Data-Driven Methods for Enhancing District Heating Network Operation

Bergsteinsson, Hjörleifur G.

Publication date:
2022

Document Version
Publisher's PDF, also known as Version of record

DTU Compute
Department of Applied Mathematics and Computer Science

Data-Driven Methods for Enhancing District Heating Network Operation

PhD Thesis

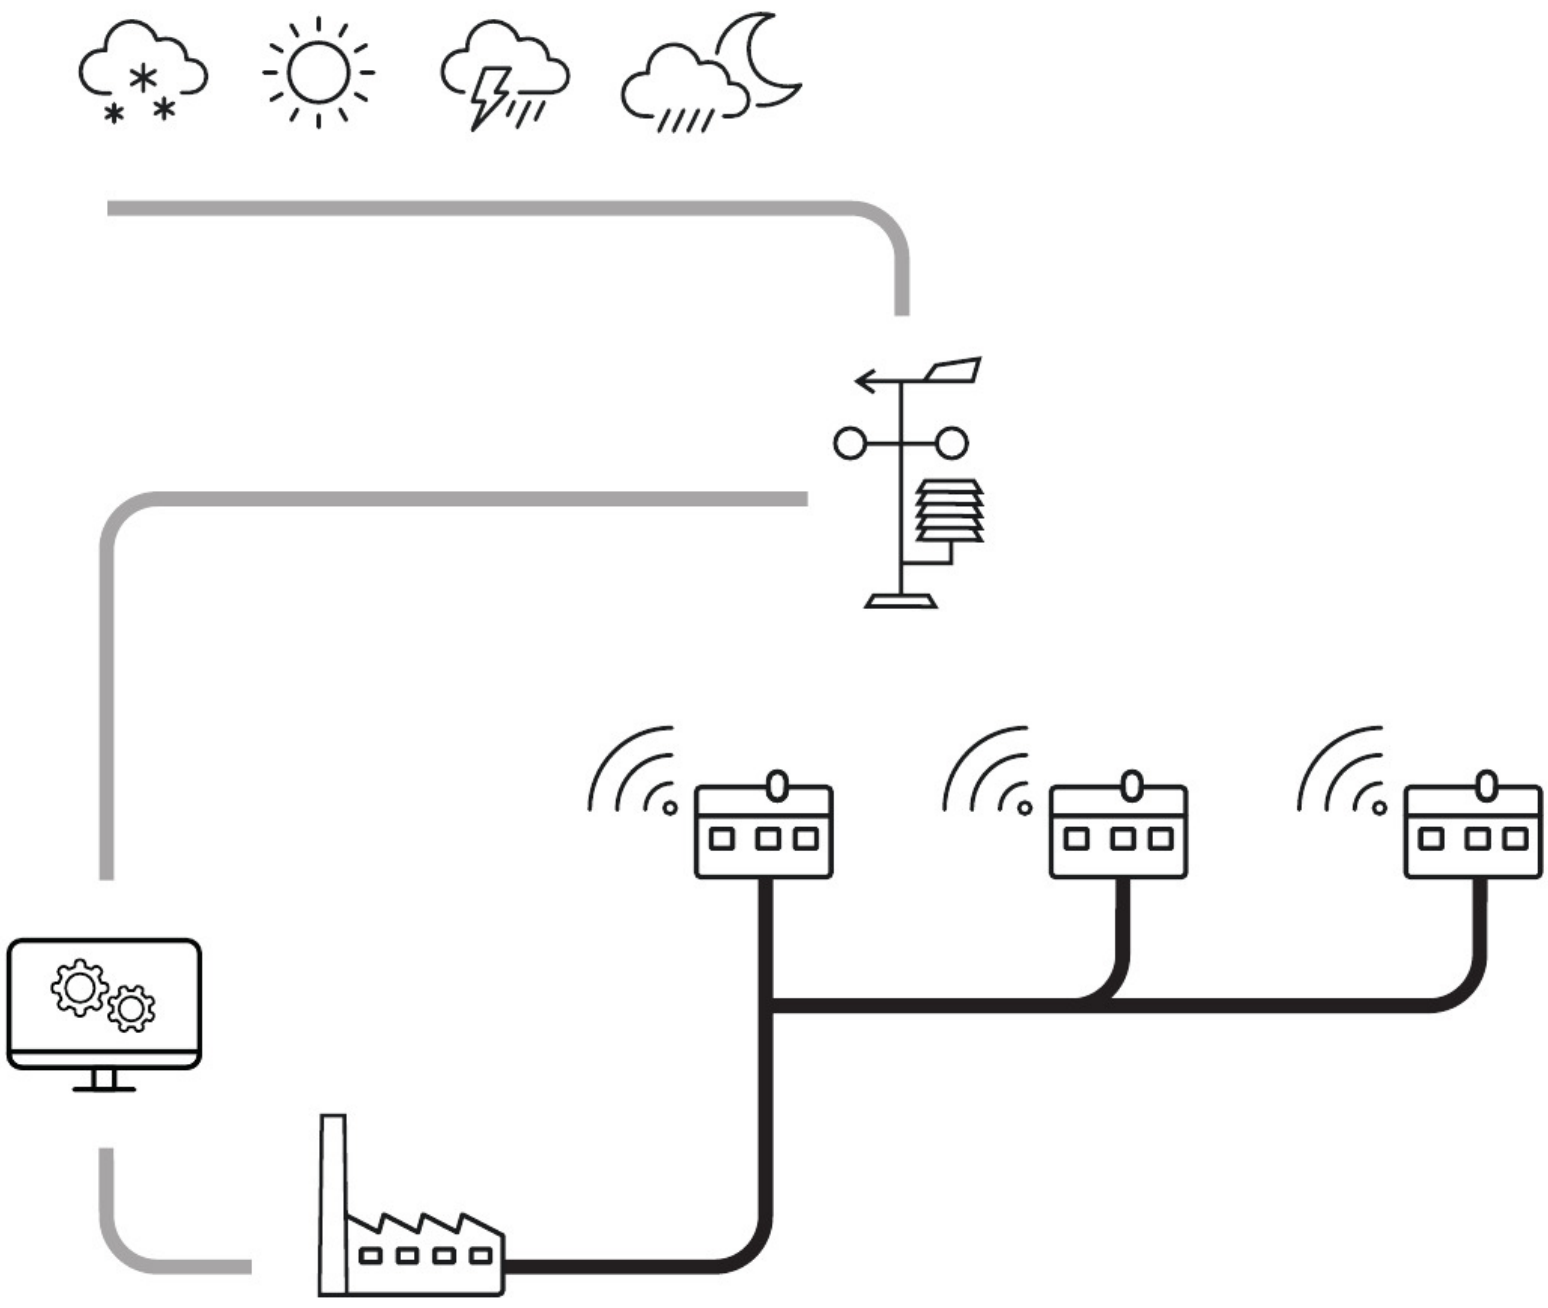

Data-Driven Methods for Enhancing District Heating Network Operation

PhD Thesis
April, 2022

By: Hjörleifur G. Bergsteinsson
Supervisor: Henrik Madsen
Co-supervisor: Jan Kloppenborg Møller
Torben Skov Nielsen

Cover photo: Bjarne Erick, 2021
Published by: DTU, Department of Applied Mathematics and Computer Science, Richard Petersens Plads, Building 324, 2800 Kgs. Lyngby Denmark
www.compute.dtu.dk
ISSN: [0909-3192]

Preface

This PhD thesis was prepared at the Department of Applied Mathematics and Computer Science at the Technical University of Denmark in fulfilment of the requirements for acquiring a PhD degree in Applied Mathematics and Computer Science.

This PhD study is supported financially by Innovation Fund Denmark through the projects HEAT 4.0 (8090-00046B), TOP-UP (9045-00017B), and FED (8090-00069B). The IDASC (18012745) project funded by Region Hovedstaden and the IEA TS4 annex on Digitalisation of District Heating: Optimised Operation (and Maintenance) of District heating and Cooling Schemes via Digital Processes Management funded by the EUDP Denmark have also supported this work.

Hjörleifur G. Bergsteinsson

.....

Signature

.....

Date

Summary (English)

This dissertation is a collection of publications that researched the operation of district heating networks. The district heating sector will contribute significantly to the future smart energy integration of renewable energy into the power system due to its high efficiency and flexibility in converting power into heat and its thermal energy storage properties. However, reaching optimal district heating operation for both individual and overall energy systems (sector coupling) involves many complex tasks, the solution of which is of great importance. One of them is to operate the district heating network efficiently. This is done by keeping the supply temperature as low as possible without violating any restrictions (e.g. maximum flow, variation of supply temperature) while at the same time meeting the heat load. This results in lower production costs and reduced heat losses in the network. Lowering the temperature in the network also leads to better investment feasibility for power-to-heat units (e.g., heat pumps), low-temperature geothermal wells, and recycling of waste heat. Without lowering the temperature, these sources would result in lower efficiency or even be disregarded.

Hence, to enhance the operational efficiency of the district heating system, the network must be operated in an optimal mode. Many district heating networks are operated by operators who use their experience and ”scarce” data to select the set points for the supply temperature. This usually leads to suboptimal operation because the network characteristics are complex, and many variables have to be taken into account, e.g. the future heat load and time delays in the network. The objective of this thesis is to propose data-driven methods that can support operators in decision-making. In particular, mathematical data-driven models will be developed based on physical knowledge and designed for use in real-time applications to increase the efficiency of operations.

Various studies have been carried out as part of the project to investigate state-of-the-art and sub-sequentially propose new methods to highlight the importance of data-driven methods to obtain optimal operation. For instance, heat load forecasting is crucial for overall district heating operations as it gives operators insight into future consumption. This gives operators the essential information to support their decisions to minimise the costs of operating the district heating system. The more accurate the forecasts, the better the operators can make their decisions. This thesis discusses the essential features for building an accurate and robust heat load forecasting model. It is shown that the localisation of the input variables and heat load has an impact on the accuracy of the heat load forecast. In addition, new methods are presented to increase the accuracy of the current operational forecasts. The proposed methods use and extend state-of-the-art methods for hierarchical forecasting. Both temporal and spatial hierarchies of district heating are considered in order to investigate the possibilities for improving today’s state-of-the-art operational forecast. It is shown that the proposed method can improve the accuracy of current operational forecasts by about 15%. These methods will be indispensable in future decentralised district heating systems as the suggested methods both increase accuracy and make them coherent across the considered temporal and spatial aggregation levels.

Methods for temperature optimisation for the district heating network are presented and discussed. A set of controllers are used in the optimisation, supply temperature controllers and flow controllers. The supply controller ensures that the temperature is sufficient at a set of selected critical points of the network. These critical points are selected such that if the temperature is sufficiently high at the critical points, then the temperature is sufficiently high everywhere. The flow controller ensures that the flow restrictions in the system are not violated, and here the time-varying electricity prices can be taken into account. Measurements at the critical point are needed to serve as temperature feedback for the controller. Typically, measurement wells are installed in areas near a group of end-users, but these measurements come with an associated cost, and the equipment needs to be maintained to ensure reliability and high precision, i.e.

temperature sensors must be finely calibrated.

This work proposes that smart meter readings can replace these measurement wells to establish the needed temperature feedback for the supply temperature controllers. Both simple and complex methods are presented. The simple method is easy to understand and can be implemented with little computational effort. On the other hand, the complex method is more robust but requires more fine-tuning. The potential savings from implemented temperature optimisation are discussed and demonstrated in a case study where the temperature is kept as low as possible. Based on real-life implementations, it is demonstrated that the precision is increased. The use of smart meters also leads to additional information about the district heating network since the network of the individual users is also taken into account. It is argued that the suggested methods give new possibilities for adaptive zonal temperature control, which could lead to further savings on heat loss and better integration of heat pumps and the use of excess heat from supermarket cooling, etc.

Finally, the potential of integrating the consumers' heating system into the network operation through the smart operation with predictive controllers that can, for example, receive signals from the district heating system to influence its heating consumption (peak shaving) is discussed.

In summary, several new methods for data-driven optimization of district heating systems are suggested. These methods can be integrated with existing methods and lead to further savings, better energy efficiency and flexibility. The suggested methods range from simple to more complex methods, but all of the methods are intended to improve online operations.

Resumé (Danish)

Denne afhandling består af en samling af publikationer, der omhandler driften af fjernvarmenetværk. Fjernvarmesektoren vil bidrage væsentligt til den fremtidige intelligente integration af vedvarende energikilder i elsystemet på grund af dens høje effektivitet og fleksibilitet i forbindelse med omdannelse af strøm til varme og dens termiske energilagringsegenskaber. At opnå optimal fjernvarmedrift for både individuelle og koblede energisystemer indebærer imidlertid mange komplekse opgaver, hvis løsning er af stor betydning. En af dem er at drive fjernvarmenettet effektivt. Dette gøres ved at holde fremløbstemperaturen så lav som muligt uden at overtræde nogen begrænsninger, så som maksimalt flow og variationer i fremløbstemperaturen, samtidig med at varmeefterspørgslen bliver opfyldt. Resultatet af dette er lavere produktionsomkostninger og mindre varmetab i nettet. En lavere temperatur i nettet fører også til bedre investeringsmuligheder for kraftvarmeenheder (f.eks. varmepumper), geotermiske brønde med lav temperatur og genanvendelse af spildvarme. Uden en temperatursænkning ville disse kilder have en lavere effektivitet eller slet ikke være rentable.

For at øge fjernvarmesystemets driftseffektivitet skal nettet derfor drives optimalt. Mange fjernvarmenetværk drives af operatører, som bruger deres erfaring og "sparsomme" data til at vælge indstillingsværdierne for fremløbstemperaturen. Dette fører normalt til suboptimal drift, fordi fjernvarmenetværk er komplekse, og der skal tages hensyn til mange variabler, f.eks. den fremtidige varmebelastning og tidsforsinkelser i nettet. Formålet med denne afhandling er at foreslå datadrevne metoder, der kan støtte operatørerne i deres beslutningstagning. Der vil navnlig blive udviklet matematiske datadrevne modeller baseret på fysisk viden og udviklet til brug i realtidsanvendelser for at øge effektiviteten af driften.

Som led i projektet er der gennemført forskellige undersøgelser for at undersøge den nyeste teknologi og efterfølgende foreslå nye metoder for at fremhæve betydningen af datadrevne metoder for at opnå optimal drift. For eksempel er prognoser for varmebelastningen afgørende for den samlede fjernvarmedrift, da de giver operatørerne indsigt i det fremtidige forbrug. Dette giver operatørerne de nødvendige oplysninger til støtte for deres beslutninger om at minimere omkostningerne ved driften af fjernvarmesystemet. Jo mere præcise prognoserne er, jo bedre kan operatørerne træffe deres beslutninger. I denne afhandling diskuteres de væsentligste elementer til opbygning af en præcis og robust prognosemodel for varmebelastning. Det påvises, at lokaliseringen af inputvariablerne og varmeefterspørgslen har betydning for nøjagtigheden af varmelastprognosen. Desuden præsenteres der nye metoder til at øge nøjagtigheden af de nuværende operationelle prognoser. De foreslåede metoder anvender og udvider de nyeste metoder til hierarkiske prognoser, og der tages hensyn til både tidsmæssige og rumlige hierarkier for fjernvarme for at undersøge mulighederne for at forbedre de nuværende operationelle prognoser. Det vises, at den foreslåede metode kan forbedre nøjagtigheden af de nuværende driftsprognoser med ca. 15%. Disse metoder vil være uundværlige i fremtidige decentraliserede fjernvarmesystemer, da de foreslåede metoder både øger nøjagtigheden og gør dem sammenhængende på tværs af de betragtede tidsmæssige og rumlige aggregeringsniveauer.

Derudover det blive præsenteret og diskuteret metoder til temperaturoptimering for fjernvarmenettet. Der anvendes et sæt regulatorer i optimeringen, fremløbstemperaturregulatorer og flowregulatorer. Forsyningsregulatoren sikrer, at temperaturen er tilstrækkelig på et sæt udvalgte kritiske punkter i nettet. Disse kritiske punkter er valgt således, at hvis temperaturen er tilstrækkelig høj på de kritiske punkter, er temperaturen tilstrækkelig høj overalt. Strømregulatoren sikrer, at strømningsrestriktionerne i systemet ikke overtrædes, og her kan der tages hensyn til de tidsvarierende elpriser. Målinger ved det kritiske punkt er nødvendige som temperaturfeedback til regulatoren. Typisk installeres der målebrønde i områder nær en gruppe slutbrugere, men disse målinger er forbundet med omkostninger, og udstyret skal vedligeholdes for at sikre pålidelighed og høj præcision, dvs. at temperaturfølerne skal være finkalibrerede.

I dette arbejde foreslås det, at aflæsninger fra intelligente målere kan erstatte disse målebrønde for at etablere den nødvendige temperaturfeedback til temperaturregulatorer for forsyningstemperaturen. Der præsenteres både enkle og komplekse metoder. Den enkle metode er let at forstå og kan gennemføres med en lille beregningsindsats. På den anden side er den komplekse metode mere robust, men kræver mere finjustering. De potentielle besparelser som følge af den gennemførte temperaturoptimering diskuteres og demonstreres i en case-undersøgelse, hvor temperaturen holdes så lav som muligt. På grundlag af virkelige implementeringer påvises det, at præcisionen øges. Brugen af intelligente målere fører også til yderligere oplysninger om fjernvarmenettet, da der også tages hensyn til nettet til de enkelte brugere. Det anføres, at de foreslåede metoder giver nye muligheder for adaptiv zonetemperaturstyring, hvilket kan føre til yderligere besparelser på varmetabet og bedre integration af varmepumper og udnyttelse af overskudsvarme fra supermarkedskøling osv.

Endelig diskuteres potentialet i at integrere forbrugernes varmesystem i netdriften gennem intelligent drift med forudsigelige styringer, der f.eks. kan modtage signaler fra fjernvarmesystemet for at påvirke dets varmeforbrug til at yde for eksempel peak shaving.

Sammenfattende foreslås der flere nye metoder til datadrevet optimering af fjernvarmesystemer. Disse metoder kan integreres med eksisterende metoder og føre til yderligere besparelser, bedre energieffektivitet og fleksibilitet. De foreslåede metoder spænder fra enkle til mere komplekse metoder, men alle metoderne har til formål at forbedre online-driften.

Acknowledgements

I want to thank my supervisors at DTU Compute, Henrik Madsen and Jan Kloppenborg Møller, for giving me the opportunity to do this work and advance my skills in statistical modelling. I would also like to thank you for the support and time you have given me to develop my ideas and theoretical knowledge of the topics covered in this thesis - and more.

I wish to thank ENFOR for allowing me to do my PhD in collaboration with the industry. This allowed me to do applied research in the field of district heating. Especially my co-supervisor, Torben Skov Nielsen, who supported me with information and theoretical knowledge about statistical modelling and district heating.

I want to give special thanks to my office colleagues in room 039 (Cities Office); Christian Ankerstjerne Thilker, Christoffer Rasmussen, Jaume Palmer Real, Dominik Dominkovic and Rune Grønborg Junker. Kudos for welcoming me to the office and for all our fruitful conversations about research and everyday life. You have definitely helped me during stressful times.

I also want to thank Phillip B. Vetter for our discussion on stochastic differential equations and modelling. Thanks to Jan Lorenz Svensen for our discussion about modelling and control theory. A big thank you also goes to Mikkel Lindstrøm Sørensen for our discussions during PhD meeting with Jan regarding topics in our thesis, especially hierarchical forecasting.

Many thanks to Peder Bacher for taking the time to discuss everything, especially forecasting and programming.

Thanks to the University of Iceland for hosting me for a month in the autumn of 2020. Special thanks to Ólafur Pétur Pálsson for mentoring me during my stay and for all the fruitful discussions about modelling and district heating.

I am very grateful to all the people I met during this PhD. All the fruitful discussions with you during the PhD were the highlights of this journey. It is a great privilege to have had the opportunity to conduct a PhD.

Finally, I am grateful and humbled for the support of my fiancée, Ína Salome, and my daughter, Hera Salome. A big thank you also goes to my family for their unwavering support.

List of Works

Publications

Peer-Reviewed

A Hjörleifur G. Bergsteinsson, Jan Kloppenborg Møller, Peter Nystrup, Ólafur Pétur Pálsson, Daniela Guericke and Henrik Madsen. "Heat Load Forecasting using Adaptive Temporal Hierarchies". *Applied Energy*, Volume 292, April 12, 2021.

B Hjörleifur G. Bergsteinsson, Torben Skov Nielsen, Jan Kloppenborg Møller, Sara Ben Amer, Dominik Franjo Dominković and Henrik Madsen."Use of Smart Meters as Feedback for District Heating Temperature Control". *Energy Reports*, Volume 7, October 28, 2021.

C Hjörleifur G. Bergsteinsson, Phillip B. Vetter, Jan Kloppenborg Møller and Henrik Madsen."Estimating Temperatures in a District Heating Network Using Smart Meter Data". *Energy Conversion and Management*, Volume 269, October 1, 2022.

D Christian Ankerstjerne Thilker, Peder Bacher, Hjörleifur G. Bergsteinsson, Rune Grønborg Junker, Davide Cali and Henrik Madsen."Non-linear Grey-Box Modelling for Heat Dynamics of Buildings". *Energy and Buildings*, Volume 252, June 28, 2021.

E Christian Ankerstjerne Thilker, Hjörleifur G. Bergsteinsson, Peder Bacher, Henrik Madsen, Davide Calì and Rune G. Junker."Non-linear Model Predictive Control for Smart Heating of Buildings". *E3S Web Conf.*, Volume 246, March 29, 2021.

Book Chapter

F Hjörleifur G. Bergsteinsson, Jan Kloppenborg Møller, Christian Ankerstjerne Thilker, Daniela Guericke, Alfred Heller, Torben Skov Nielsen, and Henrik Madsen."Handbook of Low Temperature District Heating". In *Handbook of Low Temperature District Heating*, Green Energy and Technology. Springer, 2022

Report

G Hjörleifur G. Bergsteinsson, Sara Ben Amer, Per Sieverts Nielsen and Henrik Madsen. "Digitalization of District Heating". *Technical University of Denmark*, 2021.

Submitted/To be Submitted

H Hjörleifur G. Bergsteinsson, Jan Kloppenborg Møller and Henrik Madsen."Quantification of Heat Demand Forecast Accuracy Improvements by Localized Weather Forecast". Submitted To *International Journal of Sustainable Energy*.

I Peder Bacher, Hjörleifur G. Bergsteinsson, Linde Frölke, Mikkel L. Sørensen, Julian Lemos-Vinasco, Jon Liisberg, Jan Kloppenborg Møller, Henrik Aalborg Nielsen and Henrik Madsen. "onlineforecast: An R package for adaptive and recursive forecasting", Submitted To *The R Journal*.

J Hjörleifur G. Bergsteinsson, Mikkel Lindstrøm Sørensen, Jan Kloppenborg Møller and Henrik Madsen. "Heat Load Forecast Accuracy Improvement with Spatial Hierarchies", To Be Submitted.

Presentations

Conference Presentations

K **Quantification of Heat Demand Forecast Accuracy Improvements by Localized Weather Forecast**., The 2nd International Sustainable Energy Conference, April 2022, Graz, Austria,

This conference paper was selected to be featured as a full paper in a special issue of the journal of International Journal of Sustainable Energy.

L **Use of Smart Meters as Feedback for District Heating Temperature Control.**, 17th International Symposium on District Heating and Cooling, September 2021, Nottingham, England, *This presentation and paper was awarded an IEA DHC Research Excellence Award.*

M **An Investigation of Aggregation Levels in Temporal Hierarchies for Heat Load Forecasting**, 40th International Symposium On Forecasting, October 2020, Virtual.

N **Methods for Identifying Critical Temperature for control of Low-Temperature DH Systems.**, Conference on Sustainable Development of Energy, Water and Environment Systems, October 2019, Dubrovnik, Croatia.

Workshops Presentations

O **Methods for data-intelligent operations of low temperature DH systems**, Data intelligent operation of district heating and district cooling systems, April 2019, Zagreb, Croatia.

Contents

Part I

Summary Report

1 Introduction

Methods for efficient and flexible operation of the energy sector are currently the focal point due to increasing shares of renewable energy sources across all sectors. The increase in renewables is expected to lead to sustainable energy production by enabling the phasing-out of fossil fuels and other greenhouse gas-emitting energy sources as climate change becomes more threatening due to fossil fuel energy production. However, the implementation of future weather-driven energy systems poses a challenge for every sector in terms of how to make the switch due to its intermittent nature, while a combustion plant driven by fossil fuels is a robust and steady method of generating energy on demand. Thus, efficient implementation of weather-driven energy systems calls for methods for unlocking flexibility everywhere in the energy system. Sector coupling and energy storage solutions are key elements for the needed transition. The challenge is also not identical when comparing the same sector between countries. It depends on the climate zone and the portfolio of energy sources for the particular geographical location and its energy sector.

For example, heat demand in Iceland is mainly provided by geothermal heat sources [1]. In contrast, heat generation in Denmark typically represents a more extensive portfolio of heat sources to fulfil the heat load, e.g., biomass, natural gas, oil, solar thermal or power-to-heat units [2]. Thus, district heating can either be operated with constant energy sources (e.g. geothermal production) or with variable renewable energy sources that are not always available (e.g. solar thermal) or are too expensive (heat pumps when electricity prices are high). Energy systems could also produce energy using a mixture of constant and variable energy sources. However, having variable energy sources, accurate methods must be applied to achieve optimal operation. For example, since most renewable energy sources are only available irregularly, another problem that needs to be solved in the transition of the energy sector to 100% renewable energy is how to store energy in periods of time when renewable energy sources produce more than the demand requires, i.e. when there is an energy surplus. For instance, in a country with a large installed wind power capacity, high wind power generation at night becomes a problem when there is low demand. There are several possibilities, e.g. it could be sold to another country, stored in batteries, or used to produce heat for district heating, either with electric boilers or heat pumps. This is a challenging problem, as a market player (the utility) has to evaluate which strategy is optimal to satisfy consumer demand and reduce operating costs.

At the same time, district heating systems are increasingly becoming more decentralized due to the newly available heat sources. For instance, industrial waste heat can be used either directly or with the help of heat pumps to raise the temperature before being fed into the network [3]. Solar thermal plants are designed to supply part of the entire network or are stored in seasonal TES systems [4]. Heat pumps are integrated into the network where they use an available heat source, e.g. ambient water [5]. Therefore, district heating is becoming more decentralized by integrating more local heat sources into the network. Exploiting these potentials of district heating can improve the overall efficiency of the smart energy system [6]. This can also incentivize the establishment of multi-temperature zones within the network, where different areas are operated with different supply temperatures with local heat units that can increase the supply temperature quickly by mixing. Areas with low-energy buildings can be operated with lower supply temperatures, where the primary heat source could be, e.g. heat pumps.

High energy conversion and energy storage efficiency are desirable as society becomes more reliant on electricity and requires flexibility due to the increasing share of intermittent energy sources. It is often mentioned that district heating will be one of the most crucial elements of future energy systems due to its high efficiency in converting other energies into heat and

its ability to store energy in its thermal energy storages [7]. It is a collective system that can efficiently supply heat from multiple heat units to many consumers in densely populated areas. District heating usually has several different heat generating units in its portfolio and selects the heat units to be used in order to keep production costs low depending on the future scenario, i.e. production optimization [8]. The future district heating will have even more heat units and more units that are driven by renewable sources. District heating systems offer a range of possible storage solutions, including the possibility of having several thermal energy storage (TES) systems in the network. Therefore, district heating can increase the flexibility of the overall energy system by taking advantage of fluctuating renewable energy production by generating heat with power-to-heat units when there is a surplus. Thus, when the heat demand is low, electricity can be converted to heat and efficiently stored in TES and discharged later to reduce peak load or electricity-based generation when the prices are high.

Hence, converting the surplus electric energy into heat is very attractive due to its high efficiency, e.g., using heat pumps, which can then be used directly to meet the heat load or stored as heat. Storing the surplus as heat in thermal energy storage systems is very effective, as it can be stored for a few days or long-term in large storage tanks. Hence, district heating systems offer opportunities for providing needed storage and flexibility to balance the fluctuating wind and solar power generation, especially if the network is operated optimally [9]. Optimal decision-making and operation of the district heating system are required to make this practical and as effective as possible.

Therefore, future district heating systems will become a very complex process to achieve optimal operation, increasing sector coupling and providing flexibility to the overall energy system. The system will become decentralized with the new additional heat sources and TES. Different supply temperatures will apply to different areas, depending on the ability of the areas to have low temperatures and the heat source. Residential Buildings have different energy efficiency levels and different dimensions, e.g. the radiators, and consequently different supply temperature requirements than industrial and commercial areas, e.g. hospitals. Figure 2.2 shows this type of system with a common heat production area connected to the different temperature zones via a transmission network, using heat exchangers to transfer the heat to the local distribution network. Each distribution line would have its local temperature levels zones, depending on the needed supply temperatures. For example, areas with new houses are generally more energy efficient. They therefore do not require as high temperatures as older, insufficiently insulated dwellings. Primarily, new houses with floor heating can be supplied with sufficient heat even with low supply temperatures.

District heating must be operated optimally to ensure that its flexibility potential is maximized, and a decentralized district heating system with several temperature zones requires data-driven methods for operation in order not to risk the main objective of district heating, namely to ensure that consumer heat load is met at the lowest possible cost. Accurate data-driven methods are necessary to support the district heating for making decisions or supporting the operations in decision-making. Operators need accurate forecasts (demand, prices, production plans, etc.) to take the optimal path for the operation. Therefore, the data and methods used in operations must be accurate and robust to achieve optimal operations. State-of-the-art methods must be challenged constantly, and new data-driven methods must be studied when, e.g., new sensors or data become available. However, it is also important to present the current district heating operation to establish a baseline, share knowledge and discuss the advantages and disadvantages.

This thesis presents novel data-driven methods to increase the efficiency and flexibility of the operation of district heating networks. Additional data on district heating from the digitalization transformation are presented, e.g., measurements from smart meters. It is shown how these can

be used to increase the efficiency of district heating operations. The focus will be on heat load forecasting and temperature optimization of the network. Both forecasting and temperature optimization are essential to improve network operation and exploit the flexibility potential. For instance, the efficiency of heat pump operation can be increased by lowering the supply temperature of the network using data-driven temperature optimization.

1.1 Research Questions

There are two main research questions (RQ) that form the basis of this PhD thesis:

RQ1: How can the current state-of-the-art heat load forecast be improved?

- (a) What are the potential accuracy improvements of using temporal hierarchy for heat load forecast by using forecast information at different resolutions?
- (b) What are the potential accuracy improvements of using spatial hierarchy for heat load forecast by using aggregated forecast information at different areas?

RQ2: Can the data from the digitisation of district heating systems be used to improve the operation?

- (a) Can data-driven methods be developed that use measurements from smart meters at the consumer level to reliably estimate network temperature such that it can be used in temperature control?
- (b) What are the potential savings to current network operations from using data-driven temperature optimisation to control the supply temperature?
- (c) What are the benefits of developing data-driven methods to describe thermal dynamics in buildings in a form that can be used for predictive control of the heating system?
- (d) Are local climate stations beneficial for district heating operations?

RQ1, aims to investigate how to improve the state-of-the-art operational heat load forecasts. Improving the accuracy of heat load forecasting is highly desirable in order to be able to increase the utility's operational efficiency. Forecast hierarchies were investigated to see if accuracy could be improved by using temporal or spatial hierarchies. First, additional heat load forecasts had to be made at new temporal or spatial aggregation levels. A grey box method is proposed to build a heat load forecasting model using physical knowledge of the system to improve model performance. Forecast hierarchies are proposed to improve operational heat load forecasts, which has the added benefit of imposing coherent constraints on the forecasts on various aggregation levels. For example, with temporal hierarchies, the short-term (half-hourly, hourly, etc.) and long-term (day, week, etc.) forecasts are coherent, facilitating improved decision-making for, e.g. energy production planning. For spatial hierarchies, the forecasts are coherent for the total and the sub-areas consumption, which improves decision-making in terms of production and temperature optimisation. Paper I describes an `R` package that produces online forecasts using recursive and adaptive methods, emphasising grey-box methods. Paper A proposes a new method for providing adaptive temporal hierarchies, and it is demonstrated that the adaptivity leads to significant improvements for heat load forecasts. Paper J explores the benefits of using spatial hierarchies and demonstrates significant accuracy improvements. Also, adding a new aggregation level to the hierarchy further improves accuracy and optimises the forgetting factor for each forecast horizon.

RQ2 aims to explore the possibilities arising from the digitalization of the district heating sector, e.g. smart meters installed at consumers' premises. Also, the frequently available data that is not

heat-related can be useful for operations. For instance, information from local climate stations is often available at several locations in urban areas. This information is precious to utilities as heat consumption follows the local temporal and spatial variations of the weather. Therefore, it has to be investigated how additional data to the standard data available (e.g. SCADA measurements) can be used to improve district heating operations. The potential savings from having data-driven temperature optimisation is also discussed and demonstrated in an online demo case by comparing it to a traditional temperature optimisation that is purely physically simulated driven. In Paper G, a case study is presented that investigates whether new additional data, measurements from smart meters and local climate stations, can be used to improve the operation of the network. Paper B and Paper C present novel methods for estimating network temperature using measurements from a set of smart meters. It is concluded that the estimated network temperature can be used, e.g., for temperature feedback for network temperature optimisation. Paper H suggests and demonstrates improvements for the heat load forecasting model accuracy by localising the weather forecast (the models' inputs) using climate station measurements. Paper D and Paper E propose novel methods for modelling the thermal dynamics of a building and using the model for predictive control of the building's heating system.

1.2 Overview

The thesis is structured in two parts. First, Part I introduces and summarises the s embedded in the thesis. Within this part, Chapter 2 introduces the state-of-the-art operation of district heating and explains which factors are essential for an efficient operation of a district heating network. In the next two chapters, the two main topics of this thesis are introduced and discussed with a focus on the outlined research questions. 1) Chapter 3 presents heat load forecasting and the suggested methods to improve today's state-of-the-art forecasts, 2) Chapter 4 gives an overview of data-driven temperature optimisation and presents how new additional data in district heating can improve the optimisation. Finally, Chapter 5 concludes Part I by discussing and summarising the results of this thesis. It also presents further research that is in line with the results.

Part II contains a collection of the articles included in this thesis.

2 Background

This chapter is intended to introduce the district heating system by explaining the technical aspects of the system. It will also present some of the challenges and opportunities of the system. The overview of the district heating system will be described from a Danish perspective, as district heating systems are different in different countries; it can be said that each district heating system is unique. This is due to the system characteristics, e.g. the heat units used, the pricing scheme, the characteristics of renewable generation (wind, solar, hydro, geothermal, etc.), and the social component, e.g. how the heat is consumed. However, the methods presented in this thesis are generalisable and should be suitable for any district heating system with minor adjustments to the local environment. Past, present and future district heating systems are discussed to provide a basis for the papers presented in this thesis and to demonstrate the importance of data-driven methods for district heating. First, the general district heating operation is presented in Section 2.1. Then, in Section 2.2, weather in cities is presented as it is the driving factor for heat consumption. Digitalisation in the district heating sector is introduced in Section 2.3, and it is explained why it will play an essential role in future district heating systems. Finally, Section 2.4 presents the future of the district heating system, and it is argued why data-driven methods will be a prerequisite for an efficient operation.

2.1 District Heating

The concept of district heating is to establish a centralised system for the efficient production and distribution of heat. District heating is made up of three main components, *heat generation units* (production), *network of pipes* (distribution), and *consumers* (consumption), as shown in Figure 2.1. This is the basis of district heating; this is how it was developed and how it is still operated today. The primary purpose is to deliver heat conveniently, efficiently, and safely to consumers. However, how heat is produced, delivered and consumed has changed over time and will continue to evolve to make operations more efficient and with the transition to more weather-driven energy production. District heating systems can become more efficient by either improving the currently available methods for using data (identifying problems, e.g. leaks) and using data-driven methods (e.g. temperature optimisation) or by updating their hardware by replacing old heating units with more efficient units, replacing old pipes (reducing heat losses) or adding new technologies to the system (e.g. heat pumps and storage solutions).

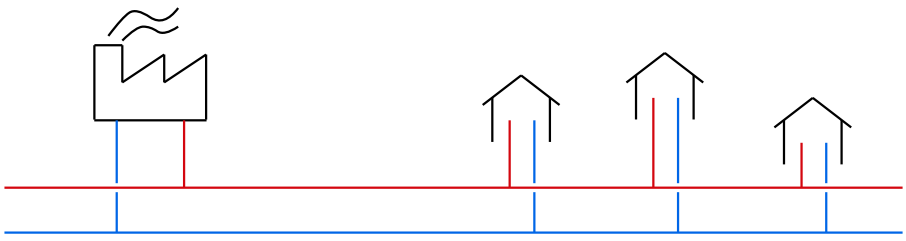

Figure 2.1: Traditional district heating system with central production, network and consumers.

District heating has undergone various generational phases in which operational efficiency has improved, as described in Lund et al. [6]. The first generation of district heating distributed heat in the form of steam generated by fossil fuels, primarily coal. This resulted in a system with low efficiency and high heat losses in the network, as the steam temperature was usually very high (up to 200℃). The second generation moved away from steam and towards pressurised hot water

with a temperature of about 100℃ (usually higher). In the third generation, the hot water temperature was lowered, usually to below 100℃. In addition, different heat units were introduced between generations by adding waste, biomass and solar thermal plants to the portfolio of fossil fuel plants. The 2nd generation also introduced combined heat and power (CHP) plants, which produce electricity and heat simultaneously. The residual steam from the electricity turbines that produce electricity is used to heat the water in the district heating network. These plants are also called cogeneration plants. During the 3rd generation, the CHP plants moved from oil due to the energy crisis (high oil price) to cheaper fuel, e.g. coal or biomass. So the general trend between generations is to lower distribution temperatures, improve generation flexibility, improve materials and expand the portfolio of heat units.

The fourth and current generation aims to keep the supply temperature below 70℃ and to continuously improve the pipe properties and the design of the supply and return pipes in the network to achieve higher efficiency. Buildings also need to make a greater contribution by increasing their energy efficiency through better-insulated envelopes and new methods of heating buildings, such as floor heating that can maintain thermal comfort with a lower supply temperature level than conventional radiators. Reducing the heating load for the individual buildings also makes it possible to have lower temperatures in the system. In addition, the term *smart* is often associated with the fourth generation as it focuses on operating the system intelligently, using data and data-driven methods to improve operations. For instance, Nielsen and Madsen [10] identifies a model that describes the heat consumption dynamics based on consumption measurements, which can be extended into a forecasting model to predict the heat load. It also drives the integration of different energy sectors so that they work together, i.e. smart energy systems by integrating electricity, gas and heat grids to achieve optimal operation for the individual and the overall system [6]. Moreover, energy systems are coupled with integrating heat with electricity supply, i.e. CHP plants or heat generation with heat, e.g. heat pumps. This type of sector coupling is becoming increasingly important. It will play an important role in providing flexibility in power consumption for balancing large fractions of wind and solar power. Future energy systems will have a high share of renewable energy systems (RES) to produce energy. The overall energy system needs to focus on flexibility to maximise the potential of intermittent energy sources, for instance, as described in Dominković et al. [11] where electricity and district heating networks are coupled to increase savings by utilising the flexibility of hot water storage in the network. Also increase flexibility and efficiency with a combination of residual resources, e.g. waste and biomass. Therefore, intelligent decision-making in system operation is required, considering several signals before making an optimal decision, e.g. a combination of historical measurements, weather forecasts, price signals and future load on the system to achieve optimal operation.

