Point-wise Network

This case is divided into two sub-cases; case Ia, with no pipes included in the model and case Ib, where pipes corresponding to approximately two hour delays have been implemented. The results are visualized in Figure 25 - Figure 33.

7.2.1 Decision Variables

The only decision variable in this case is the KVV fuel load.

7.2.2 Discrete Optimization Features

For case Ia, it can be observed in Figure 25 that the heat production follows the customers demand exactly. Based on this heat production, the KVV will produce the maximal amount of electricity according to the polyhedron at every time step. The electricity production is displayed in Figure 27. The difference between the lines represents the inaccuracy of the polyhedral KVV model compared to the detailed model used in the EDP. The higher UCP profit in Figure 29 is a consequence of the different electricity productions, as this generates a larger income.

For case Ib more heat is generated at earlier time in order to satisfy the customer demand. In Figure 26 we observe that to compensate for the heat loss, the KVV produces more heat, and to compensate the delay, the heat is produced at an earlier step in time. The return temperature dependency of the pipe heat losses can be seen in the decline in production between hour 10 and 14, corresponding to the lowering of the return temperature which happens during these hours.

7.2.3 Continuous Optimization Features

For the EDP optimization, the characteristics of the optimization results are more dependent on which sub-case that is considered than in the UCP. In Case Ia, similarly to the UCP results, the KVV production is identical to the heat demand during most of the optimization time. The small differences between the signals, which can be seen during the first hour in Figure 25, is an initial transient during which the supply temperature is optimized. The close similarity is a result of the implemented customer model, which decides the mass flow in the system so that its heat demand is fulfilled. No pipes or other means of storing energy in the network, except for a small control volume, limits the degrees of freedom in the optimization result significantly.

In Figure 31 the mass flow rate and customer inlet temperature in the district heating network, and the VK1 steam pressure, are illustrated for case Ia. From these plots it is visible that the mass flow is kept as high as possible, with the water mass flow, the VK1 steam pressure and customer inlet temperature all being limiting constraints. The drop in steam pressure between hour 6 and 20 is explained by the customer return temperature profile, which has a minimum at hour 14. The possibility to optimize the supply temperature, rather than deriving it from the outdoor temperature is an important feature of the implemented optimization method.

For case Ib the EDP optimization results clearly differ from the UCP results, as can be seen in Figure 26.

The heat production is in general smaller in the EDP results compared to the UCP. This depends on the different pipe heat loss models. In the UCP the supply temperature is assumed to be 90^{o}C, whereas the EDP has a supply temperature slightly below 60^{o}C. This results in a lower temperature difference between the water and the surrounding temperature in the EDP, which also reduces the heat loss. The difference is mostly visible for the first part of the optimization time, logically coinciding with a lower water supply temperature in the EDP results.

Usage of heat stored in the pipes is also present in case Ib. This is the explanation for the reduced heat production during the first part of the optimization interval. The optimization algorithm chooses this solution in order to reduce the supply temperature which maximizes the electricity production. The increased heat production in the end of the optimization time is the results of constraints on the pipe energy, stating that the final energy level must be the same as at the start of the optimization.

An interesting difference between the optimization solutions can be seen in the KVV production for the peak at hour 7. The explanation for the two separate production peaks around the heat demand peak is a combination of different phenomena exploited by the optimization. The dip in production leading up to the demand peak is explained by usage of heat stored in the pipes. This lowers the customer supply temperature, forcing an increase in mass flow. The mass flow constraint then gets active, forcing an increase in KVV load. The following load profile for the KVV is a function of the maximizing the mass flow, causing mostly the VK1 steam pressure constraint to be active.

The pipe delay time in the model together with the theoretical delay time based on the mass flow are displayed in Figure 33. The maximal difference is 20 minutes, coinciding with the period when the mass flow is maximized.

The total profit for case Ib, displayed in Figure 30, is lower than for case Ia. The reason for this is the pipe heat losses in case Ib, which forces a greater heat production and therefore increases the fuel cost. The difference in revenue between the UCP and EDP depends on the difference in heat loss explained above.

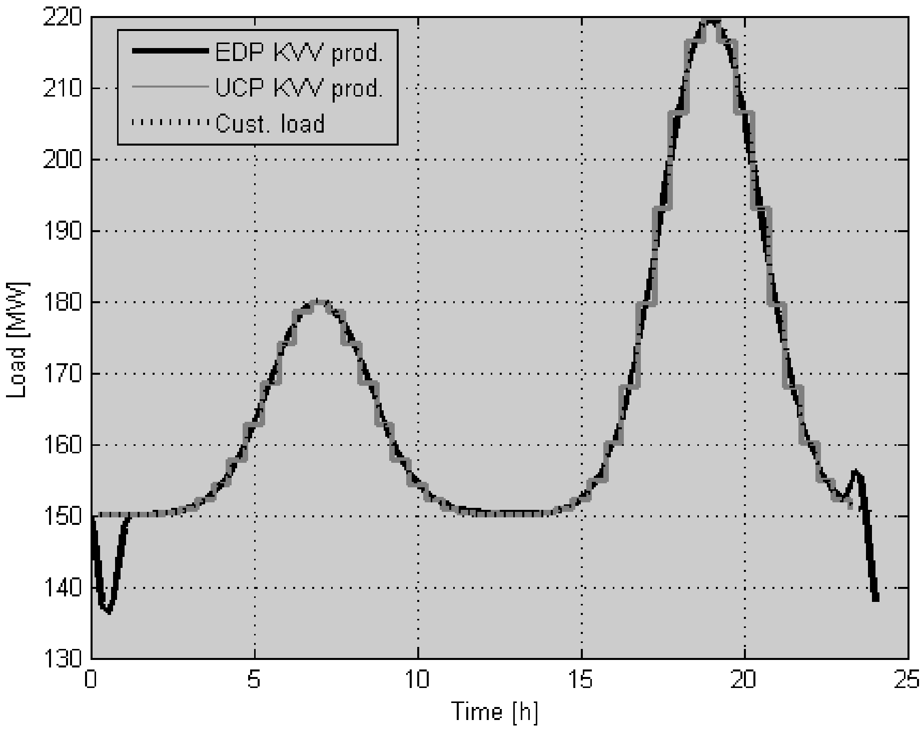

Figure 25: Case Ia KVV and customer load for the discrete and continuous optimization.

Figur 25. Fall Ia, last för KVV och kund för den diskreta och den kontinuerliga optimeringen.

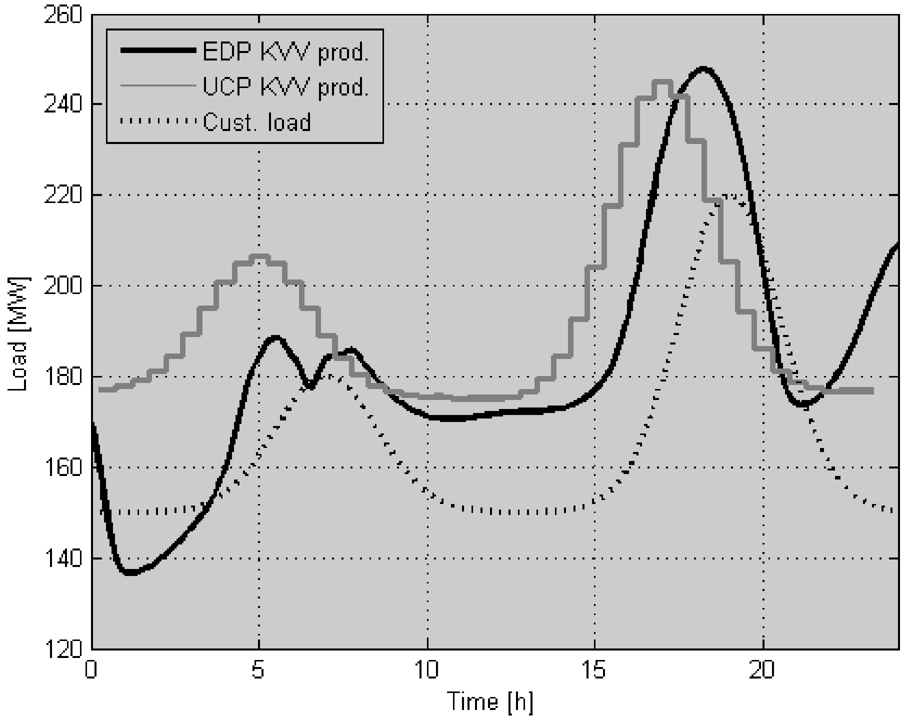

Figure 26: Case Ib KVV and customer load for the discrete and continuous optimization.

Figur 26. Fall 1b, last för KVV och kund för den diskreta och den kontinuerliga optimeringen.

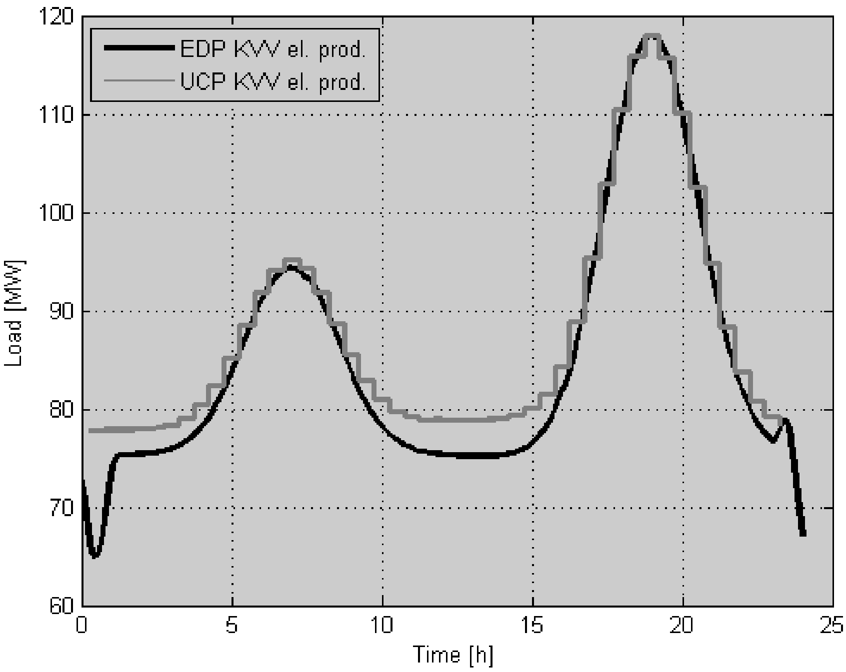

Figure 27: Case Ia electricity production for the discrete and continuous optimization.

Figur 27. Fall Ia, elproduktion för den diskreta och den kontinuerliga optimeringen.

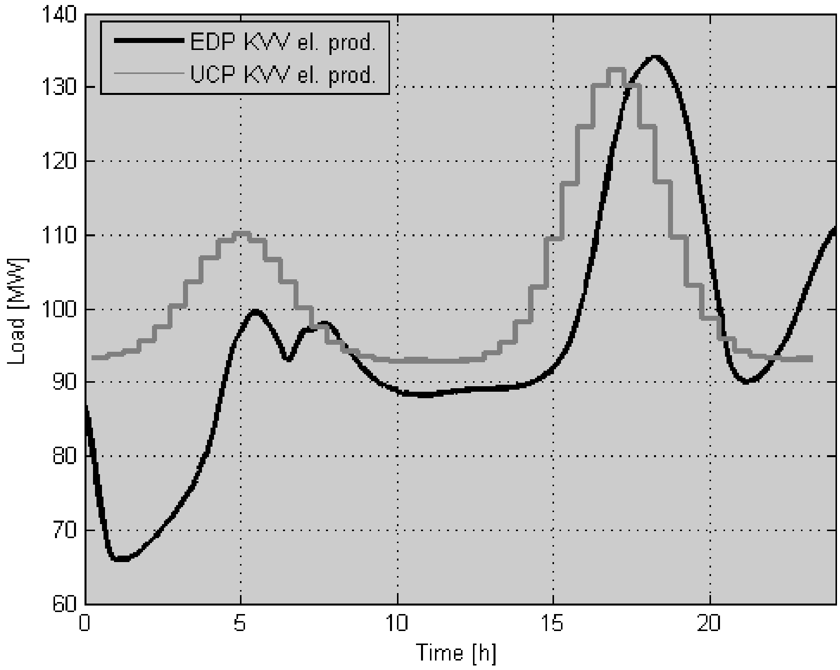

Figure 28: Case Ib electricity production for the discrete and continuous optimization.

Figur 28. Fall Ib, elproduktion för den diskreta och den kontinuerliga optimeringen.

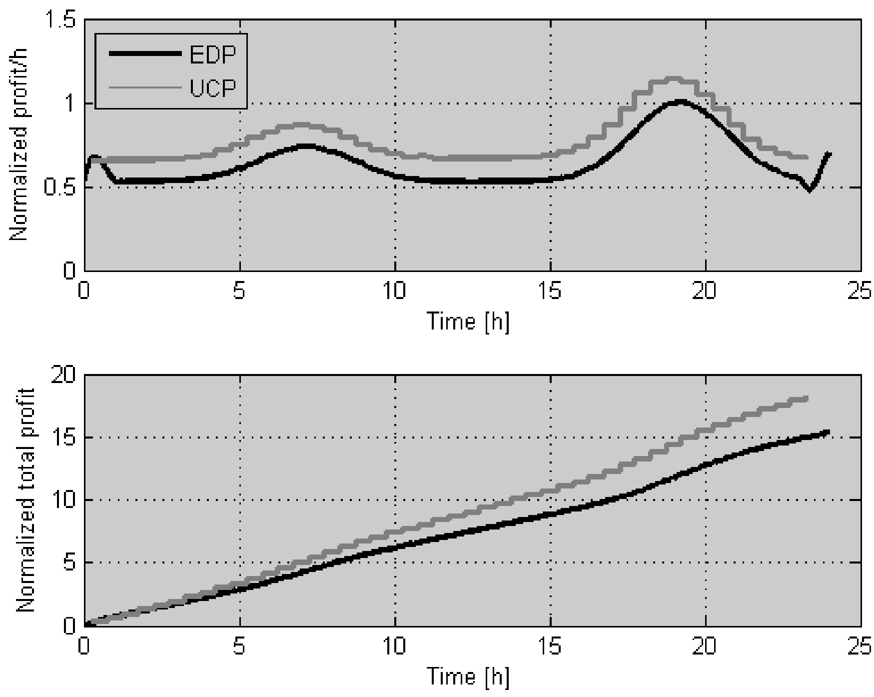

Figure 29: Case Ia normalized profit for the discrete and continuous optimization.

Figur 29. Fall Ia, normaliserad vinst för den diskreta och den kontinuerliga optimeringen.

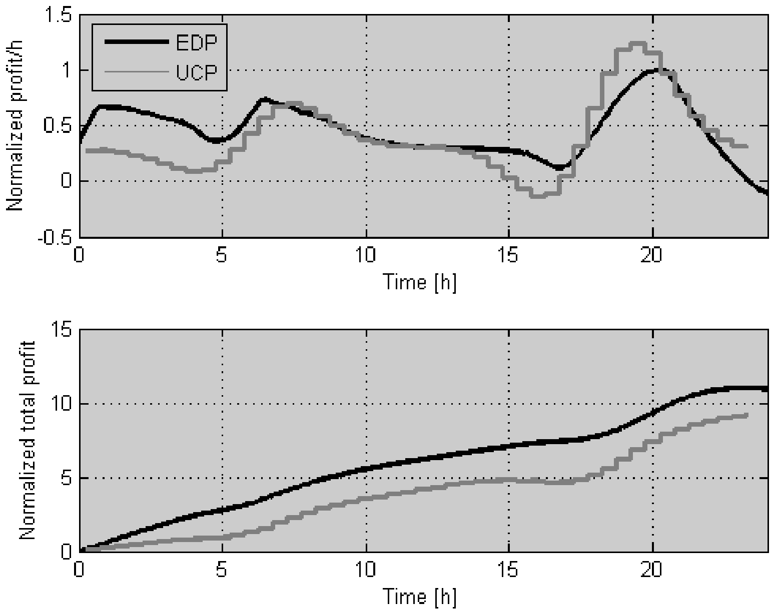

Figure 30: Case Ib normalized profit for the discrete and continuous optimization.

Figur 30. Fall Ib, normalized vinst för den diskreta och den kontinuerliga optimeringen.

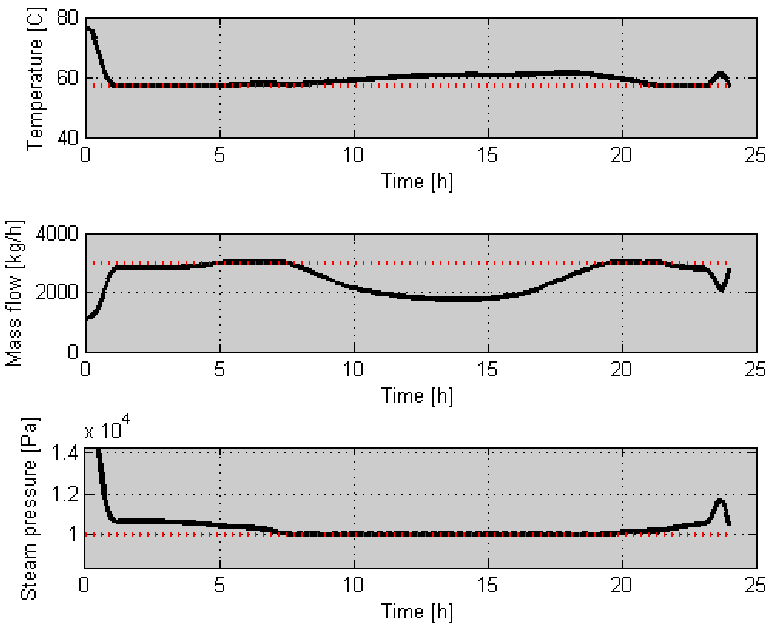

Figure 31: Case Ia supply water temperature and mass flow and VK1 steam pressure and corresponding constraints in the continuous optimization.

Figur 31. Fall Ia, framledningstemperatur och -massflöde och VK1 ångtryck tillsammans med motsvarande bivillkor i den kontinuerliga optimeringen.

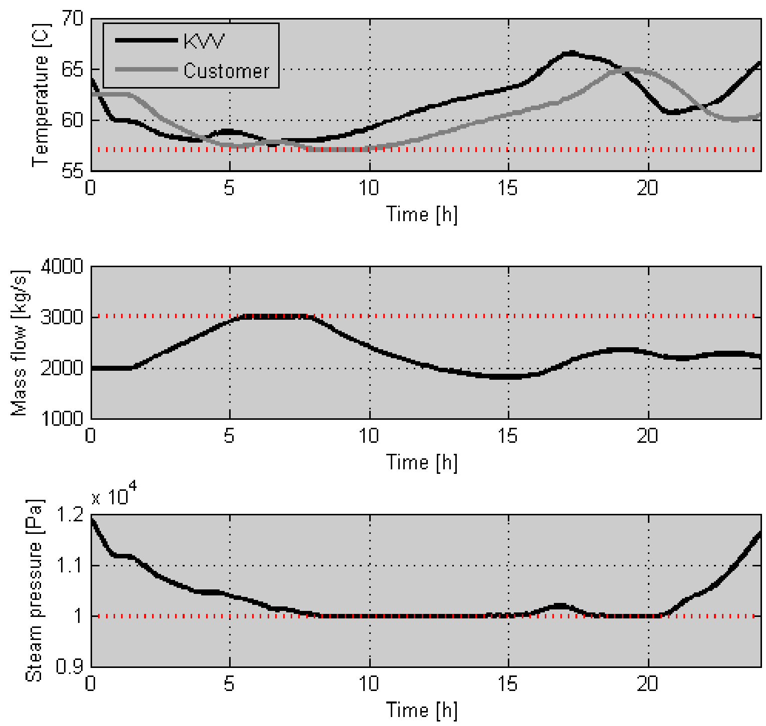

Figure 32: Case Ib supply water temperature and mass flow and VK1 steam pressure in the continuous optimization.

Figur 32. Fall Ib, framledningstemperatur och -massflöde och VK1 ångtryck tillsammans med motsvarande bivillkor i den kontinuerliga optimeringen.

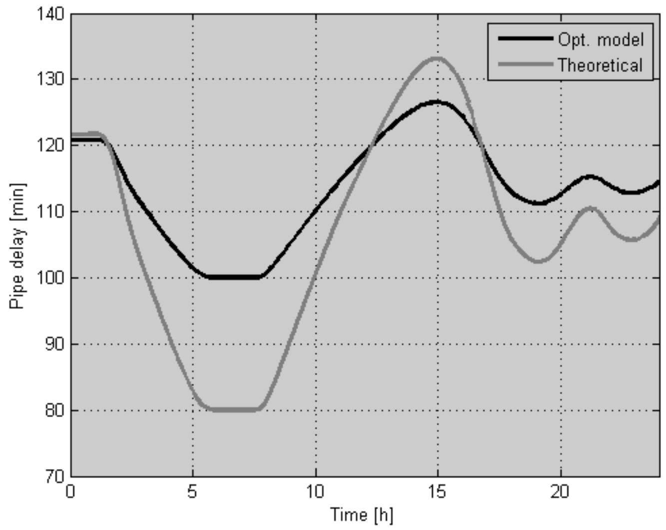

Figure 33: Case Ib pipe delay time in the continuous optimization compared to theoretical value.

Figur 33. Fall Ib, tidsfördröjning i rör i den kontinuerliga optimeringen tillsammans med teoretiskt värde.

7.3 Case II: Distributed Network

This case is similar to case I, but the load demand is divided between three different customers rather than one. The customers are distributed in a network by placing them in parallel with pipes in between, as described in section 4.4. Two sub-cases are considered; case IIa without and case IIb with pipe heat loses. The results are visualized in Figure 34 - Figure 42.

7.3.1 Decision Variables

The only decision variable in this case is the KVV fuel load.

7.3.2 Discrete Optimization Features

For both sub-cases the difference in delay time between the customers results in wider and lower peaks. The heat losses in case IIb introduce an offset in heat production compared to case IIa, in order to fulfill the customer heat demand. A comparison between case Ia and case IIa reveals that the distributed customer model reduces the production peak by approximately 11 MW in the UCP results.

7.3.3 Continuous Optimization features

Like in case I the EDP optimization results differ from the UCP results due to the possibilities to store energy in the network and use mass flow changes to handle changes in customer demand. This happens for both sub-cases.

In case IIa it is clearly visible from Figure 34 that heat stored in the pipes is used to fulfill the customer heat demand during the first part of the optimization time. This lowers the customer supply temperature which is economically beneficial. To handle the second peak the opposite happens, as the pipes are heated in advance in order to reduce the maximal load for the KVV.

In Figure 35 the load profile for case IIb is displayed. An important effect of the pipe delay can be observed here. During the first peak the customer inlet temperature is almost constant and the load peak is handled by increasing the mass flow, as can be seen in Figure 41. As the water is incompressible this means that the increased load needs to be handled without any delay in the KVV. For the second peak however, only increasing the mass flow would violate the VK1 steam pressure constraint, so the temperature is also increased. When the temperature increases to handle the peak, the delay time is visible when comparing the KVV production and the customer load. However as the mass flow also increases the delay in the EDP is less pronounced than in the UCP results.

Similarly to the first case, the mass flow in case II is in general maximized, with the supply temperature at the furthest customer station and the VK1 steam pressure being limiting constraints, as can be seen in Figure 40 and Figure 41. An important difference introduced by the distributed network in case IIb is that temperature constraints for individual customers can be handled. For the implemented network model the farthest customer will have the active constraint as the heat losses are greatest for this customer. In case IIb this is customer 3.

A clear difference between the subcases can be observed regarding the handling of the hour 19 load peak. In case IIa the production peak is significantly wider and lower compared to case IIb. The reason why this kind of production profile is not possible when pipe heat losses are present is that the heat losses in the return lines are lowering the temperature of the water entering the KVV. This reduces the VK1 steam pressure, making the constraint active for a higher production level.

The pipe heat losses in case IIb are similar to those in case Ib, as can be seen in Figure 42. The outdoor and return temperature profiles are the main contributors to the characteristics of these signals. The maximum in outdoor temperature results in the minimum in return temperature at hour 14, which creates a minimal temperature difference in the pipes at this hour and also a minimum in heat losses. The higher supply temperature during the second peak is also influencing the heat loss profiles.

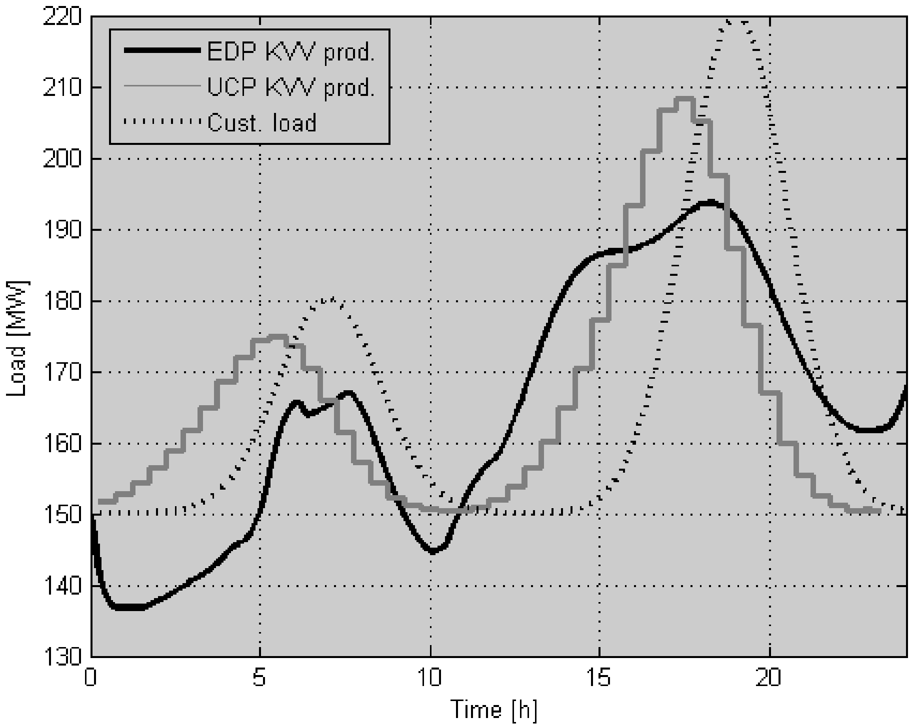

Figure 34: Case IIa KVV heat production and total customer heat demand for the discrete and continuous optimization.

Figur 34. Fall IIa, värmeproduktion för KVV och sammanlagt värmebehov för den diskreta och den kontinuerliga optimeringen.