Regardless of their generation, district heating systems are usually physically large and complex to operate. However, it can be argued that the latest generation is even more complex to operate due to the coupling of the sector with other energy sectors. The operation is driven by the heat consumption of the consumers, which is momentarily and can change quite abruptly. The system is physically large as many kilometers of pipes are used to transport hot water to consumers to satisfy the load. For instance, the total length of the district heating system in Lund, Sweden, is around 2000 km. Transporting heat to consumers is a complicated process as two variables are regulated to fulfil the heat load: the supply temperature and the flow rate (assuming the return temperature is constant). Changes in flow (and pressure) in the network are 1000 times faster than changes in temperature, but there is a limit to the maximum flow in the system. If the flow is above the maximum for too long, the consumers at the end of the network will not receive enough heat. Therefore, the time τ it takes to send one unit of hot water from point A to point B must be known so that the hot water arrives with sufficient temperatures before the flow limit

is reached. Therefore, due to the time delays, the expected future heat load at time $t + \tau$ governs the operation at time t. Accurate prediction of heat load facilitates the operation. However, production optimisation must then be done to select which heat units should be used to produce the heat to minimise production costs. Also, the operation of the pumps (adjusting the flow) needs to be considered during periods of extremely high electricity prices. Hence, electric price forecasts and bidding strategies are also required, e.g. considering the time-varying costs for the pumps and if the district heating also produces electricity.

Thus, several complex decisions have to be made to find an optimal strategy to meet the heat load at the lowest possible cost. This must be done without violating the requirements or technical constraints of the system. Due to the system's complexity, the three main systems (*production*, *network* and *consumers*) are presented separately in the following section. However, it is important to realise that the performance of each component affects the other components, and the others constrain each component. Consequently, if you improve one component without violating a constraint, another component will also benefit. For instance, if the supply temperature in the *network* is lowered, the amount of heat that needs to be generated in *production* will be reduced due to the reduced heat loss in the network. Therefore, it will also reduce the heating cost for the *consumer*.

2.1.1 Heat Production Operation

Heat production generates heat, which is then delivered to consumers using the hot water in the network. The heat can be generated with several different heat units, which can be centralised or decentralised across the network. For instance, CHP plants, solar thermal plants, peak load boilers, power-to-heat (e.g. heat pumps) or industrial waste heat can be used to generate heat. The combustion-driven heat units can use different fuels as input, e.g. coal, natural gas and wood chips. Cogeneration of electricity and heat leads to high efficiency because the steam from combustion that powers the turbine drives the generator producing electricity, which must then be condensed before it is reused in the cycle. The return water in the district heating network can therefore be used as a coolant in the condensation and reheated. As a result, the efficiency of the CHP plant can be very high (> 90%). In contrast, the combustion process without heat generation for pure electricity generation would be about 40% for coal-fired power plants and about 60% for gas-fired combined cycle power plants. However, there are some limitations for CHP plants, e.g. the supply temperature can reduce electricity generation if it needs to be high. If the return temperature is too high, the condensation process is less efficient [12]. As society becomes more and more electricity-dependent, district heating production needs to follow the unbalance in electricity generation and load due to the shares of RES so that the electricity system can take advantage of the flexibility of district heating, the ability to store heat and the high efficiency of converting electricity into heat. Therefore, district heating production needs to consider both electricity and heat load profiles, which are somewhat similar. The diurnal profile of heat load is usually low at night, with a peak in the morning and then again in the evening. At the same time, the electricity load is also low at night, rises to some level in the morning and usually stays there over the day, then peaks in the evening.

Optimisation of district heating production is necessary to minimise production costs while meeting the desired heat load and maximising flexibility potential [2]. It consists in selecting the different heating units that need to be activated to fulfil the load at a given time. There may also be TES systems that can be charged and discharged to the network. Therefore, several factors influence the production costs, e.g. the cost per unit of energy for each heating unit, the price of electricity, the bidding strategy for electricity generation and the heat load. The future weather forecast and its uncertainty are crucial factors that need to be taken into account, as it affects all the previously mentioned factors. System-related constraints must also be taken into account, e.g. the desired supply temperature, which limits electricity generation. The planning horizon

is also important for bidding on the electricity markets (day-ahead and balancing market) and the long time span of start/stop planning of heating units. In addition, due to the time lag in the system, heat production has to take place several hours in advance (in some large systems, even for days) to meet the load at a later point in time. Accurate heat load forecasts and a good understanding of the forecast's uncertainty will improve production optimisation [13].

Furthermore, the power production from CHP plants can be traded on the electricity markets, which can generate additional revenue and thus reduce overall costs. However, they also have to buy electricity for their power-to-heat plants on the market. Clearly, production optimisation is a complex process that will increase in complexity and uncertainty with a larger share of RES plants, power-to-heat and bidding on the electricity markets. In short, production optimisation is about identifying a cost function for operating a district heating system and choosing a strategy to meet the heat load of consumers while minimising costs. A further and more detailed introduction to production optimisation for district heating can be found in Guericke, Schledorn, and Madsen [2] and Blanco [14]. Also, the complexity of electricity price forecasting due to the volatility of RES can be found in Jónsson et al. [15].

2.1.2 Network Operation

A district heating network consists of a mixture of pipelines with different pipes and pumps that transfer the water through the pipes. These pipes consist of two main loops, the supply and the return water loop, which supply hot water to the consumer and return it after the heat has been transferred over to the substations. These loops are usually laid at a depth of about one meter underground. The depth and total length of the network depend on the system, but they are generally very long, hundreds to thousands of kilometres. The pipes are usually made of steel. However, for a network with very low temperatures, it has been suggested to use pipes made of plastic [16]. The network pipes are pre-insulated to minimise heat loss in the network, but there will always be some heat loss. However, with an intelligent operation, this can be kept to a minimum.

Because of the enormous length of the network, heat loss is an essential factor in operating the network. A rule of thumb states that heat loss is at least 5% of the total heat load, and more often, it is between 10% and 20%, depending on the system and operation. The heat loss is highly dependent on the supply temperature and is proportional to the difference between the water temperature to the surrounding ground temperature. So the higher the supply temperature, the higher the heat loss. The heat loss also depends on the condition of the pipes (e.g. age, insulation, etc.) and the surrounding soil (humidity and temperature). Keeping the heat loss low, the supply temperature must be kept as low as possible while varying the flow to fulfil the heat load. However, there are limitations to how low the temperature can be due to e.g. temperature restrictions and the risk of legionella. Also, lowering the temperature results in a higher flow, and the system has a maximum flow limit based on the pressure limit of the hydraulics. In practice, the maximum flow limit also depends on the pumping costs. There is also a risk of cavitation in pumps and substations. Reducing heat losses is very important for the operation of the district heating network, as it reduces the amount of fuel required. Also, for CHP plants, lowering the supply temperature results in an increase in the ratio of power to heat output, and electricity is more valuable than heat [12]. This improves the economic and environmental operation of the district heating network.

Network operators usually only adjust the supply temperature in the network, while the flow rate and return temperature depend on the consumer operation, the load and the operation of their heating system. While the supply temperature is adjusted, the flow through the substations is regulated to fulfil the load. The return temperature depends on the ability of the substations to cool down the hot water. District heating networks are usually designed to maintain a fixed

differential pressure in the network, which the pumps try to maintain by varying the flow. Due to physical limitations, changing the supply temperature at one point in the network also takes time (time delay). If the supply temperature is set too low, this can cause damage to the pumps and substations if they are operated at high flow for a long time. The time delay is therefore an important variable for the network operators. Furthermore, the time delay is also time-varying, as it depends on the flow in the system. These factors, *time delay*, *heat loss* and *pressure* (also friction losses, but these are usually negligible) affect the two main components used to meet the heat load, *temperature* and *flow*.

A district heating network operation requires an advanced method that uses physical knowledge to optimise the supply temperature to minimise costs while providing sufficient temperature to consumers. This must consider the system's physical nature (network characteristics), measurements of the network (temperature feedback) and future heat loads to find optimal future set points for the network temperature. A good understanding of the physical limitations of the network and the characteristics of the buildings to which the network supplies heat. Then, measurements can be used to estimate all the non-linear relationships between all the variables, e.g. the time delay varies with the change in flow, and the time delay affects the heat loss, or the longer it takes to transport the water, the higher the temperature loss will be. The ability to operate the network with high precision and thus keep the temperature low and not have a too high strain on the network without violating any requirements has significant advantages for the utility. In the future, district heating will be strongly integrated into smart energy systems and introduce a higher proportion of RES units, which usually have higher operating efficiency at lower temperatures in the network. The introduction of more power-to-heat and solar thermal units into the system, together with TES, will increase the complexity of temperature optimisation of the network as these units will operate in a more decentralised way in the future, e.g. district heating network with multi-temperature zones. A hierarchy of controllers will be required to perform adequate temperature optimisation of the network.

2.1.3 Consumer Operation

The consumer operation consists of a substation located at the consumer and transferring heat from the primary side (district heating) to the secondary side (building) to maintain thermal comfort (indoor air temperature) and deliver domestic hot water. There are usually two heat exchangers at each substation, one for space heating and one for domestic hot water. A hot water tank can also be installed at the consumer. Depending on the situation, the consumer may need to maintain both the house's substation and water-based heating system, e.g. pipes, radiators, floor heating, etc. From a district heating point of view, an adequate substation cools down the hot water to an acceptable level before returning it to the network. An adequate substation for the consumer maintains the indoor air temperature within the thermal comfort level and the constant temperature of the domestic hot water consumption and is reliable and inexpensive.

The consumer's heat load consists of space heating and domestic hot water, while district heating is the consumer's total heat load plus the system's heat loss. A significant part of the heat load originates from space heating to maintain thermal comfort. The heat load varies over the day (diurnal profile), the week (weekly profile) and the year (annual profile), creating a complex time-varying process. This time-varying behaviour coincides with the climate and social consumption behaviour in the place where the district heating is applied. The ambient air temperature is the most significant climate variable that drives the heat load. Heat is used to maintain certain indoor air temperatures when the ambient air temperature is low to compensate for heat loss through the building envelope. This is referred to as physical heat dependency. The social component drives the heat peaks in the diurnal profile and the different profiles between days (e.g. the difference between weekdays and weekends). Social heat consumption is driven by the area's cultural behaviour, for instance, showering before work at certain hours (hot tap water consumption)

and the perspective of thermal comfort.

A well-calibrated substation satisfies the consumer's heat load and cools the hot water on the district heating side to the accepted temperature. The lower the return temperature, the better the network will perform, e.g. by decreasing the heat loss and reducing the necessary flow for satisfying the consumer heat load. This makes it feasible to reduce the supply temperature in the network. The typical operation of substations consists of pumping the hot water through the heating system to maintain thermal comfort based on the radiator thermostat setting. In this way, the space heating is cut off when it is not needed. Sometimes a schedule management tool is also installed in the substation to use different setpoints for different periods, e.g. a night setback schedule. The night setback lowers the setpoints of the heating system during the night, as it is assumed that this will not affect thermal comfort during sleep. This results in energy savings for the buildings and lower consumer heating costs. However, this can increase the heat load during peak hours, as more heating is needed when people wake up to achieve normal thermal comfort [12]. In addition, buildings today are better insulated, which leads to higher time constants. Therefore, the indoor air temperature does not drop as quickly, so less heating is needed during the night.

Future district heating systems require buildings to become more low energy driven by constructing new energy-efficient buildings or renovating them (e.g. through better insulation) to ensure optimal operation of the low-temperature networks. Then less heat is needed to maintain thermal comfort, and the supply temperature can be lowered so that other heat units, such as heat pumps, can better be used. Also, intelligent control that uses data and advanced predictive methods are becoming more popular as awareness of inadequate and inefficient heating operations in buildings grows. In addition, more measurements are available (and cheaper), e.g. indoor air temperature, CO2 and humidity, which can be used to improve heating operations. There is always a meter in the substations to track energy consumption, which is used for billing the consumer for heating consumption. The heat consumption is calculated based on the substation's temperatures and flow measurements on the district heating side. Other sensors can be attached to the building's heating system, radiators and indoor climate to measure important variables and optimise the heating system with various cost functions, e.g. to maintain thermal comfort or reduce heating costs. There may also be some motivating tariffs from district heating, e.g. bonuses or penalties based on the return temperature. In the future, time-variable prices for district heating may also become common, like in the electricity and gas sectors.

2.2 Weather in Cities

The operation of the district heating system is highly dependent on the climate in which the consumer is located. Therefore, the climate's effects on the consumers' heat consumption must be analysed and understood to achieve optimal operation. In particular, the ambient air temperature is the most important single factor determining heat consumption. Most of the heat consumption is used to maintain the indoor air temperature within the thermal comfort of the building. However, many different aspects affect the influence of climate on heat consumption in different areas, cities, countries and continents, such as building infrastructure and social behaviour - or, as mentioned before, *Physical* and *Social* heat load.

Nielsen and Madsen [10] identify which and how climate variables can be used to model heat consumption in a district heating system. Three climate variables are proposed for use, *Ambient Air Temperature*, *Solar Radiation* and *Wind Speed*. A detailed physical description of how these climate variables influence heat consumption based on physical and steady-state considerations is proposed. The physical relationship between maintaining thermal comfort and heat transfer in buildings through walls, windows and ventilation is described. A physical derivation model is proposed and then translated into a model that can be identified with the available measurements

since the model must be structurally identified and estimated to be of any use. Nielsen and Madsen [10] propose to use a rational transfer function to estimate the effect of climate on indoor air temperature as filters. The instantaneous effect of climate values is also important. The model and the proposed method for building an online heating load prediction model are shown to perform adequately. Therefore, an important step in building a model that accurately captures heat consumption is to identify which climate variables affect heat consumption and how they affect it. It is important to recognise what drives the consumption for each district heating system as it can vary between systems, especially between countries with different climates.

A detailed understanding of the impact of climate on heat consumption at the district heating plant site is therefore essential for optimal operation. This knowledge can then be used to create a heat load forecast model that is used as input for all operations, e.g. temperature control and production planning. The temperature optimisation of the network is also strongly influenced by the climate. However, to be able to forecast the future heat load, future values of the climate variables are needed as inputs.

Therefore, a weather forecast of the essential climate variables is required, and the correlation between the weather forecast and the heat consumption must be taken into account when building forecast models. Numerical Weather Prediction (NWP) is used to forecast the weather by simulating physically based partial differential equations of atmospheric processes [17]. However, weather forecast models are usually tuned to rural areas and not urban areas where district heating is used. The difference between the climate in rural and urban areas is considerable. Research shows that the ambient air temperature in urban areas is usually higher than in rural areas [18]. This effect is termed urban heat island (UHI). A UHI is an urban area that is warmer than the surrounding rural areas due to human activities or human building infrastructure. In most cases, the NWP does not represent the UHI effect and thus the weather within cities, so a systematic bias between the NWP and the local weather stations is often observed [19]. It would therefore be beneficial to localise the weather forecast to remove bias and thus reduce the forecast uncertainty of weather forecasts for DH applications.

2.3 Digitalization in District Heating

Almost all industries have or are going through a transition where sensors are installed at key locations where measurements are required. The district heating sector is also going through this transition, where more and more data is becoming available from places that were not available before. For instance, heating-related measurements at the consumers and the city climate are measured through smart meters and local climate stations. Previous generations of heating meters were usually only used to measure energy consumption in order to bill the consumer. However, the European Union started to require that houses which are connected to a district heating network are equipped with smart meter devices where feasible [20]. Smart meters are IoT sensors that regularly send data back to a centralized database. In district heating, measurements from these smart meters are still used to generate bills based on consumers' consumption. However, these measurements are now available more frequently and often in near real-time, and in addition, the consumers also have the possibility to view them.

It has taken some time for the district heating sector to reap the benefits of frequent measurements, but recently these measurements have become attractive to research and industry. New generation measuring devices have proven helpful in identifying leaks with high precision of the leakage location [21, 22]. Fast and accurate leak detection in district heating can save the utility a lot of money and time. Intelligent control of building heating systems also needs these readings to determine optimal setpoints to minimise the cost function based on user requirements, such as reducing energy consumption or lowering the return temperature.

Weather measurements and forecasts are also part of this digitalization and are very important for the operation of district heating. More and more local climate stations are being installed as these measurements are becoming increasingly important for several different sectors. For instance, the local climate stations can be used to determine the uncertainty of the NWP (i.e. the reliability of the forecast for the area), which can be used in the heat load forecasts to improve the uncertainty of these forecasts for decision-making in temperature and production optimisation. In addition, more detailed information on the local climate will increase the knowledge of how heat consumption is related to climate. This can improve district heating operations by using this knowledge, e.g. as a more accurate input for heat load forecasting models.

IoT sensors are attractive to all sectors, especially the renewable energy sector. These sensors and measurements have the potential to improve operational efficiency and thus reduce operating costs. However, big data can be a burden for companies, handling large amounts of data and extracting valuable information from the data. Therefore, it is now important to derive useful information from these measurements to improve operations and thus reduce costs. Data-driven methods can benefit from these measurements to enrich them with additional information that makes them robust or even enables the use of data-driven methods.

2.4 Future District Heating Systems & Data-Driven Methods

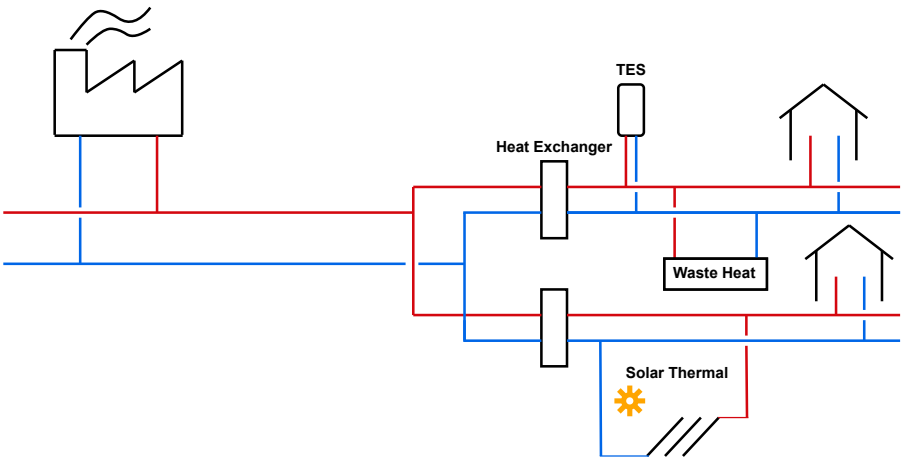

Figure 2.2: Decentralised district heating system with local heat units and TES for each area along with one main production site. Here multi-temperature zones could be used by optimising the heat exchangers with different supply temperatures.

District heating is an essential component of current and future smart energy systems due to its unique ability to store heat using thermal energy storage (TES) systems. It also has high efficiency in converting electricity into heat, e.g. using heat pumps (power-to-heat). Therefore, district heating has a high potential to make the energy system more flexible. However, the higher the temperature of the hot water in the system, the lower the efficiency, as is the case with heat pumps [5], TES [7], and CHP plants [12]. Thus, operating the district heating network with low hot water temperatures to fulfil the heat load of the consumers has significant advantages for the district heating utility and the overall energy system, as district heating maximises its flexibility potential. Also, as more data is available than ever in the district heating sector, the new data could have the potential to increase operational efficiency. Data has the unique ability to improve operations as it provides the ability to measure performance and react to changes. More importantly, it provides the ability to understand the system's dynamics to build models that predict the future and make intelligent decisions based on this knowledge. As district heating systems become an important player in the transition to the future energy system driven by

renewable energy, it is important that operational efficiency is also improved for the existing systems, opposite to purely focusing only on the new district heating systems. For instance, installing heat pumps in a current high-temperature district heating system would lead to a lower operational performance due to inefficient temperature optimisation.

Traditional district heating systems, as discussed in Section 2.1, are operated in a conservative manner without attempting to use new methods to improve efficiency. District heating is also becoming more data-rich as more IoT sensors are installed, e.g. smart meters and local climate stations. This makes it possible to use data-driven methods to enhance their operation in terms of delivering the heat load at minimal cost, thereby increasing savings. It can also improve the efficiency of operations in other energy sectors. District heating systems are changing from a traditional system with centralised operation and large production units to a system with still large production units but also new additional heat sources and thermal storage, i.e. a more decentralised system. Hence, moving from the system shown in Figure 2.1 to an example of a decentralised system with several areas and heat sources (with thermal storage) as shown in Figure 2.2. The decentralised system can have multiple temperature zones, as it can be divided into low-energy buildings, old buildings with insufficient thermal insulation or industries with a high heat load. Within these areas, different heat sources can be used to satisfy the heat demand - with the help of thermal energy storage. Also, with the mixing of temperature, the supply temperature can be increased rather fast by using e.g. residual heat from supermarkets compared to raising the temperature at the plant due to the time delay. Large production plants, such as CHP plants, can still be beneficial for the system to improve sector coupling by producing electricity when prices are profitable at low RES periods, as highlighted in David et al. [5], and using power-to-heat technologies at high RES production when electricity prices are low.

In summary, the operation of the district heating system is complex and consequently calls for data-driven methods as energy operators of district heating have the difficult task of finding the optimal strategy for the utility. First, the operational production planning of the system has to be optimised. Decision-making is complex as several things need to be considered, e.g. which heat units to produce the heat, operate the network, act on the electricity market, etc. In addition, the systems often contain TES that need a plan for when to charge and discharge. Hence, optimising production becomes a highly complex problem to solve. Secondly, the network must be efficient to improve production operations by keeping the temperature stable and as low as possible by using data-driven temperature optimisation. The district heating sector is becoming more decentralised with more heat units in the network and waste heat from industry, increasing temperature optimisation complexity. Thirdly, the heat consumption behaviour of consumers is also changing as there are more incentives to keep heat consumption low and to operate their substations optimally. Efficient operation of the consumer's heating system has become easier due to new technologies (efficient thermostats, availability of data) and greater awareness among people to operate the heating system sufficiently.

This complexity of operating a district heating system shows that more advanced tools are needed to help operators meet heating needs while minimising costs. While district heating is becoming more complex, more data is available to support decision-making for each operation. At the same time, advanced methods and a better understanding of them and the system can make decisions more accurate. However, the rapid changes in the district heating industry can be a burden for a utility. There is a need to build robust and precise production optimisation, temperature optimisation and efficient substation operation for the consumers. This thesis focuses on *heat load forecasting* and *temperature optimisation* and how new data sources can be used to improve district heating operations using data-driven methods. Accurate heat load forecasts are essential for all operations as input for the three main components of district heating in order to optimise their operation. Temperature optimisation is required to increase production opti-

misation's flexibility and ensure that consumers receive the desired heat without violating any constraints.

3 Heat load Forecast

Forecasts for the district heating sector are essential to achieve optimal operation of the system. As the district heating operation will become essential in providing flexibility for future energy systems, an optimal operation must be achieved to maximise its potential. Forecasts of the heat load are therefore necessary to enhance the decision-making for district heating systems. It prepares the operation for the expected load, and the more accurate they are, the better decisions are made. For instance, temperature optimisation needs a heat load forecast at time t for the future time to select optimal temperatures that will reach the consumers at time $t + \tau$ as it takes time to transport the hot water. Hence, the physical nature of the system requires an accurate expectation of the future such that the system operates efficiently.

Traditional district heating systems are usually operated with only one central production area with multiple heat units, e.g. CHP plant and solar thermal plant, also perhaps one or two peak boilers in the network to operate during peak hours. Thus, a forecast of the total heat load of the system would be sufficient where the inherent nature of heat load is adequately modelled, i.e. the non-linearity and non-stationarity. However, as mentioned in Section 2.1, district heating systems are becoming more complex. The systems are becoming more decentralised with areas by operating with their different temperature levels, TES systems, and local heat units, e.g. heat pumps. Thus, more localised operations are needed with multiple temperature optimisations and local heat load forecasts. Since heat load is mainly determined by dependence on the weather and the rest by the consumer's dynamics, a localised heat load forecast is needed to incorporate the local climate and social patterns. Also, continuous improvements to the current state-of-the-art forecasts are needed as better forecast accuracy will increase savings through better decision-making and increase the flexibility potential of district heating. New methods are constantly being proposed to challenge the current state-of-the-art. However, it has been shown that higher accuracy is achieved when multiple forecasts are used than compared to use only one forecasting method is used [23, 24]. A new method, hierarchical forecasting, has demonstrated potential in energy forecasting by exploiting information sharing between either temporal, spatial, or temporal-spatial hierarchies. It is a special branch of combination forecasting by reformulating it into a combination of direct forecasts using the linear coherent constraints [25].

This chapter presents heat load characteristics and forecast methods. Section 3.1 explains how the heat load models were built in this thesis by using physical knowledge of the system and use statistical methods to estimate model coefficients, i.e. grey-box modelling. The grey-box method is also used for model building and system identification using physical insights to find a relationship between the response and the input variable. Also, the model coefficients are updated recursively, and past information is down-weighted using exponential weights. Section 3.1 is based on the results from Paper F and Paper I. Section 3.2 introduces a localised forecast for district heating, where an individual model for each area is shown to be important, based on the results from Paper G. As mentioned in Section 2.2, an accurate forecast of the weather in the heat consumption areas has the potential to improve the accuracy of the heat load forecast. Therefore, in Section 3.2, the adjustment of the NWP by adapting to the local climate is presented based on the results from Paper H. Finally, hierarchical forecasting is presented in Section 3.3 with the aim of improving the state-of-the-art in operational heat load forecasts by sharing information between forecasts with temporal and spatial hierarchies based on the results in Paper A and Paper J.

3.1 Heat Load Characteristics and Forecasting

Paper F & Paper I

Figure 3.1: The figure shows heat load from one area in Brønderslev. The top plot shows the daily average of the heat load and the ambient temperature over one year, while the lower plot shows the hourly consumption over one week and the ambient temperature.

The heat load consists of two components, space heating and domestic hot water. It can also be argued that heat loss in the system should also be taken into account. Bacher et al. [26] divides the heat load in a single-family building into two components, space heating and domestic hot water usage. Space heating is described as a slow-moving component because it responds to a low-pass filtered response to ambient air temperature and solar radiation, i.e. a smooth pattern due to the thermal insulation of the envelope. Hence, the non-linearity part of the heat load as the demand does not react instantaneously as the weather changes. On the other hand, the domestic hot water component is noisy and has a high frequency. The dynamics of the heat load are therefore unique for each area as it is strongly influenced by the local climate, space heating, and the noisy component of the local consumption pattern, domestic hot water. The heat load also varies over time as the climate changes; e.g. during warm periods, space heat is unnecessary, while during cold, the heat load is governed by space heating.

Consequently, heat load dynamics vary between countries, within countries and even within cities. The within-city variation is due to differences in building infrastructure and occupancy behaviour. For example, the requirements of industrial areas, old residential buildings and low-energy houses are vastly different. Knowledge of the local climate is crucial for understanding the dynamics of space heating, as heat load is strongly correlated with climate, e.g. ambient air temperature. At the same time, domestic hot water is driven by occupancy behaviour, e.g. showering in the morning. Figure 3.1 shows in the upper plot the time-varying relationship between the heat load and the ambient temperature, where the heat load is high in cold periods and low in warm periods. The lower plot shows the hourly consumption over a week with the ambient air temperature to illustrate the daily and weekly profiles. The daily heat load profile usually has two significant peaks, but only the morning peak is significant in this example. Thus, this demonstrates the difference in heat load between areas. The weekly profile shows the difference between days, especially between weekdays and weekends. This indicates that heat load is a non-stationary (time-varying) and non-linear process driven by weather factors and social behaviour.

District heating operation is thus mainly determined by physical components. Therefore, it can be argued that applying grey-box modelling to build a model for heat load, where the components of the model are found using physical intuition, together with statistical techniques to identify other components and estimate the parameters of the model, is beneficial. Nielsen and Madsen [10] propose to use grey-box modelling to build a model for heat consumption where steady-state physical relationships (the dynamic response to changing climate conditions) are used together with calendar information and social behaviour to build the model structure. They also distinguish between the full physical model and an identifiable model where the model coefficient can be estimated from available measurements. This modelling approach for heat consumption and the proposed forecasting method is described in more detail in Nielsen and Madsen [27]. Using this physical knowledge of how the response variable reacts to changes from the input variables will improve the model accuracy in capturing the dynamics of the system, i.e. grey-box modelling. The coefficients of the models are then estimated using statistical methods, e.g. Ordinary Least Squares (OLS) or Maximum Likelihood (ML). The heat load forecasting framework must be computationally fast and adaptive, i.e. as new information becomes available, the model must be updated and produce new forecasts. The framework must also be able to handle the non-stationarity and non-linearity that is prevalent in heat load. Thus, the framework needs to be robust and simple so that it can be reconfigured when new information becomes available. In addition, the framework creates a model for each horizon, such that the models are tuned for each horizon. This thesis proposes to use recursive and adaptive regression models to produce a robust and accurate heat load forecast. The model framework described below is the framework that is used in all publications in this thesis regarding forecasts.

The proposed framework is a regression model with a dependent variable Y_t and p independent variables, $X_{1\mathrm{t}}, \ldots, X_{\mathrm{pt}}$, and its general form is written,

$$Y_\mathrm{t} = f(X_\mathrm{t}, \mathrm{t}, \theta) + \epsilon_\mathrm{t} \tag{3.1}$$

where $f(X_\mathrm{t}, \mathrm{t}, \theta)$ is a known mathematical function of the $p+1$ independent variables, $X_\mathrm{t} = (X_{1\mathrm{t}}, \ldots X_{\mathrm{pt}})^\mathrm{T}$ and t where t is the time index but with unknown coefficients $\theta = (\theta_1, \ldots, \theta_\mathrm{m})$. ϵ_t is a random variable with $\mathbb{E}[\epsilon_\mathrm{t}] = 0$, $\mathrm{Var}[\epsilon_\mathrm{t}] = \sigma_\mathrm{t}$, and $\mathrm{Cov}[\epsilon_{\mathrm{t,i}}, \epsilon_{\mathrm{t,j}}] = \sigma^2 \Sigma_{\mathrm{i,j}}$.

Regression models usually take the form of a general linear model,

$$Y_\mathrm{t} = x_t^\mathrm{T}\theta + \epsilon_\mathrm{t} \tag{3.2}$$

where the model is now a linear combination of the independent variables, and the coefficients can be estimated using maximum likelihood methods. However, suppose the assumptions of independent and/or normal distributed observation are disregarded. In that case, the estimates can still be found using least squares methods, e.g. OLS or Weighted Least Squares (WLS), and be unbiased. The estimates can thus be found,

$$\hat{\theta} = [X^\mathrm{T}\Sigma^{-1}X]^{-1}X^\mathrm{T}\Sigma^{-1}Y, \tag{3.3}$$

where it minimises the sum of squared errors,

$$S(\theta) = (Y_\mathrm{t} - X\theta)^\mathrm{T}\Sigma^{-1}(Y_\mathrm{t} - X\theta) \tag{3.4}$$

The OLS estimate has $\Sigma = I$ (i.e. unweighted), while the WLS estimate has known weights in Σ. From the estimates, the prediction of future values of Y at time $t+k$ is calculated as a function of the independent variable $x_{\mathrm{t+k}}$,

$$\hat{Y}_{\mathrm{t+k}} = \mathbb{E}[Y_{\mathrm{t+k}}|x_{\mathrm{t+k}}] = x_{t+k}^\mathrm{T}\hat{\theta} \tag{3.5}$$

Heat load forecasting models from this methodology are widely used, e.g. Dotzauer [28] creates a forecasting model from a linear combination of the relationship between ambient temperature and load and the social profile component of heat load. The coefficient is estimated using the OLS method. The proposed forecasting model is very simple; nonetheless, it demonstrates accurate predictions of heat load. In Dahl et al. [29], three machine learning methods are proposed, one of which is an OLS regression model. The OLS model performs the worst of the three methods. However, the inputs in the OLS model were treated as linear dependence on heat load due to its linear structure, while the other two methods are non-linear models.

The estimates of these methods lead to static estimates and only allow a linear combination of the independent variables for the OLS. Therefore, they are not ideal for heat load prediction as it is a non-stationary and non-linear system. Furthermore, heat load forecasts are needed for the online operation of the system, e.g. temperature optimisation is usually evaluated hourly or when new observations are available. Due to the non-stationary heat load, it is also desirable to give less weight to past observations or even remove them from the estimation procedure to obtain adaptive estimation. In addition, heat load forecasting models have been widely studied in the literature, but the operational aspect is often ignored. It is sometimes unclear whether the errors and error evaluations are performed in-sample or out-of-sample, where an in-sample should be used to tune the parameters of the model (model building), while an out-of-sample should be used to investigate the forecasting performance of the model. Also, the robustness of the proposed forecasting model, how it performs over a long period during online operation, and how much data is needed for estimation. A discussion of how robust the coefficients are or how they change over time is frequently not included. This and other information is required in order to build a robust operational online forecasting model.

Online and adaptive forecasting methods imply a framework that allows model estimates to adapt to slow changes in the system during operation. This thesis proposes to use the regression model in Eq. 3.1 estimated by adaptive recursive estimation with exponential forgetting. To recursively estimate the parameters of Eq. 3.2 so that the parameter estimates are updated as new observations become available. Start from the parameter estimate in Eq. 3.3, but rewrite it as,

$$\hat{\theta}_{\mathrm{t}} = R_t^{-1} h_{\mathrm{t}}, \tag{3.6}$$

where

$$R_{\mathrm{t}} = \sum_{i=1}^{\mathrm{t}} x_{\mathrm{i}} x_i^{\mathrm{T}}, \text{ and } h_{\mathrm{t}} = \sum_{i=1}^{\mathrm{t}} x_{\mathrm{i}} y_{\mathrm{i}}. \tag{3.7}$$

The recursive update of these two matrices becomes,

$$R_{\mathrm{t}} = X_t^{\mathrm{T}} X_{\mathrm{t}} = \sum_{i=1}^{\mathrm{t}} x_{\mathrm{i}} x_i^{\mathrm{T}} = x_{\mathrm{t}} x_t^{\mathrm{T}} + \sum_{i=1}^{\mathrm{t-1}} x_{\mathrm{i}} x_i^{\mathrm{T}} = x_{\mathrm{t}} x_t^{\mathrm{T}} + R_{\mathrm{t-1}}, \tag{3.8}$$

$$h_{\mathrm{t}} = X_t^{\mathrm{T}} Y_{\mathrm{t}} = \sum_{i}^{\mathrm{t}} x_{\mathrm{i}} y_{\mathrm{i}} = x_{\mathrm{t}} y_{\mathrm{t}} + \sum_{i}^{\mathrm{t-1}} x_{\mathrm{i}} y_{\mathrm{i}} = x_{\mathrm{t}} y_{\mathrm{t}} + h_{\mathrm{t-1}}. \tag{3.9}$$

Eq. 3.6 then becomes,

$$\hat{\theta}_{\mathrm{t}} = R_t^{-1} h_{\mathrm{t}} = R_t^{-1}(x_{\mathrm{t}} y_{\mathrm{t}} + h_{\mathrm{t-1}}) = R_t^{-1}(x_{\mathrm{t}} y_{\mathrm{t}} + R_{\mathrm{t-1}} \hat{\theta}_{\mathrm{t-1}}) \tag{3.10}$$

$$= R_t^{-1}(R_{\mathrm{t}} \hat{\theta}_{\mathrm{t-1}} - x_{\mathrm{t}} x_t^{\mathrm{T}} \hat{\theta}_{\mathrm{t-1}} + x_{\mathrm{t}} y_{\mathrm{t}}) \tag{3.11}$$

$$= \hat{\theta}_{\mathrm{t-1}} + R_t^{-1} x_{\mathrm{t}} (y_{\mathrm{t}} - x_t^{\mathrm{T}} \hat{\theta}_{\mathrm{t-1}}). \tag{3.12}$$

Thus, Recursive Least Squares (RLS) estimation. However, all observations are given the same weight; therefore, the estimation will become less adaptive as more observations become available. Ljung and Söderström [30] propose an RLS scheme with exponential forgetting, where past observations are exponentially down-weighted. Eq. 3.4 therefore becomes,

$$S_{\mathrm{t}}(\theta_{\mathrm{t}}) = \sum_{i=1}^{\mathrm{t}} \lambda^{\mathrm{t-i}}(y_{\mathrm{i}} - x_i^{\mathrm{T}}\theta_{\mathrm{t}})^2, \tag{3.13}$$

where the Σ in Eq. 3.4 is now a diagonal matrix with the diagonal given by $\mathrm{diag}(\Sigma_{\mathrm{t}}) = \lambda^{\mathrm{t-i}},\ i = 1, \ldots, t$. The estimates are then updated following the same steps in Eq. 3.6 to Eq. 3.12,

$$\hat{\theta}_{\mathrm{t}} = \hat{\theta}_{\mathrm{t-1}} + R_t^{-1} x_{\mathrm{t}}(y_{\mathrm{t}} - x_t^{\mathrm{T}}\hat{\theta}_{\mathrm{t-1}}), \tag{3.14}$$

$$R_{\mathrm{t}} = \lambda R_{\mathrm{t-1}} + x_{\mathrm{t}} x_t^{\mathrm{T}}. \tag{3.15}$$

This method thus recursively estimates the parameters and down-weights past observations, making it online and adaptive. Further information and a more detailed explanation of regression models can be seen in Madsen [31].

It is also proposed that an individual model is created for each forecasting horizon, Eq. 3.1 therefore becomes,

$$Y_{\mathrm{t+k|t}} = f(X_{\mathrm{t+k|t}}, \mathrm{t}, \theta_{\mathrm{k}}) + \epsilon_{\mathrm{t+k|t}} \tag{3.16}$$

where k indicates the forecasting horizon and the interpretation of the subscript notation t+k|t on a variable is that it is the k-step prediction calculated using only available information at time t, i.e. conditional on time t. Thus, each model tunes its coefficients using the corresponding respond variable, $Y_{\mathrm{t+k|t}}$ and input matrix, $X_{\mathrm{t+k|t}}$. Further details can be seen in Paper I.