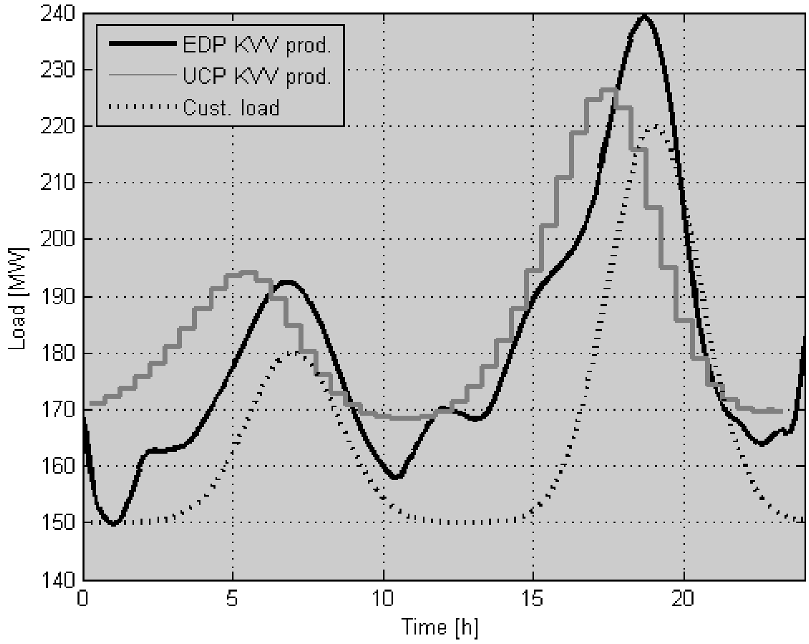

Figure 35: Case IIb KVV heat production and total customer heat demand for the discrete and continuous optimization.

Figur 35. Fall IIb, värmeproduktion för KVV och sammanlagt värmebehov för den diskreta och den kontinuerliga optimeringen.

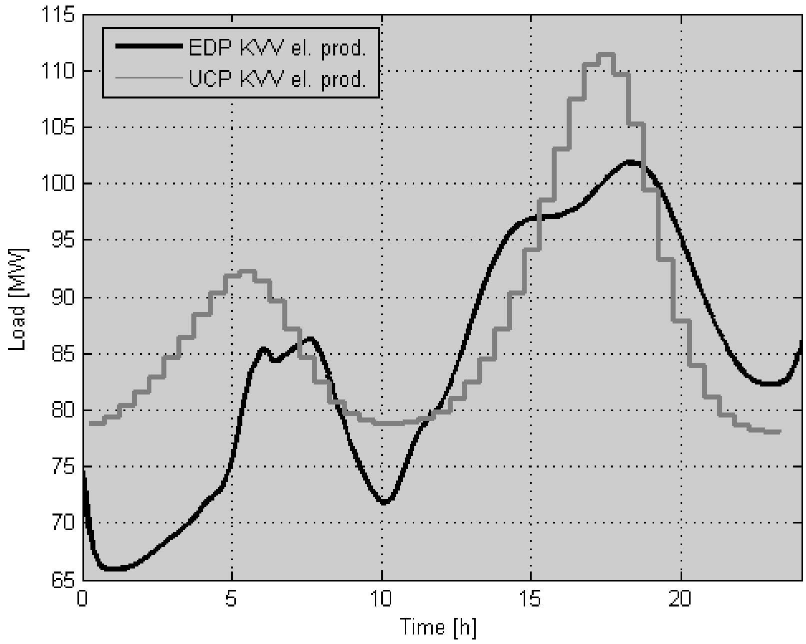

Figure 36: Case IIa electricity production for the discrete and continuous optimization.

Figur 36. Fall IIa, elproduktion för den diskreta och den kontinuerliga optimeringen.

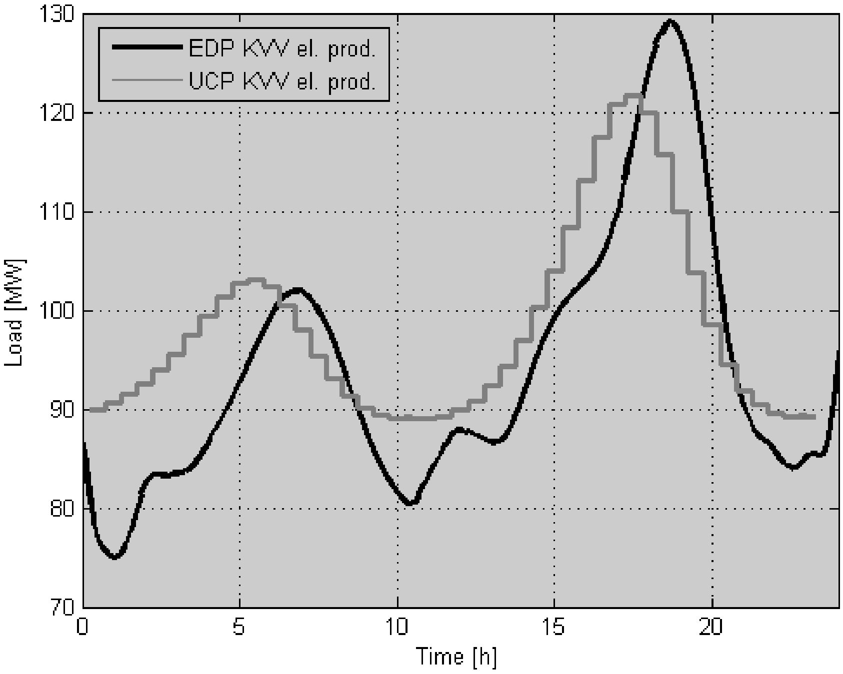

Figure 37: Case IIb Electricity production for the discrete and continuous optimization.

Figur 37. Fall IIb, elproduktion för den diskreta och den kontinuerliga optimeringen.

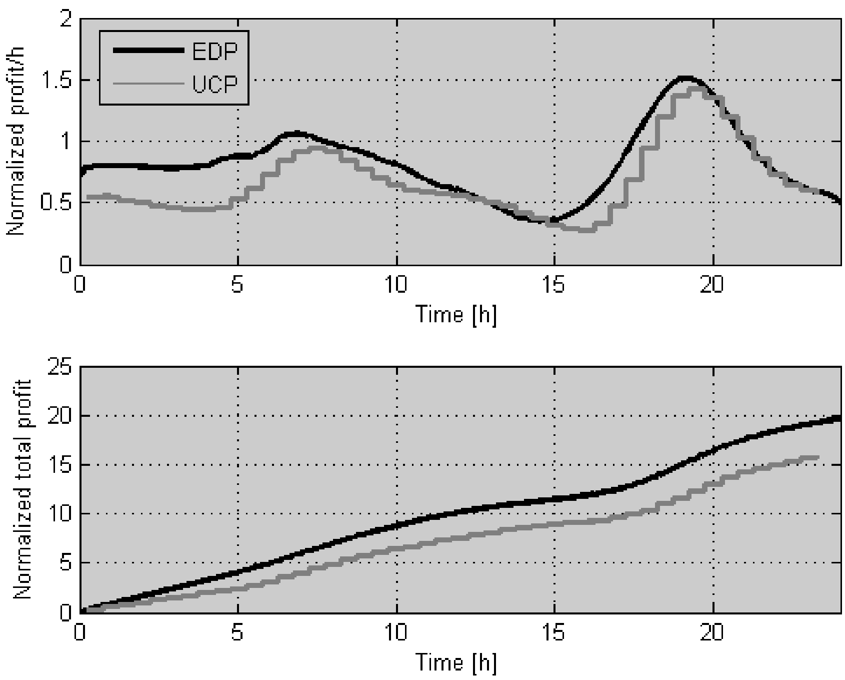

Figure 38: Case IIa normalized profit for the discrete and continuous optimization.

Figur 38. Fall IIa, normaliserad vinst för den diskreta och den kontinuerliga optimeringen.

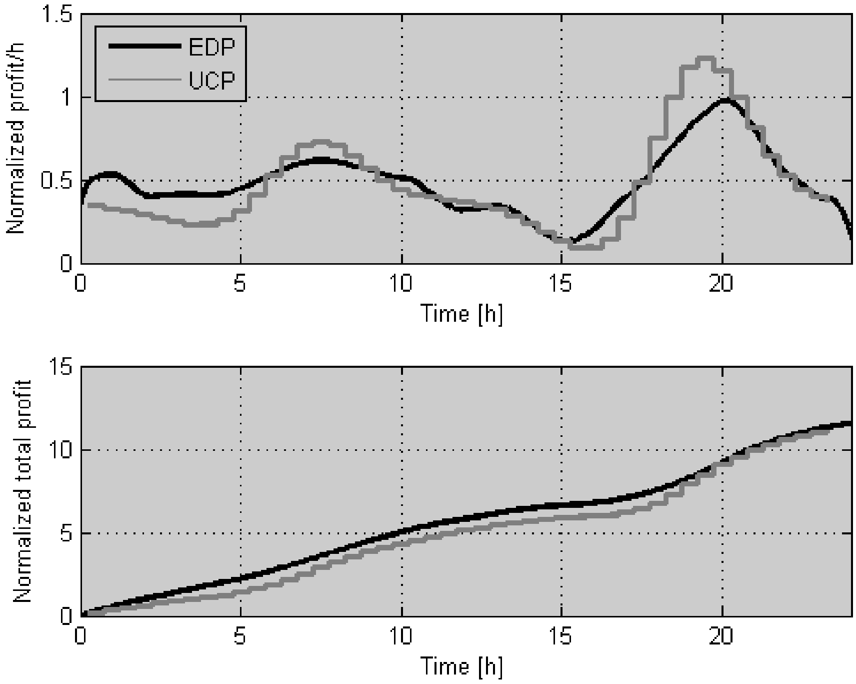

Figure 39: Case IIb normalized profit for the discrete and continuous optimization.

Figur 39. Fall IIa, normaliserad vinst för den diskreta och den kontinuerliga optimeringen.

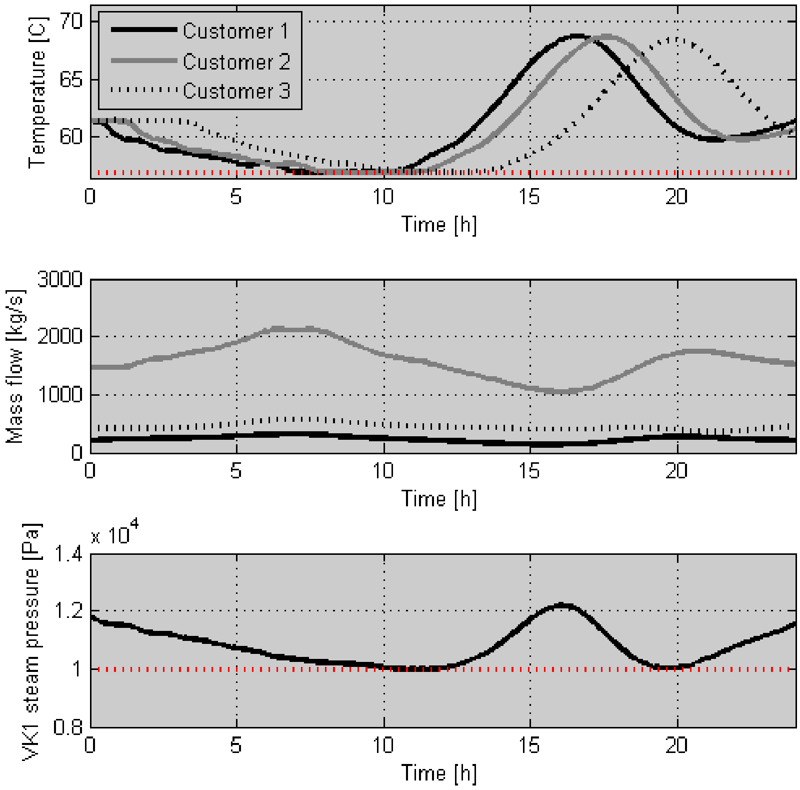

Figure 40: Case IIa customer inlet water temperatures, customer mass flows and condenser steam pressure in the continuous optimization.

Figur 40. Fall IIa, massflöde och framledningstemperatur för respektive kund, samt kondensortryck för den kontinuerliga optimeringen.

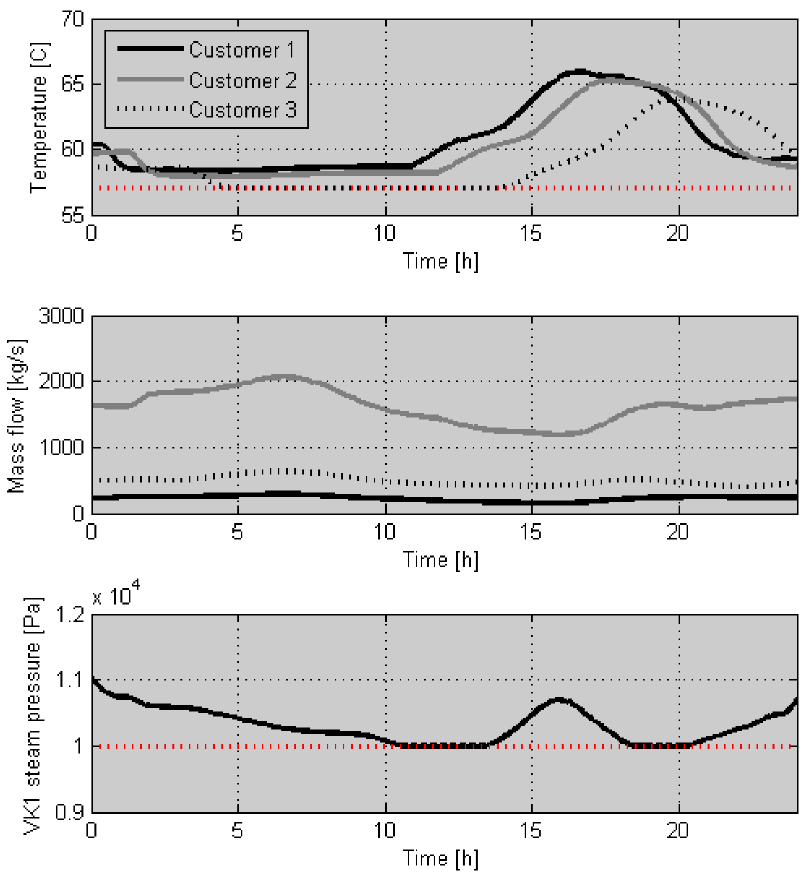

Figure 41: Case IIb customer inlet water temperatures, customer mass flows and condenser steam pressure in the continuous optimization.