RLS with an exponential forgetting factor has been used in numerous studies, usually for dynamic systems for control and forecast purposes. Bacher, Madsen, and Nielsen [32] propose to use this forecasting framework for online prediction of PV systems with 15-min resolution measurements to handle changing conditions of a PV system. Pinson et al. [33] propose to use recursive and adaptive methods to estimate the coefficient of a transfer function that describes the relationship between hot water at the plant and a point in the network to be used for temperature control. Therefore, the coefficients can vary over time as the dynamics change, including the time-delay coefficient. Palsson, Madsen, and Søgaard [34] propose a control scheme that allows the system's parameters to vary over time to control supply temperature in a district heating network.

The non-linear relationship between the heat load and the independent variables in Eq. 3.1 results in the parameters entering the model in a non-linear way. Therefore, the response cannot be written as a linear combination of the independent variables. However, if the independent variables are transformed beforehand, it can be rewritten as a linear model with a closed-form solution of the parameter estimates.

A framework is therefore proposed that uses the above methods to produce forecasts. Using the methods described above, the framework consists of two stages, *transformation* stage and *regression* stage. The framework can therefore handle both non-stationarities and non-linearities to produce accurate forecasts.

In the first stage, *transformation*, the independent variables are mapped either using a function (e.g. splines, Fourier Harmonics, etc.) or as instant effects of them. After the transformation, it should be possible to capture the dynamics between the response and the independent variables with a linear model. Therefore, in the second stage, *regression*, the parameters of the regression model are estimated using a least squares method.

In Paper I, this forecasting framework methodology has been implemented as an `R` package, the `onlineforecast`, for ease of use. Section 3.1.1 demonstrates the heat load forecasting model using the `onlineforecast` `R` package with code snippets.

3.1.1 `onlineforecast` Example

This section presents a brief modelling example to demonstrate how grey-box modelling works in practice with physical knowledge and statistical methods. The `R` package, `onlineforecast` from Paper I is used to estimate the coefficients of the models and produce online forecasts.

The measurements used are from one of the areas of the district heating utility *Brønderslev Forsyning*, and the Numerical Weather Predictions (NWPs) used as input for the forecast models were provided by MetFor™[1].

The first suggestion for a model would be to include components that model both the social and climatic parts of heat load. To model the social dynamics, it is proposed to use the Fourier series to capture the diurnal variations and forecasts for ambient air temperature. So, *Model 1* in regression form is given by,

$$\hat{Y}_{t+k|t} = \theta_{0,k} + \theta_{1,k} T^{\mathrm{a,NWP}}_{t+k|t} + \mu(t, n_{\mathrm{har}}, \boldsymbol{\alpha}_{\mathrm{diu}}), \tag{3.17}$$

where θ_0 is the intercept, $T^{\mathrm{a,NWP}}_{t+k|t}$ is the NWP prediction of the ambient air temperature, and $\mu(\cdot)$ is the Fourier series,

$$\mu(t, n_{\mathrm{har}}, \boldsymbol{\alpha}_{\mathrm{diu}}) = \sum_{n=1}^{\mathrm{n_{har}}} \alpha_{1,\mathrm{n}} \sin\left(\frac{2\pi}{P} nt\right) + \alpha_{1,2} \cos\left(\frac{2\pi}{P} nt\right). \tag{3.18}$$

P is the period of the series; in this case, the diurnal heat pattern is modelled, so it is 24, and t is the time of day in hours. n_{har} is the number of harmonics, and $\boldsymbol{\alpha}$ is the coefficient vector for the harmonics. The `onlineforecast` code snippet for *Model 1*,

```
#### Model 1
model <- forecastmodel$new()
model$output <- "Y" # The response variable, Y (heat load)
## Inputs
model$add_inputs(Ta = "Ta", # NWP of the air temperature
                 mu_tday = "fs(tday/24, nharmonics = 3)", # Harmonics
                 mu = "one()")  # Intercept
```

See the rest of the code in Appendix A where e.g. how to optimise the forgetting factor.

Using the forecasting framework proposed in Section 3.1 and Paper I, the offline parameters must be first optimised. In this model, the offline parameters consist only of the forgetting factor λ. The training period is from 2020-02-01 to 2020-12-31 and the scoring period is from 2020-03-01 to 2020-12-31, which is the period where the RMSE is to be calculated and minimised. The first month is therefore a burn-in period for initialising the model.

The residual analysis of the one-step prediction ($\hat{\epsilon} = Y_{\mathrm{t}} - \hat{Y}_{\mathrm{t|t-1}}$) of *Model 1* is performed in Figure 3.2 to validate the model. The plot on the left shows the autocorrelation (ACF) of the residuals of the one-step predictions. It shows that there are still dynamics to be modelled due to the significant correlation for all lags. To investigate whether the model captures the diurnal variation, a boxplot of the residuals for each hour of the day is shown in the middle graph. The boxplot shows that there is still some modelling to be done. It could be more harmonious or the diurnal profile between weekday and weekend needs to be modelled separately. Since it is known that the heat consumption does not react immediately when the ambient temperature

[1] https://enfor.dk/services/metfor/

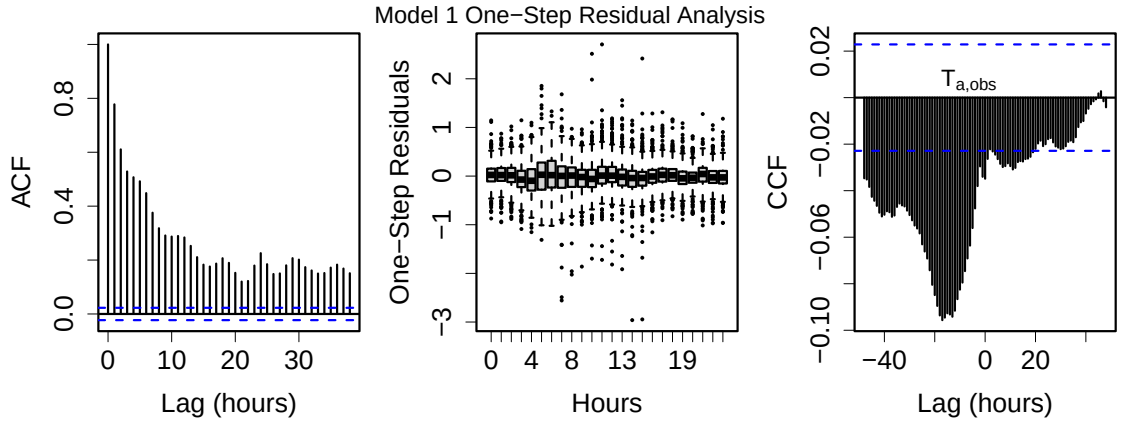

Figure 3.2: Figure shows the residual analysis of *Model 1*, where the left plot shows the ACF, the middle shows the residuals per hour of the day, and the right plot shows the CCF plot of the residuals and the observed ambient air temperature

changes as modelled in *Model 1*, the cross-correlation function (CCF) between the measured ambient temperature and the residuals is investigated to see if there is still a structure. The CCF is shown in the right plots and illustrate that there are still dynamics that need to be modelled.

From this, it can be argued that several components can be added to *Model 1* based on the residual analysis. When performing forward selection, one should try all possible additions of a component and select the one with the lowest RMSE and perform further residual analysis to continue the modelling procedure. However, this example aims to demonstrate how physical knowledge can help the model-building procedure. Thus, *Model 2*

$$\hat{Y}_{t+k|t} = \theta_{0,k} + \mu(t, n_{\text{har}}, \alpha_{\text{diu}}) + \theta_{1,k} H_a(q) T^{\text{a,NWP}}_{t+k|t}, \tag{3.19}$$

where the filter is

$$H_{a_{\text{T}_\text{a}}}(q) = \frac{1 - a_{\text{T}_\text{a}}}{1 - a_{\text{T}_\text{a}} q^{-1}}. \tag{3.20}$$

The filter models the slow reaction by smoothing the temperature with a time constant a_{T_a}. The time constant is also optimised offline with the forgetting factor. Figure 3.3 shows the same residual analysis of *Model 1* performed for *Model 2*. Note that the CCF now has a much lower correlation between the temperature and the residual. The `onlineforecast` code snippet for *Model 2*,

```
model2 <- forecastmodel$new()
model2$output <- "Y"  # The response variable, Y (heat load)
## Inputs
model2$add_inputs(Ta ="lp(Ta, a1 = 0.9)",#low pass of the NWP of the air temp.
                  mu_tday = "fs(tday/24, nharmonics = 3)", # Harmonics
                  mu = "one()") # Intercept
```

See the rest of the code in Appendix A where e.g. how to optimise the forgetting factor and the time constant for the low-pass filter.

The comparison of the two models using the RMSE over the test period for each horizon from 1 to 72 steps ahead is shown in Figure 3.4. It shows that modelling the slow response of heat load and ambient air temperature significantly improved the model accuracy. These physical insights can help make the model more appropriate by using simple methods, e.g. filters. However, the residual analysis of *Model 2* shows that there are still some dynamics to be captured, e.g. the modelling of weekends and weekdays mentioned earlier. The modelling procedure is not

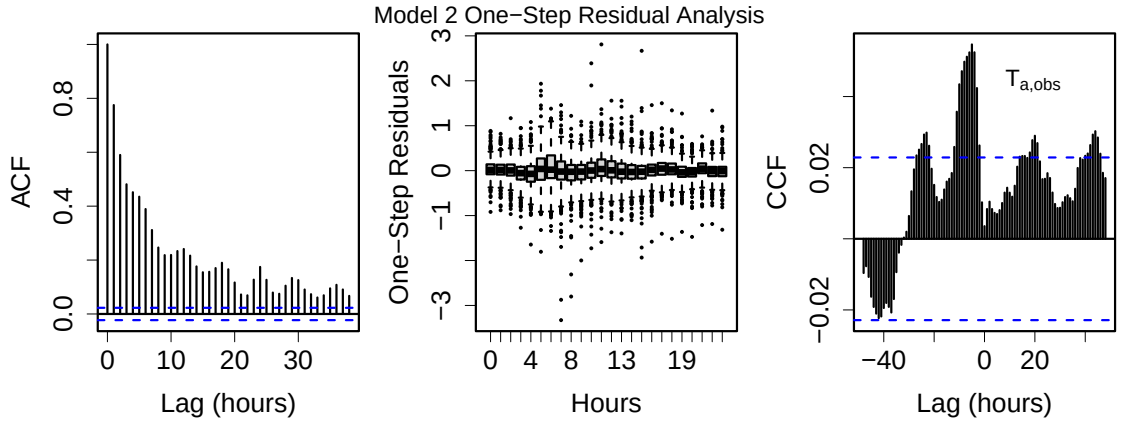

Figure 3.3: Figure shows the error analysis of *Model 2*, where the left plot shows the ACF, the middle plot shows the residuals per hour of the day, and the right plot shows the CCF plot of the residuals and the observed ambient air temperature.

further elaborated here, but a good example of the modelling and validation procedure is shown in Bacher et al. [35].

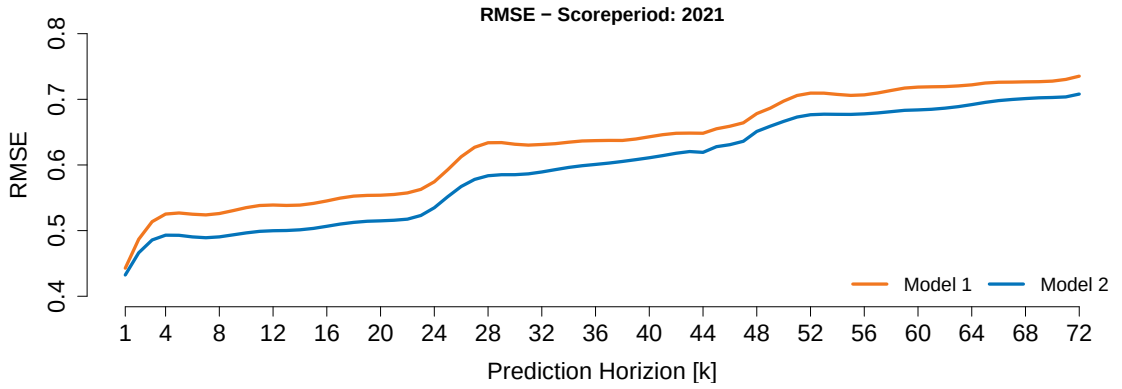

Figure 3.4: The figure shows the RMSE for the two models for each horizon from one to 72.

To illustrate the need for a recursive and adaptive forecasting framework when the response variables have non-stationary dynamics, two of the estimates (θ_0, θ_1) of both models are presented in separate plots in Figure 3.5. The Fourier series estimates are disregarded to simplify the explanations. These plots show how quickly the dynamics change over time, as the estimates frequently change except in summer. Unfortunately, measurements over a long period are missing for the summer period. However, from the available measurements, it appears that the estimates are relatively more constant than the estimates in winter, as shown in Figure 3.5. This can be explained by the fact that the heat load is more constant in summer and therefore the changes in the dynamics between heat load and input variables are small.

3.2 Localized Forecast

Paper H & Paper G

Accurate knowledge of consumer heat consumption is essential for efficient network operation and production. As mentioned earlier, consumption includes both space heating and domestic hot water usage, where space heating is strongly correlated with the climate. As space heating is used to maintain thermal comfort, it depends on the energy efficiency of the building and the local climate. Therefore, an accurate representation of space heating load is more focused

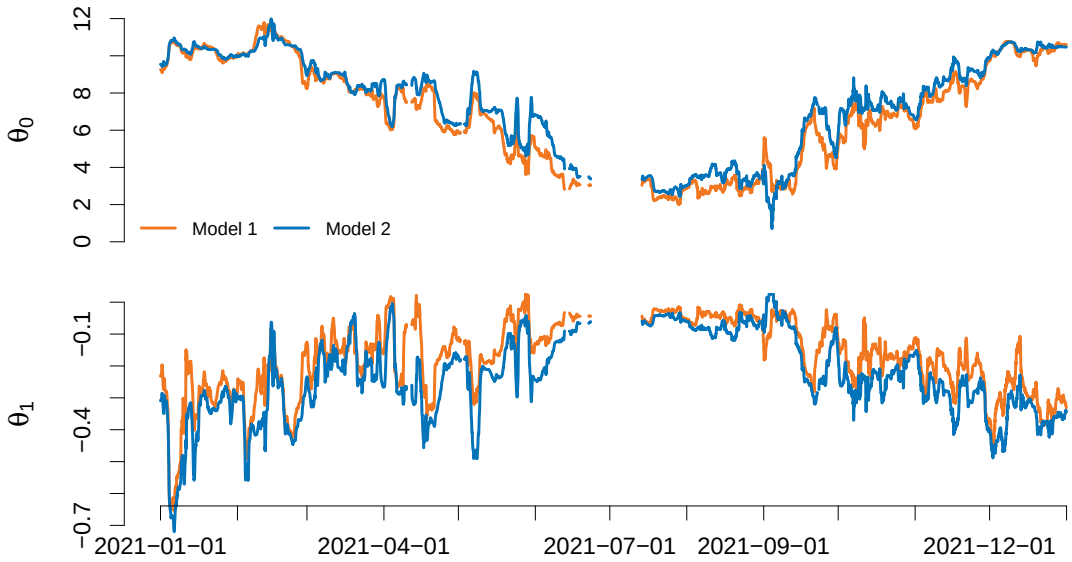

Figure 3.5: The two plots in this figure shows how the coefficient estimates vary over time as they are updated when new information is available. The top plot shows the intercept estimate, θ_0, and the bottom plot shows the ambient air temperature estimate, θ_1.

on individual houses/apartments. In this thesis, it has been assumed that aggregating individual consumption into local area consumption is reasonable. Creating and maintaining a heating load forecast for each smart meter would be too time-consuming and computationally expensive to achieve high accuracy. It can also be argued that smart meter data is too noisy to produce an accurate heat load forecast, i.e. long periods without consumption and changes in the diurnal profile. This might be too difficult to make forecasts with high accuracy that provide sufficient information to district heating operators. However, the aggregated heat load for a specific area, e.g. an area controlled by a heat exchanger which supplies heat from the transmission network to the distribution network, is sufficient to create an accurate forecast model to operate the network. Therefore, a forecast model is needed that represents the heat load in a local area, i.e. a localised forecast model.

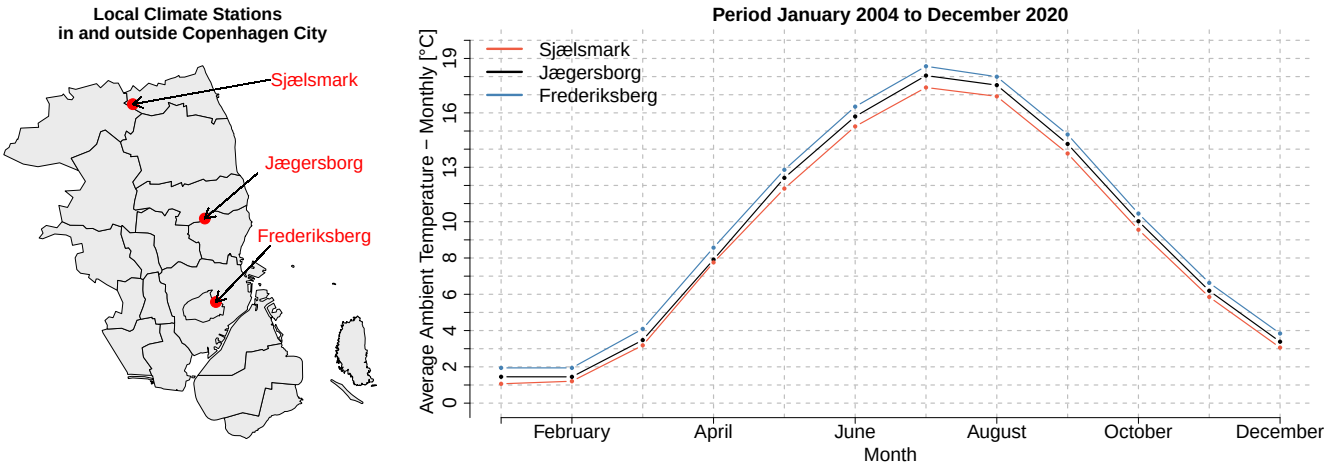

Figure 3.6: Urban Heat Island: Copenhagen demonstrates using three climate stations located in different proximity to the centre of Copenhagen: in the centre (Frederiksberg), in the outskirts (Jægersborg) and in a rural area (Sjælsmark), as shown in the map and the plot showing the difference in monthly average temperature for the three stations. Figure from Paper G.

Furthermore, not only does the difference in building mass affect how much heat load is needed to maintain thermal comfort but there is also research that has shown that ambient air temperature can vary within a large city, see, e.g. Steeneveld et al. [18]. Therefore, measuring the temperature in the area where the district heating network is located is important. Traditionally, the ambient air temperature has only been monitored in rural areas, but with digitalization using cheap IoT sensors, these temperatures are also measured within cities. The Danish Meteorological Institute, for example, has recently started giving the public access to its climate stations located all over Denmark. However, it is important to know that the air temperature measured at airports may differ from the temperature in cities, where the air temperature is exposed to human activities and the built environment. Research shows that the ambient air temperature in urban areas is usually higher than in rural areas [18]. This effect is termed urban heat island (UHI). Research related to UHI's has recently gained more attention due to concerns that climate change, with an average temperature increase of 2 to 3 K, will cause more severe heat waves in the future and place a significant burden on urban populations. However, this problem is mostly relevant for hotter climates with no district heating or at least outside the heating season. It is, however, appropriate to the energy sector as a whole, as there is a significant demand for cooling during a heat wave. Unfortunately, there is not the same interest in studying temperature differences between urban and rural areas in winter that would be relevant for the district heating sector.

Figure 3.6 shows the temperature differences of three different climate stations in Denmark, of which 1) a climate station is located in Frederiksberg, which is close to the city centre; a densely populated area, 2) The climate station Jægersborg is located in the outskirts of Copenhagen and 3) the climate station Sjælsmark located in a rural area north of Copenhagen. The data were taken from the Open Meteorological Data provided by the Danish Meteorological Institute (DMI) [36]. For each climate station, the past hourly mean temperature from January 2004 to December 2020 was extracted, and then the monthly average ambient air temperature for the same period was calculated. The locations of the climate stations are shown in the left plot in Figure 3.6, while the monthly average temperature for the climate stations is visualised on the right plot. The monthly average temperature shows the UHI effect as the climate station in the city has the highest average temperature, then the climate station between urban and rural areas and finally, the climate station in a rural area has the lowest average temperature. This shows that different temperatures in the city need to be considered for optimal heat load forecasting, temperature and production optimisation.

Therefore, a localised forecasting model of district heating can enable more precise operation. The localised forecast model uses observations of the heat load in a specific area and input variables representing the area, e.g. a localised weather forecast. The localised weather forecast represents the climate in that particular area. However, these local weather forecasts are usually unavailable because NWPs are designed for rural areas, as mentioned in Section 2.2. Therefore, they do not provide an optimal representation of the local climate. Thus, the measurements from the local climate stations can be used to adjust the NWP to provide accurate information about the climate.

The motivation for the following subsection was to investigate both the performance of heat load forecasts scaled from larger areas to smaller areas and the improvement of the accuracy of heat load forecasts by localising weather forecasts with measurements from a climate station. The following methods and results are from Paper G and from Paper H. Both papers use the same case study as part of the IDASC project. The project investigated the operation of a district heating network in Copenhagen. The Tingbjerg area was chosen because it can be considered a district heating island within the whole system. It has only one heat exchanger between the Tingbjerg area and the transmission lines to deliver heat to the consumers. Furthermore, it is a small area. For a more detailed summary, see Paper G.

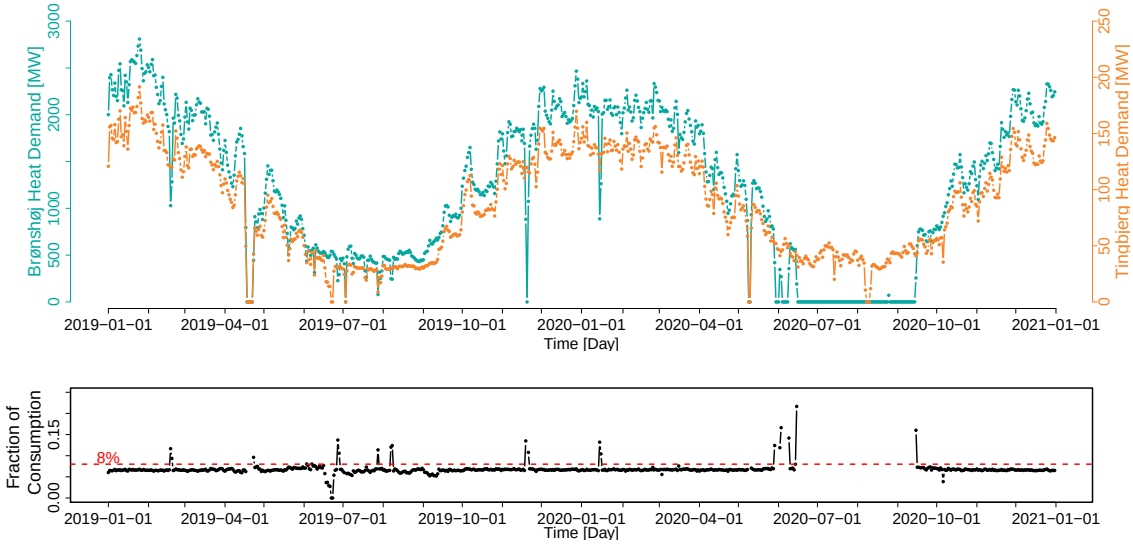

Figure 3.7: The daily heat load of Brønshøj and Tingbjerg, a small area within Brønshøj. The fraction of Tingbjerg's daily consumption compared to Brønshøj's consumption is shown in the plot below. From Paper G.

3.2.1 Scaling Heat Load Forecast

Sometimes heat load forecasts are scaled from total consumption to smaller areas within the network. This is usually done when the district heating system has a transmission and distribution system where heat exchangers transfer heat between the transmission and distribution pipes. Therefore, temperature optimisation on the secondary side is required, and heat load forecasts are needed to find the optimal supply temperatures needed to satisfy demand while minimising costs.

In Paper G, a case study was carried out to investigate the operation of a small distribution network in Copenhagen inside an area which is called Tingbjerg. The network had a heat exchanger to transfer the heat. However, the heat load forecasts used for the heat exchanger operation were scaled from a larger area, namely Brønshøj. The forecast was scaled down using a fixed fraction between the historical Tingbjerg consumption and the historical Brønshøj consumption. The ratio estimated by the district heating utility was that Tingbjerg consumption was 8% of Brønshøj consumption. Figure 3.7 shows the historical load of the two areas over two years in the upper plot and the computed fraction between them in the lower plot. The fraction from this historical data is closer to 7% than 8%. Since Tingbjerg is a small area with low heat consumption, the measurements are noisier and therefore more difficult to predict.

The compared the previous operational heat load forecast with the new localised forecast, which uses Tingbjerg's consumption measurements and localised inputs to predict future load. This was done to illustrate the importance of a localised forecast, especially to improve temperature optimisation. The new localised forecast was provided by the HEATFOR™[2] software. The new forecasting model is tailored to social consumption patterns and the relationship between building mass and local climate, as shown in Figure 3.8. The RMSE for each horizon was then calculated using a scaled forecast using the 8% and 7% ratios and the new localised heat load forecasts. The results are shown in Figure 3.9. The top plot shows that the localised forecast is significantly more accurate than the scaled forecasts. The lower plot highlights how the localised forecast performs significantly better for prediction horizons up to eight steps ahead. These short horizons

[2] https://enfor.dk/services/heatfor/

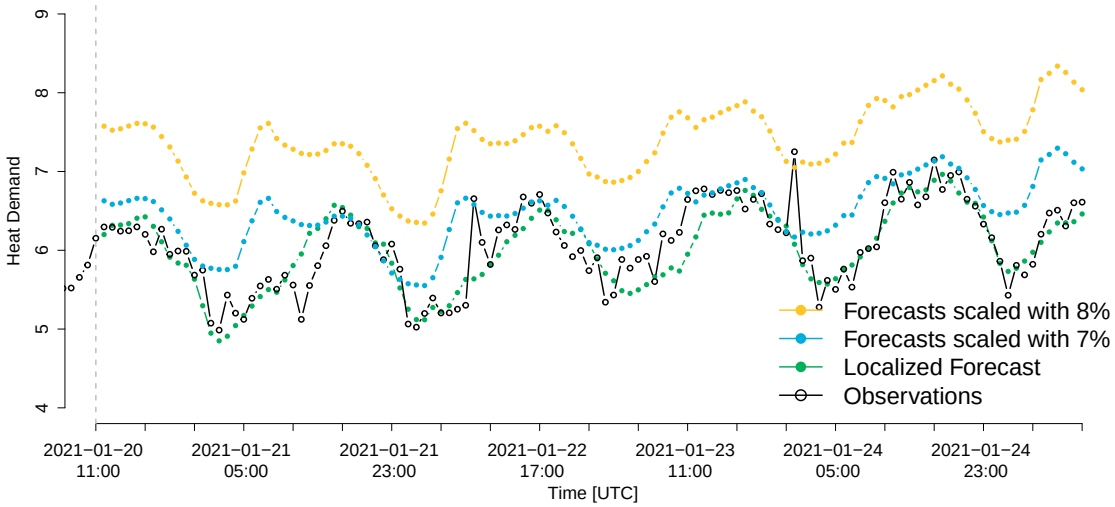

Figure 3.8: Figure demonstrates the performance of three forecast models, two scaled with ratios 7% and 8% and the localised forecast. It shows the one- to 72h steps ahead forecasts when generated at 2021-01-20 11:00. From Paper G.

are the most crucial since they match the longest time delay for water to reach the consumers, which is essential for temperature optimisation. Depending on the heat load prediction at time t for the future time t+τ, the temperature optimisation ensures that the hot water temperature leaving the plant at time t will be sufficient when it reaches the consumer at time t+τ.

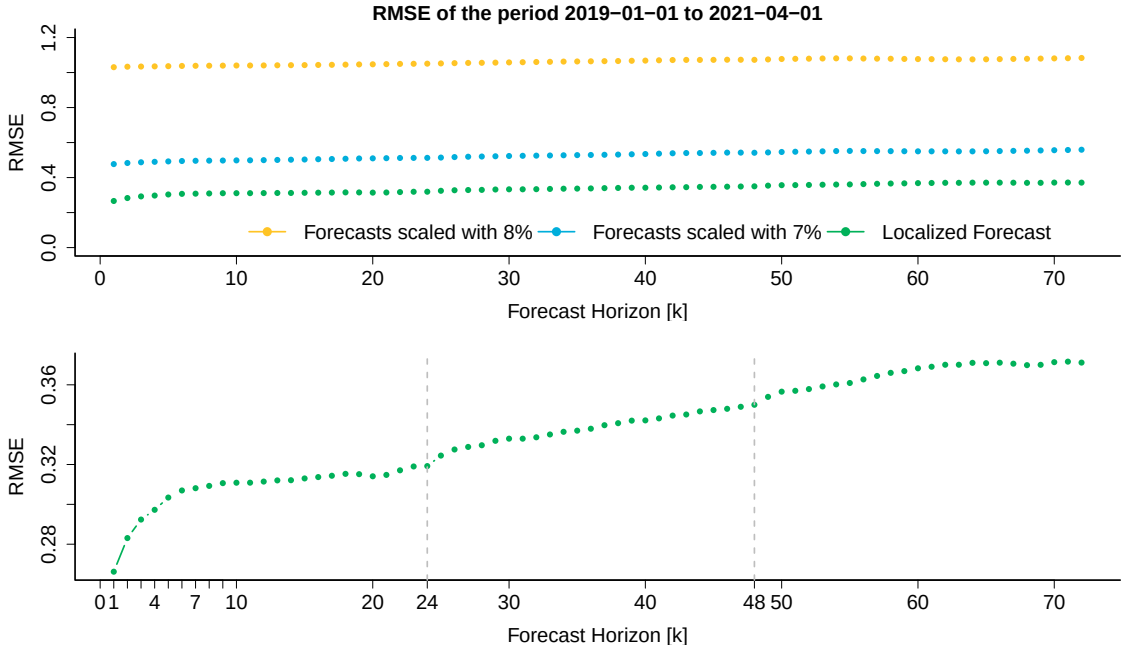

Figure 3.9: Performance of forecasts compared with RMSE for one to 72 steps ahead. The top plot visualises the results, and the bottom plot shows the localised forecast to demonstrate the effect of a local climate station to improve the short-term forecasts, i.e. the forecast horizon between 1 and 8 hours. Figure from Paper G.

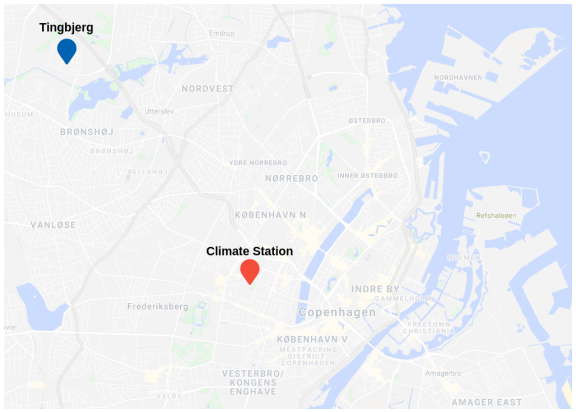

Figure 3.10: Figure shows the location of the climate station and Tingbjerg, where the district heating is located. The distance between the two places is approximately 6 km apart. Figure from Paper H.

3.2.2 Localize Inputs

As mentioned earlier, the UHI requires the operation of the district heating network to consider the local climate for each area where district heating is operated. In Paper H, it is proposed to adjust the NWP using a local climate station in Copenhagen, which is approximately 6 km away from Tingbjerg, as shown in Figure 3.10. The proposed method is simple and will recursively and adaptively update the coefficient estimates of the adjustment model to adapt to the current climate in the area.

The raw NWP of the ambient air temperature does not consider the local climate in the area where the heat load forecast is needed. However, by using a local climate station, it can be adjusted to the local climate. In Paper H, the Model Output Statistics method proposed by Glahn and Lowry [37] is used to adjust the NWP through a regression model,

$$y_{\mathrm{t}} = \beta_0 + \beta_1 T_a^{\mathrm{nwp}} + \epsilon_{\mathrm{t}}, \tag{3.21}$$

where the dependent variable is the observed weather variable y at the local weather station, and the explanatory variable is the weather forecast variable T_a^{nwp}. The forecasts of this regression model are then used as inputs to the heat load forecast model proposed in Paper H, or

$$\begin{aligned} \hat{y}_{\mathrm{t+k|t}} =& \beta_{0,\mathrm{k}} + \beta_{1,\mathrm{k}} y_{\mathrm{t}} + \mu_{\mathrm{k}}(t, n_{\mathrm{har}}, \alpha_{\mathrm{diu}}) + \\ & \beta_{3,\mathrm{k}} H(q) T_{a,t+k|t}^{\mathrm{obs,nwp}} + \beta_{4,\mathrm{k}} H(q) W_{s,t+k|t}^{\mathrm{nwp}} + \beta_{5,\mathrm{k}} H(q) G_{t+k|t}^{\mathrm{nwp}}, \end{aligned} \tag{3.22}$$

where the coefficients of the model are, the NWP of wind speed (m/s), $W_{s,t+k|t}^{\mathrm{nwp}}$, the NWP of global radiation (W/m2), $G_{t+k|t}^{\mathrm{nwp}}$ and $\mu(t, n_{\mathrm{har}}, \alpha_{\mathrm{diu}})$ which describe the diurnal curve using Fourier harmonic series, where t is the time of day (hour), n_{har} is the number of harmonics, α_{diu} is a vector consisting of the coefficients for the included harmonics, and finally $T_{a,t+k|t}^{\mathrm{obs,nwp}}$ is a combined sequence of measured and forecasted ambient air temperatures (℃), including observed measurements and NWP of the ambient air temperature, i.e.,

$$T_{a,t+k|t}^{\mathrm{obs,nwp}} = \{\ldots, T_{a,t-1}^{\mathrm{obs}}, T_{a,t}^{\mathrm{obs}}, T_{a,t+1|t}^{\mathrm{nwp}}, T_{a,t+2|t}^{\mathrm{nwp}}, \ldots, T_{a,t+k|t}^{\mathrm{nwp}}\}. \tag{3.23}$$

$H(q)$ in Eq. 3.22 is a transfer function acting as a low-pass filter with a stationary gain equal to one,

$$H(q) = \frac{1-a}{1-aq^{-1}}, \tag{3.24}$$

where a is the time constant describing how the buildings are affected by changes in the corresponding climate variables, e.g. ambient air temperature.

In Paper H, three different localised NWPs were created using Eq. 3.21,

- Parameters estimated using OLS estimation, i.e. constant parameters (*lm*)
- Parameters estimated using RLS with forgetting factor ($\lambda = 0.994$) (*rls*)
- Parameters estimated using RLS with forgetting factor ($\lambda = 0.998$), i.e. whereby the parameters change more slowly over time than in case 2 (*rls2*)

The proposed method offers the possibility to adjust the NWP to the local climate by updating it with local measurements, providing better information for predicting heat load. It also offers the possibility to estimate the uncertainty of the NWP for that location[19].

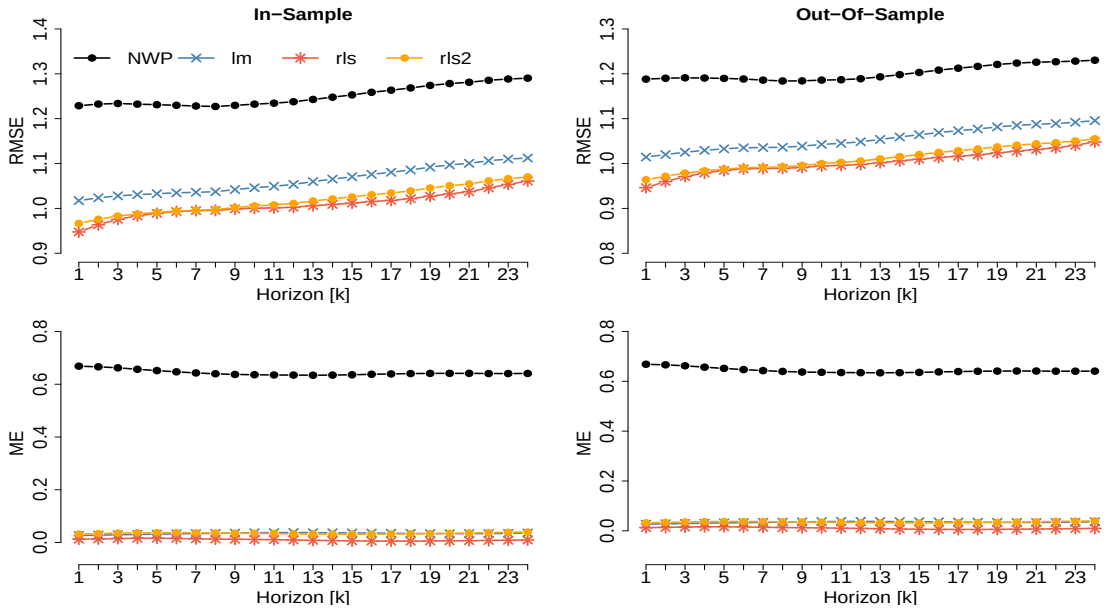

Figure 3.11: Figure shows the accuracy of the NWPs in Tingbjerg, with the in-sample in the left plots and the out-of-sample in the right plots. The RMSE is shown in the top and ME in the bottom plots. Figure from Paper H.

Figure 3.11 shows the result of localising the NWP to a local climate station. The left plots show the results from the in-sample period, while the right plots show the results of the out-of-sample period. Two error scores are used: the RMSE in the upper plots and the Mean Error (ME) in the lower plots. The plots show that the proposed method for locally adjusted forecasts performs much better than the forecasts based directly on the NWP. From the ME, it can be seen that the NWP is biased; it underestimates the temperature in the city, as expected, due to the UHI effect. The *rls* method with the optimised forgetting factor shows the best performance.

These three ambient air temperature-adjusted NWPs and the raw NWP are then used in Eq. 3.23 to predict the heat load in Tingbjerg. To benchmark the performance of using different inputs into the same forecasting model, the relative root mean squared error (RRMSE) is used,

$$\text{RRMSE} = \frac{\text{RMSE}}{\text{RMSE}_{\text{base}}} - 1. \tag{3.25}$$

The RRMSE demonstrates either improvement or decline in performance compared to the base forecast, where negative values correspond to improvements in accuracy over the base forecast.

The RMSE is when the adjusted NWP is used as input to the forecasting model, while $\text{RMSE}_{\text{base}}$ is when the raw NWP is used. The results presented in Figure 3.12 show that all three localised predictions improved the accuracy of the heat load predictions compared to the raw NWP model. The *rls* method, which adapts faster, has the best performance overall, especially for the horizons up to 12 hours ahead. As mentioned earlier, it is a very important horizon to improve due to the nature of temperature optimisation. For the last horizons, *rls* performs similarly to the other localised predictions but still shows about 1.5% higher accuracy than using the raw NWP as input.