Figur 41. Fall IIb, massflöde och framledningstemperatur för respektive kund, samt kondensortryck för den kontinuerliga optimeringen.

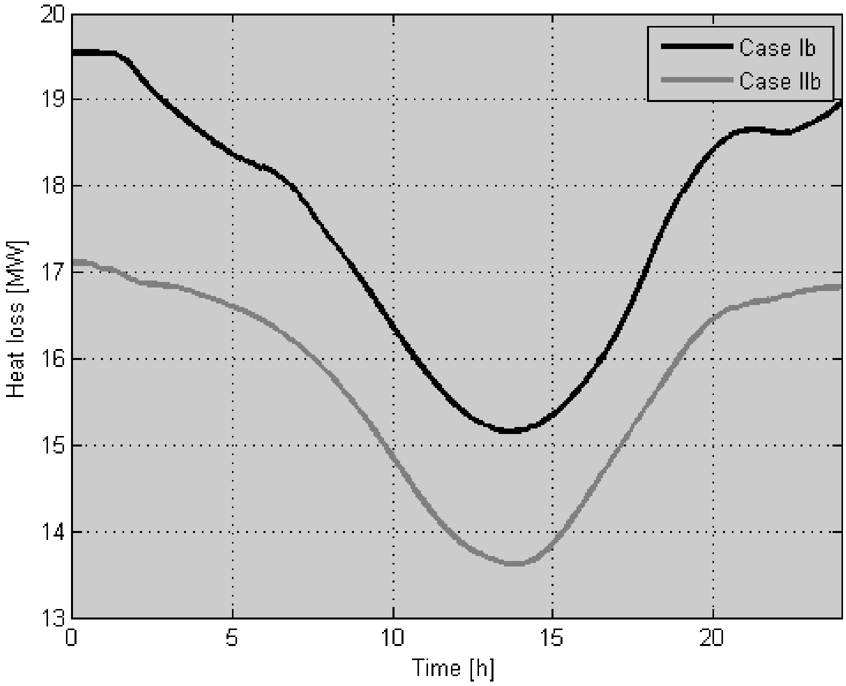

Figure 42: Total heat losses for case Ib and IIb in the continuous optimization.

Figur 42. Sammanlagda värmeförluster för fall Ib och IIb i den kontinuerliga optimeringen.

7.4 Case III: Optimization over Several Days

The following optimization case was created in order to investigate the performance of the implemented production planning strategy in a more complex and realistic scenario. Therefore a UCP optimization horizon over four days is considered and three heat production units are used. Two of these units, the KVV and the AFA, are run at all times, while the status of the last one, the Husbyborg oil boiler, is a decision variable in the UCP. The EDP is solved for five separate 20 hour windows for case IIIa, and six windows for case IIIb, with the Husbyborg starting time and accumulator energy level as input. However, the last hours of each EDP solution are disregarded in the production planning to avoid transient behaviors at the end of the optimization interval. In case IIIa the last two hours are disregarded, while the last five hours are disregarded in case IIIb. The starting point of each 20 hour optimization scenario is correspondingly determined by the status in the network 18 and 15 hours into the previous optimization result for the respective case. The disregarded section needs to be longer in the second sub-case due to the slower dynamics introduced by the pipes.

As oil is much more expensive than the peat which is used in the KVV, or the waste used in the AFA, it is desirable to delay the start-up of the oil boiler for as long as possible. The start-up cost and maintenance costs are additional reasons for avoiding starting extra units.

The two sub-cases have different network topologies. In case IIIa a point-wise net with no pipes added to the model is used. This corresponds to the typical network implementation in today's standard methods for unit scheduling.

In case IIIb a distributed network with three customers and pipes between is used. Here the Husbyborg plant is placed the same distance away from Boländerna as the second customer, in order to mimic the topology of the real network. This unit can only be used to heat customer 3, as it otherwise would be necessary to introduce reversible pipe mass flows. For this reason an extra constraint is added in the UCP formulation. As the plant is connected in parallel with the customers a fixed mass flow through the unit have been assumed in the EDP formulation in this case.

The models are initialized with the Husbyborg plant turned off, the AFA running at full load and the KVV load being 50%. As there is a negative cost on the AFA fuel this unit is kept running at maximal load throughout the optimization time.

The results are visualized in Figure 44 - Figure 52.

7.4.1 Decision Variables

The decision variables in this case are the KVV fuel load, the Husbyborg fuel load and the accumulator pump speed.

7.4.2 Discrete Optimization Features

In both sub-cases the accumulator is used to handle the peaks in the heat demand profile, as seen in Figure 43 and Figure 44, where the load profiles are displayed and Figure 49 and Figure 50, which display the accumulator usage. As the heat demand is increasing the Husbyborg oil boiler must eventually be started. In the first sub-case this happens during the first peak in day 2, after 30.5 hours, while in case IIIb the start-up happens after 39.5 hours. In this sub-case the distributed customer model decreases the peak at the main production site which is what makes the delay in the start-up of the oil-boiler possible. The difference in start-up time is an important result of the distributed network model as running extra units is expensive.

As the heat demand increases further the load of the oil boiler is increased in steps. As the oil is more expensive than the heat it produces, the increase in load reflects on the hourly profit in Figure 47 and Figure 48, which decreases in corresponding steps.

7.4.3 Continuous Optimization Features

As in previous cases, the finer discretization and more complex modeling in the EDP formulation compared to the UCP result in some differences between the optimization results. The EDP signals are in general less smooth and involve more control actions and transient behaviors. However, the general characteristics are similar for the two solutions, especially for the first sub-case. Features such as optimization of supply temperature and better constraint handling should though guarantee that the quality of the EDP results is higher than the quality of the UCP results.

When the load profiles in Figure 43 and Figure 44 are examined, a difference in the modeling of the KVV between the UCP and the EDP can be observed, as the EDP model allows for a slightly higher heat production at maximal load. This forces a reduced production from the KVV in the EDP during the first days of the optimization, as the accumulator constraints and Husbyborg start-up time are based on a lower maximal load.

In Figure 45 and Figure 46 the electricity productions are visualized. The greater variations in the electricity production in the EDP results compared to the UCP depends on the more detailed KVV and district heating water description in the continuous formulation. The on average higher electricity production in the EDP is a consequence of differences in the KVV modeling.

In Figure 47 and Figure 48 one can see that the profit has a great similarity both between the UCP and EDP and between the two sub-cases. Due to the additional fixed costs of starting and running the Husbyborg oil boiler, and the added fixed income from running the AFA, the relative difference is very small. The start-up cost of the Husbyborg plant is visible in the UCP results as a small dip in the hourly revenue.

In case IIIa the EDP accumulator signal is very similar to the corresponding signal in the UCP, as can be observed in Figure 49. The only difference is that parts of the load peaks are also handled by the production units in the EDP, first by the KVV and later, when the oil boiler is started and the KVV is running at maximum load, the Husbyborg plant. For case IIIb the difference is greater, which is shown in Figure 50. This depends on the addition of the pipes to the model, which increases the degrees of freedom in the EDP compared to the UCP, and the shorter optimization windows in the second sub-case. Having more similar KVV models and improving the implementation of the accumulator energy constraint would therefore be desirable in order to improve the integration.

Like in previous cases the district heating water mass flow is kept close to the maximal constraint. In case IIIa this constraint is present for both the customer and the KVV mass flows, these signals are displayed together with the constraint in Figure 51. In Figure 52 the supply temperatures and mass flow for each customer in case IIIb are showed. As there are additional production units heating the district heating water leaving the KVV, the customer supply temperature is clearly higher in case III compared to previous cases, and the lower customer supply temperature constraint does therefore not get active in either of the sub-cases. In case IIIb the inlet temperature for customer 3 has a different profile than the other customers. The reason for this is the extra mass flow that is coming from the Husbyborg plant. Before the plant is started the temperature is lower as the supply water is mixed with return water due to this mass flow. When the plant is started, it increases the water temperature more than the rest of the supply water, resulting in a higher supply temperature for customer 3 than the other customers.

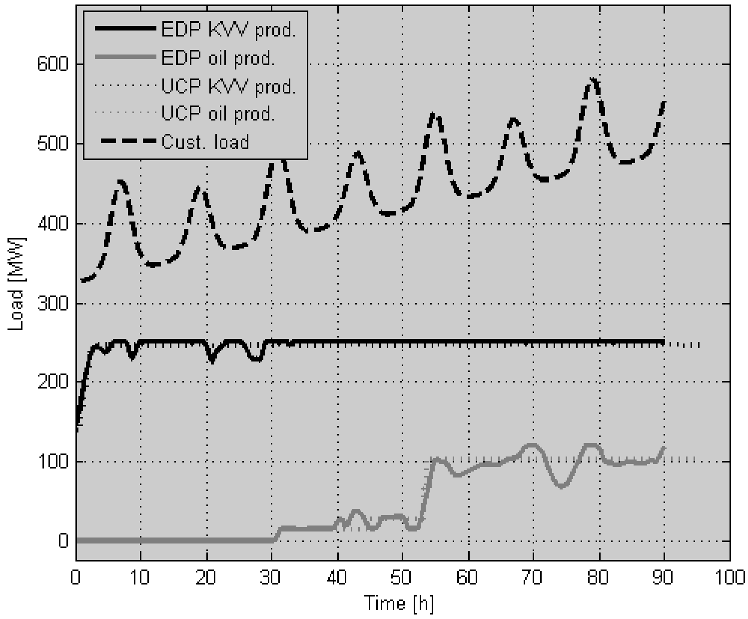

Figure 43: Case IIIa heat production and customer load for the discrete and continuous optimization.

Figur 43. Fall IIIa, värmeproduktion och värmebehov för den diskreta och den kontinuerliga optimeringen.

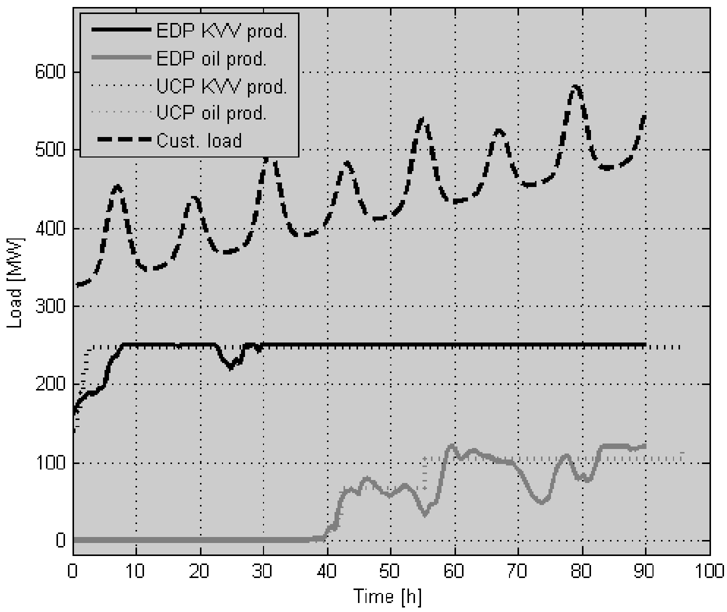

Figure 44: Case IIIb heat production and total customer load for the discrete and continuous optimization.

Figur 44. Fall IIIb, värmeproduktion och sammanlagt värmebehov för den diskreta och den kontinuerliga optimeringen.

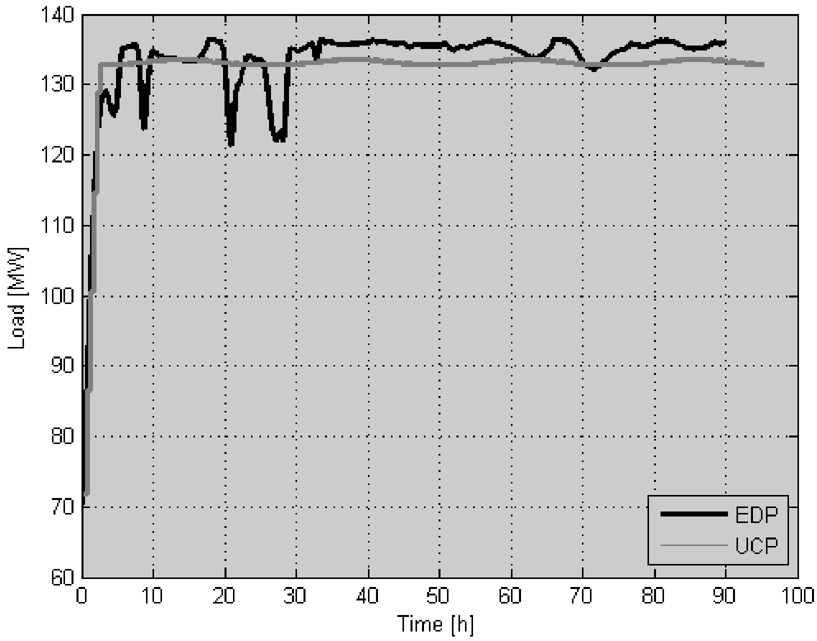

Figure 45: Case IIIa electricity production for the discrete and continuous optimization.

Figur 45. Fall IIIa, elproduktion för den diskreta och den kontinuerliga optimeringen.

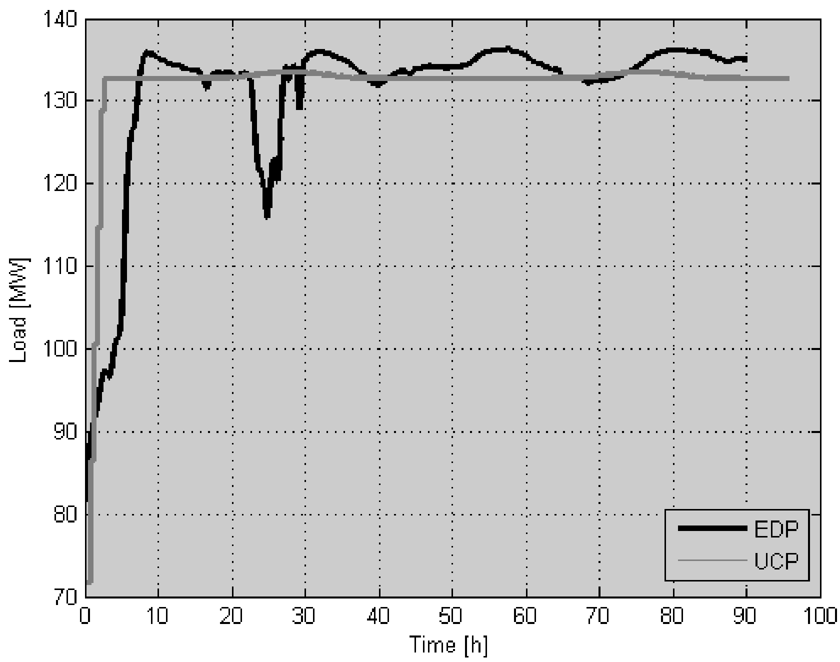

Figure 46: Case IIIb electricity production for the discrete and continuous optimization.

Figur 46. Fall III3, elproduktion för den diskreta och den kontinuerliga optimeringen.

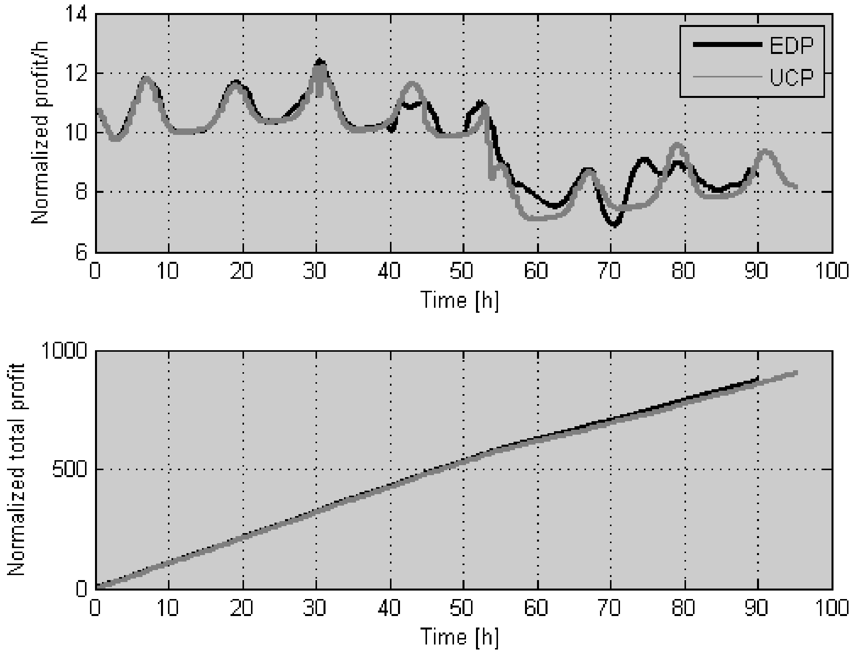

Figure 47: Case IIIa normalized profit for the discrete and continuous optimization.

Figur 47. Fall IIIa, normaliserad vinst för den diskreta och den kontinuerliga optimeringen.

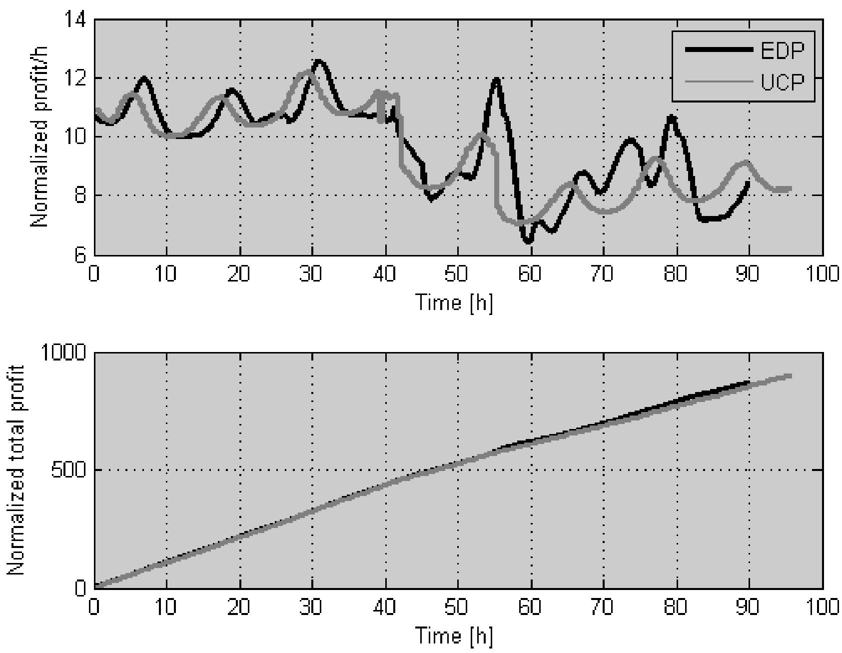

Figure 48: Case IIIb normalized profit for the discrete and continuous optimization.

Figur 48. Fall IIIb, normaliserad vinst för den diskreta och den kontinuerliga optimeringen.

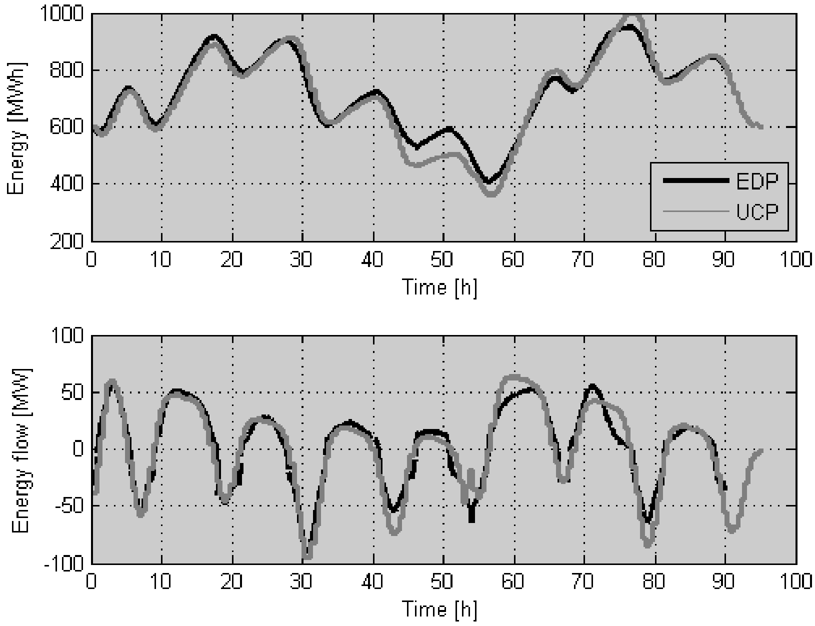

Figure 49: Case IIIa accumulator energy content and flow for the discrete and continuous optimization.

Figur 49. Fall IIa, energiinnehåll och -flöde för ackumulatorn i den diskreta och den kontinuerliga optimeringen.

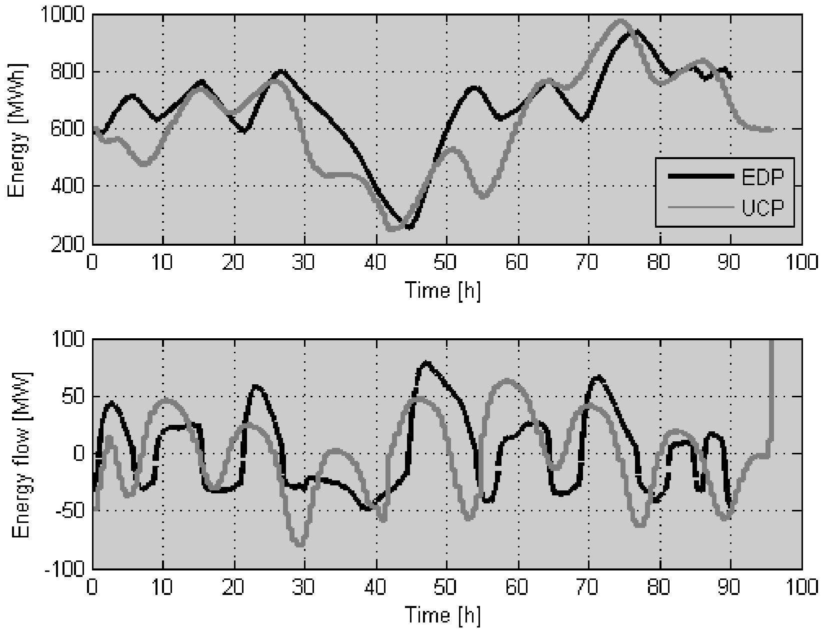

Figure 50: Case IIIb accumulator energy content and flow for the discrete and continuous optimization.

Figur 50. Fall IIb, energiinnehåll och -flöde för ackumulatorn i den diskreta och den kontinuerliga optimeringen.

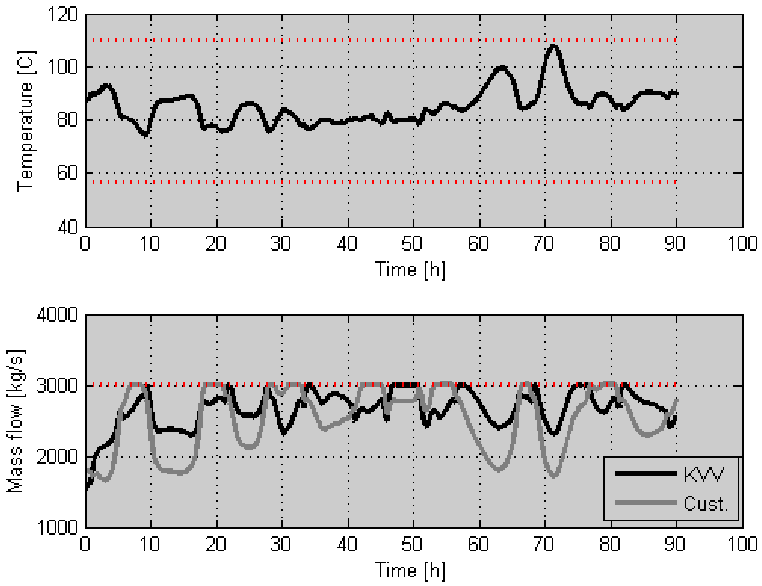

Figure 51: Case IIIa customer supply temperature and mass flows in the continuous optimization.

Figur 51. Fall IIIa, framledningstemperatur och massflöde för kunden i den kontinuerliga optimeringen.

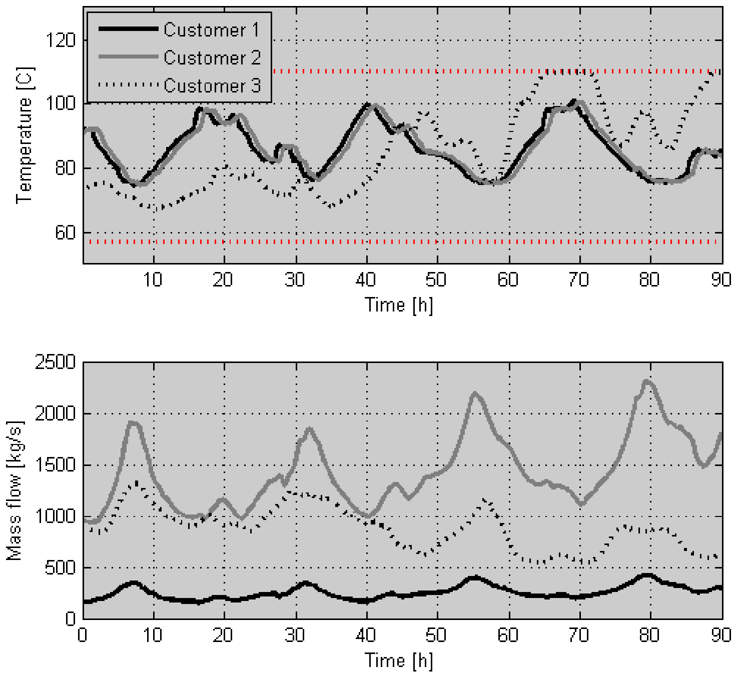

Figure 52: Case IIIb customer supply temperatures compared to the temperature constraints and mass flows in the EDP.

Figur 52. Fall IIb, framledningstemperatur och massflöde för kunderna i den kontinuerliga optimeringen.