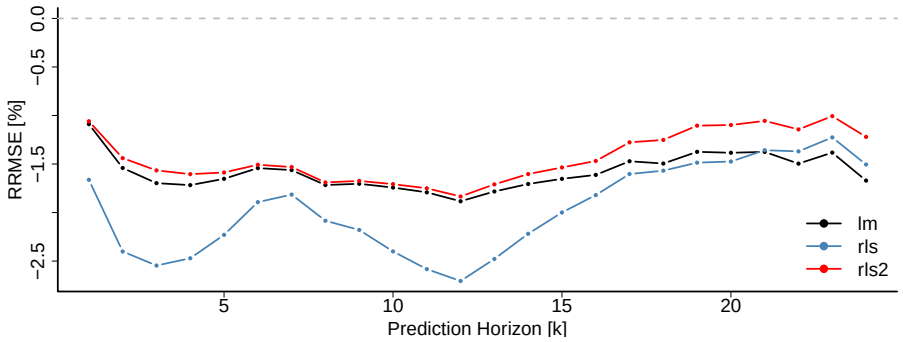

Figure 3.12: The plot shows the RRMSE of the proposed method using three different methods to estimate the coefficient of the adjusting model in Eq. 3.21 for the same forecasting model using the raw NWP as inputs. Figure from Paper H.

3.2.3 Main Findings

Based on the results of Paper G and Paper H, and highlighted in the above subsections, the importance of accurately representing local climate and knowledge of social behaviour is emphasised. First, it was shown that using scaled heat load forecasts from a large area does not accurately represent the dynamics in a smaller area. Social behaviour is different, the daily heat load profile is not the same, and the fraction used to scale the forecast varies over time, e.g. different seasonal heating effects in the area. It has also been shown that a good understanding of the local climate leads to higher accuracy of the heat load forecast by comparing the performance of using raw NWP as inputs versus adjusting the raw NWP to a local climate station within the city and using them as input to the same forecasting model. The higher accuracy is due to the fact that the NWPs were developed for rural areas and not for urban areas, so the NWPs do not consider the urban climate. Therefore, the heat load forecast benefits from localisation and shows an improvement in accuracy for all forecast horizons. There is a significant improvement in accuracy for the shorter horizons, e.g. from one-hour to eight hours ahead. The short horizon is crucial for the operation of the district heating network, as data-driven temperature optimisation becomes more efficient when the consumption can be predicted with high accuracy. The optimisation can deliver the desired hot water temperature at time t and reach the consumer at time $t + \tau$. τ is the time delay of sending one unit of water from the plant to the consumer to fulfil the heat load at the future time using the heat load forecast. Temperature optimisation will be able to select future set-points with high precision and minimise costs by using more accurate heat load forecasts.

Accurate localised heat load and weather forecasts are important for the efficient operation of district heating. The benefits come from the potentially higher precision in temperature optimisation, as the future load can be predicted more accurately, thus keeping the supply temperature as low as possible. This reduces both costs and heat losses in the system. The accuracy improvements for heat load forecast will also have benefits for production optimisation as the planning of

heat units can be done with more precision. District heating systems that are decentralised with heating units and multiple temperature optimisations in the network need this to operate each area optimally. Digitalisation will also play a major role in enabling this, as more climate stations will be available and more detailed information on how heat consumption reacts to changes, e.g. through smart meters.

3.3 Hierarchical Forecast

Paper A & Paper J

A hierarchy is the definition of a certain rank structure. In the world of data (in this context, time series), this can be one of two structures: the spatial (cross-sectional) or the temporal hierarchy. The spatial hierarchy is the structure of groups or aggregations, e.g. from the heat load in a building in an area (aggregation of all heat load in buildings in this area) to the entire district heating system (aggregation of areas). The temporal aspect is the resolution of the data, e.g. the heat load in one hour, in 12 hours (aggregation of the 12 one-hour loads) and for the whole day (aggregation of the two 12-hour periods of a day). So the commonality is that there is a bottom level and a predefined structure for the upper levels and that they must be coherent, i.e. the summation of the lower group must be consistent with the next level in the hierarchy.

Forecasts usually have to be produced in a hierarchical structure and are typically independent, so they are not necessarily coherent. For example, heat load forecasts for different areas of a district heating network and the total system are usually prepared separately. This is sub-optimal for the district heating as the forecasts are not coherent, and the operation will be operated on information that is not aligned. However, hierarchy forecasting has been proposed to make individual forecasts coherent and has demonstrated promising results in improving individual forecast accuracy by sharing information through the hierarchy.

The original proposal for using hierarchical forecasts was to use Bottom-Up, Top-Down, or middle-out to make them coherent. These proposals have drawbacks, and one of the biggest is using only individual forecasts at one level of the hierarchy to distribute to other levels. Hyndman et al. [38] propose to make individual forecasts coherent at all levels using the least squares method. These individual forecasts are called base forecasts in the reconciliation process. They can be made by any method for any group at any level of the hierarchy as long as they produce numerical values used in the reconciliation. This method of estimating the weights in the reconciliations is referred to as optimal because it creates an optimal combination of the base forecasts. In Hollyman, Petropoulos, and Tipping [25], it is shown that forecast reconciliation is a special case of forecast combinations by reformulating it into a combination of direct forecasts using linear coherent constraints. It is also discussed that a forecast combination tends to have higher accuracy than using only a single forecasting method.

As mentioned, Hyndman et al. [38] propose a regression approach to estimate the mapping matrix $\boldsymbol{G}$ which maps the base forecast to a vector of bottom forecasts using generalised least squares estimation in which the coherence errors are used to estimate the covariance matrix. The base forecasts are then written in the regression form,

$$\hat{Y}_{t+k|t} = \boldsymbol{S}\beta(k) + \epsilon(k), \tag{3.26}$$

where $\beta(k) = \mathbf{E}[Y_{\ell,t+k}|Y = y_1, \ldots, y_t]$ is the unknown conditional mean of the future values of the most disaggregated observed series, i.e. the reconciled forecasts. $\boldsymbol{S}$ is the summation matrix

which is used to describe the structure of the hierarchy; see, e.g. Figure 3.13, which results in

$$\boldsymbol{S} = \begin{bmatrix} 1 & 1 & 1 & 1 \\ 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 \end{bmatrix}. \tag{3.27}$$

The $\epsilon(k)$ represents the error between the base forecasts and their expected value, the coherence error $\hat{Y} - \tilde{Y}$. It is assumed that the error $\epsilon(k)$ has a mean of zero and a covariance matrix Σ. Therefore, it is possible to use the generalised least squares estimation to estimate $\beta(k)$ in Eq. (3.26). If Σ is assumed to be known and the base forecasts are unbiased, the reconciled forecasts can be estimated as

$$\tilde{y} = \boldsymbol{S}(\boldsymbol{S}^{\mathrm{T}}\Sigma^{-1}\boldsymbol{S})^{-1}\boldsymbol{S}^{\mathrm{T}}\Sigma^{-1}\hat{y}, \tag{3.28}$$

where the mapping matrix for the reconciliation forecasts is $\boldsymbol{G} = (\boldsymbol{S}^{\mathrm{T}}\Sigma^{-1}\boldsymbol{S})^{-1}\boldsymbol{S}^{\mathrm{T}}\Sigma^{-1}$. Hence, then the multiplication of the summation matrix and mapping matrix is the projection matrix, $\boldsymbol{P} = \boldsymbol{SG}$

However, estimating the covariance matrix from coherency errors is impossible because it is not identifiable, as shown in Wickramasuriya, Athanasopoulos, and Hyndman [39]. Therefore, it has been proposed to use the errors from the base forecast to estimate a covariance matrix to be used. Several suggestions for the covariance estimator have been proposed. Hyndman et al. [38] uses OLS to estimate the covariance estimator, i.e. an identity matrix that leads to an equal weighting of all base forecasts. In Hyndman, Lee, and Wang [40], the OLS method is then extended to WLS, where the base forecasts are weighted according to the variance of their residuals, i.e. replacing the identity matrix with the corresponding variance. Three different WLS methods are proposed in Athanasopoulos et al. [41]. However, these are naive methods as they disregard any cross-correlation between levels. In Wickramasuriya, Athanasopoulos, and Hyndman [39] and Nystrup et al. [42], it is shown that higher accuracy is obtained when the full covariance matrix of the base forecast errors is used compared to using only the variance.

Therefore, a covariance matrix must be estimated for hierarchical forecasts to produce optimal coherent forecasts. It is proposed to use the base forecast errors to estimate the covariance matrix, as it is not identifiable from the coherency errors. Several different methods for estimating the covariance matrix have been proposed, with varying degrees of success. However, it has been shown that using the full matrix leads to the highest improvements in accuracy. Nevertheless, all the proposed methods have shown an improvement in accuracy. Thus, the optimal estimation method must be found for each process to achieve optimal improvements.

Several recent studies have demonstrated accuracy improvements and the benefits of the reconciliation process to produce coherent forecasts in the energy sector [43]. For instance, Nystrup et al. [42] demonstrates the improvement for short-term electricity load forecasts, and Jeon, Panagiotelis, and Petropoulos [44] demonstrates the reconciliation process for probabilistic forecasting of wind power and electric load to ensure coherence, resulting in higher accuracy. Improvements in power forecast for PV systems have also been demonstrated in Yagli, Yang, and Srinivasan [45] where spatial-temporal hierarchy is imposed.

3.3.1 Hierarchical Forecast in District Heating

Hierarchies could be a beneficial tool for district heating to improve its forecasts and thereby the operation. For instance, temporal hierarchies can be used to improve forecasts for energy operator operators and, above all, make them more coherent and, hopefully, more accurate. These

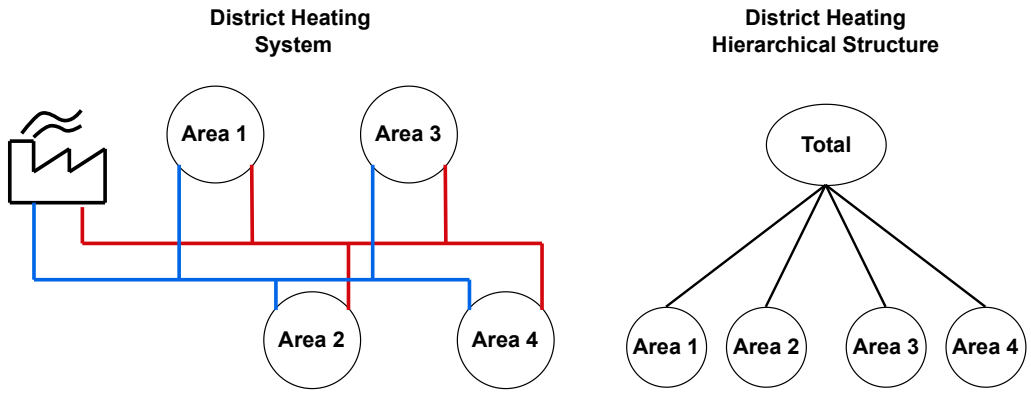

Figure 3.13: Example of the structure of a spatial hierarchy within district heating. Figure from Paper J.

operators must consider decisions for the next hour, the next day, the next weeks, and extended periods. They use forecasts for these different periods, usually divided into short, medium and long-term, and it could be beneficial to link them together and make them coherent. This will lead to optimal decision-making instead of making decisions based on forecasts that are not coherent. As district heating becomes increasingly decentralised, areas will need localised heat load forecasts to optimally operate the area's network. The heat load forecasts are then made for the total heat load and each area, which are therefore not necessarily coherent. Figure 3.13 demonstrates the natural spatial hierarchy of a district heating system. Hierarchical forecasting is ideal for production and network operators as the reconciled forecasts will be coherent for their optimisation. Furthermore, hierarchical forecasting has demonstrated that they usually improve accuracy, which is an important advantage.

However, there is little or no discussion in the literature about the empirical covariance matrix for non-stationary data, which must be robust and stable in order to produce an accurate, reconciled forecast. The covariance matrix for non-stationary processes must be able to adapt to any changes in the system in order to react accordingly. This is the same argument as when producing heat load forecasts using the method in Section 3.1, as these are produced recursively, and past information is exponentially down-weighted. Two papers have been made in this thesis where it is proposed to estimate the covariance matrix based on non-stationary requirements and demonstrate that the method results in high accuracy improvements for both papers. The first paper, Paper A, explores the use of a temporal hierarchy constraint for the 24-hour forecast with a full natural structure between the hourly and daily forecast levels, i.e. the levels in the hierarchy is thus $\ell = \{1, 2, 3, 4, 6, 8, 12, 24\}$. In Paper J, the spatial hierarchy was studied using two different district heating systems. One has few areas, and the other has many to demonstrate the potential of the hierarchy framework. In both papers, the heat load is predicted over long periods of time, and the measurements are therefore not stationary. Therefore, in Paper A and Paper J, it is proposed to use a recursive and adaptive covariance estimator estimated from the in-sample base forecast errors. Three different estimators are proposed in Paper A,

$$\text{Expanding Window:} \qquad \hat{\Sigma}_t = \frac{1}{t} e_t e_t^{\mathrm{T}} + \frac{t-1}{t} \hat{\Sigma}_{t-1}, \tag{3.29}$$

$$\text{Rolling Window:} \qquad \hat{\Sigma}_t = \frac{1}{t-j} \sum_{i=t-j}^{t} e_i e_i^{\mathrm{T}}, \tag{3.30}$$

$$\text{Exponential Smoothing:} \qquad \hat{\Sigma}_t = \lambda \hat{\Sigma}_{t-1} + (1-\lambda) e_t e_t^{\mathrm{T}}. \tag{3.31}$$

The first method is *expanding window*, where all available past errors are used to estimate the

estimator. The second is *rolling window*, where a fixed window of some of the available errors is used, and the third is *exponential smoothing*, where past information is exponentially down-weighted using a forgetting factor. It has also been suggested to shrink the covariance matrix to ensure that it is well-defined and does not become singular, for example. This has been done in some other hierarchy studies, e.g. Wickramasuriya, Athanasopoulos, and Hyndman [39] and Nystrup et al. [42]. Thus, a shrinkage method is also proposed, where the covariance estimator is shrunk (e.g. using the three methods mentioned above) by using a shrinkage method proposed by Ledoit and Wolf [46]. This method (optimal shrinkage intensity) tends to shift the extreme values towards central values, which reduces the estimation error,

$$\text{Shrinkage: } \hat{\Sigma}_t^* = \lambda^*_{\text{shrink}} \hat{\Sigma}_t^{\text{d}} + (1 - \lambda^*_{\text{shrink}}) \hat{\Sigma}_t, \tag{3.32}$$

where $\hat{\Sigma}_t^{\text{d}}$ are the diagonal entries from the covariance estimator. Thus, as λ^*_{shrink} increases, the off-diagonal elements are shrunk towards zero since the shrinkage target is the diagonal variance of the levels in the hierarchy. The shrinkage intensity parameter λ^*_{shrink} has a closed-form solution where the mean squared error is minimised,

$$\lambda^*_{\text{shrink}} = \frac{\sum_{\text{i}\neq\text{j}} \widehat{\text{Var}}(\hat{\sigma}_{ij})}{\sum_{\text{i}\neq\text{j}} \hat{\sigma}_{ij}^2}, \tag{3.33}$$

where $\hat{\sigma}_{\text{ij}}$ is the i, j'th element of the covariance estimator.

In Paper A, recursive and adaptive estimation of the variance of the covariance estimator is proposed,

$$\widehat{\text{Var}}(\hat{\Sigma}_{\text{t}}) = \lambda(1-\lambda)^2 (e_t^2 (e_t^2)^{\text{T}} - \hat{\Sigma}_t^2) + \lambda^2 \widehat{\text{Var}}(\hat{\Sigma}_{\text{t}-1}). \tag{3.34}$$

The shrinkage parameter can then be updated easily at each time step,

$$\lambda^*_{\text{shrink},t} = \frac{\sum_{\text{i}\neq\text{j}} \widehat{\text{Var}}(\hat{\Sigma}_t)_{ij}}{\sum_{\text{i}\neq\text{j}} (\hat{\Sigma}_t)_{ij}^2}. \tag{3.35}$$

Thus, the shrinkage intensity parameter can be re-estimated whenever new information becomes available, and the past information is down-weighted using exponential weights with the forgetting factor λ. This is done to replace the variance estimation of the covariance estimator as proposed by Schäfer and Strimmer [47] in Appendix A.

3.3.2 Workflow of the online reconciliation process

To get a better overview of hierarchical forecasting, Figure 3.14 depicts a workflow of the reconciliation process. It shows a generalised overview of the process and is described in detail below. However, before the reconciliation process is started, the base forecast models and the covariance matrix have to be initialised.

1. Generate base forecasts. Base forecasts can be generated from the desired model forecast for any level and horizon. The only requirement is that you have a numerical value for each level and forecast horizon. Then, a base forecast vector at time t, $\hat{Y}_{t+1|t}$, is created containing each aggregation level forecast with the corresponding forecast horizon. An example of this vector is a case for a temporal hierarchy where the top level is the 24-hour, then the 12-hour, and at the bottom the 4-hour forecasts,

$$\hat{Y}_{t+1|t} = \begin{bmatrix} \hat{Y}^{\text{24h}}_{t+k|t} = [\hat{y}^{\text{24h}}_{t+1|t}]^{\text{T}} \\ \hat{Y}^{\text{12h}}_{t+k|t} = [\hat{y}^{\text{12h}}_{t+1|t},\ \hat{y}^{\text{12h}}_{t+2|t}]^{\text{T}} \\ \hat{Y}^{\text{4h}}_{t+k|t} = [\hat{y}^{\text{4h}}_{t+1|t},\ \hat{y}^{\text{4h}}_{t+2|t},\ \ldots,\ \hat{y}^{\text{4h}}_{t+6|t}]^{\text{T}} \end{bmatrix} \tag{3.36}$$

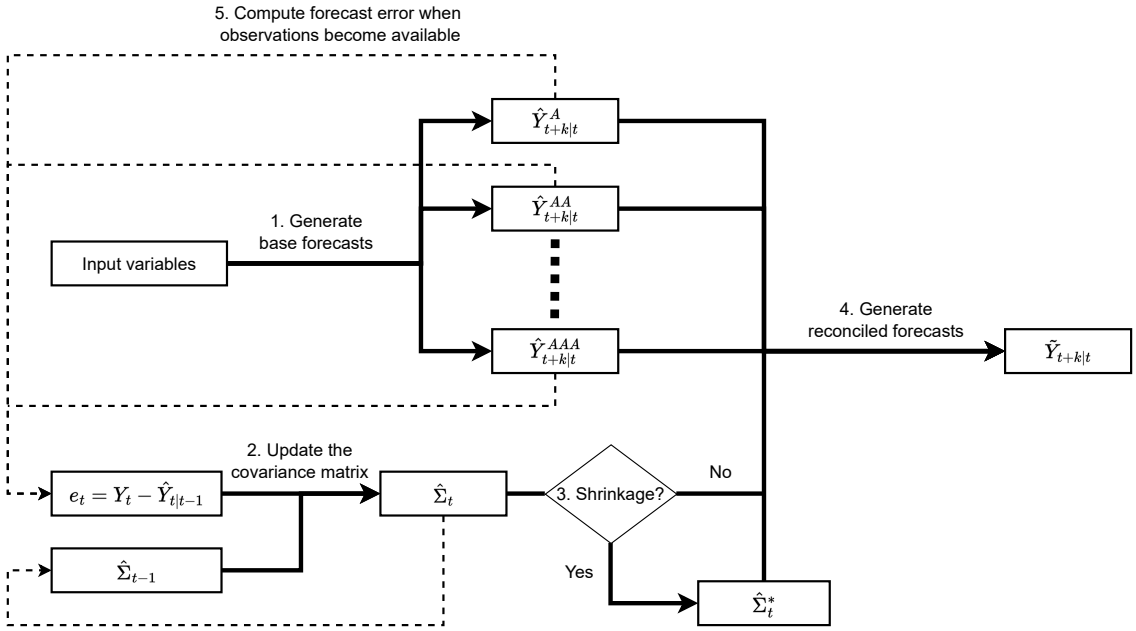

Figure 3.14: Overview of the reconciliation process. The first iteration of this figure is from Paper A.

*2. **Update the covariance matrix**.* A covariance matrix is needed to combine the base forecasts for the reconciliation process. It is proposed that for online forecasts, the covariance matrix is updated recursively using either Eq. 3.29 or Eq. 3.31. Since the covariance matrix is updated recursively, the previous covariance matrix must be retained for updating when a new observation becomes available. However, other methods could be used if the estimator does not need to be updated recursively, e.g covariance in Eq. 3.30.

*3. **Shrinkage**.* If the covariance matrix is to be shrunk before the reconciliation process, then this should be done after the matrix has been recursively updated. In Paper A, it is suggested that if there is a shrinkage for online forecasts, it should then be the recursive shrinkage method as suggested in Eq. 3.34 and Eq. 3.35.

*4. **Generate reconciled forecasts**.* When the base forecasts and the covariance matrix are ready, the reconciled forecasts are calculated using Eq. (3.28).

*5. **Compute forecast error when observations become available**.* As new observations become available, the base forecast residuals $e_t = Y_t - \hat{Y}_{t|t-1}$ are calculated, and the covariance matrix is updated. The updated covariance matrix is then used to generate new reconciled forecasts.

***Optimization of hyperparameters**.* If the reconciliation procedure has hyperparameters, then these must be optimised to make the reconciliation forecast as accurate as possible. Therefore, the hyperparameters are trained over the training period to find the parameters that minimise the RMSE. For instance, the forgetting factor for the covariance estimator in Eq. 3.31 must be optimised.

Therefore, the proposed reconciliation prediction workflow must have a low computation time for all steps, and the pipeline must be aware of any delays. First, the input variables can arrive at any hour (since the resolution of the forecast is hourly). Second, the calculation time for the base forecasts. Third, when new observations become available for updating the covariance matrix. Fourth, the calculation time for updating the covariance matrix and, finally, the calculation time

for producing the reconciliation forecasts.

This proposed framework for the online reconciliation process is ideal for real-world operations as not much information needs to be stored. A new reconciliation forecast can be generated when new information is available in a short time, as the computation time is short. It is also suitable for a non-stationary process as the covariance estimator can be adaptive.

3.3.3 Results: Main Findings

Two separate papers, Paper A and Paper J explore the possibility of improving the heat load forecasts using temporal and spatial hierarchy. In both papers, the prediction horizon ranged from 1-step to 24-steps ahead, with hourly resolution. The objective was to improve the state-of-the-art operational heat load forecast from the HeatFor™[3] software.

Paper A investigates the improvements potential of using temporal hierarchies for heat load forecasts for the heat production planner for the district heating system in Greater Copenhagen, *Varmelast*. The temporal hierarchies improved the heat load forecast by about 15% on average over three years by sharing information between different temporal resolution levels. Three different covariance estimators were compared, with all estimators showing an improvement in accuracy over the base forecast. The adaptive estimator provided the most significant improvements. The improvements result either from the truncation in the AR process or from the different levels of aggregation of the forecast using different memories when producing the base forecasts. For example, rapid changes in the heat load or ambient air temperature do not immediately affect heat load forecast if they have high memory, while having a lower memory would have a more significant effect. Thus, in the reconciliation process, it can react faster to rapid changes as the weights in the reconciliation will rely more on forecasts that have lower memory and can adapt faster to the changes in heat consumption.

Paper J, the improvements of spatial hierarchies are investigated using heat load forecasts from two different district heating utilities in Denmark. First, the utility *Brønderslev Fjernvarme* system is used, as it has only a few areas, namely only three. The second utility is *Fjernvarme Fyn*, which has several smaller areas, 12 in total. Therefore, two different hierarchy structures are used, and it is also proposed to add a new additional level of aggregation to the hierarchy in *Fjernvarme Fyn*. This study also used a state-of-the-art hourly operational heat load forecast from the same software as in Paper A. It was found that the accuracy could be significantly improved depending on the forecast horizon, ranging from 2% to 15%. The higher the horizon, the higher the accuracy improvements were achieved for both utilities. In addition, adding a new level of aggregation to the hierarchy was shown to increase accuracy improvements. The improvements in this study are due to the small forgetting factor used in the covariance estimator so that it can respond quickly to any rapid changes in the system. This allows the weights in the reconciliation to be changed quickly if needed. It was also demonstrated using a simulation study to show the effects of using a small forgetting factor. It is also discussed that the optimal forgetting factor for operational reconciliation prediction might be too low due to possible very large prediction errors due to overshoot/undershoot, as shown in the simulation study. Therefore, more robust results could be obtained by choosing a slightly higher forgetting factor but with a smaller improvement in accuracy. The higher and more robust forgetting factor reduces the probability of high prediction errors.

Figure 3.15 shows the accuracy improvements for the total heat load in the spatial hierarchy forecast at *Fjernvarme Fyn* utility from Paper J to highlight the potential accuracy improvements from hierarchical forecasting. The *Oper. Base* is the base forecast from the software HeatFor™, *Oper. Rec* and *Oper. Agg Rec* are the reconciliation forecasting using two different hierarchical

[3] https://enfor.dk/services/heatfor/

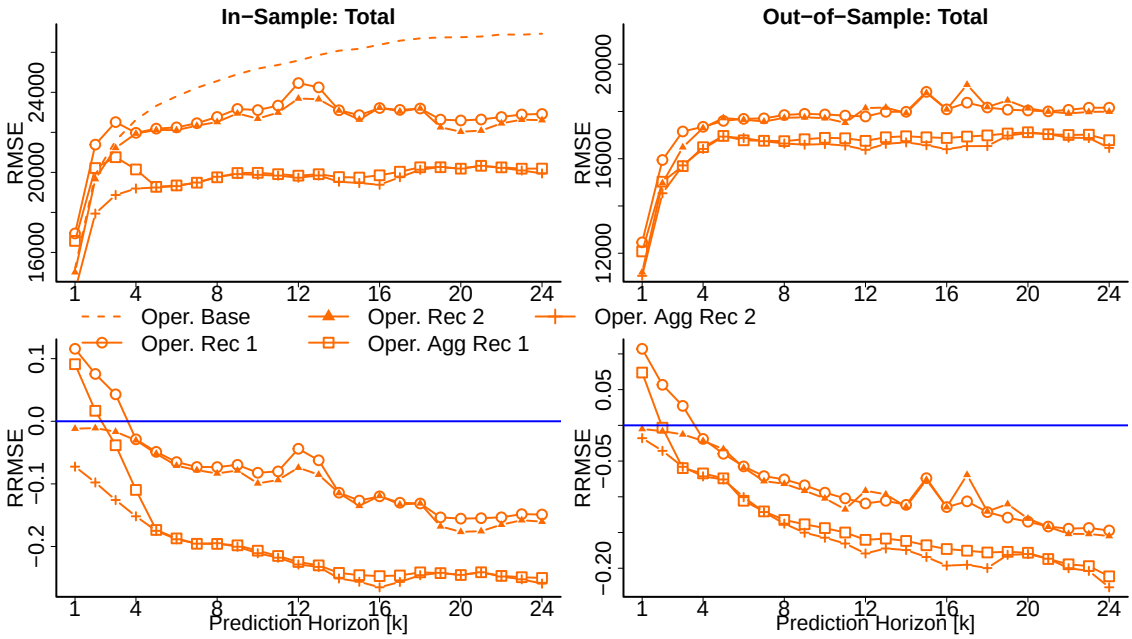

Figure 3.15: The plots show the accuracy improvements of the base forecast in *Fjernvarme Fyn*, where two different hierarchies are investigated. The top plots show the RMSE, and the bottom plots show the improvements using the RRMSE, while the left and right plots show the result in the in-sample and out-of-sample periods. Figure is from Paper J.

structures. *Oper. Rec* uses the simplest form of hierarchy with the total heat load and, at the bottom, the heat load from all 12 areas, while *Oper. Agg Rec* hierarchy has added an additional level with three new aggregation groups, which aggregate four bottom areas. Finally, the number 1 or 2 in the legend for the reconciliation forecasts refers to whether the forgetting factor was the same for all horizons or whether it was optimised for each horizon. The plots demonstrate that the spatial hierarchy can improve forecast accuracy however to ensure always improvements, the forgetting factors must be optimised for all horizons. Further improvement in accuracy is also possible by adding more levels to the hierarchy, as shown in Figure 3.15 by comparing the results of the two hierarchy structures used. In this case, the improvements are almost doubled by adding an additional level.

It is shown that significant accuracy improvements in heat load forecast can be achieved through temporal and spatial hierarchy forecasting. Thus, sharing all available information through the hierarchy increases forecast accuracy. Accurate heat load forecasts have a high potential to improve all operations and decision-making in district heating, reducing costs and increasing the safety of fulfilling the consumer's heat load. In addition, the reconciliation process makes the forecast coherent in the hierarchy, which is also beneficial for operators, especially production operators and energy planners. Thus, both main objectives of operating a district heating system are improved by hierarchical forecasting. Also, it increases the feasibility of operating district heating decentralised, e.g. multi-temperature level zones.

In summary, using hierarchies and reconciliation with an adaptive covariance estimator is shown to improve state-of-the-art operational heat load forecasting in three different cases when either a temporal or a spatial hierarchy is used. The significant improvements will positively impact reducing operating costs and meeting the heat load of consumers through more efficient district heating operations.

4 Temperature Optimisation in District Heating Networks

The main objective of district heating is to fulfil the heating needs of consumers at all times while minimising costs. The supply temperature is an important variable to minimise the costs, as keeping it low as possible lowers costs and reduces heat losses while also having positive benefits for other parts. For instance, efficient operation of the network temperature is essential to maximise the flexibility potential of district heating. The lower the supply temperature, the more likely it is that heat from new sources can be used, e.g. excess heat from supermarkets that otherwise could not be used due to too low temperatures entering the network. The operation of heat pumps also becomes more efficient, making it a better investment opportunity for them. Therefore, efficient operation of the supply temperature in district heating is crucial to reduce costs and maximise flexibility. There are robust methods for operating the network temperature, i.e. temperature optimisation. Temperature optimisation determines the optimal set points for the supply temperature that simultaneously satisfy the heat load, do not violate any system requirements and reduce heat production costs and heat losses in the network. Precision is also an important variable for optimisation so that temperatures are operated as close as possible to the constraints without violating them and that the temperatures are not fluctuating too much.

In this chapter, data-driven temperature optimisation is introduced and discussed. It also introduces new methods developed in this PhD thesis to make temperature optimisation more feasible for district heating networks. First, Section 4.1 briefly presents how the supply temperature could be optimised and what needs to be considered while creating data-driven temperature optimisation. For a detailed explanation of data-driven temperature optimisation, see Section 4.2 in Paper F, which goes into more detail with references to studies. However, in this PhD study, no temperature optimisation was created. The motivation was to investigate the savings from state-of-the-art temperature optimisation and compare it to a more traditional control method. It was also to investigate whether the new additional data in district heating that has become available through digitalization (e.g. smart meters) can be used for optimisation. Section 4.1 presents the results of Paper B and Paper C describing new methods to estimate the temperature feedback of the network. Savings from temperature optimisation and comparing different operations in district heating are discussed in Section 4.2, where the results of using a data-driven method and a non-data-driven method are compared using results from Paper G and Paper B. Finally, Section 4.3 presents the results from Paper D and Paper E, which address grey-modelling of a school building and control of the heating system to achieve efficient heating operation. Intelligent control of the consumer heating system is also discussed in the context of increasing the flexibility of the network.

4.1 Temperature Control in DH

Paper F & Paper G

Section 2.1.2 describes the operation of district heating networks; this knowledge can be used to create accurate data-driven temperature optimisation. First, by investigating the physical relationship between the supply temperature $[T_\mathrm{s}]$, the return temperature $[T_\mathrm{r}]$, the flow rate $[\dot{m}]$, and the heat load Q is

$$Q = \dot{m} c_\mathrm{p} (T_\mathrm{s} - T_\mathrm{r}), \tag{4.1}$$

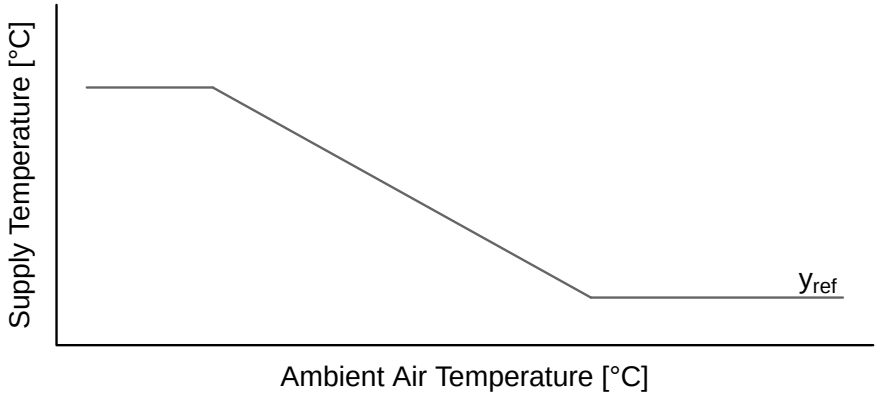

Figure 4.1: An example of a reference curve used to control the network's supply temperature for the given ambient air temperature. From Paper C.

with the specific heat constant of the water $[c_p]$ to identify which variable can be used as a control variable. The supply temperature, return temperature, and flow rate can be varied to satisfy the heat load in the district heating network, i.e. possible control variables. However, the supply temperature is usually varied, while the flow rate and return temperature depend on the consumer's operation, the load and the heating system's operation. Other physical influences must also be taken into consideration. For instance, the network is usually physically large, and the time delays are usually more than three hours, so the temperature optimisation needs must also consider the time delay. Furthermore, since the time delay varies depending on the flow, the optimisation becomes even more complex [33]. The temperature optimisation must also have a heat load forecast to ensure sufficient temperatures in the network without violating constraints, e.g. too high flow rate due to too low temperatures. Heat load forecasting is also a highly complex task because, as discussed in Section 3.1, it is driven by the physical and social components [10]. Climate and social behaviour are the driving factors of heat consumption in the places where district heating is used, e.g. how the ambient air temperature influences the space heating to maintain thermal comfort (physical) and when people use hot tap water (social). Therefore, temperature optimisation is a complex task, and advanced data-driven methods are needed to make it accurate and robust to maximise cost savings and ensure consumers receive sufficient heat.

Traditionally, supply temperatures have been controlled based on rules that depend on the ambient air temperature, i.e. a reference curve [48]. Figure 4.1 shows an example of a reference curve. Alternatively, it was controlled based on the operator's experience. However, a combination of both is usually used. A reference curve is a method in which the supply temperature is determined as a function of the current ambient air temperature or as a constraint that the supply temperature must not be lower than the reference curve at the given ambient air temperature. The supply temperature is kept low and constant at high ambient temperatures in order to provide sufficient temperature for domestic hot water usage during non-heating periods. The temperature is kept high enough to eliminate any risk of bacteria. When the ambient air temperature drops below a certain value, the supply temperature starts to increase linearly with the falling ambient air temperature until the maximum supply temperature of the system is reached. This control scheme is designed to ensure that the consumers receive sufficient heat, as it models the relationship between supply temperature and ambient air temperature as a worst-case scenario, i.e. the temperature should not be below the reference curve. The curve also takes into account that the supply temperature has a sufficiently large safety margin for the given ambient

air temperature. The flow rate is then determined based on the current supply temperature in the network and the heat consumption. However, a more optimal approach would vary the flow until it approaches the maximum flow of the system before increasing the temperature of the network, as increasing the temperature is more expensive than the electricity needed to increase the flow rate. The precision of the traditional method is low as it does not use temperature feedback of the system in the temperature optimisation, only as validation that is not violating the reference curve. It is also not a predictive controller and only considers one climate variable to vary the supply temperature. Due to not having feedback from the system, there is no direct constraint that ensures that the supply temperature in the network is not below the reference curve; see Figure 4.8 in Section 4.2 for demonstration.

Using a single reference curve in the system usually results in a high supply temperature, enough to disregard any other known factors that influence heat consumption. This scheme is therefore a conservative estimate as it only takes one variable into account to ensure that the heat supply is sufficient at all times, which often leads to an unnecessarily high supply temperature. This leads to higher costs and more heat losses in the system. The curve is therefore not an optimal strategy, as it does not take into account other climate variables known to influence heat consumption, such as wind speed, wind direction and solar radiation. It also ignores the social behaviour of heat load and the time-varying relationship. One of the important time-varying relationships is the time lag between when the water is produced and when it reaches the consumers, the time delay. It is also not a predictive control, i.e. it does not look into the future when choosing the supply temperature. Also, the effects of weather on heat load do not occur immediately, as buildings are known to have a slow thermal response. The slow thermal response in a single building is demonstrated in Madsen and Holst [49]. Preferably, the controller should take into account predictions for the heat load and the network characteristics of the system (time delay, heat loss). When all of the above factors are included in a control scheme using predictive methods, new opportunities arise to lower the supply temperature without violating any constraints and increase the supply temperature's precision while lowering costs and reducing heat losses. It is also important to note that the hydraulic limitation of the DHN imposes some restrictions on the minimum supply temperature to ensure that the flow is below the maximum limit with a specific safety margin. Hence, the supply temperature should be decided based on these points mentioned above, e.g., flow, time delay, heat load, and weather forecast, where the supply temperature at the plant is the control variable, and the network temperature is the response variable. Therefore, a predictive method is required to select optimal supply temperature set points to satisfy the consumer heat load without violating any constraints and increase the savings. Otherwise, the consumers at the end of the network will not be supplied with sufficient heat.

The DHS usually also has restrictions based on the above factors on the operation of the network due to physical limitations and additional constraints made by the utility. Nielsen [50] describes the following usual restrictions:

- *A maximum allowable flow rate in the system:* The restrictions in the flow rate are due to the (always) limited pumping capacity, the risk of cavitation in heat exchangers and difficulties maintaining a sufficiently high differential pressure in the remote parts of the network during periods with high flow rates.
- *A minimum guaranteed inlet temperature at the consumers:* This restriction is due to limitations in the consumer installations as well as minimum domestic water usage temperature requirements imposed by hygienic concerns.
- *A maximum allowable supply temperature:* This restriction is put on the systems in order not to damage pipelines and consumer installations.

- *Limited short-term variation in the supply temperature:* The stresses inflicted on the network by large and frequent fluctuations in the supply temperature dictate that the short-term variations in supply temperature should be limited.
- *Maximum allowable diurnal variations of the supply temperature:* In some systems, the size of the expansion tanks may impose limitations on the allowable diurnal variation of the supply temperature.

Thus, there are many constraints and physical limitations in the context of temperature control, so several aspects must be considered when operating a network to ensure optimal operation. Optimal operation of the network is achieved by minimising production costs without compromising the safe operation of the system, negatively affecting the system's maintenance costs or compromising consumer satisfaction. Note that the physical descriptions and system limitations listed here do not apply to every system. Each network is unique and subject to different physical constraints or limitations. However, they must be considered before implementing temperature optimisation to reduce the risk of failure and achieve optimal network operation.

Madsen et al. [51] propose a data-driven temperature control scheme that uses measurements of network temperature, heat load forecasts, and return temperature forecasts to select the optimal future set points of temperature. Any heat load forecasting model can be used, but accuracy is critical for optimal control. For instance, the forecasting model framework proposed in Chapter 3 is an appropriate choice as it is an adaptive and recursive method where new information is weighted higher than older information when updating the forecast. Hence, newer measurements are given more weight when selecting the temperature set points. A model of the temperature relationship between the production and the critical points is important because it can be used to determine the characteristics of the network, e.g. time-constant and temperature loss. Thus, knowing that the supply temperature leaving the plant at time t arrives at the critical point at time $t + \tau$ and satisfies the reference curve (the guaranteed temperature given by the utility) and heat load condition. The heat load condition is that the supply temperature is high enough so that the flow rate is not above the maximum limit. It also uses probability constraint to ensure that the supply temperature does not violate the reference curve too frequently and that the variation of the supply temperature over time is kept as small as possible [52]. The same reference curve is used as described above, but it is used at the critical points by using temperature feedback from the critical points for validation. Intelligent methods are also used to minimise the temperature variation at the critical point, keeping it close as possible to the reference curve.

Søgaard [53] and Madsen et al. [51] propose to use a statistical transfer function model to model the relationship between temperatures at the plant and critical points. For instance, a single-input single-output AutoRegressive-eXtraneous (ARX) structure with time-varying coefficients can be used to model the relationship. Here, the coefficients of the model are estimated recursively. In addition, the time delay is also recursively estimated, which is essential as it frequently changes due to its dependencies on the flow. Then, a controller can be created using the transfer function to select the optimal set points of the temperature to ensure sufficient temperature at the critical point, e.g. an eXtended Generalised Predictive Controller (XGPC) as proposed in [34] and [52]. The XGPC was proposed as an extension of the GPC controller because the controller needs to handle the time-varying process. The time-varying process is needed due to the time-varying relationship between the temperature at the plant and netpoint due to the time delay caused by flow rate changes, which the GPC cannot handle.

A schematic view of the proposed controller from Madsen et al. [51] is shown in Figure 4.2 with the addition of the NWP and load forecast, which are also required. The temperature set point at the plant is selected by the overall controller (OC), which selects the highest set point from the sub-controllers, the temperature sub-controllers (SC) and the flow sub-controller (FSC). The sub-

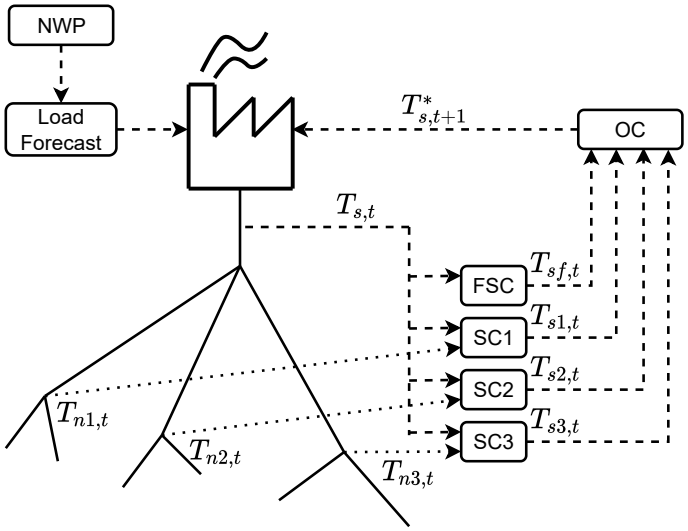

Figure 4.2: Figure shows a schematic view of the proposed temperature control from Madsen et al. [51]. It shows the two sub-controllers, the temperature sub-controller (SC), which models the relationship between the production and the network temperatures. The flow sub-controller (FSC) estimates the supply temperature based on the heat load forecasts and the flow limits. The overall controller (OC) then selects the highest set point supply temperature from the sub-controllers to use.

controller uses the information and methods described above to find the optimal future set points. One of the key elements here is the temperature feedback (the response variable) to model the network's time-varying characteristics and have a closed-loop control. Therefore, measurements at a critical point in the network are one of the most important parts of efficient data-driven temperature optimisation as it is vital to ensure sufficient temperature for the consumer at the same time lowering costs.

Usually, these critical points are typically close to consumers where the largest temperature loss occurs in the system. The location of these points can change over time as a consequence of the diurnal variations, leakage, the age of the pipes, etc. Therefore, several (e.g. five) critical points are typically used. A measurement well is installed at these points to measure the data and send it back to production. These wells are time-consuming to plan, install and maintain and very expensive.

4.1.1 Temperature Feedback

Paper G & Paper C & Paper B

Temperature feedback for data-driven temperature optimisation is necessary because it provides important information about network characteristics, such as losses and time delays between production and specific points in the network. In addition, this information can be used to ensure that consumers receive a temperature of hot water acceptable for the given operation. These measurement points that provide temperature feedback from the network can be referred to as critical points, as they are usually located where the system operators believe the lowest network temperature occurs, i.e. the highest temperature loss in the system. Therefore, if the temperatures at these critical points are sufficient, the temperatures at all other points in the network should also be sufficient. In the following, these locations are referred to as critical points. Operators can therefore use this information to adjust their set points for the supply temperature in the system. However, it would be more advantageous to have a closed-loop controller of the temperature in the plant by using the network temperature as feedback.

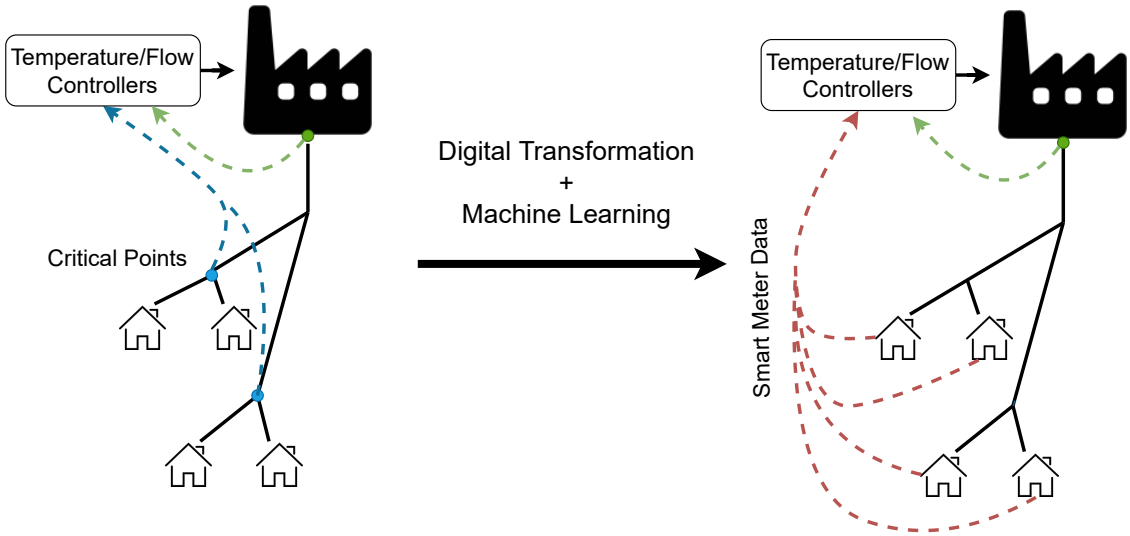

Figure 4.3: Illustration of moving away from traditional temperature feedback through the use of measurement wells at critical points to the use of smart meter readings at consumers to provide temperature feedback from the network. Figure from Paper G.

Usually, a measurement well is installed where the network's critical points are suspected. Multiple critical points are also desirable, as the location of the points may change over time as a consequence of the diurnal pattern. Therefore, multiple (e.g. three) critical points are normally used. The network can also change, e.g. pipes get older, pipes are replaced, new areas are connected to the network, buildings are renovated, etc. Therefore, it would be beneficial if the system feedback could be adapted by receiving feedback from any desired particular location in the network. As mentioned in Section 2.3, smart meters will be available in almost every house connected to the district heating network. Therefore, measurements from multiple locations in the network are available and could be used to establish temperature feedback of the network. This approach also eliminates the need to install temperature sensors in measurement wells. Figure 4.3 illustrates the proposal to use the digital transformation of multiple smart meters and use the measurements from these meters with machine learning to estimate a robust network temperature that can be used as feedback, replacing the measurement well.

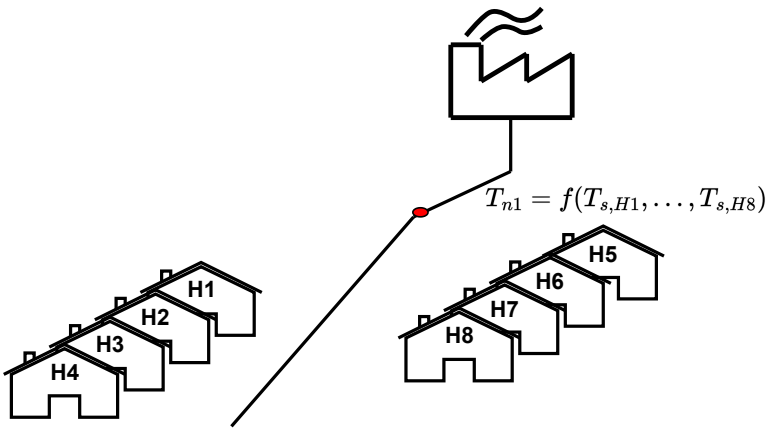

Figure 4.4: Groups of houses connected to the same distribution pipe can be used to establish netpoint temperature using smart meter readings from these houses.

In Paper B, a method is proposed that uses smart meter measurements to estimate the network

temperature. Two cases are distinguished: first, the use of smart meters in single-family homes to determine the network temperature (as illustrated in Figure 4.4), and second, the use of smart meters in large apartment buildings as direct feedback for an online temperature optimisation study. This is also discussed in Paper G, which discusses smart meters and temperature optimisation of an online case study and presents the results.

Algorithm 1 Simple algorithm to establish network temperature from smart meter measurements

Input

T matrix $[n \times N]$ of supply temperature measurements between time $t - 15$ min and $t + 15$ min

Q matrix $[n \times N]$ of flow measurements between time $t - 30$ min and $t + 30$ min

n Number of smart meters

N Number of time steps

mQ Numerical value of minimum flow allowed

q Numerical value of the desired quantile

Output

Ts Estimate network temperature at time t

procedure ESTIMATE NETPOINT TEMPERATURE AT TIME T

$Tvec \leftarrow$ *vector*

for $i = 1, 2, \ldots, N$ **do**

for $j = 1, 2, \ldots, n$ **do**

if $Q[i, j] <= mQ$ **then** ▷ If the flow is lower than the minimum, then the temperature value is replaced with NA

$T[i, j] \leftarrow NA$

end if

end for

$Tvec[i] \leftarrow median(T[i, :]$, na.remove = TRUE) ▷ Median of all temperature readings for each smart meter for the time t is computed

end for

$Ts \leftarrow quantile(Tvec, q)$ ▷ Netpoint temperature estimated from the median temperatures from all smart meters at time t using the quantile

end procedure

The proposed algorithm uses the measurements of a group of smart meters at the consumer, which are close to each other, to estimate the network temperature in the distribution pipe in the street to which the consumer is connected, see Figure 4.4. The algorithm is demonstrated in Algorithm 1, where the input is the flow rate, the supply temperature and the timestamp of the readings from the smart meters. Unfortunately, the readings from the smart meters do not necessarily come at the same time, and the readings may differ by 30 minutes between the smart meters. Figure 4.5 illustrates this by showing the discrepancy in the data from a group of smart meters where the frequency and resolution between the meters are entirely different. So before estimating the network temperature, the readings need to be cleaned up. It is suggested to aggregate the readings within the same 30 minutes by taking the median of the temperature readings from the same meter within the time interval. However, first temperatures with low flow readings are removed as they do not represent the temperature in the network when the water in the service pipe is still. Then the median temperatures at the new time stamp for each smart meter are used to estimate the network temperature. The network temperature is estimated by calculating the corresponding quantile of the new temperatures.

Algorithm 1 receives measurements of temperature T and flow Q from all smart meters that

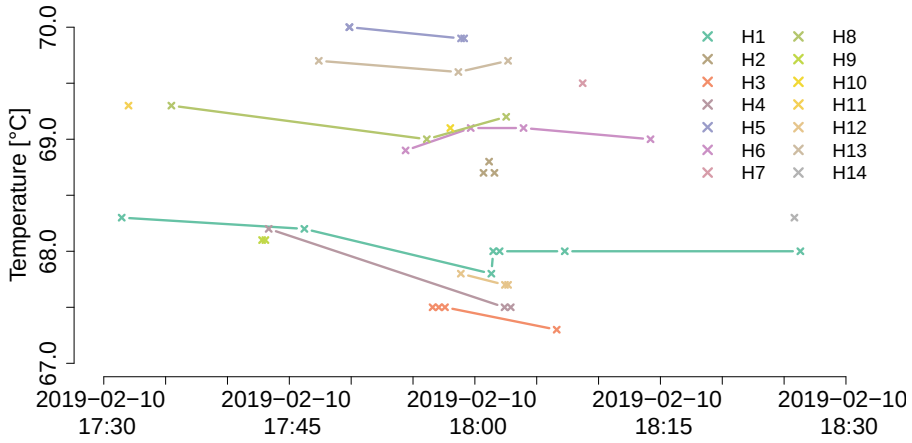

Figure 4.5: The plot demonstrates an interval of one hour. It shows the discrepancy in frequency and resolution of smart meter data. From Paper C.

have readings in the period $[\mathrm{t}-30\ \mathrm{min}, \mathrm{t}+30\ \mathrm{min}]$. n is the number of smart meters, and N is the number of readings. The input variables mQ and q are the minimum accepted flow measurement constraint and the desired quantile to be used when estimating the network temperature at time t from a vector of the median temperature of all smart meters during the period $[\mathrm{t}-30\ \mathrm{min}, \mathrm{t}+30\ \mathrm{min}]$.

The proposed algorithm determines a network temperature when measurements from the smart meters are available. Therefore, the proposed method can only estimate the temperature in the past. Usually, smart meters log values with high resolution (e.g. hourly), but they are usually sent once a day with readings from the last 24 hours. As a result, the estimated network temperature is only available for the past day when the data arrives. Also, the estimated temperature is unreliable if the quality of the measurements is low (e.g. if there are only a few temperature measurements or most of the measurements have a low flow).

In Paper C, a new method was proposed to make a more robust estimate of the network temperature when the measurements are of low quality. The grey-box method was applied by creating stochastic differential equations of the system, as shown in Figure 4.6. Here, physical knowledge was used to estimate the network temperature in the street $[T_t^{(\mathrm{s})}]$ by starting from measurements in smart meters and describing the heat dynamics in the service pipe with heat losses to the surroundings. The following partial differential equation (PDE) describes the heat transfer and associated heat losses to the surroundings in a pipe as proposed by van der Heijde et al. [54],

$$\underbrace{\frac{\partial(\rho c_{\mathrm{v}} T A)}{\partial t}}_{\substack{\text{time}\\ \text{derivative}}} + \underbrace{\frac{\partial(\rho v(c_{\mathrm{v}} T + p/\rho)A)}{\partial x}}_{\substack{\text{spatial}\\ \text{derivative}}} = \underbrace{vA\frac{\partial p}{\partial x}}_{\substack{\text{pressure}\\ \text{difference}\\ \text{energy}}} + \underbrace{\frac{1}{2}\rho v^2 |v| f_{\mathrm{D}} S}_{\substack{\text{wall friction}\\ \text{dissipation}}} + \underbrace{\frac{\partial}{\partial x}(kA\frac{\partial T}{\partial x})}_{\substack{\text{axial heat}\\ \text{diffusion}}} - \underbrace{\dot{q}_e}_{\substack{\text{heat}\\ \text{loss}}}, \tag{4.2}$$

where $\rho\ [kg/m^3]$ is the mass density of the fluid in the pipe, $c_{\mathrm{v}}\ [kJ/(kgK)]$ is the specific heat of the fluid in the pipe, $A\ [m^2]$ is the cross-sectional area of the pipe, $v\ [m/s]$ is the flow velocity,

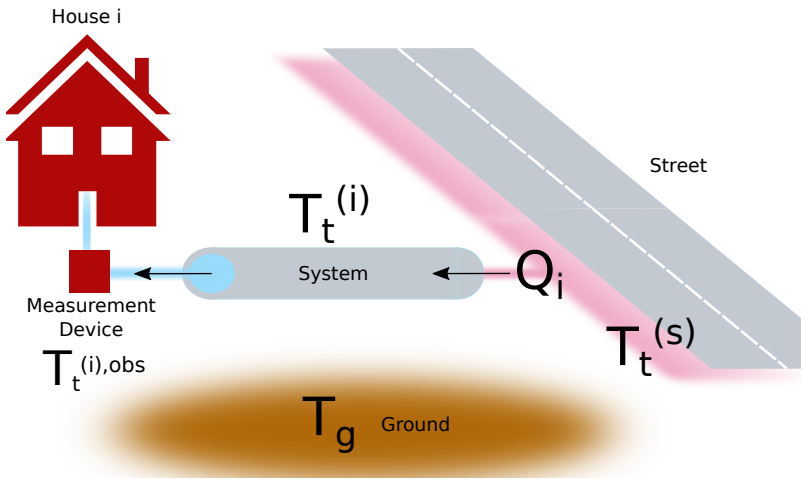

Figure 4.6: Schematic of the layout of the system in which the smart meter is located at the house, measuring all the necessary information. A service pipe transports the hot water between the distribution pipe, which is normally located on the street in front of the house, and the consumer. There is a heat loss between the service pipe and the surroundings, which depends on the temperature in the pipe and in the ground.

p $[kg/m^3]$ is the absolute pressure, f_{D} $[-]$ is the Darcy friction coefficient, S $[m]$ is the pipe circumference, k $[W/(mK)]$ is the thermal conductivity, T $[\mathrm{C}]$ is the temperature inside the pipe, and $\dot{q}_{\mathrm{e}}$ $[W/m]$ is the heat loss per unit length.

From the PDE in Eq. 4.2 and a few assumptions as shown in Paper C, it is possible to formulate a stochastic differential equation that describes the heat dynamics over a pipe,

$$\mathrm{d}T_{\mathrm{t}} = C^{-1}\left(c_{\mathrm{v}}Q_t(T_t^{(\mathrm{s})} - T_t) - R^{-1}(T_t - T_{\mathrm{g}})\right)\mathrm{d}t + \sigma\,\mathrm{d}\omega_t. \tag{4.3}$$

Jointly with the observation equation, this type of model is also referred to as a grey-box model in the literature; see, for instance, Madsen and Holst [49].

Eq. 4.3 models the heat dynamics over the service pipe, i.e. the temperature loss from the network pipe to the smart meter. However, there is no information about the temperature at the network pipe, $T_t^{(\mathrm{s})}$, and the only information available is the temperature measured at a group of smart meters that are located close together. The street temperature must, in principle, be greater or equal to the highest observed house temperature. Using this fact, the street temperature dynamic is modelled by a random walk process. Hence, no real drivers for the street temperature. Combing this with the model of the service pipe gives the system of equations

$$\mathrm{d}T_t^{(\mathrm{i})} = C_i^{-1}\left(c_{\mathrm{v}}Q_t^{(\mathrm{i})}(T_t^{(\mathrm{s})} - T_t^{(\mathrm{i})}) - R_i^{-1}(T_t^{(\mathrm{i})} - T_t^{(\mathrm{g})})\right)\mathrm{d}t + \sigma_{\mathrm{i}}\,\mathrm{d}\omega_t^{(\mathrm{i})}, \tag{4.4a}$$

$$\mathrm{d}T_t^{(\mathrm{s})} = \sigma_{\mathrm{s}}\,\mathrm{d}\omega_t^{(\mathrm{s})}, \tag{4.4b}$$

as illustrated in Figure 4.6.

The individual house temperatures $T_t^{(\mathrm{i})}$ are directly observed by the meter $T_t^{(\mathrm{i,obs})}$ and it is assumed that the uncertainty of these measurements can be approximated by Gaussian noise,

$$T_t^{(\mathrm{i,obs})} = T_t^{(\mathrm{i})} + e_t^{(\mathrm{i})}, \quad e_t^{(\mathrm{i})} \sim N(0, V_t^i), \tag{4.5}$$

Since the quality of the temperature measurement depends on the flow, a flow-dependent variance construction was made using the logistic function S,

$$V_t^i = \mathbb{V}\left[e_t^{(\mathrm{i})}\right] = \sigma_{\mathrm{obs}}^2 + S(-Q_{\mathrm{i}}(t) + b) = \sigma_{\mathrm{obs}}^2 + \frac{K}{1 + \mathrm{e}^{a(Q_{\mathrm{i}}(t) - b)}}, \tag{4.6}$$

This variance construction aims to reduce the weight of observations collected under low flow conditions where the observations contain very limited information about the street temperature. The parameters of the logistic function are a is the curve's steepness, b is the midpoint, and K is the maximum value of the function. The flow-dependent variance constructor proved necessary when houses shut off the heating for long periods of time, as the temperature quickly converged to a low-temperature value. This affected the network temperature estimate and introduced more uncertainty into the process. The values of the variance construction should be estimated separately for each cluster of smart meters, as the low flow threshold can change from area to area.

In Paper C, the parameter estimation was done using the maximum likelihood approach with two different approaches. The first method uses the `R` package Template Model Builder (`TMB`) [55] methodology, which uses a generalised mixed effects model using the Laplace approximation to integrate the states in the system. The latter method uses a discrete Kalman filter. However, the filter was implemented with the `TMB` package to take advantage of the automatic differentiation of the package, which reduces the computation time by using the gradient and the Hessian with respect to the parameters of the likelihood function.

The computation time for parameter estimation must be fast because the system is large and scales with the number of smart meters used. Therefore, for each smart meter, a new equation from Eq. 4.3 is added to the system, adding three new parameters to the system equation and one to the observation equation. However, the observation variance was proposed to be fixed in Paper C as it is quite stable and reduces the computation time. Therefore, the number of parameters was $3M + 1$, where M is the number of smart meters, and one represents the variance of the random walk process. Thus, 46 parameters had to be estimated for a group of 15 houses and data with a resolution of 5 minutes for one month was used. The two methods were compared. The Kalman filter method using the automatic differentiation from `TMB` package was about ten times faster and gave almost the same estimate of the parameters.

4.1.2 Main Findings

In Paper B and Paper C, two different methods were proposed to estimate the temperature of the network, considered as the critical point for the temperature feedback of the network. In Paper B, an additional method is presented using three smart meters as the critical point so that the temperature readings are directly used as feedback. Measurements from smart meters in buildings that consume a lot of heat (e.g. hospitals, industry, large apartment buildings, etc.) are ideal for use as temperature feedback because they are robust. The flow is often high, so the temperature measurements should be of good quality as they represent the temperature in the distribution pipe in the network adequately. However, these types of consumers are not present at all points in the network, and smart meters need to be used from single-family buildings that consume less heat. The measurement from buildings that consume heat irregularly is not optimal for temperature feedback, as the hot water in the service line to the house becomes still, and the temperature starts to drop. Therefore, low supply temperature measurements do not give an accurate signal of the hot water temperature in the distribution pipe in the street.

The two proposed methods were developed to use irregular and low-quality measurements from smart meters in single-family buildings to create accurate temperature feedback of the hot water temperature in the distribution pipe. The first method from Paper B is simple and can be easily

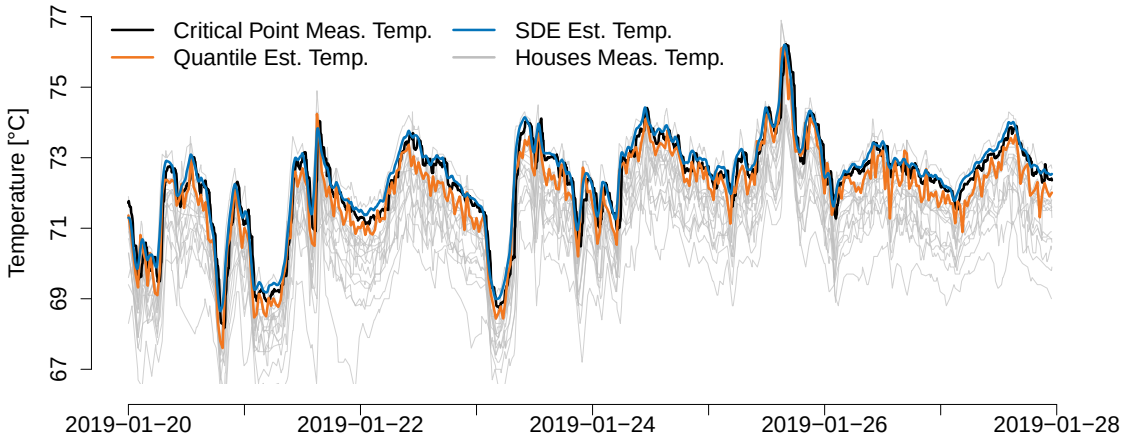

Figure 4.7: Figure shows the measured house temperatures (grey lines) together with the measured critical point temperatures (black line). The estimated temperatures were calculated using either the method proposed in Paper B (orange line) or the method proposed in Paper C. The measurements used here are from Area 1 from Paper C.

implemented to create temperature feedback from a group of smart meters. The method estimates the temperature by taking, for example, the 90th percentile of the temperature readings from the smart meters that have been cleaned as described above. The second method from Paper C is a more robust method that uses SDEs and Kalman filter to estimate the temperatures from the smart meters. However, this method is complex and requires more computation time. The first method is also a non-parameter method with one ”parameter” that needs to be adjusted, namely the quantile. In contrast, the second method has multiple parameters and scales with the number of smart meters used.

Due to the complexity and advance of the second method, it leads to a more robust and precise temperature estimate, as shown in Figure 4.7. The temperature estimation of the two methods is shown; the ”Quantile Est. Temp.” is from the first method, and the ”SDE Est. Temp.” is from the second method. The estimated temperatures can then be compared to the measured temperature at a critical point. The critical point, in this case, is just outside the area where the buildings are located that provide the smart meter readings. The temperature readings from the smart meters are also visualised to show the difference between the two methods. The estimated temperature from the second method (”Quantile Est. Temp.”) is more accurate and precise as it does not fluctuate as much as the estimated temperature from the first method. While the estimated temperatures from the first method (Quantile Est. Temp.) follow the smart meter readings reasonably closely and are therefore more affected by the irregular heat consumption.

The estimated temperature from these two methods can be used as temperature feedback, which can either be used only to determine whether the temperature in the network is high enough or can be used for more intelligent operation, e.g. temperature control. This will give network operators a better understanding of the network and its characteristics. It also increases the feasibility of a district heating network with multiple temperature zones, as it is possible to get temperature feedback from anywhere in the network, i.e. a decentralised network. Finally, using smart meters instead of installing and maintaining a measurement well will increase the utility’s savings, as installing and operating these chambers are time and cost-consuming.

4.2 Savings

Paper G

The most significant savings from data-driven temperature optimisation to control the supply temperature in the network come from lowering the supply temperature. This reduces production costs as not as much heat needs to be generated, as lowering the supply temperature means less heat loss in the system and an increase in the ratio of electricity to heat output in CHP-powered district heating systems [12]. This is beneficial for the utility as electricity is more valuable than heat, thus allowing for more profitable operation. As mentioned in Chapter 2, lowering the supply temperature in the network also increases the flexibility of the system, e.g. the operation of heat pumps becomes more feasible [5].

However, lowering the supply temperature results in a higher flow, as the heat load remains the same, so pumping costs increase when the supply temperature is lowered. This is reasonable as the electricity cost of running the pumps is negligible compared to the cost of producing a higher supply temperature. In Nielsen and Madsen [56], an online operation of temperature optimisation in the district heating network in Roskilde, Denmark, was carried out. An economic analysis was done to reduce the supply temperature and increase the flow, which showed that large savings could be achieved through data-driven temperature optimisation. Although pumping costs increased, they were negligible compared to the savings from the reduction in supply temperature[1].

An additional advantage of data-driven temperature optimisation, which uses feedback from the critical points, is that it uses a reference curve as a constraint that the temperature at these points must not be below the reference (with a probability of, e.g. 99%) [51]. This ensures that consumers get what they were promised. In addition, the load on the system is improved because the fluctuations in the supply temperature in the system are limited, i.e. it must not fluctuate too much. Suppose the supply temperature fluctuates too much over long periods of time. In that case, this increases the maintenance costs of the network, as the pipes are subjected to too much stress as they expand and contract with temperature fluctuations Nielsen et al. [57].

The Paper G and Paper B were produced as part of a small project on data-driven temperature optimisation of a small district heating network. One of the tasks was to compare traditional control operations and data-driven temperature optimisation operations in an online setting. Thus, the data-driven temperature optimisation controlled the temperature of a network in a real-world case study for a heating season that started on 1 November 2020 and lasted until 1 April 2021. One of the main tasks of this project was to investigate whether smart meters could be used as feedback from the network for control. Smart meter measurements were investigated to see if a group of smart meters could be used to estimate network temperature as shown in Section 4.1.1. However, the buildings inside the district heating network are mainly large apartment buildings therefore temperature readings from their smart meters could be used directly. Large apartment buildings usually have constant heat consumption; therefore, temperature readings are quite reliable as the hot water does not become still as for single-family buildings. In this project, three apartment buildings were selected that had the most reliable measurements in the past to be used as temperature feedback for data-driven temperature optimisation.

The previous operation used an open-loop temperature optimisation based on a purely physical simulated-driven model (i.e., white-box model) to estimate the supply temperature set points at the heat exchanger supplying heat to Tingbjerg. The open-loop refers to that no feedback

[1]Note: The electricity price influences pumping costs, and at the time of writing, the electricity price market is historically high and very volatile. An additional signal (electricity price) for the controller might be needed to compare the economic costs of running the pumps at high flow or reducing the supply temperature. This also strongly depends on which heat production units are available.

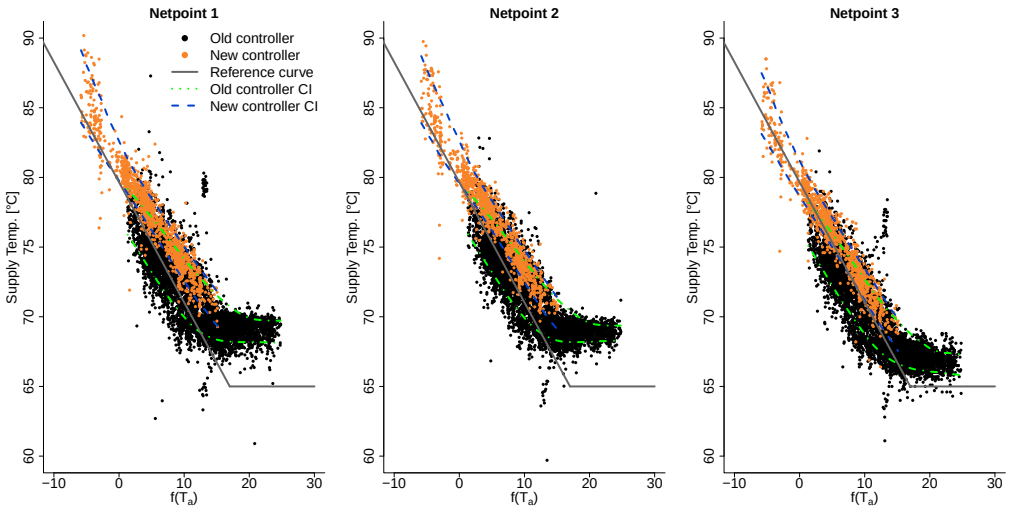

Figure 4.8: Figure compares the operation between the two controllers at the three critical points in the Tingbjerg district heating network by plotting the supply temperature against the rolling average of the ambient air temperature for the last 24 hours. The reference curve and the estimated confidence interval of the operation are shown. Figure from Paper B.

from the system was used, i.e. no temperature feedback. The white-box model is based on the knowledge of the network in the area (combined with the operator's experience) using the scaled heat load forecast from Section 3.2.1 and the ambient air temperature as inputs to estimate the set points. In the new operation, state-of-the-art temperature optimisation was used, HEATTO™[2] to compare the advantage of using a data-driven method. Figure 4.8 shows the supply temperature at the three critical points for both operating periods compared to the rolling average of the ambient air temperature for the last 24 hours. The reference curve used for both operations is also plotted with the points along the confidence intervals of the points obtained using the non-parametric quantile regression with the 10th and 90th percentiles. Analysis of this figure shows that the data-driven operation with feedback operates with higher precision and rarely violates the reference constraint, while the previous operation violates it quite frequently. The dispersion is significantly larger during the previous operation, which is evident from the data points and the comparison of the confidence intervals.

Figure 4.8 shows that the previous operation was allowed to violate the reference curve without receiving complaints from consumers, so the reference curve is too high. Therefore, a "what if" scenario was created where the data-driven operation was adjusted downwards by 5℃, and a new reference curve was proposed that fits as a constraint based on the previous operation. The "what if" scenario is shown in Figure 4.9. The figure shows that the temperature can be lowered significantly using a data-driven method if the constraints are similar to those under which the previous operation could be performed without complaint.

Since it is impossible to run two online operation controllers simultaneously, the comparison between the two operations must be made between the different heating seasons. Usually, degree days are used to compare the two operations. Degree days are often used for calculations on the effect of ambient air temperature on energy consumption. In this case, degree days compare the supply temperature between heating seasons. Degree days, T^{dd}, are calculated as the positive difference between the average ambient temperature ($\bar{T}_a$) over a day and a cut-off of heating

[2] https://enfor.dk/services/heatto/

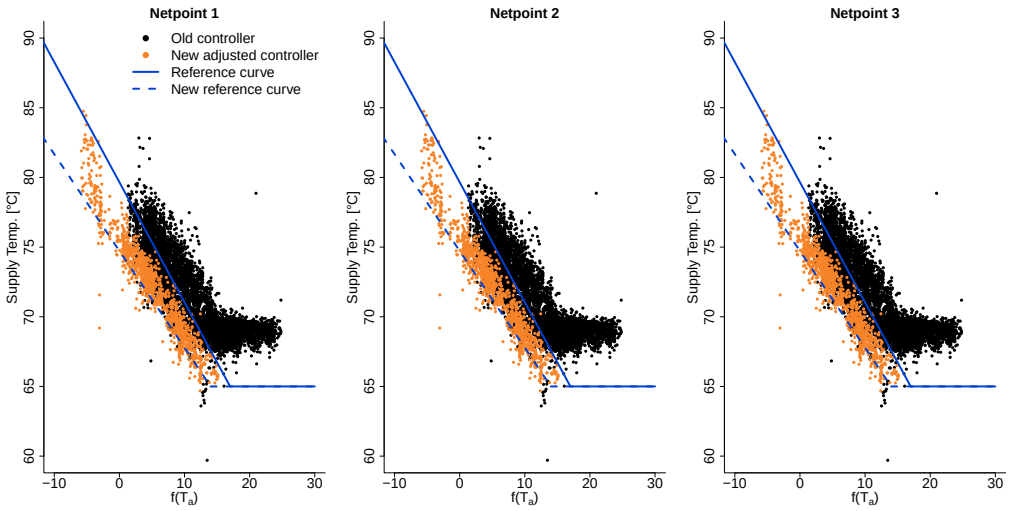

Figure 4.9: Figure demonstrate the ”what if” scenario by adjusting the data-driven operation by 5℃ to show that perfroamnce of the operation when is not constrained by a reference curve that is too high. Figure from Paper B.

demand from buildings (17℃ is used here), i.e.

$$T^{dd} = \max(0, 17 - \bar{T}_a). \tag{4.7}$$

Therefore, it is then possible to compare the operation using the network's average daily supply temperature as a degree day function. Figure 4.10 shows the performance for both operations and the ”what-if” scenario in the right-hand plots for both the supply and return temperatures. The data-driven operation (the *new controller*) performs worse than the previous operation (*old controller*), as expected because the data-driven operation is subject to higher constraints as it does not violate the reference curve. To compare the operations more accurately, a regression model was fitted for each operation using Ordinary Least Squares (OLS) to estimate the parameters of the model with an intercept and slope, as shown in Figure 4.10 and rewritten here below,

$$\text{New controller: } T_{supply} = 68.48 + 0.71T^{dd}, \tag{4.8}$$

$$\text{Old controller: } T_{supply} = 62.61 + 0.97T^{dd}, \tag{4.9}$$

$$\text{New adjusted controller: } T_{supply} = 63.48 + 0.71T^{dd}. \tag{4.10}$$

The slope of the two operations shows us that the data-driven operation gives lower temperatures on days with higher degrees. However, there is a large difference in the intercept due to the reference curve. The data-driven operation may not fall below it, while the previous operation is scattered around it. It can therefore be assumed that the data-driven operation would perform better if the reference curve were lower than the current one, matching what appears to be the ”real” reference curve from the previous operation. This is also evident from the model for the adjusted controller (the ”what-if” scenario) and the data points of the supply temperature as a function of degree days in Figure 4.10.

Figure 4.10 also shows the controllers' precision in the plots on the left, where the ambient air temperature and the supply temperature are shown together over the operating periods. The lower plot on the left shows the difference series (current value minus past value) of the supply temperature for both operations over time. These plots show that the data-driven operation is more precise, as the supply temperature has a significantly lower variance in the difference series.

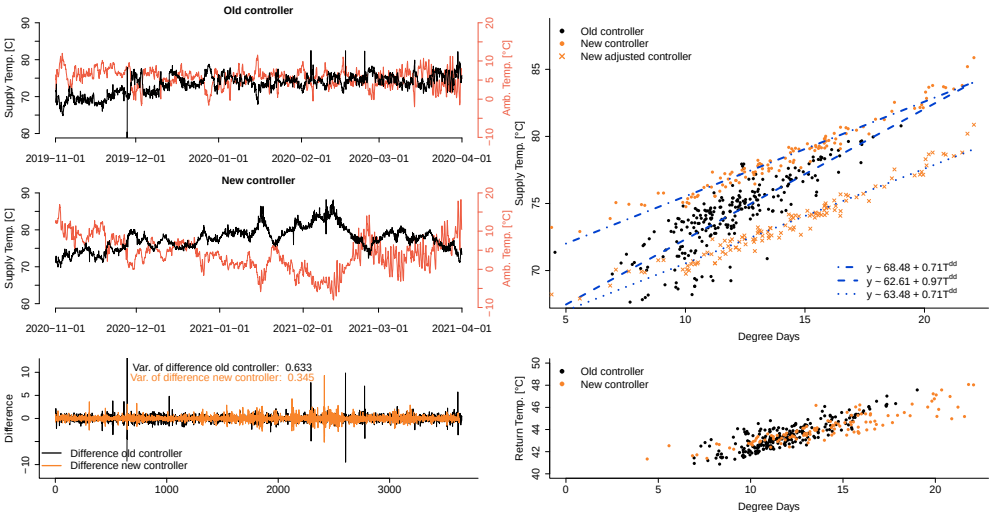

Figure 4.10: Figure shows different plots comparing the two operations in Tingbjerg. The plots on the left compare the robustness of the operations, while the plot on the right shows the performance of the operations with respect to supply and return temperatures. The plots on the right also show the "what if" scenario, where all supply temperatures are reduced by 5℃, as discussed previously. Figure from Paper B.