7.5 Case IV: Stochastic Programming

In this section an optimization case with an uncertain heat demand profile is investigated. The methods presented in section 6.2.3 are used to solve the production planning problem.

Three production units with different characteristics are present in the setup; the KVV, the Husbyborg oil boiler and a solid fuel boiler (SFB) that is not present in the real Uppsala network. Initially only the KVV is running, but depending on which demand profile that is realized the start-up of additional units might be necessary to fulfill the customer heat demand. In order to highlight the impact of stochastic programming on the UCP, the additional units have been chosen to have clearly different behaviors. The Husbyborg oil boiler has a short start-up time, but the oil makes it expensive to use. The SFB has a much lower fuel cost, but takes significantly longer to start, representing the characteristics of a cold start of the boiler. The specific parameter values of this unit are presented in Table 19.

Table 19: SFB characteristics. The start-up cost is normalized with the corresponding cost for the Husbyborg oil boiler.

Tabell 19. Karakteristik för fastbränslepannan. Uppstartskostnaden är normaliserad mot motsvarande kostnad för Husbyborgsverket.

Parameter	Value
Max capacity [MW]	120
Min capacity [MW]	15
Fuel cost (normalized)	0.52
Fixed cost (normalized)	2.02
Variable cost (normalized)	0.017
Start-up cost (normalized)	2.67
Start-up time [h]	7
Efficiency	0.896

The same type of demand profile is used as in the previous cases, with two demand peaks on a 24 hour optimization time. However, the linear offset used in this case is different as it represents the difference between the scenarios during stage two, which is starting after 12 hours. In order to prevent infeasibility problems, the demand profiles of stage two are however identical for the first time step of this stage. Otherwise the production at the last time step of stage 1 would be impossible to determine as there is a delay of one time step between production and consumption. Four different scenarios are included in the model, with demand profiles and probabilities displayed in Figure 53.

The KVV modeling is simplified compared to the deterministic cases, as only a constant return temperature is considered and the polyhedral representation of the plant is therefore replaced with a polygon.

The heat production for each unit and scenario for the different production planning strategies is presented in Figure 54 - Figure 56.

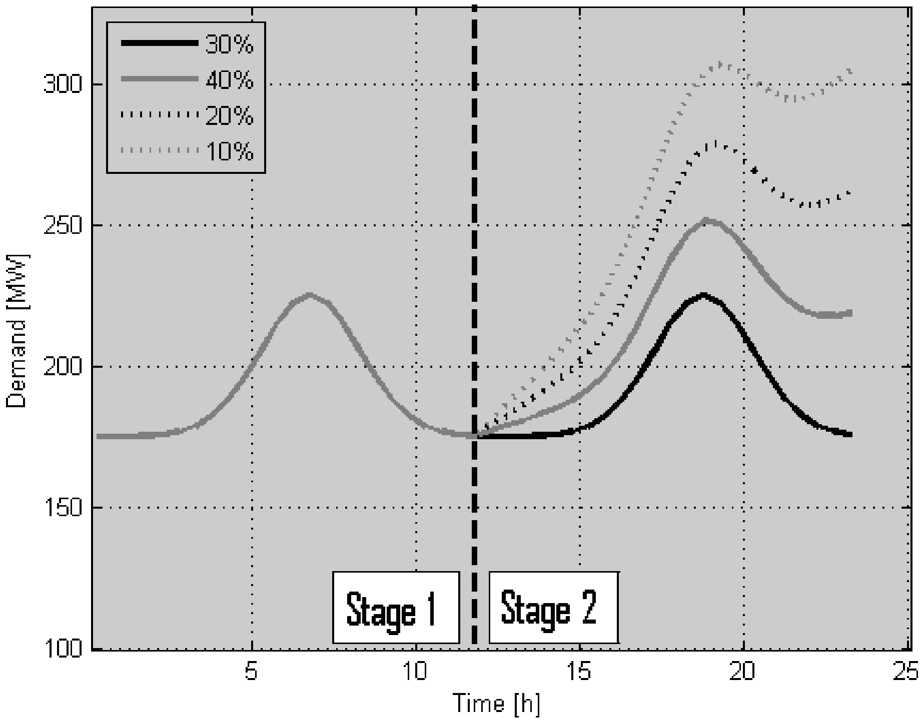

Figure 53: Case IV total customer heat demand profiles during stage 1 and 2.

Figur 53. Fall IV, kundernas totala värmebehov under de två stegen.

7.5.1 Result Features

The total profit for each of the scenarios is presented in Table 20, normalized with the perfect information result.

Table 20: Expected profit from the three planning approaches.

Tabell 20. Förväntad vinst från de tre planeringsmetoderna.

Approach	Expected profit (normalized)
Perfect information	1
Stochastic programming	0.972
Wait-and-see	0.948

As expected the perfect information strategy results in the highest profit. The plans derived for this case are interesting as they provide a reference when the other results are analyzed. In the perfect information strategy no additional units are started for the first scenario and only the Husbyborg oil boiler is started in scenario 2. For scenario 3 and 4 the SFB is used to handle the increased load instead, since the higher load makes it more profitable to run this unit, even though the start-up time is longer. In these scenarios the start-up of the SFB is initialized before the end of phase 1.

When stochastic programming is used to solve the optimization problem no additional unit is started during phase 1. This implies that the results coincide with the perfect information results for scenario 1 and 2. In scenario 3 and 4 the SFB is used to handle the increased heat demand. However, the start-up is delayed compared to the perfect information results, making it necessary to also start the Husbyborg plant. This involves a greater cost compared to the ideal strategy in these scenarios, reducing the total profit.

The wait-and-see approach forces an early start-up of the SFB, as this is the ideal strategy in scenario 4. This is close to the perfect strategy in scenario 3 as well, whereas the perfect results indicate that using the oil boiler instead would be more beneficial in scenario 2. In scenario 1 the biggest drawback of the robust planning is visible, as the SFB is started unnecessarily. Even though the unit is shut down as soon as possible, the start-up reduces the profit in this scenario considerably, as it involves additional costs and forces a reduced heat and electricity production from the KVV.

The results from this stochastic example shows that it is possible to achieve a higher profit by considering the probability of different future scenarios when the production planning strategy is developed, compared to only planning for handling the worst case scenario. This is especially true in situations when there is a choice of which unit to start or stop and the available units have different operating characteristics.

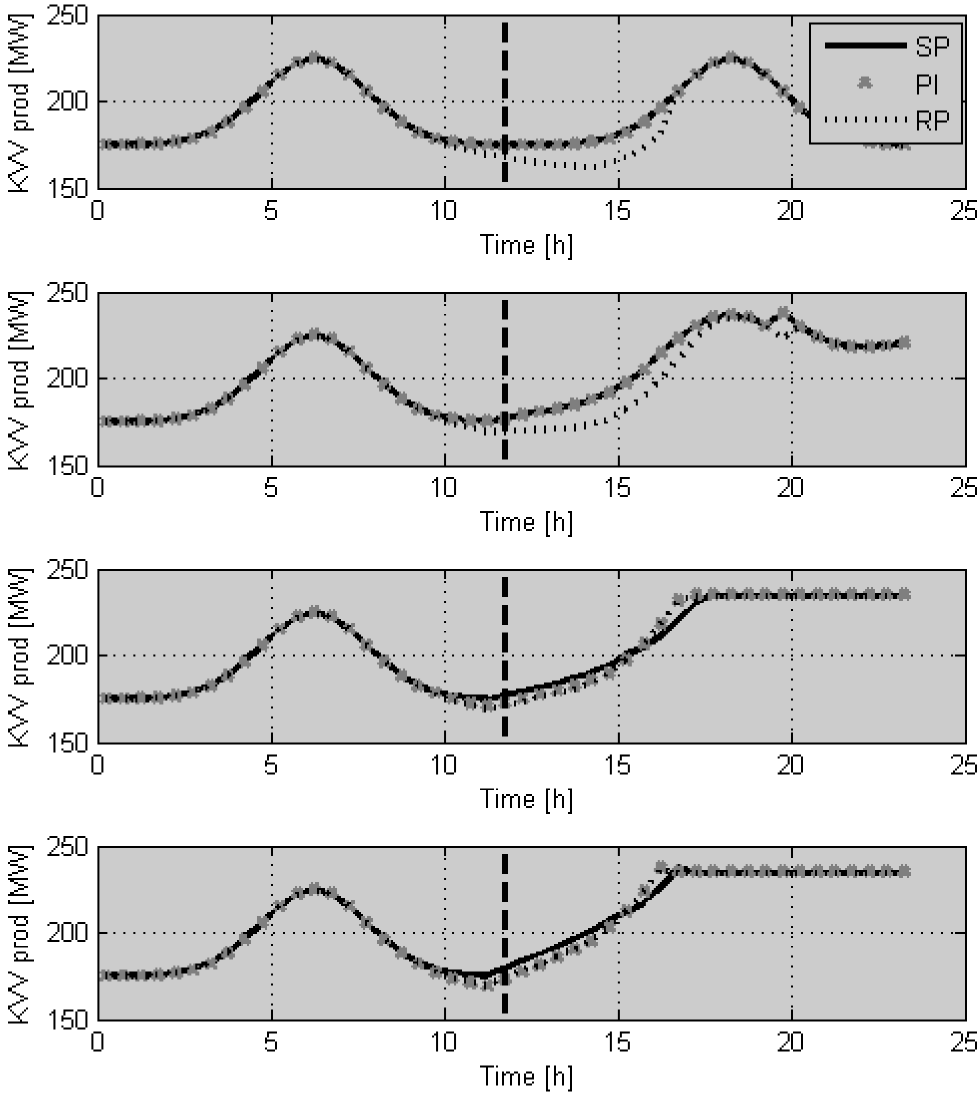

Figure 54: Case IV KVV production in different scenarios, with increasing heat demand from top to bottom. SP is the stochastic programming results, PI is the perfect information results, and RP is the wait-and-see results.

Figur 54. Fall IV, värmeproduktion från KVV för olika scenario, med ökande värmebehov uppifrån och ner. SP, PI och RP är resultaten för stokastisk programmering, perfekt information respektive wait-and-see metoden.

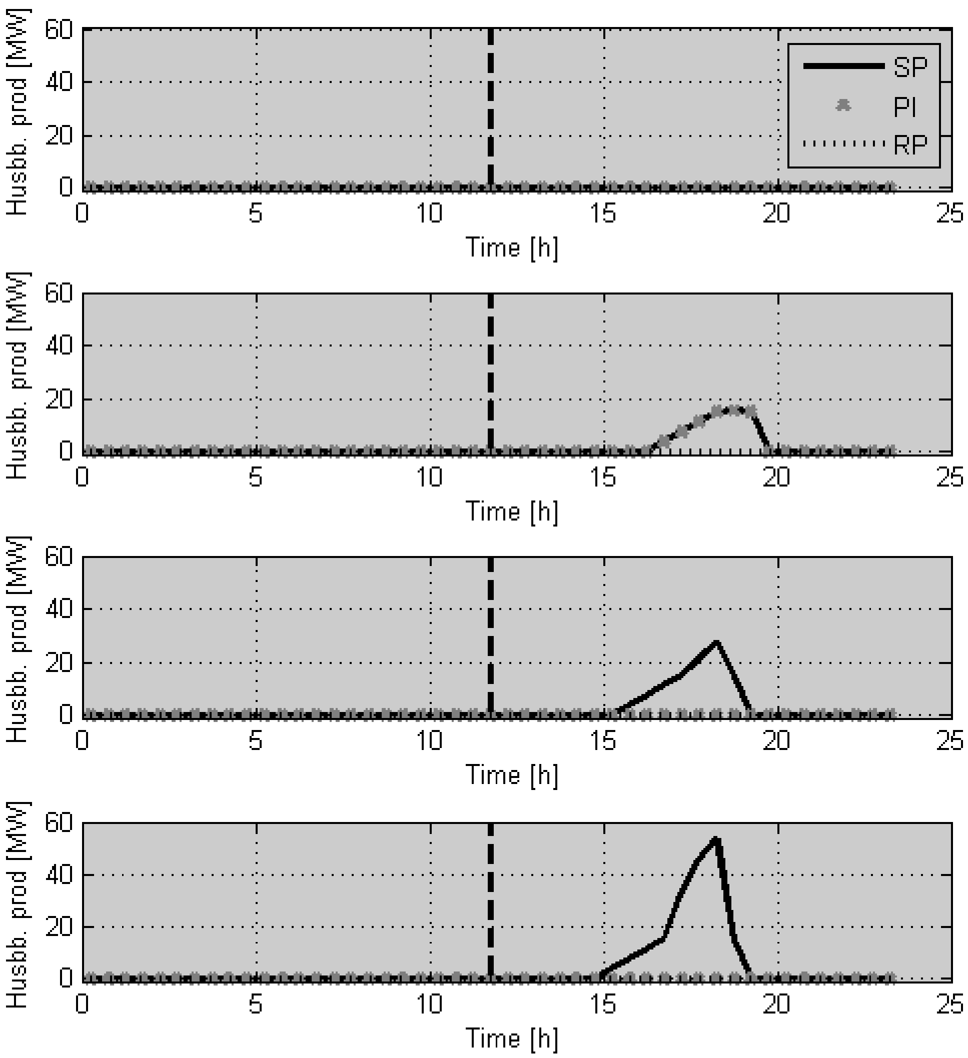

Figure 55: Case IV Husbyborg production in different scenarios, with increasing heat demand from top to bottom. SP is the stochastic programming results, PI is the perfect information results, and RP is the wait-and-see results.

Figur 55. Fall IV, värmeproduktion från Husbyborg för olika scenario, med ökande värmebehov uppifrån och ner. SP, PI och RP är resultaten för stokastisk programmering, perfekt information respektive wait-and-see metoden.