This can also be seen visually, although it was significantly colder and more volatile during the data-driven operation period.

The savings estimate from lowering the supply temperature between two heating seasons is an approximation if only supply temperature and ambient air temperature data are used. However, it gives a good idea of the potential savings, as lowering the supply temperature increases the utility's savings. Madsen et al. [58] suggests a rule of thumb for the savings that result from lowering the supply temperature in the CHP plant; For each degree lowered, the saving for heat loss in the network is 0.5%, and the saving from more efficient production is 1%, so the savings can be calculated as follows

$$\text{Savings} = (\text{Cost}_{\text{before}} * x\ [^{\circ}C] * 0.5\%) + \text{Shares}_{Production}(\text{Cost}_{\text{before}} * x\ [^{\circ}C] * 1\%) \quad (4.11)$$

where x is the supply temperature difference between the operations to evaluate the savings

Estimating the savings in Tingbjerg using data-driven temperature optimisation compared to the previous operation would not result in any savings, as Figure 4.10 shows that the supply temperature has increased on average. However, analysing the "what if" scenario, it can be assumed that the supply temperature could be reduced by about 3℃ on average. By using the rule of thumb, the savings would be about 4.5%.

4.2.1 Main Findings

In Paper G, it was shown that it is possible to use smart meters as feedback of the temperature network. To demonstrate this, data-driven temperature optimisation using smart meter measurements as feedback was performed in an online experiment. Savings and improvements for the operation of a district heating network by reducing the supply temperature were also discussed. A comparison between open-loop (old controller) and closed-loop (new controller) controllers is carried out to demonstrate the potential of a closed-loop controller where the network characteristics are taken into account when calculating the future set points for the supply temperature.

Unfortunately, the reference curve imposed on the data-driven operation was too high compared to the previous operation. Due to not having feedback on the critical points in the previous op-

eration, it was not recognised that the reference curve constraint was frequently violated. This led to difficulties in comparing the two operations because the data-driven operation uses the reference curve as a constraint that must not be violated. Therefore, comparing the supply temperature would not be an accurate performance measurement. However, it can be seen that the data-driven operation is more efficient and accurate when comparing the supply temperatures between the operations, as it rarely violates the reference curve while keeping the temperatures very close to it. This positively impacts operations as frequent and rapid changes in supply temperature are reduced. This reduces maintenance costs for the utility and increases the certainty that the correct temperature is delivered without too much uncertainty.

In the Tingbjerg study, no savings were demonstrated because the supply temperature was higher for the data-driven operation. However, other studies have demonstrated this. For instance, a study at Brønderslev Forsyning showed a temperature lower by about 7℃ for a degree day of 10, as shown in Paper F in Figure 5, and the savings for a 7℃ lower temperature would be 10.5%. Another study was carried out at Svebølle Viskinge Fjernvarmeselskab, where the temperature was lowered by 8℃ on average, and the savings would then be 12%, as shown in Paper G in Figure A.1.

4.3 Smart Network

Paper D & Paper E & Paper F

The optimal operation of district heating systems also requires ”smart” consumers to achieve optimal operation. Consumers can make their buildings more energy efficient by better thermal insulation (i.e. storing heat more efficiently) or by using smart control of their heating, or optimally would be doing both. Smart control reduces energy consumption and/or increases the flexibility of the buildings by shifting the load by preheating the houses before peak hours. Buildings can be considered as heat reservoirs that consume more heat and increase the indoor air temperature to the maximum thermal comfort when the flexibility signal from the district heating indicates it. The flexibility signal from the district heating could be a price signal or just a signal indicating a surplus of heat in the system. For instance, the indoor air temperature can be raised to a maximum, e.g. one or two hours before people wake up and start to use the domestic hot water, so there is no heating during these peak times in the morning, i.e. peak shaving. This can be done through the use of intelligent control, using multiple input signals to maximize the flexibility of district heating while maintaining a comfortable indoor environment for consumers. However, flexibility and peak shaving do not reduce the overall heat load; it reduces the peak load, but the base load is increased [9]. Using heat intelligently increases the possibility of lowering the supply temperature in the network, thus making the operation more efficient by reducing operating costs and heat losses. This also means that additional heat sources for the area are feasible and thus more operationally efficient. Another advantage of smart controlling the heat consumption is that it can lower the return temperatures, which makes the operation of the heat units, especially the CHP unit, more efficient [12].

Thus, if buildings are operated more intelligently and made more energy efficient (e.g. better thermal insulation), the possibility of decentralising district heating increases as the supply temperature of the network can be lowered. It will also increase the efficiency of sector coupling. Thus, additional heat sources, such as heat pumps, can be used efficiently to cover the heat load. In Paper D, it is proposed to use a non-linear grey-box model with stochastic differential equations to describe the heat dynamics of a school building in Denmark. Heating is provided by a hydraulic heating system that delivers heat to each room via radiators. IoT sensors are used to collect data to estimate the model parameters, e.g. the indoor temperature of the room is measured and controlled thermostatic radiators are used to change the indoor temperature. The model was identified using data from an experiment conducted to generate suitable data for iden-

tification. The experiment is described in detail in Paper D. For simplicity, the mean indoor air temperature of the rooms is modelled, not that of individual rooms (the control of individual rooms is done in Thilker et al. [59]). The variables that should be modelled to predict the building's heating system are 1) the mean indoor air temperature (T_t^{i}), 2) the heat load of the building (ϕ_t^{h}), 3) the temperature of the return water of the heating system (T_t^{ret}).

The proposed model was found by splitting the system into two parts: 1) identifying the model that predicts only the indoor air temperature based on the measurements of the delivered heat from the primary side of the building's heat exchanger (i.e. the district heating side), 2) identifying the model that predicts the heat load from the primary side (district heating) to the secondary side (building heating) based on the measured indoor air temperature and set points while keeping the parameters fixed that were estimated in 1). After identifying a suitable model for these two objectives, it is possible to combine these two models into one and estimate the coefficients of the combined model using the maximum likelihood estimation procedure with the continuous-discrete extended Kalman filter to update and predict the states of the system.

An additional part of the model that needs to be highlighted is the way in which the radiator flow is modelled for thermostatic control. Thus, the amount of heat the radiator emits to regulate the indoor air temperature results from the relationship between the flow rate and the temperature difference across the radiator. In Paper D, it is proposed to estimate the flow by modelling the state of the thermostatic valve between 0 (closed, no flow) and 1 (fully open, maximum flow) using the sigmoid function and set point and the measured indoor air temperature. An offset is also proposed, as the indoor air temperature is not measured directly next to the radiators. It follows that,

$$f_t^{\text{valve}} = \frac{1}{1 + \exp\left(-\alpha(T_t^{\text{set}} + T_{\text{offset}} - T_t^i)\right)}\,, \tag{4.12}$$

where α is the slope that determines how fast the heating systems turn on and off.

The final proposed combined model,

$$\mathrm{d}T_t^i = \frac{1}{C_i}\left(\frac{1}{R_{ih}}\left(T_t^h - T_t^i\right) + \frac{1}{R_{iw}}\left(T_t^w - T_t^i\right) + A_w\phi_t^s\right)\mathrm{d}t + \sigma_1\mathrm{d}\omega_t^1\,, \tag{4.13a}$$

$$\mathrm{d}T_t^w = \frac{1}{C_w}\left(\frac{1}{R_{iw}}\left(T_t^i - T_t^w\right) + \frac{1}{R_{wa}}\left(T_t^a - T_t^w\right)\right)\mathrm{d}t + \sigma_2\mathrm{d}\omega_t^2\,, \tag{4.13b}$$

$$\mathrm{d}\Phi_t = \frac{1}{C_f}\left(\Phi_{\max} f_{\text{valve}} - \Phi_t\right)\mathrm{d}t + \sigma_3\mathrm{d}\omega_t^3\,, \tag{4.13c}$$

$$\mathrm{d}T_t^h = \frac{1}{C_h}\left(\Phi_t c_{p,w}\left(T_t^{\text{for}} - T_t^h\right) + \frac{1}{R_{ih}}\left(T_t^i - T_t^h\right)\right)\mathrm{d}t + \sigma_4\mathrm{d}\omega_t^4\,, \tag{4.13d}$$

$$\mathrm{d}T_t^{\text{ret}} = \frac{1}{C_h}\left(\frac{1}{R_{fr}}\left(T_t^h - T_t^{\text{ret}}\right)\right)\mathrm{d}t + \sigma_t\mathrm{d}\omega_t^5\,. \tag{4.13e}$$

where Φ is the flow rate of the water on the building site. The observation equations are

$$y_k^i = T_{t_k}^i + v_1\,, \qquad v_1 \sim N_{iid}(0, R_1)\,, \tag{4.14a}$$

$$y_k^h = \Phi_t c_{p,w}\left(T_t^{\text{for}} - T_t^{\text{ret}}\right) + v_2\,, \qquad v_2 \sim N_{iid}(0, R_2)\,, \tag{4.14b}$$

$$y_k^{\text{ret}} = T_{t_k}^{\text{ret}} + v_3\,, \qquad v_3 \sim N_{iid}(0, R_3)\,. \tag{4.14c}$$

Thus, the proposed model is used to describe the heating dynamics of school buildings using a physically inspired nonlinear SDE model to predict the indoor air temperature, heat load and return temperature of space heating. This model can then be used for smart building control, e.g.

to reduce energy costs, increase the flexibility of the district heating network (peak shaving), reduce CO2 emissions or create better thermal comfort in the building.

This can be done with the help of model predictive control (MPC), which uses the proposed model to achieve one of these goals. Paper E demonstrates the potential of detailed modelling of the school building and MPC to reduce the operating costs and maximise the flexibility potential. The reduction in operating costs was demonstrated by simulating the control strategy with the aim of reducing the return temperature based on the district heating utility requirement that the consumer must pay an additional charge if the return temperature is too high. The analysis showed that if the return temperature is above 40℃, the heat prices increase by 2% per additional degree of return temperature. The analysis showed that economic savings of 10% can be achieved by using smart control. The flexibility study of the smart control demo is demonstrated by simulating a heating price signal from the district heating utility. The results show that the smart control can unlock the flexibility of the building by preheating the building before the peak hours (defined as the most expensive heating hours) without violating the thermal comfort constraints of the indoor air temperature, both before the peak hours and during.

4.3.1 Main Findings

Paper D and Paper E show how buildings can become smart by modelling heat dynamics using physically driven equations and installing IoT sensors to measure the desired variables. It will also show how to set up the desired control cost function to achieve optimal heating operation. Smart buildings can then be used to increase network operation efficiency by receiving a signal from the utility (e.g. a price signal, a penalty signal, etc.) and allowing the building's control system to respond based on this signal, i.e. increase the flexibility potential of district heating. However, the building's individual control system should regulate the heating to meet the consumer's wishes in terms of thermal comfort and economy. There must therefore be some kind of coordination between the consumer and the district heating utility in order to achieve optimal cooperation between the two systems.

Hence, consumers are important actors on the way to decentralised district heating systems and maximise their flexible potential, as the higher energy efficiency of buildings leads to lower heat load and peak shaving. Therefore, more local heat sources can be used, and the utility's operating costs can be reduced.

5 Conclusion and Perspectives

This PhD thesis aims to investigate the value and benefits of new additional data sources and to apply new data-driven methods to enhance the operation of a district heating network. The focus was on improving the current state-of-the-art methods for heat load forecasting and temperature optimisation using new methodologies and new data.

A data-driven heat load forecasting model is presented, and a framework for dealing with the non-stationary and non-linear characteristics of the heat load is introduced. Also, an individual forecasting model is created for each desired forecast horizon. The framework consists of two stages, the transformation stage and the regression stage. In the transformation stage, the independent variables are modified to account for the non-linearity between them and the response variable, while in the regression stage, the parameters of the model are estimated using a least squares method. A recursive least squares method with forgetting is proposed to allow the parameters of the model to adapt over time to account for the non-stationarity in the heat load. As district heating becomes increasingly decentralised, with sub-networks having their own temperature optimisation and heat units, the importance of localised heat load forecasting is introduced. Its importance is demonstrated by localising the forecast model using area measurements and the weather input variables by adjusting them to the area's climate using local climate measurements. It is shown that the proposed localisation can increase the accuracy of the heat load forecast. Furthermore, the value of accurate heat load forecasts is discussed with regard to the operation of a district heating network, e.g. when optimising the supply temperature provided by a heat exchanger.

Finally, operational state-of-the-art heat load forecast improvements are demonstrated using hierarchical forecasting techniques where individual forecasts share information across the hierarchy. Both temporal and spatial hierarchies are suggested. A recursive and adaptive covariance estimator is proposed for the reconciliation process to achieve optimal improvements due to the non-stationary nature of the heat load. With spatial and temporal hierarchies, significant improvements in heat load forecast accuracy can be achieved, on average, about 15%. For future energy systems, accurate forecasts will be very valuable, especially for district heating and unlocking the potential flexibility to maximise energy system efficiency. More granulated forecasts for local heat load will also be needed as district heating systems become more decentralised and include more local heat units. Therefore, more accurate forecasts will benefit the optimisation of both temperature levels and production.

Data-driven temperature optimisation is introduced and demonstrated through a comparison with traditional optimisation. Traditional optimisation is based on a purely physical simulation-based model, where no network feedback is considered. In contrast, data-driven optimisation uses physical insights to operate the system using measurements to estimate the parameters of the model using statistical methods and exploiting the temperature feedback of the network to gain knowledge of the network characteristics. It is shown that data-driven optimisation can significantly reduce the supply temperature and increase precision, hence reducing operating costs. The potential savings from data-driven temperature optimisation and lowering the supply temperature for district heating networks were discussed. It is shown that network temperature feedback is an essential part of temperature control to increase accuracy by actively learning about the characteristics of the network and ensuring that consumers receive sufficient temperatures. Nowadays, there are smart meters in every house connected to a district heating network, which offers the possibility to determine the network temperature using the measurement from the meters and use it as feedback. Two methods for estimating network temperatures from a

group of smart meters from single-family houses are presented to increase the potential of data-driven temperature optimisation and the possibility of zones with multiple temperature levels. Since the network is dynamic, the critical points may change over time, so the location of the feedback needs to be changed, which can be done by selecting a new group of smart meters. One of the two proposed methods can be implemented easily where the temperature measurements are cleaned, and the network temperature is estimated by computing the 90th percentile of smart meter temperatures at the same time. While the other method is more complex, using the grey-box modelling approach to estimate the network temperatures. This is a more accurate and robust approach to estimating the network temperature due to the quality of smart meter measurements, as temperature readings are not accurate when the heat consumption is low. It is therefore proposed and demonstrated that smart meter measurements can be used as temperature feedback for temperature optimisation.

This thesis demonstrated that meter data could be used as the source for feedback for temperature control through an online case study where smart meter measurements from users with high heat consumption (large apartment buildings) were used directly as feedback. The precision of temperature optimisation using smart meters as the feedback has been shown to be higher than the previously used simulation-based approach. It is also worth mentioning that the percentile method for estimating the network temperature from the smart meter measurements has been integrated into the HEATTO temperature optimisation software and has been in operation for more than a year with promising results.

Finally, it is shown that smart control of the district heating consumer's heating system can improve its operation, either by improving thermal comfort or economic aspects, by creating a grey-box model of the thermal dynamics of the buildings. In addition, the consumers can play an important role in improving the operation of the network, e.g. by delivering flexibility for reducing the load during peak hours through interaction with the district heating. This can be done e.g. with time-varying heat prices. The future operation of district heating with the decentralised system (with local heat units and TES systems) and the interaction with the consumer buildings will be complex and challenging to operate all of its components (*production*, *network*, *consumer*) optimally. This complexity calls for the widespread utilisation of data-driven methods.

The results of this thesis can be summarised as follows: Data-driven methods and physical knowledge of the system are essential for improving network operation, and digitalisation is playing an important role in the future weather-driven energy system. The efficiency of district heating network operation is crucial for the energy transition, where the share of renewable energy sources is increasing. District heating is recognised as an essential player in the energy transition due to its energy flexibility through efficient heat storage and conversion of electricity into heat. However, robust and accurate methods are needed to maximise the flexibility potential of district heating by improving decision-making and operations. The methods presented here have proven to be robust and accurate in real-life tests and implementations. As they exploit the physical nature of the system, they are considered suitable (reliable) for operating the system in a real-life online environment.

Bibliography

[1] Gudni Axelsson et al. “Low-temperature geothermal utilization in Iceland–Decades of experience”. In: *Geothermics* 39.4 (2010), pp. 329–338.

[2] Daniela Guericke, Amos Schledorn, and Henrik Madsen. “Optimization of Heat Production for Electricity Market Participation”. In: *Handbook of Low Temperature District Heating*. Springer, 2022, pp. 179–193.

[3] Meng Yuan et al. “District heating in 100% renewable energy systems: Combining industrial excess heat and heat pumps”. In: *Energy Conversion and Management* 244 (2021), p. 114527.

[4] Daniel Tschopp et al. “Large-scale solar thermal systems in leading countries: A review and comparative study of Denmark, China, Germany and Austria”. In: *Applied energy* 270 (2020), p. 114997.

[5] Andrei David et al. “Heat Roadmap Europe: Large-Scale Electric Heat Pumps in District Heating Systems”. In: *Energies* 10.4 (2017).

[6] Henrik Lund et al. “4th Generation District Heating (4GDH): Integrating smart thermal grids into future sustainable energy systems”. In: *Energy* 68 (2014), pp. 1–11.

[7] Elisa Guelpa and Vittorio Verda. “Thermal energy storage in district heating and cooling systems: A review”. In: *Applied Energy* 252 (2019), p. 113474.

[8] Ignacio Blanco et al. “Operational planning and bidding for district heating systems with uncertain renewable energy production”. In: *Energies* 11.12 (2018), p. 3310.

[9] Annelies Vandermeulen, Bram van der Heijde, and Lieve Helsen. “Controlling district heating and cooling networks to unlock flexibility: A review”. In: *Energy* 151 (2018), pp. 103–115.

[10] Henrik Aalborg Nielsen and Henrik Madsen. “Modelling the heat consumption in district heating systems using a grey-box approach”. In: *Energy and buildings* 38.1 (2006), pp. 63–71.

[11] Dominik Franjo Dominković et al. “Implementing flexibility into energy planning models: Soft-linking of a high-level energy planning model and a short-term operational model”. In: *Applied Energy* 260 (2020), p. 114292.

[12] Lars Arvastson. “Stochastic Modeling and Operational Optimization in District Heating Systems”. PhD thesis. Centre for Mathematical Sciences, Lund University, 2001.

[13] Ruud Egging-Bratseth et al. “Seasonal storage and demand side management in district heating systems with demand uncertainty”. In: *Applied Energy* 285 (2021), p. 116392.

[14] Ignacio Blanco. “Decision-making Under Uncertainty for the Operation of Integrated Energy Systems”. PhD thesis. 2019.

[15] Tryggvi Jónsson et al. “Forecasting electricity spot prices accounting for wind power predictions”. In: *IEEE Transactions on Sustainable Energy* 4.1 (2012), pp. 210–218.

[16] Alessandro Dalla Rosa et al. *IEA DHC Annex X report: Toward 4th Generation District Heating: Experience and Potential of Low-Temperature District Heating*. 2014.

[17] Peter Bauer, Alan Thorpe, and Gilbert Brunet. “The quiet revolution of numerical weather prediction”. In: *Nature* 525.7567 (2015), pp. 47–55.

[18] G.J. Steeneveld et al. “Quantifying urban heat island effects and human comfort for cities of variable size and urban morphology in the Netherlands”. In: *Journal of Geophysical Research Atmospheres* 116.20 (2011).

[19] Philippe Crochet. “Adaptive Kalman filtering of 2-metre temperature and 10-metre windspeed forecasts in Iceland”. In: *Meteorological Applications* 11.2 (2004), pp. 173–187.

[20] European Parliament and Council of the European Union. *Directive 2012/27/EU*. Accessed on 28 March 2021. Brussels, 2012.
[21] Jonas Kjeld Kirstein. "Data-driven water distribution system analysis – exploring challenges and potentials from smart meters and beyond". PhD thesis. 2020.
[22] Kamstrup. *Leak Detector*. https://www.kamstrup.com/en-en/water-solutions/water-analytics/leak-detector. [Online; accessed 6-November-2022]. 2022.
[23] Robert T Clemen. "Combining forecasts: A review and annotated bibliography". In: *International journal of forecasting* 5.4 (1989), pp. 559–583.
[24] Tao Hong et al. "Energy forecasting: A review and outlook". In: *IEEE Open Access Journal of Power and Energy* 7 (2020), pp. 376–388.
[25] Ross Hollyman, Fotios Petropoulos, and Michael E. Tipping. "Understanding forecast reconciliation". In: *European Journal of Operational Research* 294.1 (2021), pp. 149–160.
[26] Peder Bacher et al. "Non-parametric method for separating domestic hot water heating spikes and space heating". In: *Energy and Buildings* 130 (2016), pp. 107–112.
[27] Henrik Aalborg Nielsen and Henrik Madsen. *Predicting the Heat Consumption in District Heating Systems using Meteorological Forecasts*. Informatics and Mathematical Modelling, Technical University of Denmark, 2000.
[28] E. Dotzauer. "Simple model for prediction of loads in district - heating systems". In: *Applied Energy* 73.3-4 (2002), pp. 277–284.
[29] Magnus Dahl et al. "Improving short-term heat load forecasts with calendar and holiday data". In: *Energies* 11.7 (2018), p. 1678.
[30] Lennart Ljung and Torsten Söderström. *Theory and practice of recursive identification*. Vol. 4. The MIT press series in signal processing, optimization, and control. MIT Press, 1983.
[31] Henrik Madsen. *Time series analysis*. Chapman & Hall, 2007.
[32] Peder Bacher, Henrik Madsen, and Henrik Aalborg Nielsen. "Online short-term solar power forecasting". In: *Solar energy* 83.10 (2009), pp. 1772–1783.
[33] Pierre Pinson et al. "Temperature prediction at critical points in district heating systems". In: *European Journal of Operational Research* 194.1 (2009), pp. 163–176.
[34] Olafur P. Palsson, Henrik Madsen, and Henning T. Søgaard. "Generalized predictive control for non-stationary systems". In: *Automatica* 30.12 (1994), pp. 1991–1997.
[35] Peder Bacher et al. "Short-term heat load forecasting for single family houses". In: *Energy and buildings* 65 (2013), pp. 101–112.
[36] Danish Meteorological Institute (DMI). *Meteorological Observation (metObs)*. https://confluence.govcloud.dk/display/FDAPI/Danish+Meteorological+Institute+-+Open+Data. Accessed on 24.05.2021. 2021.
[37] Harry R. Glahn and Dale A. Lowry. "The Use of Model Output Statistics (MOS) in Objective Weather Forecasting". In: *Journal of Applied Meteorology (1962-1982)* 11.8 (1972), pp. 1203–1211.
[38] R.J. Hyndman et al. "Optimal combination forecasts for hierarchical time series". In: *Computational Statistics and Data Analysis* 55.9 (2011), pp. 2579–2589.
[39] Shanika L. Wickramasuriya, George Athanasopoulos, and Rob J. Hyndman. "Optimal Forecast Reconciliation for Hierarchical and Grouped Time Series Through Trace Minimization". In: *Journal of the American Statistical Association* 114.526 (2019), pp. 804–819.
[40] Rob J. Hyndman, Alan J. Lee, and Earo Wang. "Fast computation of reconciled forecasts for hierarchical and grouped time series". In: *Computational Statistics and Data Analysis* 97 (2016), pp. 16–32.
[41] G. Athanasopoulos et al. "Forecasting with temporal hierarchies". In: *European Journal of Operational Research* 262.1 (2017), pp. 60–74.
[42] P. Nystrup et al. "Temporal hierarchies with autocorrelation for load forecasting". In: *European Journal of Operational Research* 280.3 (2020), pp. 876–888.

[43] Mikkel L Sørensen et al. “Recent developments in multivariate wind and solar power forecasting”. In: *Wiley Interdisciplinary Reviews: Energy and Environment* (2022), e465.
[44] J. Jeon, A. Panagiotelis, and F. Petropoulos. “Probabilistic forecast reconciliation with applications to wind power and electric load”. In: *European Journal of Operational Research* 279.2 (2019), pp. 364–379.
[45] Gokhan Mert Yagli, Dazhi Yang, and Dipti Srinivasan. “Reconciling solar forecasts: Sequential reconciliation”. In: *Solar Energy* 179 (2019), pp. 391–397.
[46] Olivier Ledoit and Michael Wolf. “Improved estimation of the covariance matrix of stock returns with an application to portfolio selection”. In: *Journal of Empirical Finance* 10.5 (2003), pp. 603–621.
[47] J. Schäfer and K. Strimmer. “A shrinkage approach to large-scale covariance matrix estimation and implications for functional genomics”. In: *Statistical Applications in Genetics and Molecular Biology* 4.1 (2005), pp. 1–30.
[48] I. Oliker. “Steam Turbines for Cogeneration Power Plants”. In: *Journal of Engineering for Power* 102.2 (Apr. 1980), pp. 482–485.
[49] H. Madsen and J. Holst. “Estimation of continuous-time models for the heat dynamics of a building”. In: *Energy and Buildings* 22.1 (1995), pp. 67–79.
[50] T. S. Nielsen. “Online prediction and control in nonlinear stochastic systems”. PhD thesis. Informatics and Mathematical Modelling, Technical University of Denmark, DTU, 2002.
[51] Henrik Madsen et al. “On flow and supply temperature control in district heating systems”. In: *Heat Recovery Systems and CHP* 14.6 (1994), pp. 613–620.
[52] O.P. Palsson, H. Madsen, and H.T. Søgaard. “Application of predictive control in district heating systems”. In: *Proceedings of the Institution of Mechanical Engineers, Part A: Journal of Power and Energy* 207.3 (1993), pp. 157–163.
[53] H. T. Søgaard. “Stochastic systems with embedded parameter variations - applications to district heating”. PhD thesis. Technical University of Denmark, Department of Applied Mathematics and Computer Science, 1993.
[54] B. van der Heijde et al. “Dynamic equation-based thermo-hydraulic pipe model for district heating and cooling systems”. In: *Energy Conversion and Management* 151 (2017), pp. 158–169.
[55] Kasper Kristensen et al. “TMB: Automatic Differentiation and Laplace Approximation”. In: *Journal of Statistical Software* 70.5 (2016), pp. 1–21.
[56] Torben Skov Nielsen and Henrik Madsen. “Control of supply temperature in district heating systems”. In: *Proceedings of the 8th International Symposium on District heating and Cooling*. 2002.
[57] Torben Skov Nielsen et al. *Predictive control of supply temperature in district heating systems*. Informatics and Mathematical Modelling, Technical University of Denmark, 2002.
[58] Henrik Madsen et al. *Models and Methods for Optimization of District Heating Systems.: Part II: Models and Control Methods*. Informatics and Mathematical Modelling, Technical University of Denmark, 1992.
[59] Christian Ankerstjerne Thilker et al. “Identification of non-linear autoregressive models with exogenous inputs for room air temperature modelling”. In: *Energy and AI* (2022), p. 100165.

Part II

Publications

Paper A

HEAT LOAD FORECASTING USING ADAPTIVE TEMPORAL HIERARCHIES

Authors:
Hjörleifur G. Bergsteinsson, Jan Kloppenborg Møller, Peter Nystrup, Ólafur Pétur Pálsson, Daniela Guericke and Henrik Madsen.

Published in:
Applied Energy.

Heat Load Forecasting Using Adaptive Temporal Hierarchies

Hjörleifur G. Bergsteinsson[a,*], Jan Kloppenborg Møller[a], Peter Nystrup[a,b], Ólafur Pétur Pálsson[c], Daniela Guericke[a], Henrik Madsen[a,d]

[a]*Technical University of Denmark, Denmark*
[b]*Lund University, Sweden*
[c]*University of Iceland, Iceland*
[d]*Norwegian University of Science and Technology, Norway*

Abstract

Heat load forecasts are crucial for energy operators in order to optimize the energy production at district heating plants for the coming hours. Furthermore, forecasts of heat load are needed for optimized control of the district heating network since a lower temperature reduces the heat loss, but the required heat supply at the end-users puts a lower limit on the temperature level. Consequently, improving the accuracy of heat load forecasts leads to savings and reduced heat loss by enabling improved control of the network and an optimized production schedule at the plant. This paper proposes the use of temporal hierarchies to enhance the accuracy of heat load forecasts in district heating. Usually, forecasts are only made at the temporal aggregation level that is the most important for the system. However, forecasts for multiple aggregation levels can be reconciled and lead to more accurate forecasts at essentially all aggregation levels. Here it is important that the auto- and cross-covariance between forecast errors at the different aggregation levels are taken into account. This paper suggests a novel framework using temporal hierarchies and adaptive estimation to improve heat load forecast accuracy by optimally combining forecasts from multiple aggregation levels using a reconciliation process. The weights for the reconciliation are computed using an adaptively estimated covariance matrix with a full structure, enabling the process to share time-varying information both within and between aggregation levels. The case study shows that the proposed framework improves the heat load forecast accuracy by 15% compared to commercial state-of-the-art operational forecasts.

Keywords: Heat load forecast, Adaptive forecasting, Temporal hierarchies, Forecast reconciliation, Adaptive estimator, Recursive shrinkage estimator

1. Introduction

Energy forecasting has become an essential method in the green transition due to the increased complexity of energy systems required to achieve high energy efficiency. This is highlighted in Hong et al. [1], who give an extensive historical overview of energy forecasting, including current trends. The authors point to the significant growth of research related to renewable energy forecasts during the past decade. Energy systems that rely on renewable energy sources, e.g., wind and solar, to produce either electricity or heat need an accurate prediction of the energy demand for the upcoming days to ensure that they can fulfill the demand and optimize their operation accordingly. Planning needs to take into account that wind and solar as energy sources can not always provide the required energy demand due to their weather dependency. Therefore, they are integrated with other energy sources to ensure that the energy demand is met. However, wind and solar are an important part of the green transition. The share of renewable energy in the European Union is expected to increase to 70% by 2050 [2]. Therefore, accurate forecasts are becoming increasingly valuable for the energy sector.

Many different forecasting models have been proposed in recent years for the energy sector. From classical methods, e.g. Taylor [3] using double-seasonal Holt-Winters exponential smoothing to forecast electricity demand, to more complicated models, e.g. Nielsen and Madsen [4] using a grey-box model to forecast heat consumption. However, there is probably not one unique forecasting model that can provide the best forecast in all situations for every energy system. Therefore, frequent studies suggest using multiple models, and combining forecasts typically gives better overall forecasts than using only one model [1, 5]. There has also been a significant increase in the number of studies using temporal hierarchies to improve forecast accuracy by utilizing the information between different aggregation levels and by optimally combining them with reconciliation [6, 7]. Using temporal hierarchies could be a beneficial approach to improve the accuracy of energy forecasts because they use both short-term and long-term forecasts to improve forecast accuracy and consequently give coherent forecasts on all aggregation levels. For example, district heating plants need short-term forecasts for operational optimization, such as supply temperature control, while long-term forecasts are used for planning, e.g. biomass supply planning. The district heating sector can gain heavily from the temporal hierarchies, as they need accurate and coherent short-term and long-term forecasts for control and planning.

*Corresponding Author
Email address: hgbe@dtu.dk
Postal Address: Anker Engelunds vej 1, Building 101A, 2800 Kongens Lyngby, Denmark

1.1. District Heating

District heating is an efficient way to provide heat to buildings in densely populated areas. Due to its flexibility, district heating has become a crucial part of the agenda to reach a renewable, non-fossil heat supply in the future. Additionally, district heating increases the flexibility of the overall integrated energy system by storing energy. For example, excess wind power during low-electricity-demand periods can be used to heat water with heat pumps or electric boilers that can be stored to reduce pressure on the heating system during peak hours [8]. The change from traditional fossil fuels to renewable district heating embedded in a smart energy system is referred to as 4^{th} generation district heating [9]. To pave the way, district heating systems are currently undergoing a digital transformation. Along with this transformation, the district heating systems are continuously improving the efficiency of their operations. An important prerequisite for operating a district heating system efficiently is the accuracy of heat load forecasts. Being able to look accurately into the future demand enables the operators to run the system with more precision, resulting in a more efficient operation, hence lowering the system costs.

Lowering the supply temperature in the district heating network using control methods has substantial potential for cost savings, since lower supply temperatures lead to lower heat production costs as well as reduced heat losses in the transmission and distribution network [10, 11]. The supply temperature from the plant and flow in the network are controlled by using feedback from the critical points in the network. The feedback is typically based on an adaptively estimated model describing the time-varying time-delay in the network, as well as the heat loss between the plant and the critical points in the network [4, 10, 12]. The controller uses predictions of the future heat load to adjust the supply temperature in order to reach the lowest possible temperature which still ensures that the temperatures are as required at the critical points. The flow is typically regulated to match the required heat demand while keeping the supply temperature as low as possible [10, 13].

The European Union has set a target of 5% for the share of solar thermal production in district heating systems in 2050 [14]. Tschopp et al. [15] provide an extensive review of the performance and future of solar thermal production in district heating. Current developments in the performance of collectors and control strategies make the use of solar thermal units more attractive for the district heating sector. District heating systems generally design the size of solar farms to cover the heat demand in summer periods without production from other units. By using thermal storage and appropriate control strategies at a solar farm, district heating operators have the flexibility to store heat from surplus production and use this stored heat in periods with low solar radiation [16]. The control strategies for a solar farm require accurate heat load forecasts for optimal utilization of solar heat.

Furthermore, the overall production optimization of district heating systems is highly dependent on forecasts of the heat load [17]. Since production is optimized for the heat load forecast, accurate forecasts result in reliable optimal production planning for the system, with a larger potential for savings and utilization of green energy. Without an accurate forecast, the determined production schedule will be inefficient, costly, and maybe even infeasible. To summarize, the presented scenarios on how heat load forecasts are used to operate district heating more efficiently exemplify why improved heat load forecasts are highly desired and beneficial for the district heating sector.

1.2. Heat Load Forecasting

Heat load forecasting has been frequently studied, and several methods have been proposed along with a selection of variables that influence heat demand as inputs to the forecasting model. Dotzauer [18] suggest using a simple regression model including forecasts of outdoor temperature and heat demand of consumers. This simple regression model gives an accurate prediction of the future heat load by creating a piecewise linear function linking the heat demand to the outdoor temperature and using seasonal profiles of the heat demand. Dahl et al. [11] use the weather forecast uncertainty and an ARX (autoregressive exogenous) model to forecast the heat load using outside temperature, wind speed, and solar irradiance as inputs. The input variables are determined using model selection of different input variables. Their results show that the time-varying uncertainties improve the supply temperature control of a heat exchanger, thereby lowering the supply temperature and increasing savings by reducing heat losses to the surroundings in the transportation pipes.

Nielsen and Madsen [4] model heat consumption using the grey-box modeling approach to take advantage of combining physical and statistical modeling of district heating system. They also propose which input variables to use for heat load forecasting, and in some cases they suggest filtering the input variables due to e.g. the thermal inertia of buildings. The forecasting models presented in this paper are based on the methodology proposed in Bacher et al. [19]. Bacher et al. [19] propose an adaptive and recursive method to forecast heat load for single-family houses. Their method uses recursive least squares (RLS) to estimate the coefficients in the forecasting model. This allows the model coefficients to change over time and adapt to changes. The model is therefore self-calibrating, which is important because heat load is a nonlinear process, highly dependent on weather conditions. Therefore, the coefficients of the model change as the heat load dynamics change over time along with the weather. Bacher et al. [20] use the same RLS scheme for solar-power forecasting of a PV system, and Rasmussen et al. [21] use it to forecast electrical load for supermarket refrigeration after demonstrating how to incorporate nonlinearity into the linear regression model.

1.3. Temporal Hierarchies

Temporal hierarchies have not been applied to heat load forecasts, although such hierarchies have shown promising accuracy improvements in other areas, such as tourism and electrical load forecasting [6, 7]. In temporal hierarchies, the hierarchy levels are different, non-overlapping temporal aggregations. For example, the total heat load of a single-family house over a week can be disaggregated into the total heat load for each day of the week. The heat load for each day can then be

disaggregated into the demand per hour of the day. Here, the top level of the temporal hierarchy is the total weekly heat load, which is the aggregation of the daily heat load for the days in the week, which themselves are aggregations of the hourly heat load for each hour of the day.

Athanasopoulos et al. [6] state that forecasting different temporal aggregation levels reveals different information in the data. They demonstrate that on a lower resolution frequency, e.g. the weekly heat load, the trend in the load could be more dominant, while on a higher frequency, e.g. hourly heat load, seasonality can be more dominant. Hierarchical prediction models on different aggregation levels can capture this different behavior in the data. The information at each level is shared between levels using forecast reconciliation. Forecast reconciliation is the process of optimally combining hierarchical forecasts to yield coherent forecasts. A coherent forecast fulfills the constraints defined by the temporal hierarchy framework proposed by Athanasopoulos et al. [6], i.e. any forecast at an aggregate level is equal to the sum of the respective subaggregate forecasts from the level below.

The bottom-up method is a simple approach to generate coherent forecasts where the forecasts from the lowest aggregation level are aggregated to fit the hierarchical structure created beforehand. The top-down method disaggregates forecasts from the highest aggregation level using probabilities or past experiences to divide the forecasts into lower levels. Neither of these methods takes the relationship between aggregation levels into consideration. Therefore, information is lost between the levels, leading to a sub-optimal result. A detailed summary of traditional hierarchical forecasting is given in Athanasopoulos et al. [22].

Recent studies propose different methods to optimally reconcile forecasts in a temporal hierarchy. Hyndman et al. [23] create an independent forecast for each level, which is referred to as the base forecast and can be created from any model. The base forecasts are then reconciled to yield coherent forecasts. The reconciled forecasts are linear combinations of the base forecasts computed using generalized least squares (GLS) based on the covariance matrix Σ. Wickramasuriya et al. [24] show that the covariance matrix estimated from the coherency errors is nonidentifiable and therefore impossible to estimate. Hyndman et al. [25] propose using a diagonal covariance matrix estimated from the base forecast errors leading to a weighted least squares estimate of the reconciled forecasts. Wickramasuriya et al. [24] provide theoretical justification for using the empirical covariance matrix for the base forecast errors, $\hat{\Sigma}$, as an estimator for Σ. Nystrup et al. [7] demonstrate that using the full covariance matrix for the base forecast errors results in significantly more accurate reconciled forecasts than assuming no auto- and cross-correlation between aggregation levels.

Temporal hierarchies have been applied for several types of energy forecasting. Nystrup et al. [7] use temporal hierarchies to improve short-term electricity load forecasts. Jeon et al. [26] apply the temporal hierarchy framework to ensure coherence of probabilistic forecasts of wind power production, using a cross-validation method to find the weights in the reconciliation process. Yagli et al. [27] and Yang et al. [28] use temporal hierarchies to improve the accuracy of solar power production forecasts for PV plants. These results demonstrate a promising potential of applying a temporal hierarchy to improve heat load forecasts.

The temporal hierarchy frameworks in literature frequently do not take the covariance into account and thereby result in lower accuracy improvements when dismissing the connection between aggregation levels [7]. There have also been no attempts to extend the framework to have adaptive and recursive updates of the covariance matrix. In this paper, we suggest a method to overcome these issues.

1.4. Contribution

This paper proposes a novel framework based on temporal hierarchies to improve the accuracy of heat load forecasts. This is done by combining forecasts from multiple temporal aggregation levels using an adaptive forecast reconciliation method. We propose three different approaches to update the covariance matrix, and hence the weights in the reconciliation process change over time to handle the time-varying dynamics of the heat load. The first method recursively updates the covariance matrix without any forgetting, the second method uses a rolling window with fixed width and equal weights to forget past information, and the third method uses exponentially decaying weights on past information to increase the importance of the most recent observations. All three covariance matrices use a shrinkage estimate of the full matrix structure based on the in-sample prediction errors. This enables the framework to share information both within and between aggregation levels.

We demonstrate the usefulness of the framework by improving the accuracy of state-of-the-art operational heat load forecasts for the Danish district heating planner, *Varmelast*. The operational forecasts were provided to *Varmelast* by a commercial forecast provider. The objective is to increase the accuracy of the hourly heat load forecasts (i.e., the operational heat load forecasts) for day-ahead operational planning using the proposed framework. Forecasts of the heat load for all aggregation levels for the next 24 hours are issued at 23:00 every night. The base forecasts are used in the reconciliation process to generate reconciled forecasts. These forecasts are then used to demonstrate the accuracy improvements for the *Varmelast* case study. The usage of the commercial state-of-the-art forecasts is discussed by comparing them to our own simple forecasting model and the benefit of improving the forecasts using the proposed framework.

The contribution of this article is fourfold: 1) To the best of our knowledge, this is the first article to apply a temporal hierarchy to improve heat load forecasts. 2) We propose three adaptive and recursive methods to estimate the covariance matrix, which is used to reconcile the forecasts across the hierarchy. By allowing for time-varying weights, the approach is able to handle non-stationary processes by adapting to changing dynamics when required. This improves the potential for achieving more accurate forecasts in practical applications where stationary is often an issue. 3) We shrink the covariance matrix before reconciling the heat load forecasts. In order to do so, we extend the Ledoit and Wolf [29] closed-form solution for the

optimal shrinkage intensity parameter by deriving a recursive estimation method. 4) This is the first article to include a commercial state-of-the-art operational forecast in the hierarchy and demonstrate accuracy improvements while using simple forecasts on the other temporal aggregation levels. Our results show that the suggested approach is able to improve state-of-the-art forecasts that have a direct influence on the operation of the system through optimal control of the network supply temperature. In addition to improving the current state-of-the-art forecast, our method yields coherent forecasts at multiple temporal aggregation levels which are useful for planning, e.g. purchases of biomass for a combined heat and power plant [7, 30].

The remainder of this article is organized as follows. The data is presented in Section 2. In Section 3, the base forecast models used to generate forecasts for the aggregation levels above the operational forecast are presented. Section 4 outlines the theory of forecast reconciliation and proposes three different covariance estimators. The results are presented in Section 5 and discussed in Section 6. The paper is concluded in Section 7 with a summary.

2. Data

The data for this study is the heat load in the Greater Copenhagen area in Denmark. The data was provided by *Varmelast*, who deliver heat production planning for the district heating plants within this area. The data is the hourly heat load from 1 January 2016 to 31 December 2019, denoted by

$$\{Y_t; t = 1, \dots, N\}, \tag{1}$$

with the total number of observations $N = 35064$. The measurements are visualized in Figure 1 and the seasonal dynamics of the heat load can clearly be seen. The top plot visualizes the heat load over the four years, showing the yearly seasonality. The yearly dynamics are the result of increased heating demand in the winter when the ambient temperature decreases and decreased heating demand in the summer when the ambient temperature increases. The lower plot shows one winter week to demonstrate the weekly and daily seasonality in the heat load. The weekly and daily patterns can be explained by consumer behaviour, as the demand peaks in the morning and evening when consumers leave for work and return home, respectively. The weekends have a different pattern than the weekdays as the morning peaks disappear, which can be explained by fewer people going to work on the weekend.

2.1. Numerical Weather Prediction

The numerical weather predictions (NWPs) used as input to the forecasting models were provided by MetFor$^{\text{TM}}$ from the commercial forecast provider.[1] The NWPs consist of climate variables with an hourly resolution that are updated every hour and are available to forecast the heat load, $\hat{Y}_{t+k|t}$.

An example of an NWP for the kth forecast horizon is the predicted ambient temperature [in °C] denoted by

$$\left\{T^{\text{a,NWP}}_{t+k|t}, t = 1, \dots, N,\ k = 1, \dots, 24\right\}. \tag{2}$$

[1] `https://enfor.dk/services/metfor/`

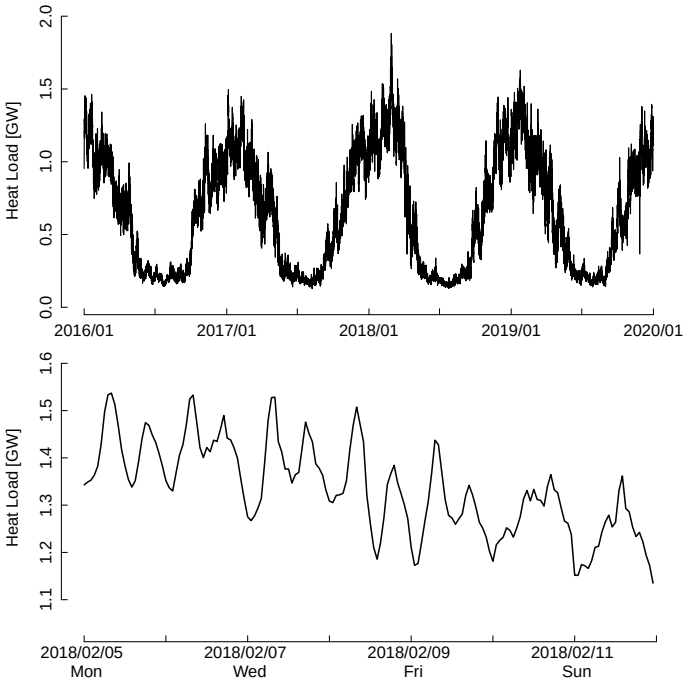

Figure 1: Time series plot of the heat load over the four-year period and of one week to visualise the seasonal dynamics in the data.

2.2. Operational Heat Load Forecasts

Varmelast uses forecasts with an hourly resolution for multiple steps ahead to create the operational plan for the district heating plants. It is crucial to increase the accuracy of these forecasts, since this will allow more accurate planning. The commercial forecast provider supplies the hourly heat load forecasts to *Varmelast*. The heat load forecasts are supplied by HeatFor$^{\text{TM}}$.[2] The heat load forecast, $\hat{Y}_{t+k|t}$, is updated every hour for k-steps ahead, as shown below,

$$\left\{\hat{Y}_{t+k|t}, t = 1, \dots, N,\ k = 1, \dots, 24\right\}. \tag{3}$$

The objective of the paper is to improve the accuracy of the day-ahead hourly heat load forecasts using forecast reconciliation. The forecasts made at 23:00 each night for the next 24 hours will be used to see if the proposed method can improve the accuracy of the hourly forecasts. Thus, the predictions horizon of interest in improving the hourly forecast accuracy is $k = 1, 2, \dots, 24$.

3. Base Forecasts

Base forecasts for each aggregation level are required for the proposed method. We will create base forecasts using the method proposed in Bacher et al. [19], i.e., recursive and adaptive regression-based models, and use some of the proposed input variables from Nielsen and Madsen [4], including their suggestions for filtering of some input variables. The filtering compensates, for instance, the thermal inertia of buildings, where

[2] `https://enfor.dk/services/heatfor/`

only slow variations in the outdoor temperature are reflected in the heat needed to maintain a particular indoor temperature. The hourly heat load forecasts are supplied by the commercial forecast provider, but forecasts on the other temporal aggregation levels are needed. As focus is on the day-ahead operation, the temporal hierarchy will span from the daily level at the top level to the hourly level at the bottom. The aggregation levels are $\ell = \{24, 12, 8, 6, 4, 3, 2, 1\}$. A total of eight different base forecast models will be created, including a simple model for the hourly level to demonstrate the difference between forecasts from a simple model and the commercial state-of-the-art forecasts.

The models will use regression form, for example

$$Y_{t+k} = \theta_0 + \theta_1 T^{\mathrm{a,NWP}}_{t+k|t} + \epsilon_{t+k}, \tag{4}$$

where Y is the heat load, θ are the coefficients, and ϵ is the residual. The predictors in this model are the intercept and the ambient temperature from the NWP. Therefore, the regressor at time t is

$$\boldsymbol{X}_t^{\mathrm{T}} = (1, T^{\mathrm{a,NWP}}_{t+k|t}), \tag{5}$$

and the parameter vector,

$$\boldsymbol{\theta}^T = (\theta_0, \theta_1). \tag{6}$$

Hence, the model can be written as

$$Y_{t+k} = \boldsymbol{X}_t^{\mathrm{T}} \boldsymbol{\theta}_{t+k} + \epsilon_{t+k}. \tag{7}$$

The base forecasts are generated using adaptive methods where the coefficients of the model are time-varying, i.e. they are updated every time a new observation becomes available. In addition, the model uses a forgetting factor to discount old information and increase the importance of the newest observations. This type of forecasting model is referred to as a recursive least squares model with a forgetting factor λ [31]. The solution at time t can be written in a recursive form. The RLS updates the coefficients when new information becomes available using:

$$\hat{\boldsymbol{\theta}}_{k,t} = \hat{\boldsymbol{\theta}}_{k,t-1} + \boldsymbol{R}^{-1}_{k,t} \boldsymbol{X}_{t-k} [Y_t - \boldsymbol{X}^{\mathrm{T}}_{t-k} \hat{\boldsymbol{\theta}}_{k,t-1}], \tag{8}$$

$$\boldsymbol{R}_{k,t} = \lambda \boldsymbol{R}_{k,t-1} + \boldsymbol{X}_{t-k} \boldsymbol{X}^{\mathrm{T}}_{t-k}. \tag{9}$$

That is, the coefficients are recursively updated by weighted least squares estimation with the weights exponentially decaying over time. The rate of decay is determined by the forgetting factor, λ. $\boldsymbol{X}_{t-k}$ is the regressor vector, $\hat{\boldsymbol{\theta}}_{t,k}$ is the coefficient vector, and Y_t is the dependent variable. The subscript k represents the unique coefficient estimates for each k-step prediction model. The regressor vector only uses information available at time $t - k$ to forecast the dependent variable at time t when updating the coefficients.

Each aggregation level has its own unique model that is estimated independently of the other aggregation levels. This is due to the fact that each aggregation level has unique dynamics. Thereby, each model structure is made in the regression form with the k-step prediction at time t being

$$\hat{Y}^{\ell}_{t+k|t} = \boldsymbol{X}_t^{\ell^{\mathrm{T}}} \hat{\boldsymbol{\theta}}^{\ell}_{k,t}, \tag{10}$$

Table 1: The base forecast models for all aggregation levels and their corresponding inputs and steps ahead. Note that the hourly model is the simple model and not the forecast provider model.

Model Aggregation level [ℓ]	Inputs	Steps ahead
Daily	$T^{\mathrm{a,NWP}}_{t,k}$, θ_0, AR(1)	$k = \{1\}$
Twelve-hourly	$T^{\mathrm{a,NWP}}_{t,k}$, θ_0, AR(1)	$k = \{1, 2\}$
Eight-hourly	$T^{\mathrm{a,NWP}}_{t,k}$, θ_0	$k = \{1, 2, 3\}$
Six-hourly	$T^{\mathrm{a,NWP}}_{t,k}$, θ_0, $\mu(t, n_{\mathrm{har}}, \alpha_{\mathrm{diu}})$	$k = \{1, 2, 3, 4\}$
Four-hourly	$T^{\mathrm{a,NWP}}_{t,k}$, θ_0, $\mu(t, n_{\mathrm{har}}, \alpha_{\mathrm{diu}})$	$k = \{1, \ldots, 6\}$
Three-hourly	$T^{\mathrm{a,NWP}}_{t,k}$, θ_0, $\mu(t, n_{\mathrm{har}}, \alpha_{\mathrm{diu}})$	$k = \{1, \ldots, 8\}$
Two-hourly	$T^{\mathrm{a,NWP}}_{t,k}$, θ_0, $\mu(t, n_{\mathrm{har}}, \alpha_{\mathrm{diu}})$	$k = \{1, \ldots, 12\}$
Hourly	$T^{\mathrm{a,NWP}}_{t,k}$, θ_0, $\mu(t, n_{\mathrm{har}}, \alpha_{\mathrm{diu}})$	$k = \{1, \ldots, 24\}$

where the ℓ superscription refers to the target aggregation level of the forecast.

The regression vector, $\boldsymbol{X}$, is created using filtered and non-filtered inputs to describe the heat consumption dynamics adequately for each aggregation level. Table 1 shows the models for each aggregation level with the inputs used and the prediction steps. The ambient temperature forecast, $T^{\mathrm{a,NWP}}_{t+k|t}$, is filtered using a low-pass filter with a stationary gain of one. The systematic heat load peaks in the data are modeled using a diurnal curve model from a harmonic function using a Fourier series, $\mu(t, n_{\mathrm{har}}, \alpha_{\mathrm{diu}})$. The index t is the time of the day in hours, n_{har} is the number of harmonics, and α_{diu} is a vector consisting of the coefficients for the harmonics. An auto-regressive (AR) term is included in the model when needed to remove auto-correlation in the error. Finally, θ_0 is the intercept in the model.

In total, eight different models are created, and each of them has a different number of prediction steps for the next day. The prediction steps for each aggregation level are shown in the last column in Table 1. For example, the top level is the daily aggregation and the prediction is therefore only one step ahead. The other, lower aggregation levels need to deliver forecasts that cover the entire day. For example, at the eight-hourly level, forecasting three steps ahead is necessary to match the 24-hour cycle. We give the daily model in Eq. (11) as an example of a model:

$$\hat{Y}_{\ell=24,t+k|t} = \theta_{0,k} + \theta_{1,k} Y_{\ell=24,t} + \theta_{2,k} H_{a_{\mathrm{T_a}}}(q) T^{\mathrm{a,NWP}}_{\ell=24,t+k|t}, \tag{11}$$

where the filter

$$H_{a_{\mathrm{T_a}}}(q) = \frac{1 - a_{\mathrm{T_a}}}{1 - a_{\mathrm{T_a}} q^{-1}}. \tag{12}$$

The θs are the coefficients of the regression model and $H_{a_{\mathrm{T_a}}}(q)$ is the first order low-pass filter of the ambient temperature forecast, $T^{\mathrm{a,NWP}}_{t,k}$. The q^{-1} is the backward shift operator, i.e. $q^{-1} x_t = x_{t-1}$, and $a_{\mathrm{T_a}} \in [0, 1]$ is a parameter which corresponds to the time constant between the heat load and ambient temperature.

The forgetting factor, λ, and the time constant, $a_{\mathrm{T_a}}$, are parameters that are estimated in an offline setting for each aggregation level by minimizing the root mean square error (RMSE) using the `optim()` function in `R` [32]. The model coefficients,

$\boldsymbol{\theta}$, are updated online whenever a new observation is available using Eq. (8) and Eq. (9).

4. Forecast Reconciliation

The reconciled forecasts are computed from the base forecasts using a temporal hierarchy. In a temporal hierarchy, each aggregation level is the total heat load for the given time resolution, i.e. Y_t^ℓ is the total heat load for the period $(t-\ell, t]$ where ℓ is the aggregation level. As the objective of this paper is to improve the hourly day-ahead heat load forecasts, the hierarchy structure was chosen as the full natural structure with all aggregation levels between the hourly aggregation at the bottom and the daily level at the top, i.e., $\ell = \{24, 12, 8, 6, 4, 3, 2, 1\}$. If an aggregation level is removed from the hierarchy it can result in lower accuracy when the true dynamics of the system are unknown. We consider the full natural hierarchy rather than only a subset of the aggregation levels, as this leads to the largest improvements in accuracy.

The number of aggregation levels could be increased by adding half-hourly or quarter-hourly forecasts; however, the NWPs that are used as inputs to the base forecast models have an hourly resolution, which is also the granularity of interest for the district heating operation. The NWPs would therefore need to be interpolated and it is not clear that this would lead to further improvements. It would also increase the dimension of the hierarchy significantly, which could become an issue. See Nystrup et al. [33] for a discussion of the trade-off between dimensionality and accuracy. We leave it for future work to consider the optimal number of aggregation levels in greater detail.

To demonstrate the concept of a temporal hierarchy, we use an example of a temporal hierarchy with only three levels, as shown in Figure 2. The lowest aggregation level is $\ell = 4h$, the second-lowest is $\ell = 12h$, and the top level is $\ell = 24h$. The top level is the sum of the two 12-hour periods, which are the sum of three different four-hour periods.

Figure 2 can be seen as the summation matrix, $\boldsymbol{S}$, for the hierarchy. The summation matrix corresponding to Figure 2 is given by

$$\boldsymbol{S} = \begin{bmatrix} 1 & 1 & 1 & 1 & 1 & 1 \\ 1 & 1 & 1 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 & 1 & 1 \\ 1 & 0 & 0 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 0 & 0 & 1 \end{bmatrix} \tag{13}$$

The general definition of the summation matrix is given in Nystrup et al. [7] as

$$\boldsymbol{S} = \begin{bmatrix} I_{\mathrm{m}/\ell_1} \otimes 1_{\ell_1}^{\mathrm{T}} \\ \vdots \\ I_{\mathrm{m}/\ell_{\mathrm{L}}} \otimes 1_{\ell_{\mathrm{L}}}^{\mathrm{T}} \end{bmatrix} \tag{14}$$

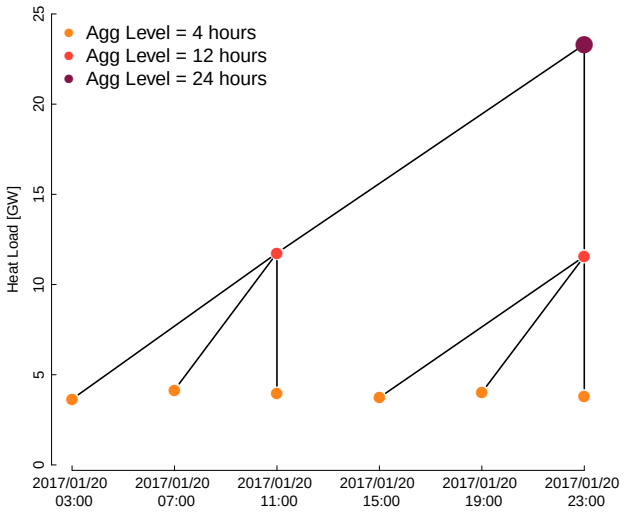

Figure 2: Example of a temporal hierarchy structure. The plot shows three different levels, $\ell = \{24, 12, 4\}$ and how they are aggregated.

where $\otimes$ denotes the Kronecker product, $I_{m/\ell}$ is an identity matrix of order m/ℓ, and 1_ℓ is an $\ell-$vector of ones. The summation matrix is used to reconcile the forecasts in order to fulfill the coherency constraints of the temporal hierarchy. The aggregation levels are a factor of m, which is the sampling frequency of the lowest level. In the example above, $\ell_1 = m$, $\ell_L = 1$, and m/ℓ is the number of observations at the aggregation level ℓ. Using Figure 2 to illustrate this, the hierarchy has aggregation levels $\ell_1 = 6$, $\ell_2 = 2$, and $\ell_3 = 1$ with $m = 6$ and the number of nodes in the structure is $n = 9$.

The base forecast models are created independently for each aggregation level. They are defined as $\hat{Y}_{t+k|t}^\ell$ for the temporal hierarchy. The number of steps ahead, k, of the prediction is the vector of $1, \ldots, m/\ell$. The forecasts are generated at time t for the hierarchy. The levels $\ell = \{24, 12, 4\}$ would result in Figure 2. Thus, we have nine base forecasts in total: six for $\ell = 4h$, two for $\ell = 12h$, and one for $\ell = 24h$ at these timestamps. The base forecasts are not necessarily coherent, i.e. the sum of the first three four-hour forecasts is not necessarily equal to the first 12-hour forecast. However, the reconciled forecasts, $\tilde{Y}_{t+k|t}^\ell$, are coherent.

The vector $\hat{Y}_{t+k|t}$ consists of all base forecasts for all aggregation levels. An example of this using the same hierarchy from Figure 2 is

$$\hat{Y}_{t+1|t} = \begin{bmatrix} \hat{Y}_{t+k|t}^{24\mathrm{h}} = [\hat{y}_{t+1|t}^{24\mathrm{h}}]^{\mathrm{T}} \\ \hat{Y}_{t+k|t}^{12\mathrm{h}} = [\hat{y}_{t+1|t}^{12\mathrm{h}},\ \hat{y}_{t+2|t}^{12\mathrm{h}}]^{\mathrm{T}} \\ \hat{Y}_{t+k|t}^{4\mathrm{h}} = [\hat{y}_{t+1|t}^{4\mathrm{h}},\ \hat{y}_{t+2|t}^{4\mathrm{h}},\ \ldots,\ \hat{y}_{t+6|t}^{4\mathrm{h}}]^{\mathrm{T}} \end{bmatrix} \tag{15}$$

which contains predictions from the three aggregation levels. The order of the vector is that the top row is the highest aggregation level and the corresponding prediction steps; in this case, the daily and the one-step ahead prediction. Then come the lower aggregation levels with their k predictions until the lowest level, which is the four-hourly aggregation with the one-step to six-steps ahead predictions. The same holds for

the reconciled forecasts, $\tilde{Y}_{t+k|t}$. In this paper, the reconciled forecasts are only computed daily at 23:00 for the next day to demonstrate the accuracy improvements. Therefore, the base forecasts generated at 23:00 are used in the reconciliation process. Consequently, $\hat{Y}_{t+k|t}$ and $\tilde{Y}_{t+k|t}$ are updated daily when the base forecasts are generated and reconciled.

There are multiple ways to make hierarchical forecasts coherent. The simplest ways are the bottom-up and top-down methods, where forecasts are made either at the lowest or highest aggregation level, respectively. The bottom-up method is defined as

$$\tilde{Y} = \boldsymbol{SG}\hat{Y}, \tag{16}$$

$$\boldsymbol{G} = \left[0_{m\times(n-m)}|I_m\right], \tag{17}$$

where $\hat{Y}$ is the base forecast made at the lowest level, $\boldsymbol{S}$ is the summation matrix matching the hierarchy structure, and $\boldsymbol{G}$ is a matrix of order $m \times n$, which extracts the m bottom-level forecasts, transforming it to the correct form for the summation matrix. The matrix $\boldsymbol{SG}$ can be seen as the projection matrix. This method does not use any information between different aggregation levels, but can be used to demonstrate the reconciliation process. The process, as shown in Eq. (16), is a linear combination of the base forecasts where $\boldsymbol{G}$ extracts and combines the base forecasts into a vector of size m of disaggregated forecasts. The summation matrix, $\boldsymbol{S}$, then creates the reconciled forecasts from the disaggregated vector. In an example using the hierarchical structure in Figure 2, the bottom-up method would only extract the four-hourly base forecasts from $\hat{Y}$. The reconciled forecasts, $\tilde{Y}$, would then be created using the summation matrix from the four-hourly forecasts. Base forecasts from the other two aggregation levels would therefore not be used, and the reconciliation process would not use information from the higher levels. However, this can be changed by modifying the projection matrix. Hence, every reconciliation process can be written using this form with different $\boldsymbol{G}$.

Van Erven and Cugliari [34] propose a game-theoretical approach to estimate the optimal reconciled forecasts in two independent steps. They prove that their reconciliation method improves any base forecast or is at least as good as the base forecast. The steps are 1) creating the best possible predictions without coherency constraints and 2) mapping them to new predictions that are coherent. They formulate the problem of computing the reconciled forecasts as a minimax optimization problem that can be solved using convex optimization. Unfortunately, the problem does not have a closed-form solution. Hyndman et al. [23] and Athanasopoulos et al. [6] propose a closed-form solution using linear regression to estimate the reconciled forecasts for structural and temporal hierarchies using generalized least squares estimation. The base forecasts are written on the regression form,

$$\hat{Y}_{t+k|t} = \boldsymbol{S}\beta(k) + \epsilon(k), \tag{18}$$

where $\boldsymbol{S}$ is the appropriate summation matrix; $\beta(k) = \mathbf{E}[Y_{\ell,t+k}|Y = y_1, \ldots, y_t]$ is the unknown conditional mean of the future values of the most disaggregated observed series, i.e. the reconciled forecasts; and $\epsilon(k)$ represents the error between the base forecasts and their expected value, i.e. the coherency error $\hat{Y} - \tilde{Y}$. The error $\epsilon(k)$ is assumed to have zero mean and covariance matrix, Σ. Hence, the generalized least squares estimation of $\beta(k)$ in Eq. (18). If Σ is assumed to be known and the base forecasts are unbiased, the reconciled forecasts can be estimated by

$$\tilde{y} = S(S^{\mathrm{T}}\Sigma^{-1}S)^{-1}S^{\mathrm{T}}\Sigma^{-1}\hat{y}, \tag{19}$$

where the matrix $\boldsymbol{G} = (S^{\mathrm{T}}\Sigma^{-1}S)^{-1}S^{\mathrm{T}}\Sigma^{-1}$.

It has been shown that Σ is never known nor identifiable [24]. Therefore, it is proposed to estimate it from the in-sample base forecast errors while imposing additional structure on the matrix. Numerous different methods have been proposed to use the $\hat{\Sigma}$ estimator instead of Σ. Hyndman et al. [23] suggest using ordinary least squares (OLS), i.e. an identity matrix, which results in the equal weighting of all base forecasts. In Hyndman et al. [25], weighted least squares is used as an estimator by using the variance of the base forecasts to create the weights. In Athanasopoulos et al. [6], three different weighted least squares estimators, *Hierarchy Variance Scaling*, *Series Variance Scaling*, and *Structural Scaling* are presented. These are naive methods for creating the estimator as the off-diagonal entries are zero, thus dismissing any cross-correlation between levels. Wickramasuriya et al. [24] propose the minimum trace (MinT) estimator, which uses all information between levels by using a full covariance matrix of the base forecast errors. Nystrup et al. [7] present four covariance estimators that account for the cross-correlation between levels.

In this study, the full covariance matrix will be used to generate the reconciled forecasts. The matrix shares information within levels and between levels. Recent results have shown that using the auto- and cross-covariance leads to significant improvements in forecast accuracy [7, 24]. The covariance matrix is estimated from the past base forecast errors,

$$\hat{\Sigma}_t = \frac{1}{t}\sum_{i=1}^{\mathrm{t}} e_i e_i^{\mathrm{T}}, \tag{20}$$

where e is the vector of the base forecast error for each aggregation level and step-ahead prediction at each time step. It is assumed that the forecast errors are unbiased, i.e., $\mathbf{E}[e_i] = 0$.

It is known that heat load follows seasonal patterns as the ambient temperature changes and space heating increases or decreases. The heat load is also more stationary over the summer period as there is no space heating. Therefore, two of the covariance estimators proposed include a forgetting factor. This allows them to forget past errors as the dynamics of the heat load changes between seasons and throughout the years. The third estimator is the full covariance matrix that uses all of the past available base forecast errors.

4.1. Method 1: Expanding Window Covariance Matrix

The first method proposed is the approach where the covariance matrix estimator is based on all past available errors. The

estimator includes both the auto- and cross-covariance between all aggregation levels. Nystrup et al. [7] applied this estimator of the covariance matrix and concluded that using all information between levels improves the forecast accuracy on all aggregation levels.

This estimator will be recursively updated every day when new observations are available and the error of the forecast can be computed,

$$\hat{\Sigma}_t = \frac{1}{t} e_t e_t^{\mathrm{T}} + \frac{t-1}{t} \hat{\Sigma}_{t-1}. \tag{21}$$

Initially, this will update the covariance matrix quickly as new information becomes available, and the updating will be slowed down later.

4.2. Method 2: Rolling Window Covariance Matrix

The second proposal is to estimate the covariance matrix using a rolling window. The estimation is performed on a fixed window of past errors where they all have equal weight. The rolling window adds new errors to the estimation when they become available and removes the oldest errors from the window. The length of the window or the memory is optimized by finding which memory yields the highest accuracy improvements of the operational heat load forecasts over the in-sample period.

The rolling window estimator is given by

$$\hat{\Sigma}_t = \frac{1}{t-j} \sum_{i=t-j}^{t} e_i e_i^{\mathrm{T}}, \tag{22}$$

where the index j is the length of the rolling window, i.e., the memory.

4.3. Method 3: Exponential Smoothing Covariance Matrix

The exponential smoothing covariance matrix is a recursive and adaptive estimator. It is updated when a new observation is available, and every observation is weighted differently using a forgetting factor, λ. The weights are normalized to ensure that the sum of the weights is one [31]. The exponential smoothing estimator is given by

$$\hat{\Sigma}_t = \lambda \hat{\Sigma}_{t-1} + (1-\lambda) e_t e_t^{\mathrm{T}}. \tag{23}$$

When $\lambda \to 1$, the covariance is not updated, and when $\lambda \to 0$, the covariance is highly influenced by the newest observations. The initial covariance matrix for $\hat{\Sigma}_0$ is usually initialized by computing the sample covariance over an initialization period in an offline setting. The forgetting factor and the initial period for the covariance are optimized to maximize the improvement in the accuracy of the operational heat load forecast as for the rolling window method.

4.4. Shrinkage

The fourth proposal is to shrink the estimated covariance matrices. Shrinking the covariance estimate has been shown to improve forecast accuracy considerably [7, 24]. The shrinkage method used is the scale and location invariant shrinkage proposed by Ledoit and Wolf [29]:

$$\hat{\Sigma}_t^* = \lambda_{\text{shrink}}^* \hat{\Sigma}_t^{\mathrm{d}} + (1 - \lambda_{\text{shrink}}^*) \hat{\Sigma}_t, \tag{24}$$

where $\hat{\Sigma}_t^{\mathrm{d}}$ is the diagonal entries from the covariance matrix $\hat{\Sigma}_t$. The shrinkage target is the diagonal variance of the levels since the off-diagonal elements of the covariance matrix are shrunk towards zero as $\lambda_{\text{shrink}}^*$ increases. Ledoit and Wolf [29] derived a closed-form solution for the optimal value of $\lambda_{\text{shrink}}^*$ by minimizing the mean squared error. This shrinkage estimator of the covariance matrix is ideal for a small number of data points with a large number of parameters. If the variance is assumed to be constant, then the optimal shrinkage parameter is given by

$$\lambda_{\text{shrink}}^* = \frac{\sum_{\mathrm{i} \neq \mathrm{j}} \widehat{\mathrm{Var}}(\hat{\sigma}_{ij})}{\sum_{\mathrm{i} \neq \mathrm{j}} \hat{\sigma}_{ij}^2}, \tag{25}$$

where $\hat{\sigma}_{ij}$ is the ijth element of the covariance matrix from the base forecast errors. The variance of the estimated covariance, $\hat{\sigma}_{ij}$, from the covariance matrix, $\hat{\Sigma}$, is computed as shown in Appendix A in Schäfer and Strimmer [35].

We extend the work by Ledoit and Wolf [29] and Schäfer and Strimmer [35] such that the variance of the empirical covariance matrix is estimated recursively. A recursive estimate of the variance of the covariance matrix is found to be

$$\widehat{\mathrm{Var}}(\hat{\Sigma}_{\mathrm{t}}) = \lambda(1-\lambda)^2 (e_t^2 (e_t^2)^{\mathrm{T}} - \hat{\Sigma}_t^2) + \lambda^2 \widehat{\mathrm{Var}}(\hat{\Sigma}_{\mathrm{t-1}}). \tag{26}$$

The shrinkage parameter can then be recursively updated at each time step using

$$\lambda_{\text{shrink},t}^* = \frac{\sum_{\mathrm{i} \neq \mathrm{j}} \widehat{\mathrm{Var}}(\hat{\Sigma}_t)_{ij}}{\sum_{\mathrm{i} \neq \mathrm{j}} (\hat{\Sigma}_t)_{ij}^2}. \tag{27}$$

The proof is given in Appendix A and the algorithmic solution is used in a recursive scheme. The recursive shrinkage will be imposed only on the exponential smoothing estimator. The other two estimators will be shrunk using Eq. (25). Hence, the estimate will be updated at each time step and shrunk before being used to reconcile the forecasts.

4.5. Optimization of Memory Parameter

For two of the covariance estimation methods, the rolling window and exponential smoothing, the memory parameter needs to be optimized to give the largest improvements. The expanding window estimator does not have a memory parameter. The initialization period for the sample covariance matrix also needs to be determined.

The initialization period and the memory parameter will be optimized to maximize the accuracy improvements of the operational heat load forecasts. The metric used to measure the accuracy improvements is the relative root mean squared error (RRMSE) that is frequently used to compare accuracy improvements of reconciled forecasts [6, 7]. It was recommended by Hyndman and Koehler [36] due to the interpretability of the relative measure when considering accuracy improvements. The RRMSE is defined as

$$\mathrm{RRMSE} = \frac{\mathrm{RMSE}}{\mathrm{RMSE}^{\mathrm{base}}} - 1, \tag{28}$$

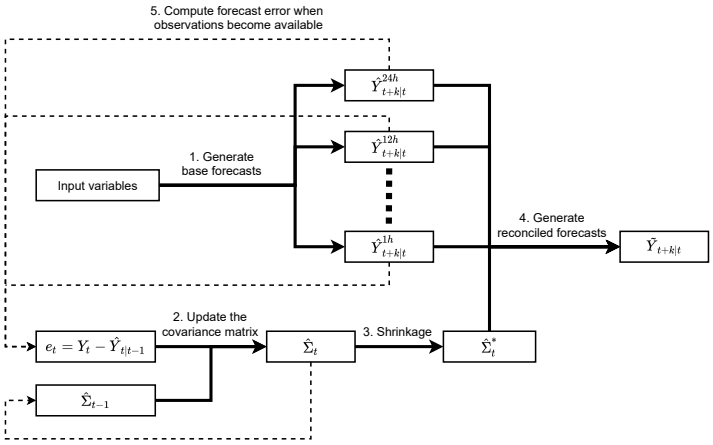

Figure 3: Overview of the workflow of the proposed method.

where negative values describe a percentage improvement of the reconciled forecasts over the base forecasts. In this paper, the RRMSE is only computed for each aggregation level. The RMSE is computed as the average error for all prediction steps for the corresponding aggregation level:

$$\text{RMSE} = \sqrt{\frac{1}{T}\sum_{t=1}^{T}\frac{1}{m/\ell}\sum_{k=1}^{m/\ell}(Y_{t+k} - \hat{Y}_{t+k|t})^2} \tag{29}$$

where T is the number of days forecasts are generated, m is the number of forecasts at the bottom level, ℓ is the aggregation level, k is the prediction horizon, Y is the heat load observation and, $\hat{Y}$ is the heat load forecast.

For the rolling window, the memory is defined as the number of past days. These past days will be used to estimate the covariance at each time step using the error from those days. The window for the rolling window has a fixed length where the errors have equal weights. The memory for the exponential smoothing is computed from the forgetting factor λ as

$$N_{\text{eff}} = \frac{1}{1-\lambda}. \tag{30}$$

where N_{eff} is the effective memory for the exponential smoothing. The memory for the exponential smoothing is not a fixed-length window as for the rolling window due to exponential decay.

The initial sample covariance matrix needs to be invertible such that the reconciled forecasts can be estimated as shown in Eq. (19). Thus, an optimal initial period needs to be defined to make this feasible. The shrinkage will also ensure that the covariance is invertible when $\lambda^*_{\text{shrink}} > 0$ [29].

4.6. Workflow of Forecast Reconcilation

The workflow of the method proposed in this paper is illustrated in Figure 3 and described in the following. Before the workflow is started, the base forecast models and the covariance matrix are initialized.

1. Generate base forecasts. The base forecasts in this paper are generated from the RLS regression models as introduced in Section 3 using the corresponding input variables for each aggregation level. A base forecast vector at time t, $\hat{Y}_{t+1|t}$, is created which includes every aggregation level forecast with the corresponding forecast horizon as shown in Eq. (15). When the commercial state-of-the-art forecasts are used, they are combined with the higher aggregation base forecasts, and the authors' hourly base forecasts are removed before the reconciliation process.

2. Update the covariance matrix. A covariance matrix is needed for the reconciliation process to combine the base forecasts. In this paper, the covariance matrix is updated recursively using one of the three methods proposed. As the covariance matrix is recursively updated, the previous covariance matrix needs to be kept for the update when a new observation becomes available.

3. Shrinkage. The covariance matrix is shrunk before the reconciliation process using either of the two shrinkage methods suggested in Section 4.4. One of the methods recursively updates the shrinkage parameter using Eq. (26) and Eq. (27) while the other estimates it directly from the current covariance matrix using Eq. (24).

4. Generate reconciled forecasts. When the base forecasts and the covariance matrix are ready, the reconciled forecasts are computed using Eq. (19).

5. Compute forecast error when observations become available. When new observations become available, the base forecast errors, $e_t = Y_t - \hat{Y}_{t|t-1}$, are computed and the covariance matrix is updated. The updated covariance matrix is then used to generate new reconciled forecasts.

5. Results

In this section, the results from the proposed method, using three different covariance estimators, are presented and evaluated. The results will be discussed in terms of improving the accuracy of the operational forecast using the temporal hierarchy. The benefits of having a state-of-the-art forecast in the proposed method will be discussed by comparing the accuracy result with the result of using the authors' simple model proposed in this article.

The optimization of the hyperparameters for the covariance estimators is presented in Section 5.1. The estimators are then used with the selected hyperparameters to create reconciled forecasts from the base forecast for the *Varmelast* case study. The results for the reconciled forecasts are presented in Section 5.2.

5.1. Hyperparameter Optimization

As described above, the memory and initialization period for the covariance estimators need to be optimized. They are determined by minimizing the RRMSE. The year 2016 will be used as initialization period and years 2017 and 2018 as in-sample training period to estimate the optimal initial period and memory. Data from year 2016 is only available from 17 January until the end of the year, as the first 16 days are used for initialization of the base forecasts.