.

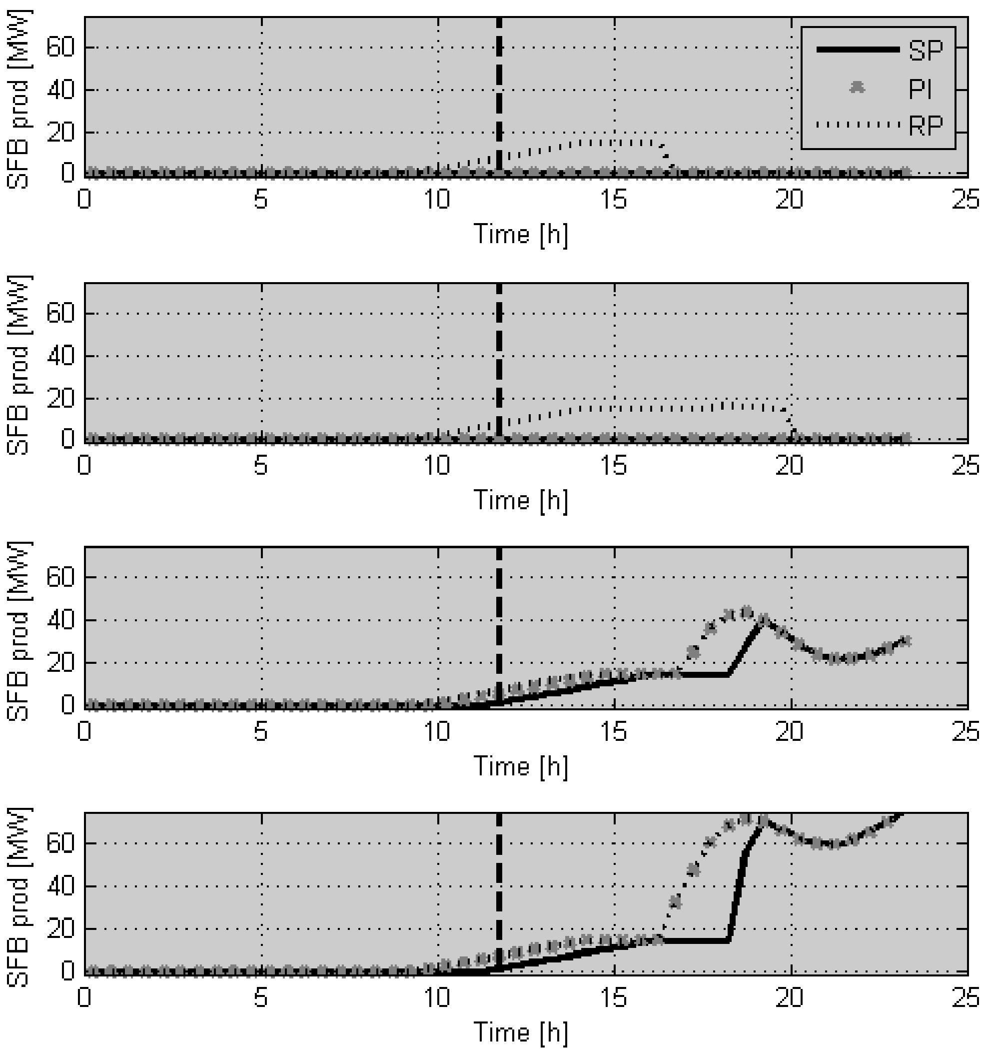

Figure 56: Case IV SFB production in different scenarios, with increasing heat demand from top to bottom. SP is the stochastic programming results, PI is the perfect information results, and RP is the wait-and-see results.

Figur 56. Fall IV, värmeproduktion från fastbränslepannan för olika scenario, med ökande värmebehov uppifrån och ner. SP, PI och RP är resultaten för stokastisk programmering, perfekt information respektive wait-and-see metoden.

7.6 Case V: Stochastic programming with uncertain electricity price

In this section, we formulate a stochastic programming model in which variations in both electricity price and demand are considered. Using this model, comparisons with a replanning approach are made on a small sample case study using mock-up data. The objective of the study is not to accurately model electricity price and demand variation, but only to determine whether the stochastic nature of electricity price and heat demand could theoretically have an impact on the profitability of running a district heating plant.

7.6.1 Problem Data

In this study the unknown demand and price in stage two is discretized into three possible outcomes. The problem setup is as follows:

- Eight time steps, four in the first stage and four in the second stage.
- Two production units: one KVV and one additional unit denominated HVC.
- A single customer, connected with pipes which incur a one-unit delay.
- Possibility to overproduce heat, since the heat demand is unknown at the time of the production. The excess heat returns to the KVV from the customer after one time unit and is limited to at most 65MW.
- No start-up/shut-down ramping.
- No heat losses.

The characteristics of the HVC is displayed in Table 21.

Table 21: HVC characteristics.

Tabell 21. Karakteristik för HVC.

Parameter	Value
Max capacity [MW]	30
Min capacity [MW]	100
Fuel cost (normalized)	0.52
Fixed cost (normalized)	2.02
Variable cost (normalized)	0.017
Efficiency	0.92

The three outcomes contain data as outlined in Table 22. As can be seen, at time 5 the heat demand spikes at 325 MW in outcome 1, 300 MW in outcome 2 and 275 MW in outcome 3.

Table 22: Stochastic outcomes in stage two.

Tabell 22. Stokastiska utfall i steg 2.

Stage	Time [sample]	Demand			El. Price [normalized]		
1	1		150			1	
	2		150			1	
	3		199			1	
	4		266			1	
2	5	325	300	275	0.75	1	1.25
	6	291	266	241	0.75	1	1.25
	7	224	199	174	0.75	1	1.25
	8	175	150	150	0.75	1	1.25
	Probability:	0.4	0.3	0.3	0.4	0.3	0.3

7.6.2 Results

The results of the study are shown in Table 23 below.

Table 23: Expected profit from the three planning approaches.

Tabell 23. Förväntad vinst från de tre planeringsmetoderna.

Approach	**Expected profit (normalized)**
Perfect information	1
Stochastic programming	0.9673
Wait-and-see	0.9490

The perfect information approach gives an indication on what the possible gains are from improving forecasting technology and gives an upper bound on the profit for all correct methods using the defined setup, but is otherwise not useful for planning in practice. The wait-and-see approach corresponds to robust planning, and has been considered as a practical alternative for a full stochastic programming approach. This method gives a lower bound on the profits of an optimal stochastic programming method. Considering stochastic programming, the results show that it was, for this particular setup, possible to reach a marginally higher profit by explicitly considering the possible outcomes in stage 2, and plan for them in stage 1. The difference between stochastic programming and robust planning is approximately 1.9 % on this particular setup. The difference can be compared to the theoretical maximum given by the perfect information approach, which is 5.4 % higher than the expected profit from wait-and-see.

The individual profiles for the three approaches are shown in Figure 57, Figure 58 and Figure 59. Comparing the results for the stochastic programming approach in Figure 58 to the results for the wait-and-see approach in Figure 59, we can see that in the latter case, the HVC starts up already at time 3, while this does not happen in the former case. Instead, in the first scenario, using the stochastic programming approach, the HVC runs at time 5 in stage 2, which does not happen in the wait-and-see approach. Starting the HVC at time 4 instead of 3 makes the return temperature lower at time 5. This allows the KVV to produce at a higher level (thus allowing the production of more electricity) in scenario 3 when the price is high. As the expected price in stage 2 is only 195, this opportunity is not taken into account in stage 1 in the wait-and-see approach, and as preparations have not been made, the opportunity is therefore lost when replanning is finally made at time 5. This illustrates the benefit of a stochastic approach.

One clear drawback of this optimization case is the dependency between the heat demand and the electricity price in the considered scenarios. In order to create a more realistic scenario these signals should be uncorrelated. Such a formulation would however increase the complexity of the optimization problem considerably. The effect of increasing the complexity of a stochastic formulation will be investigated in the next section.

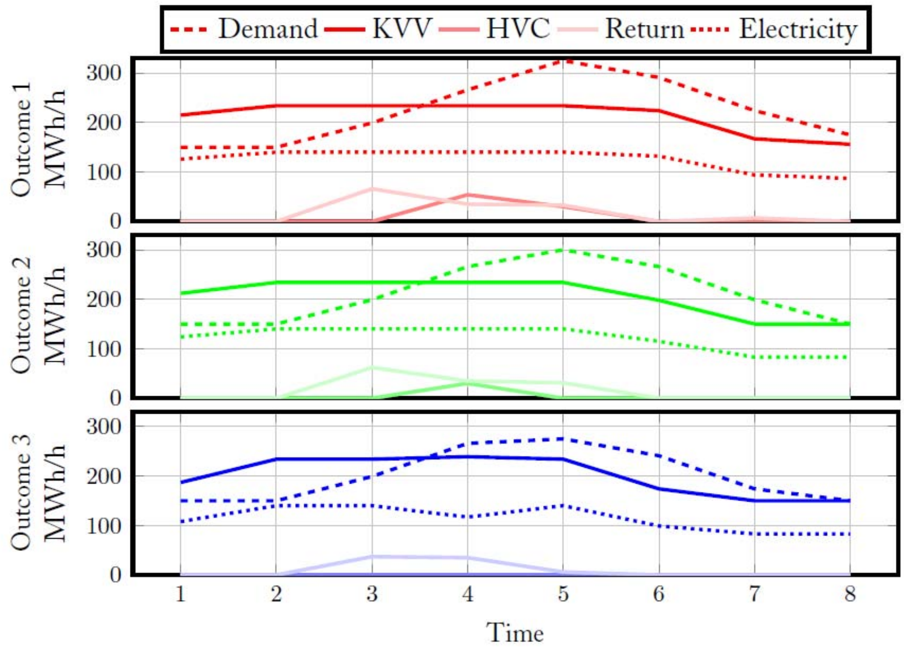

Figure 57: Perfect prognostics results.

Figur 57. Resultat vid perfekt prognostisering.

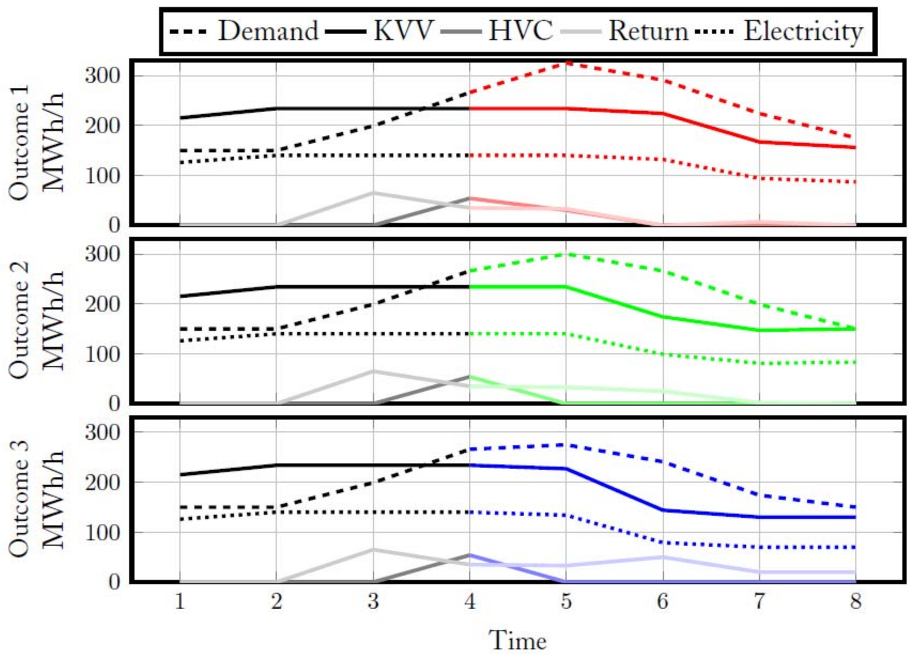

Figure 58: Stochastic programming results.

Figur 58. Resultat för stokastisk programmering.

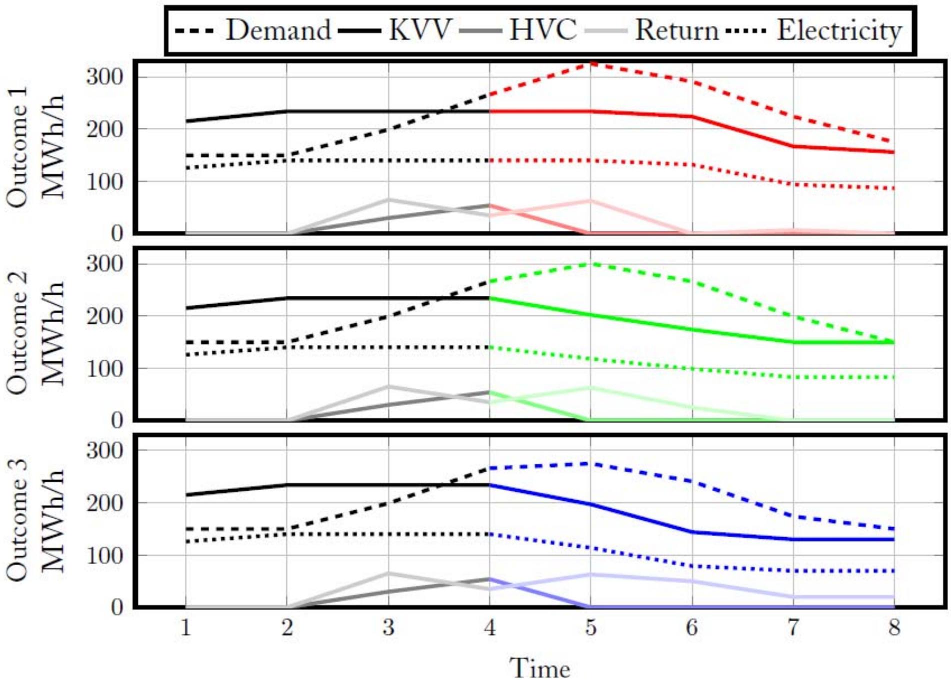

Figure 59: Wait-and-see approach results.

Figur 59. Resultat för wait-and-see metoden.

7.7 Scaling Test

In this section a comparison of the performance of the two MILP solvers used in this project, GLPK and Gurobi, is conducted. For this purpose an optimization setup identical to the one in case IV is used, except for the probability distribution of the scenarios and the number of scenarios. The heat demand scenarios are for simplicity assumed to be equally probable. The complexity of the optimization problem increases rapidly with an increased number of scenarios, making this a useful parameter in a scaling test.

Based on the computation time of the open-source solver the number of scenarios have been varied between two and seven. The calculation time for each solver when solving the stochastic programming formulation for different numbers of scenarios is presented in Table 24.

The results show an exponential dependency between the number of scenarios and the calculation time for the open-source solver, whereas the time is increasing more linearly for the commercial solver. For seven scenarios the Gurobi optimizer finds the optimal solution approximately 160 times faster than GLPK finds the same one.

Table 24: Computation times for open-source and commercial MILP solver.

Tabell 24. Beräkningstider för open-source-, respektive kommersiell lösare.

Number of scenarios	GLPK time [s]	Gurobi time [s]
2	1.6	1.5
3	7.9	1.4
4	7.4	2.0
5	56.4	2.2
6	404.1	2.7
7	517.5	3.2

For the setups used for the UCP in this project the open source solver has proven sufficient, but based on the results from the scaling test it seems to be beneficial to use a commercial solver when more complex MILP problems are investigated. Using a commercial tool would seem especially useful if more complex stochastic formulations are considered, as the open-source solver struggles in this case.

8 Summary

This section summarizes the main contributions of the project.

8.1 General

In general, the work has shown that

- It is possible to:
  - Model distributed consumption and production using a distribution network in both UCP and EDP.
  - Optimize production planning schedules based on models with distributed consumption and production in both UCP and EDP.
- Modeling of a distribution network, as compared to point-wise consumption and production, has impact on optimized production planning: reduction of costly production peaks, heat losses, pipe heat storage.
- Optimization horizons have been extended in both UCP (several days) and EDP (24h) compared to [9].
- Stochastic optimization is able to utilize load demand prediction probabilities in production planning.

8.2 Modeling

As the optimizations utilize the models in any way possible, it is important that the models describe the reality as closely as possible.

Main results for the physics-based modeling for the district heat distribution and customers are

- Both production and consumption of heat are distributed. This means that
  - Total heat demand is distributed among customers.
  - Delay times are individual for each customer.
  - Analysis on effects of distribution, compared to point-wise production and consumption, can and has been performed.
- The distribution net is built using physics-based pipes, giving the following properties
  - Time delay of a pipe is supply/return flow dependent.
  - Heat loss of a pipe depends on supply/return water temperature, pipe length, and outdoor temperature.
  - Distribution net can be used as accumulator as the pipes are capable of storing energy.
- A customer model that
  - Utilizes a return temperature model based on outdoor temperature.
  - Calculates its individually needed water supply flow rate based on supply and return water temperature.
  - Determines overall distribution hydraulics.

Main results for the physics-based modeling for the production units are:

- Kraft- och värmeverk (KVV) has generation of heat and electricity that depend on return temperature and flow as well as load level, and has been validated with measurement data with low relative error (~5%).
- The physics-based KVV model has been used as reference for incorporating return temperature dependence in the KVV UCP model by means of a polytope in return temperature-heat-electricity space.

8.3 Optimization

All resulting planning schedules have been optimized for maximized production profit, taking into consideration heat, electricity and fuel prices as well as start-up and maintenance costs for the production units.

Physics-based modeling of the distribution network has impacted optimization results, and the main characteristics due to this are

- Costly production peaks can be decreased as each customer has an individual time delay and heat demand, yielding production spread out over a longer time period.
  - In case III, the peak reduction delayed the start of an expensive boiler.
- Distribution net is used as an accumulator, controlled by supply temperature.
- Heat loss in distribution net, dependent on outdoor temperature, can be compensated for by higher heat production.

Other characteristics of the resulting optimized planning schedules are

- Production profit over optimization horizon is maximized.
- Balance between heat production and heat consumption.
- Electricity production is maximized due to high sell price.
- Supply temperature is minimized yielding lower boiler load and supply flow (active constraints, see below).
- Heat demand peaks can be handled by increased supply flow, supply temperature, or both simultaneously.

The physics-based model approach in Modelica together with the optimization in JModelica.org allows setting constraints on physically relevant signals. The main active constraints in the planning schedule are

- Individual temperature constraints at each customer provides safety towards bacterial growth; active in e.g., case I/IIab.
- A decreased supply temperature requires higher supply flow, which in shown cases, e.g., IIIa, reaches limit due to pump flow rate constraint.
- A high return flow can give too low condenser pressure, which has an active constraint in case I.
- Production units have start-up/stop trajectories as well as minimum and maximum loads that have been active.

The examples with stochastic optimization has shown that

- It is possible to formulate a stochastic optimization problem based on heat demand prediction probabilities.
- Taking heat demand prediction probabilities into account, it is possible to increase the expected profit.

A simple scaling test has revealed that

- The commercial Gurobi solver performs considerably better than the open-source package GLPK when MILP problems containing multiple stochastic scenarios are solved.

9 Future Work

The work in [9] focused mainly on the production units, having a simple net and customer, while the work in this report has built upon [9] and focused more on the distribution net and distribution of production and consumption. For implementing the approach as a decision support tool, future work could aim at the customer. However, improvements on production units and distribution net can be done as well. This section presents some possible improvements and ideas.

9.1 Modeling

The distribution net considered in this work is one-dimensional and the water media model used in the distribution net is a fixed-pressure media. To increase the physics-based modelling of the net, the following improvements are required:

- General structures of the distribution net:
  - Several branches and junctions (not one-dimensional) are needed to model general district heating water nets. In turn, this may require a finer distribution of the total heat load prediction.
  - Pipes of the distribution net need to support reversing flow.
- Including pressure calculations for the district heating water will make it possible to model the pressure drop along the net and at customers. This would also make it possible to have constraints on pumps for controlling the pressure difference across critical customers.

Customers have requirements on the district heating water system and affect it at the same time. An improved customer model that reflects these properties could further improve the optimized planning schedule. Improvements include:

- A pressure drop model, making it possible to set constraint on desired pressure of the district heating water.
- Customer grouping into different types such as residential, commercial and industry. In turn requiring improved heat demand prediction models for the different types.
- Inclusion of different demand side management (DSM) strategies at customers to reduce load peaks and avoid start-ups of expensive boilers.

The overall fidelity of the UCP models could be improved such that the model consistency towards the physics-based EDP models is higher. Advancements in these models can include

- Dependence on supply and return temperature in units, distribution net and customer. Such studies have been made, see e.g., [7].
- Better parameterization of operating range, i.e., not excluding any extreme points when generating the UCP models from the EDP models.

9.2 Optimization

The generality of the optimization formulation could be increased by

- Adding handling of time varying electricity, heat and fuel prices, which was seen as interesting in [9].
- Adding start cost dependence on unit status history, e.g., lower start cost at warm starts than at cold starts, requiring temperature in the UCP unit model.
- Removing or handling effects of limited optimization horizons such as
    - Accumulator energy at optimization horizon end.
    - Pipe energy at optimization horizon end.

  This is valid for both UCP and EDP optimizations as well as in the integration step.

Stochastic programming showed a benefit when having uncertainties in heat demand predictions. Future work include therefore

- Construction of more advanced and reality like scenarios with integration with EDP.
- Analysis of different stochastic programming formulations.
- Construction of a heat demand prediction model that provides e.g., a nominal heat demand prediction and a probability distribution of the heat demand as a function of time.

Frameworks for implementing and solving stochastic mixed-integer linear programs are available in e.g., PySP, an open source modeling and solver library, see [28].

Improvements can be made in the dynamic optimization framework JModelica.org, mainly targeting

- Memory consumption by implementing a BLT algorithm.
- Implementing warm-start for supported solvers if using previous optimization results as initial guess is desired.
- Evaluating different solvers, such as WORHP [29].

10 Bibliography

[1] A. Pruessner and M. R. Bussieck, "Mixed-Integer Nonlinear Programming," GAMSWorld, 2009.

[2] D. Häggståhl and E. Dotzauer, "Produktionsplanering under osäkerhet - Simulatorbaserad produktionsplanering av medelstora kraftvärmeanläggningar," Värmeforsk report 898, 2004.

[3] Vitec, "Vitec - Välkommen till Vitec," Vitec, 2014. [Online]. Available: http://www.vitec.se/. [Accessed 27 10 2014].

[4] Energy Opticon, "Energy Opticon," Energy Opticon, 2014. [Online]. Available: http://www.energyopticon.com/. [Accessed 27 10 2014].

[5] ABB, "ABB," ABB, 2014. [Online]. Available: http://new.abb.com/se. [Accessed 27 10 2014].

[6] E. Dotzauer, "Algorithms for Short-Term Production planning of Cogeneration Plant," Lic. Thesis, Linköping University, 1997.

[7] E. Dotzauer, "Produktionsplanering av el och värme - Matematiska modeller och metoder," Forskningsrapport MDH ISt 2002:2, Inst. för Samhällsteknik, Mälarladelens Högskola, 2002.

[8] J. Kvarnström, E. Dotzauer and E. Dahlquist, "Produktions- och distributionsplanering av fjärrvärme," Värmeforsk, 2006.

[9] S. Velut, P.-O. Larsson, J. Windahl, K. Boman and L. Saarinen, "Non-linear and Dynamic Optimization for Short-term Production Planning," Värmeforsk, 2013.

[10] L. Saarinen, "Model-based control of district heating supply temperature," Värmeforsk rapport 1555, Stockholm, 2010.

[11] A. Cervantes and L. Biegler, "Optimization Strategies for Dynamic Systems," in *Encyclopedia of Optimization*, Springer, 2009, pp. 2847-2858.

[12] J. Åkesson, C. Laird, K. Lavedan, K. Prölss, H. Tummescheit, S. Velut and Y. Zhu, "Nonlinear Model Predictive Control of a CO2 Post-Combustion Unit," *Chemical Engineering Technology,* vol. 35, no. 3, pp. 445-454, 2011.

[13] E. Sällberg and A. Lind, "Start-up Optimization of a Combined Cycle Power Plant," in *9th International Modelica Conference*, Munich, Germany, 2012.

[14] L. Saarinen and K. Boman, "Optimized district heating supply temperature for large networks," Värmeforsk, 2012.

[15] Vattenfall AB, *Översikt Uppsalaanläggningen,* Uppsala: Vattenfall AB, 2014.

[16] L. Saarinen, "Modeling and control of a district heating system," 2008.

[17] B. Sundén, Värmeöverföring, Studentlitteratur, 2006.

[18] S. Velut and H. Tummescheit, "Implementation of a transmission line model for fast simulation of fluid flow dynamics," in *Modelica Conference*, Dresden, 2011.

[19] Sandia National Laboratories, "Pyomo - Pyomo," 2014. [Online]. Available: https://software.sandia.gov/trac/pyomo/wiki/Pyomo. [Accessed 22 10 2014].

[20] Gurobi Optimization, "Gurobi - The overall fastest and best supported solver available," Gurobi Optimization, 2014. [Online]. Available:

http://www.gurobi.com. [Accessed 22 10 2014].

[21] GNU Project, "GLPK - GNU Project - Free Software Foundation (FSF)," GNU Project, 2014. [Online]. Available: https://www.gnu.org/software/glpk/. [Accessed 22 10 2014].

[22] "www.jmodelica.org," Modelon AB, 2014. [Online]. Available: www.jmodelica.org. [Accessed 2014].

[23] J. Åkesson, K.-E. Årzén, M. Gäfvert, T. Bergdahl and H. Tummescheit, "Modeling and Optimization with Optimica and JModelica.org—Languages and Tools for Solving Large - Scale Dynamic Optimization Problems," *Computers and Chemical Engineering,* vol. 34, no. 11, pp. 1737-1749, 2010.

[24] J. Andersson, PhD Thesis: A General Purpose Software Framework for Dynamic Optimization, leuven: Arenberg Doctoral School, KU Leuven, 2013.

[25] A. Wächter and L. T. Biegler, "On the implementation of an interior-point filter line-search algorithm for large-scale nonlinear programming," *Mathematical Programming,* vol. 196, no. 1, pp. 25-68, 2006.

[26] J. Arroyo and A. Conejo, "Modeling of start-up and shut-down power trajectories of thermal units," *IEEE Transactions on power systems,* vol. 19, no. 3, 2004.

[27] T. Shiina and J. R. Birge, "Stochastic unit commitment problem," *International Transactions in Operational Research,* vol. 11, no. 1, 2004.

[28] Sandia National Laboratories, "PySP," 2013. [Online]. Available: https://software.sandia.gov/trac/coopr/wiki/PySP. [Accessed 14 August 2013].

[29] C. Büskens, D. Wassel, S. Geffken, R. Schäfer and M. Gerdts, "WORHP - Large-scale Sparse Nonlinear Solver," 2014. [Online]. Available: http://www.worhp.de/. [Accessed 30 10 2014].