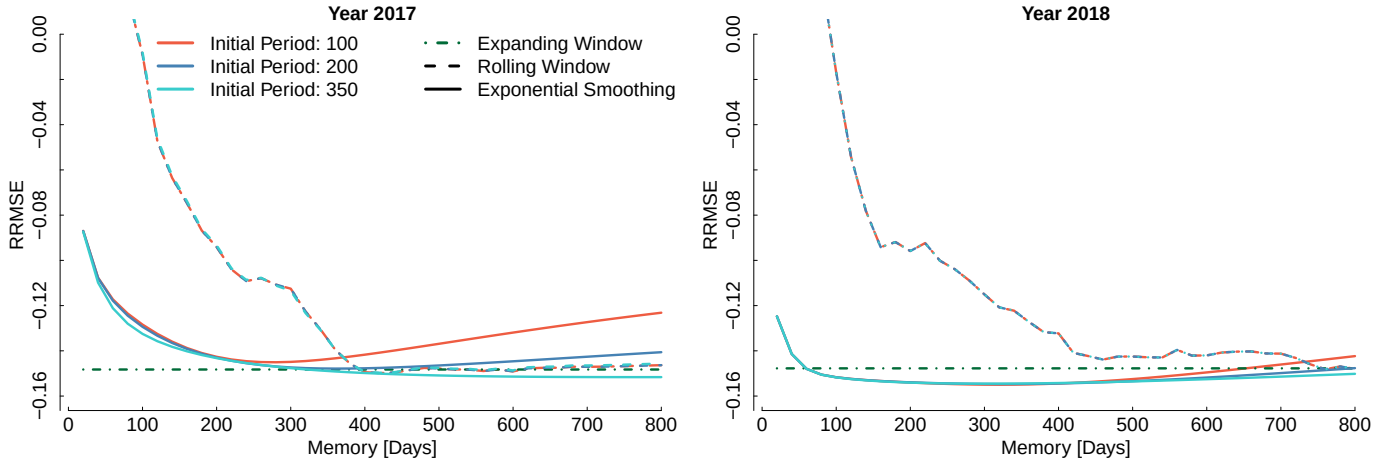

Figure 4: Optimizing the memory for the Exponential Smoothing and Rolling Window estimators using three different initial periods for the sample covariance for two different in-sample years. The optimization is done on the hourly aggregation level.

Figure 4 shows the RRMSE of the operational heat load forecasts for different memories and initial periods for all three methods in the two in-sample years. The rolling window and exponential smoothing methods test three different initialization periods of 100, 200, and 350 days. The initialization period does not seem to affect the results for the rolling window, as it results in the same accuracy improvements for the three initial periods. However, the initialization period affects the exponential smoothing, as the plot for 2017 demonstrates. When comparing the result between the two years, it can be seen that the results are sensitive to the initialization period and the memory length. The 2017 test year is closer to the initial sample covariance, $\hat{\Sigma}_0$, which results in three different curves based on the initial period while the curves for 2018 are more similar. When the memory is longer than what is optimal, the improvements seem to decrease again. The rate of deterioration after the optimal point depends on the initial period for which the sample covariance is estimated. Figure 5 demonstrates the effects between the memory and the initialization period for the exponential smoothing method. The plot shows the influence of the older observations, depending on the memory. The longer the memory is, the more influence older observations have. In other words, the initial sample covariance matrix will influence the current estimator, and the magnitude of the influence depends on the memory. Hence, when the memory is long and the initial period is 100 days, then the initial sample covariance has a considerable impact on the covariance matrix many months later. For example, the influence of the initial sample covariance on the current covariance after one year is close to 0.16 with a memory of 200 days, as the green curve in Figure 5 shows.

The rolling window does not outperform the other two methods as the memory increases. Figure 4 shows that the rolling window converges to the same accuracy improvements as the expanding window method with increasing memory size. This is not the case for the exponential smoothing method, as it seems to have an optimal point where it yields the maximum improvement. Based on these results, the initial period was set

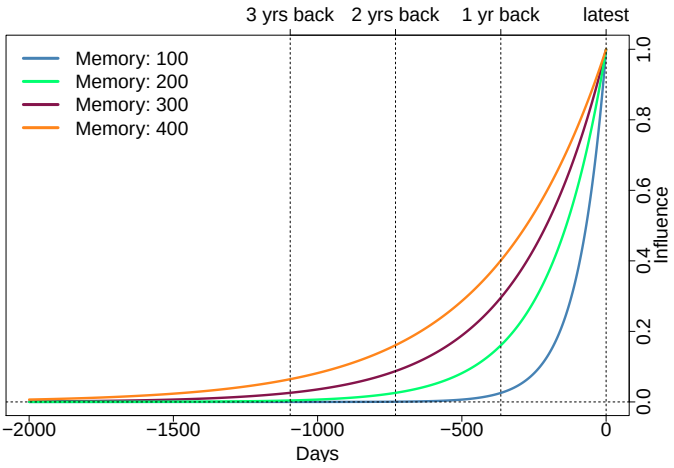

Figure 5: Influence of older observation for different memory lengths for exponential smoothing. The weights are normalized by $1 - \lambda$.

to the first 350 days, i.e. the entire year of 2016 after removing the first 16 days for the base forecast initialization to compute the initial sample covariance. The memory for the exponential smoothing was set to 365 days, i.e. a full year, as these parameters are hyperparameters and have a flat curve around the optimal point. The memory for the rolling window is the same as for exponential smoothing.

5.2. Empirical Results

The results for the three years, 2017, 2018, and 2019 are presented in Table 2 along with the result for all three years combined. The 2017 and 2018 data were used to optimize the hyperparameters as shown in Section 5.1, while out-of-sample data from 2019 is used to validate the results. The hourly results in the table are for the commercial state-of-the-art forecasts. The hourly improvements are highlighted in Table 2 to emphasize that they are of most importance in this case. They are the operational heat load forecasts used to operate the district heating. The higher aggregation levels are included in the result for completeness. The RMSE of the base forecasts is shown in the first

Table 2: RMSE for the base forecasts and RRMSE for the reconciled forecasts for daily heat load in the Greater Copenhagen area. Results are shown for three different years and the whole period from 2017 to 2019. The years 2017 and 2018 were used to optimize the hyperparameters while 2019 is the out-of-sample year. The result from the hourly forecasts is highlighted using grey-background to emphasize that the objective is to improve the hourly forecasts.

	2017				2018			
	Base RMSE	Expanding Window	Rolling Window	Exp. Smoothing	Base RMSE	Expanding Window	Rolling Window	Exp. Smoothing
Daily	0.5585	-17.88	-17.7	**-18.1**	0.6218	**-23.02**	-19.56	-22.94
Twelve-hourly	0.3151	-17	-16.72	**-17.19**	0.3766	-25	-22.25	**-25.01**
Eight-hourly	0.3333	-39.78	-39.61	**-39.9**	0.3508	-40.24	-38.56	**-40.48**
Six-hourly	0.2628	-41.07	-40.77	**-41.17**	0.2876	-42.16	-40.21	**-42.23**
Four-hourly	0.1715	-35.24	-34.86	**-35.34**	0.1725	-31.81	-30.16	**-32.32**
Three-hourly	0.1273	-31.98	-31.62	**-32.09**	0.1315	-30.34	-28.62	**-30.75**
Two-hourly	0.0846	-29.07	-28.64	**-29.16**	0.088	-27.99	-26.33	**-28.51**
Hourly	0.0372	-14.83	-14.26	**-14.92**	0.0389	-14.77	-12.91	**-15.44**
	2019				2017–2019			
	Base RMSE	Expanding Window	Rolling Window	Exp. Smoothing	Base RMSE	Expanding Window	Rolling Window	Exp. Smoothing
Daily	0.6058	-29.6	-30.21	**-30.41**	0.596	-23.75	-22.49	**-23.93**
Twelve-hourly	0.3602	-28.91	-28.54	**-29.1**	0.3516	-24.08	-22.83	**-24.2**
Eight-hourly	0.3759	-49.82	-49.36	**-50.01**	0.3538	-43.51	-42.72	**-43.69**
Six-hourly	0.3106	-49.67	-49.36	**-49.87**	0.2876	-44.64	-43.75	**-44.76**
Four-hourly	0.1852	-40.74	-40.22	**-41.05**	0.1765	-36.06	-35.19	**-36.37**
Three-hourly	0.141	-36.36	-35.56	**-36.54**	0.1334	-33.03	-32.05	**-33.26**
Two-hourly	0.0924	-32.94	-32.02	**-33.13**	0.0884	-30.09	-29.07	**-30.36**
Hourly	0.0389	-14.67	-13.3	**-14.83**	0.383	-14.75	-13.46	**-15.07**

Bold values represent the best performance for each aggregation.

column for each year followed by the RRMSE for the three proposed covariance estimators: expanding window, rolling window, and exponential smoothing. The table shows that, for all aggregation levels and for both the in-sample and out-of-sample years, the reconciled forecasts using any of the estimation methods improve the base forecasts. The improvement of the operational heat load forecasts is consistently around 15% in all years. The improvement of the other levels varies more between years and the improvements are even higher in the out-of-sample year. The reconciliation approach performs at least as well out-of-sample as in-sample, which highlights its robustness. Each level has an improvement ranging from 10% to 50%, with the six- and eight-hourly aggregation levels having the largest improvements. These large improvements compared to the base forecasts could indicate that the model is not able to capture the seasonal behavior in the data.

Based on the results in Table 2, we can summarize the following for three estimators:

Expanding Window. The covariance matrix grows over time and always improves the base forecasts at all aggregation levels. The expanding covariance matrix does not require optimization of the hyperparameters, which makes it favorable when few data points are available.

Rolling Window. The covariance matrix with fixed window size and equal weighting performs the worst out of the three methods. Optimization of the hyperparameters showed that the rolling window will converge to the same result as the expanding window with increasing memory, but it never outperforms the expanding window method.

Exponential Smoothing. The adaptive and recursively updated covariance that exponentially weights past errors yields the largest accuracy improvements. Having an adaptive and recursive covariance matrix results in the best performing estimator, as it always gives the largest accuracy improvements for the operational heat load forecasts. It also gives the largest improvements on the other levels in nearly all cases. Across all three years, it outperforms on all aggregation levels.

Figure 6 shows the heat load observations, base forecasts, and reconciled forecasts using the exponential smoothing estimator for four of the aggregation levels. The grey dashed vertical lines represent the hour 23:00 every day when the models are updated and forecasts are generated for the next day. The plots show that the reconciled forecasts are frequently better than the base forecasts. There are, however, a few time periods where the reconciled forecasts are worse. Yet, overall the reconciled forecasts result in higher accuracy, as the RMSE in the lowest plot shows for the operational level. It is difficult to understand where the improvements at the lowest aggregation level come from. The improvements could come from truncation in the AR process and different memory in the forecasting models for the higher aggregation levels. The higher the aggregation level, the shorter the memory, and therefore it can react more rapidly. For example, rapid changes in the heat load or the outdoor temperature are shared with the lower levels.

The eight-hourly plot shows that the base forecast is only able to capture the trend, not the daily seasonality. However, the reconciled forecasts capture both the trend and seasonality. The reconciled forecasts are able to share information with the lower aggregation levels, where the seasonality is more dominant. The forecasts for the higher aggregation levels benefit from sharing information with the lower levels when they are not able to capture the seasonality. The lower aggregation levels benefit from sharing information with higher levels as they can capture the trend more accurately. Thus, temporal hierarchies using the covariance matrix with auto- and cross-covariance, which shares information both within each level and between levels in the reconciliation process, significantly improves the accuracy of all forecasts, including the operational heat load forecasts.

Figure 7 shows four different error measures for each prediction horizon for the hourly day-ahead forecasts. The top plot shows the RMSE in 2019 for each prediction step. Reconciling

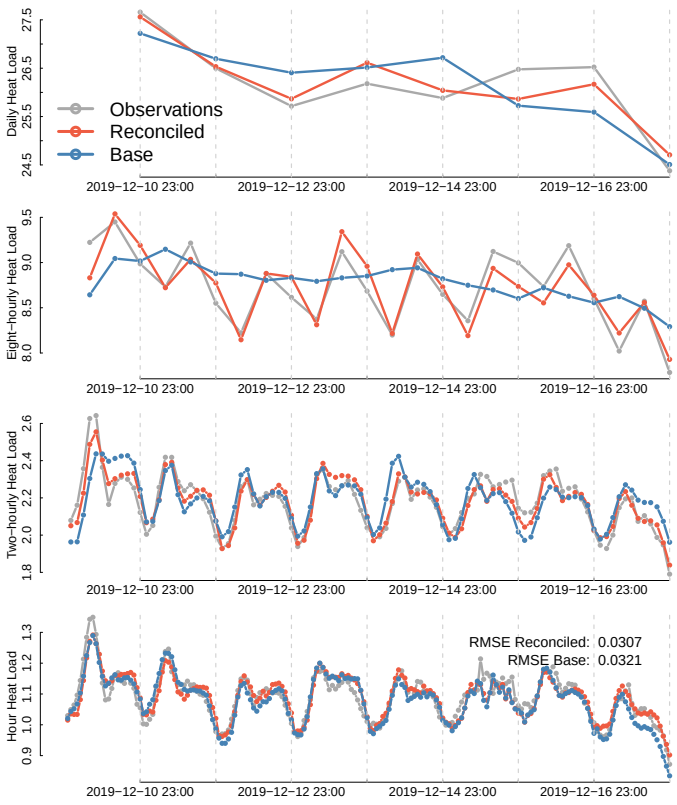

Figure 6: Plots of the observations, base and reconciled forecasts showing the improvements of the reconciled forecasts.

the forecasts using the proposed exponential smoothing covariance improves accuracy for each step. The second plot shows the difference in RMSE between the base and reconciled forecasts for each prediction horizon. The plot shows that the reconciled forecasts improve the base forecast for all horizons. In general, the differences are quite similar in size except for the peak around $k = 6$. The third plot shows the *root mean scaled squared error* (RMSSE),

$$\text{RMSSE} = \frac{\text{RMSE}}{\sqrt{\frac{1}{N-1}\sum_{t=2}^{N}(y_t - y_{t-1})^2}}, \tag{31}$$

for each prediction step. Similar results are seen when comparing the RMSSE and RMSE plots. The reconciled forecasts improve accuracy for all prediction horizons, as shown in the bottom plot, where the RRMSE is plotted for each step k. The plot also demonstrates that the accuracy improvements using the proposed method are consistent over the three years. Generally, the error should increase with the forecast horizon. This is not always the case, as the two top plots show some variation over the horizons due to the diurnal variation of heat load. This is can be seen in the RMSE and RMSSE plots, where the shape of the error follows the daily heat load profile, i.e. the error peaks in the mornings and afternoons when the heat load starts to increase. The horizon accuracy from the NWPs used as inputs to the models could also be a contributing factor here, combined with the daily profile. Comparing the RRMSE of the $k = 1$ horizon between the years shows that accuracy is improved the least in 2019. For the first six steps ahead, the RRMSE varies between the years, while the RRMSE for the longer horizons is more consistent across the years. A possible explanation for this is that the first six hourly forecasts are the heat load between 23:00 and five in the morning where the heat load is relatively stationary.

The heating season can be split into three different periods. The first period is the first four months of the year, when the heat load is quite high and the transition from winter to summer occurs. The second period is the summer months from May through the end of August, when the heating demand in the Greater Copenhagen area is low. The third period is the last four months of the year, when the heat load starts to increase again as outdoor temperature decreases. Figure 8 shows the cumulative squared error of the operational heat load forecasts for the year 2019 split into these three periods. The three plots show that the temporal hierarchy improves the accuracy in all three periods. Hence, the methods can perform even though the heat demand changes concurrently with the outdoor temperature. The largest accuracy improvements occur around the times when the heat load changes the most. For example, after October the heat load starts to increase as shown in Figure 1, and that is also when a significant change in accuracy improvement happens. The same occurs in the transition period in the spring. In the summer period, when the heat load is stationary, there are only small improvements using the proposed method. All of the three covariance estimators outperform the commercial state-of-the-art base forecast in all three periods, with the exponential smoothing method performing the best.

To demonstrate how powerful temporal hierarchies are, we compare reconciled forecasts based on our own simple base forecasts from Section 3 for all aggregation levels to the state-of-the-art base forecast from the commercial provider. The simple model is the hourly model from Table 1. Note that this model uses accurate weather forecasts provided by the forecast provider as input. Figure 9 shows that the state-of-the-art base forecast is significantly better than our own hourly base forecast. Yet, when applying the proposed adaptive and recursive covariance estimator to our base forecasts, the resulting reconciled forecasts outperform the commercial base forecasts slightly. Thus, creating forecasts from simple models at multiple aggregation levels can compete with state-of-the-art operational heat load forecasts in terms of accuracy. As expected, the reconciled forecasts based on the commercial state-of-the-art base forecasts are even more accurate than those based entirely on our own base forecasts. In other words, using more accurate base forecasts results in more accurate reconciled forecasts.

6. Discussion

In this work, we have suggested recursive and adaptive methods to update the covariance matrix used in forecast reconciliation, and we have shown that the suggested methods lead to significant improvements of the accuracy of operational heat load forecasts. The covariance estimator in the reconciliation process uses the auto- and cross-covariance to connect the

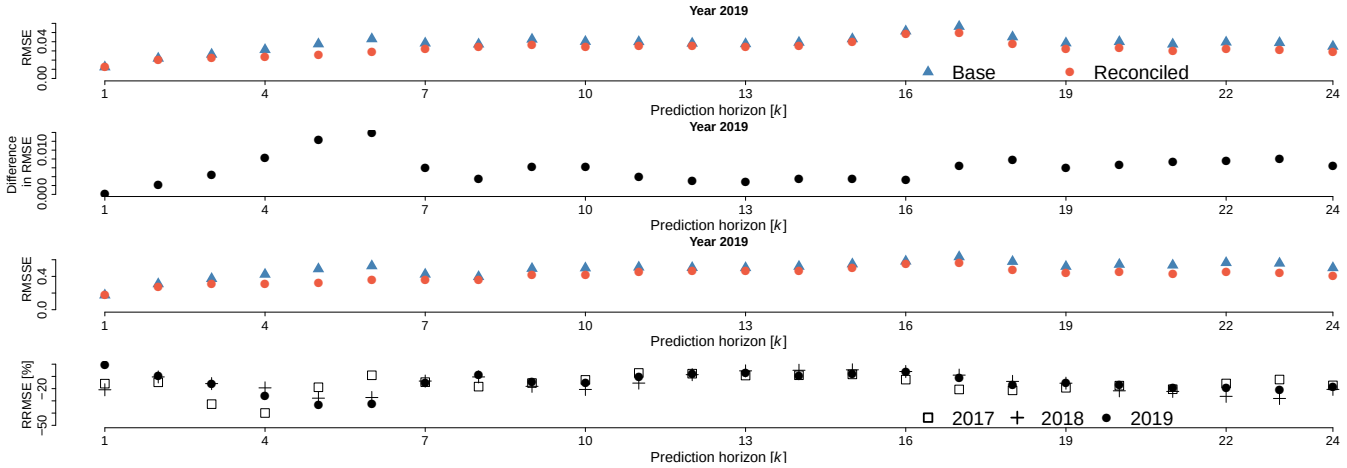

Figure 7: Error measures for the different prediction horizons from commercial hourly aggregation forecasts are illustrated. The plots show four different measures, the RMSE, the difference in RMSE between base and reconciled forecasts, RMSSE, and RRMSE for each prediction horizon. The reconciled forecast improves the accuracy for all horizons in all three years.

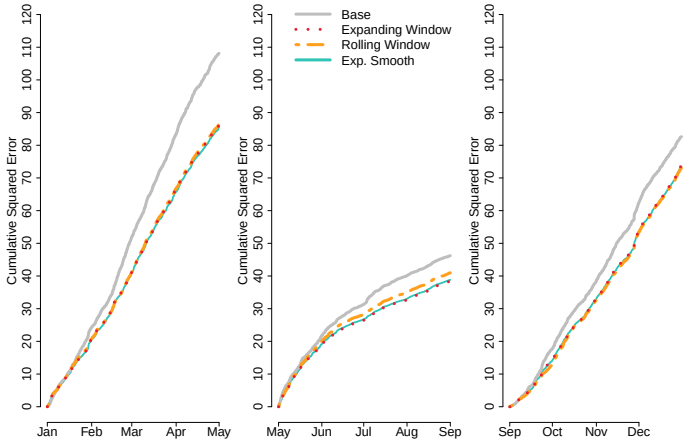

Figure 8: Cumulative squared error of the base forecasts and the reconciled forecasts from the three covariance estimators. The three plots divide the year 2019 into three different periods with respect to heating demand and transition period.

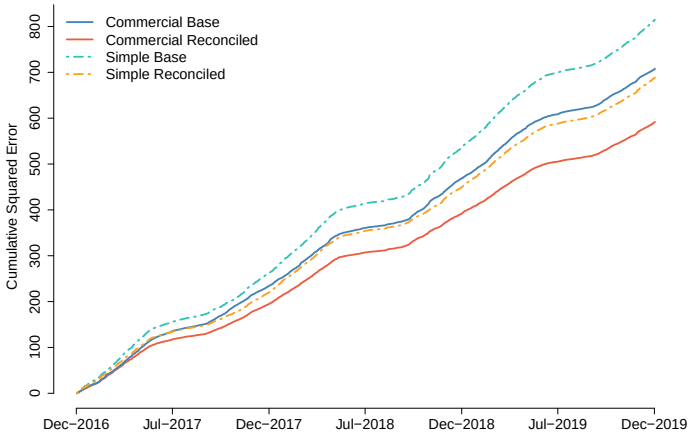

Figure 9: Cumulative squared error comparing state-of-the-art hourly forecasts from the commercial provider and forecasts created using the authors' simple model for the operational level. More accurate base forecasts lead to more accurate reconciled forecasts.

base forecasts from all aggregation levels, and this improves the forecasts by sharing information within and between the hierarchy levels. The covariance matrix was made recursive and adaptive as the heat load is a time-varying process due to changes in the outdoor temperature. The results show that the covariance method needs an entire year of in-sample data for initialization and a long memory to yield the largest improvements. However, it only needs to store the previous covariance matrix and current errors for the reconciliation procedure.

We focused solely on improving the base forecasts for the next day that were generated at 23:00 every night. The proposed method could easily be applied to recursively reconcile new base forecasts generated every hour for the next 24 hours using the same hierarchy. If only interested in one-hour ahead forecasts a different temporal hierarchy would be more beneficial. The strength of the proposed method is its adaptive property for the covariance matrix and the proposed recursive estimation of the shrinkage parameter. The weights for the reconciliation process and the shrinkage parameter are able to change over time when new information becomes available. This is highly relevant for systems that are governed by non-stationary processes.

Regression models were used to generate base forecasts for all aggregation levels above the operational hourly level for district heating. However, any forecast model could have been used to generate the base forecasts. Information is shared between the base forecasts through the covariance matrix, and different levels contribute with different information, e.g. the trend from higher levels and seasonality from the lower levels, to achieve accuracy improvements. Therefore, the temporal hierarchy structure needs to be chosen based on the application and the dynamics of the system. Commercial base forecasts were used to demonstrate the ability of the proposed method to achieve large accuracy improvements, even when combining

state-of-the-art forecasts with simple forecasts on higher levels.

Similar findings of accuracy improvements when forecasting for different temporal levels and using auto- and cross-covariance matrices to compute the reconciled forecasts are reported in other applications. Wickramasuriya et al. [24] propose using the full covariance matrix estimated from the base forecast errors and the same shrinkage as used in this paper in the forecast reconciliation process. They benchmark different structures of the covariance matrix to demonstrate the importance of using all of the information in the covariance matrix by sharing it within and between aggregation levels. Nystrup et al. [7] has suggested various structures and regularization for the auto- and cross-correlation to maximize the accuracy improvements when reconciling forecasts in a temporal hierarchy.

Improving commercial state-of-the-art operational forecasts will result in more reliable operational planning of the district heating system. It will help to find the optimal production schedule of the units and improve control of the heat delivered to the network. Blanco et al. [17] propose a method that optimizes the production and creates bids for the day-ahead and balancing electricity markets using stochastic programming that needs an accurate heat load forecast. Thus, by improving operational heat load forecasts, the system can be optimized more accurately and the overall system cost can be reduced by profiting from the electricity market.

7. Conclusion

We demonstrate how temporal hierarchies can be used to improve the accuracy of operational heat load forecasts. The heat load forecasts are generated each day at 23:00 for the next 24 hours. State-of-the-art heat load forecasts from a commercial energy forecasting provider are included to see if accuracy improvements are possible.

The empirical results are based on four years of heat load data. The first year is used to initialize the method and the next two years to find the optimal hyperparameters for the covariance estimators. The last year is used to show the accuracy improvement on out-of-sample data.

Three covariance estimators are presented based on the base forecast errors from the different aggregation levels. The first estimator is recursively updated and uses an expanding window, meaning that all past errors are used with the same weight. The second estimator uses a rolling window with a fixed number of past errors having the same weight. The third estimator uses exponential smoothing to put weights on past errors. This is the estimator that we propose for the reconciliation process to improve the accuracy of heat load forecasts. The covariance matrix estimate is shrunk before it is used to compute the reconciled forecasts. The shrinkage parameter is recursively updated using the method proposed.

The hyperparameters for the estimators are selected by minimizing the RRMSE over two years. The three different estimators are tested and compared using RRMSE. The exponential smoothing estimator yields the largest accuracy improvements in the reconciliation process in our case study. This estimator improves the accuracy of the commercial state-of-the-art operational heat load forecasts by, on average, 15% over the three years. The proposed method improves the base forecasts at all aggregation levels and for all forecast horizons considered. When using only simple base forecasts as inputs at all aggregation levels, the reconciled forecasts have similar accuracy to the commercial state-of-the-art base forecast. However, using the state-of-the-art base forecast as input, the reconciled operational heat load forecast becomes significantly better. Hence, it is necessary to include a state-of-the-art forecast to get maximum accuracy.

There are several possible directions for future research on improving heat load forecast accuracy using temporal hierarchies. One possibility is to investigate the optimal number of aggregation levels in greater detail, including an analysis of how longer-term forecasts (e.g., one-week ahead) would perform. Another option is to improve the covariance estimator used in the reconciliation process, e.g. through dimensionality reduction or by imposing additional structure on it. Finally, it is possible that the accuracy of the reconciled operational heat load forecasts could be further improved by improving the accuracy of the base forecasts.

8. Acknowledgement

This work is funded by Innovation Fund Denmark through the projects Heat 4.0 (8090-00046B), CITIES (1305-00027B), Flexible Energy Denmark (8090-00069B), the IDASC (18012745) project funded by Region H, and the Norwegian FME-ZEN project funded by the ZEN partners and the Research Council of Norway.

We thank the energy forecasting provider ENFOR, especially Henrik Aalborg Nielsen and Torben Skov Nielsen, for their support and for supplying the forecasts & information. We would also like to thank Varmelast for making their heat load data available.

Appendix A. Recursive Shrinkage

Computing the optimal estimated shrinkage intensity, $\hat{\lambda}^*_{\text{shrink}}$, as Schäfer and Strimmer [35] demonstrate in Appendix A, while doing recursive and adaptive estimation of the covariance matrix is infeasible. Here we show how to recursively estimate the empirical variance and covariance of the *individual entries* of the recursive and adaptive covariance matrix Σ_t.

The recursive and adaptive covariance matrix is defined as

$$\hat{\Sigma}_t = (1-\lambda)\sum_{i=0}^{t} \lambda^i e_{t-i} e_{t-i}^{\mathrm{T}}, \tag{A.1}$$

where λ is the forgetting factor and e_{t-i} is the base forecast error vector at time $t-i$. We will assume that the error is normally distributed with zero mean and covariance, V. Then, we can assume that Eq. (A.1) comes from the Wishart distribution. Therefore, the matrix from the Wishart distribution computes the variance of the individual entries.

The variance and covariance of the individual entries of Σ_t are computed by taking the variance of each term in Eq. (A.1),

$$X_{t-i,kj} = e_{t-i,k} e_{t-i,j}^{\mathrm{T}}, \tag{A.2}$$

$$\mathrm{Var}(\Sigma_{t,kj}) = \mathrm{Var}\left[(1-\lambda)\sum_{i=0}^{t}\lambda^{i} X_{t-i,kj}\right], \tag{A.3}$$

assuming independence,

$$\mathrm{Var}(\Sigma_{t,kj}) = (1-\lambda)^2 \sum_{i=0}^{t}\lambda^{2\mathrm{i}}\mathrm{Var}[X_{t-i,kj}], \tag{A.4}$$

$$\mathrm{Var}(X_{t-i,kj}) = \mathbb{E}[(X_{t-i,kj} - \mathbb{E}[X_{t-i,kj}])^2], \tag{A.5}$$

$$= \mathbb{E}[X_{t,kj}^2 - 2X_{t-i,kj}\mathbb{E}[X_{t-i,kj}] + \mathbb{E}[X_{t-i,kj}]^2], \tag{A.6}$$

$$= \mathbb{E}[X_{t,kj}^2] - 2\mathbb{E}[X_{t-i,kj}]\mathbb{E}[X_{t-i,kj}] + \mathbb{E}[X_{t-i,kj}]^2, \tag{A.7}$$

$$= \mathbb{E}[X_{t,kj}^2] - \mathbb{E}[X_{t-i,kj}]^2, \tag{A.8}$$

$$= \mathbb{E}[X_{t,kj}^2] - \sigma_{t,kj}^2, \tag{A.9}$$

$$\mathrm{Var}(\Sigma_{t,kj}) = (1-\lambda)^2 \sum_{i=0}^{t}\lambda^{2i}(\mathbb{E}[X_{t,kj}^2] - \sigma_{t,kj}^2). \tag{A.10}$$

With this we have an analytical solution to estimate the variance and covariance of the indivudual entries of Σ_t. We also present an algorithmic solution of Eq. (A.10) where the objective is to have the result in a matrix. The method is shown here

$$\widehat{\mathrm{Var}}(\Sigma_t) = (1-\lambda)^2 \sum_{i=0}^{t}(\lambda^{2i} e_{t-i}^2 (e_{t-i}^2)^{\mathrm{T}} - \hat{\Sigma}_t^2), \tag{A.11}$$

where e_t is the vector of all base forecast errors at time t and the squared values are element-wise. We can now rewrite Eq. (A.11) to get a recursive estimate of variance and covariance of the individual entries of Σ_t, as demonstrated below:

$$\widehat{\mathrm{Var}}(\Sigma_t) = (1-\lambda)^2\lambda(e_t^2(e_t^2)^{\mathrm{T}} - \hat{\Sigma}_t^2) + \lambda^2(1-\lambda)^2 \sum_{i=0}^{t-1}\lambda^{2i}(e_{t-i}^2(e_t^2)^{\mathrm{T}} - \hat{\Sigma}_t^2), \tag{A.12}$$

$$\widehat{\mathrm{Var}}(\Sigma_{t-1}) = (1-\lambda)^2 \sum_{i=0}^{t-1}\lambda^{2i}(e_t^2(e_t^2)^{\mathrm{T}} - \hat{\sigma}_{t,kj}^2), \tag{A.13}$$

$$\widehat{\mathrm{Var}}(\Sigma_t) = (1-\lambda)^2\lambda(e_t^2(e_t^2)^{\mathrm{T}} - \hat{\Sigma}_t^2) + \lambda^2\widehat{\mathrm{Var}}(\Sigma_{t-1}) \tag{A.14}$$

References

[1] T. Hong, P. Pinson, Y. Wang, R. Weron, D. Yang, H. Zareipour, Energy forecasting: A review and outlook, IEEE Open Access Journal of Power and Energy 7 (2020) 376–388. doi:10.1109/OAJPE.2020.3029979.

[2] International Renewable Energy Agency (IRENA), Global energy transformation: A roadmap to 2050, 2018.

[3] J. Taylor, Short-term electricity demand forecasting using double seasonal exponential smoothing, Journal of the Operational Research Society 54 (2003) 799–805. doi:10.1057/palgrave.jors.2601589.

[4] H. Nielsen, H. Madsen, Modelling the heat consumption in district heating systems using a grey-box approach, Energy and Buildings 38 (2006) 63–71. doi:10.1016/j.enbuild.2005.05.002.

[5] R. T. Clemen, Combining forecasts: A review and annotated bibliography, International Journal of Forecasting 5 (1989) 559 – 583. doi:https://doi.org/10.1016/0169-2070(89)90012-5.

[6] G. Athanasopoulos, R. Hyndman, N. Kourentzes, F. Petropoulos, Forecasting with temporal hierarchies, European Journal of Operational Research 262 (2017) 60–74. doi:10.1016/j.ejor.2017.02.046.

[7] P. Nystrup, E. Lindström, P. Pinson, H. Madsen, Temporal hierarchies with autocorrelation for load forecasting, European Journal of Operational Research 280 (2020) 876–888. doi:10.1016/j.ejor.2019.07.061.

[8] D. Dominković, R. Junker, K. Lindberg, H. Madsen, Implementing flexibility into energy planning models: Soft-linking of a high-level energy planning model and a short-term operational model, Applied Energy 260 (2020). doi:10.1016/j.apenergy.2019.114292.

[9] H. Lund, S. Werner, R. Wiltshire, S. Svendsen, J. E. Thorsen, F. Hvelplund, B. V. Mathiesen, 4th generation district heating (4gdh): Integrating smart thermal grids into future sustainable energy systems, Energy 68 (2014) 1 – 11. doi:https://doi.org/10.1016/j.energy.2014.02.089.

[10] H. Madsen, K. Sejling, H. T. Søgaard, O. P. Palsson, On flow and supply temperature control in district heating systems, Heat Recovery Systems and CHP 14 (1994) 613 – 620. doi:https://doi.org/10.1016/0890-4332(94)90031-0.

[11] M. Dahl, A. Brun, G. Andresen, Using ensemble weather predictions in district heating operation and load forecasting, Applied Energy 193 (2017) 455–465. doi:10.1016/j.apenergy.2017.02.066.

[12] T. S. Nielsen, H. Madsen, L. Gottlieb, Control of supply temperature in district heating systems with multiple supply points, Ecos 2005 - Proceedings of the 18th International Conference on Efficiency, Cost, Optimization, Simulation, and Environmental Impact of Energy Systems (2005) 1071–1079.

[13] O. P. Palsson, H. Madsen, H. T. Søgaard, Application of predictive control in district heating systems, Proceedings of the Institution of Mechanical Engineers, Part A: Journal of Power and Energy 207 (1993) 157–163. doi:10.1243/PIME_PROC_1993_207_029_02.

[14] M. Schubert, C. Holter, R. Soell, Solar district heating (sdh): Technologies used in large scale sdh plants in graz - operational experiences and further developments, 2010, pp. 140–142.

[15] D. Tschopp, Z. Tian, M. Berberich, J. Fan, B. Perers, S. Furbo, Large-scale solar thermal systems in leading countries: A review and comparative study of denmark, china, germany and austria, Applied Energy 270 (2020). doi:10.1016/j.apenergy.2020.114997.

[16] R. Halvgaard, P. Bacher, B. Perers, E. Andersen, S. Furbo, J. Jørgensen, N. Poulsen, H. Madsen, Model predictive control for a smart solar tank based on weather and consumption forecasts, Energy Procedia 30 (2012) 270–278. doi:10.1016/j.egypro.2012.11.032.

[17] I. Blanco, D. Guericke, A. Andersen, H. Madsen, Operational planning and bidding for district heating systems with uncertain renewable energy production, Energies 11 (2018). doi:10.3390/en11123310.

[18] E. Dotzauer, Simple model for prediction of loads in district - heating systems, Applied Energy 73 (2002) 277–284. doi:10.1016/S0306-2619(02)00078-8.

[19] P. Bacher, H. Madsen, H. Nielsen, B. Perers, Short-term heat load forecasting for single family houses, Energy and Buildings 65 (2013) 101–112. doi:10.1016/j.enbuild.2013.04.022.

[20] P. Bacher, H. Madsen, H. A. Nielsen, Online short-term solar power forecasting, Solar Energy 83 (2009) 1772 – 1783. doi:https://doi.org/10.1016/j.solener.2009.05.016.

[21] L. B. Rasmussen, P. Bacher, H. Madsen, H. A. Nielsen, C. Heerup, T. Green, Load forecasting of supermarket refrigeration, Applied Energy 163 (2016) 32 – 40. doi:https://doi.org/10.1016/j.apenergy.2015.10.046.

[22] G. Athanasopoulos, R. Ahmed, R. Hyndman, Hierarchical forecasts for australian domestic tourism, International Journal of Forecasting 25 (2009) 146–166. doi:10.1016/j.ijforecast.2008.07.004.

[23] R. Hyndman, R. Ahmed, G. Athanasopoulos, H. Shang, Optimal combination forecasts for hierarchical time series, Computational Statistics and Data Analysis 55 (2011) 2579–2589. doi:10.1016/j.csda.2011.03.

006.

[24] S. L. Wickramasuriya, G. Athanasopoulos, R. J. Hyndman, Optimal forecast reconciliation for hierarchical and grouped time series through trace minimization, Journal of the American Statistical Association 114 (2019) 804–819. doi:10.1080/01621459.2018.1448825.

[25] R. J. Hyndman, A. J. Lee, E. Wang, Fast computation of reconciled forecasts for hierarchical and grouped time series, Computational Statistics and Data Analysis 97 (2016) 16 – 32. doi:https://doi.org/10.1016/j.csda.2015.11.007.

[26] J. Jeon, A. Panagiotelis, F. Petropoulos, Probabilistic forecast reconciliation with applications to wind power and electric load, European Journal of Operational Research 279 (2019) 364–379. doi:10.1016/j.ejor.2019.05.020.

[27] G. Yagli, D. Yang, D. Srinivasan, Reconciling solar forecasts: Sequential reconciliation, Solar Energy 179 (2019) 391–397. doi:10.1016/j.solener.2018.12.075.

[28] D. Yang, H. Quan, V. Disfani, C. Rodríguez-Gallegos, Reconciling solar forecasts: Temporal hierarchy, Solar Energy 158 (2017) 332–346. doi:10.1016/j.solener.2017.09.055.

[29] O. Ledoit, M. Wolf, Improved estimation of the covariance matrix of stock returns with an application to portfolio selection, Journal of Empirical Finance 10 (2003) 603–621. doi:10.1016/S0927-5398(03)00007-0.

[30] D. Guericke, I. Blanco, J. M. Morales, H. Madsen, A two-phase stochastic programming approach to biomass supply planning for combined heat and power plants, Or Spectrum 42 (2020) 863–900.

[31] H. Madsen, Time series analysis, CRC Press, 2007.

[32] R Core Team, R: A Language and Environment for Statistical Computing, R Foundation for Statistical Computing, Vienna, Austria, 2019.

[33] P. Nystrup, E. Lindström, J. K. Møller, H. Madsen, Dimensionality reduction in forecasting with temporal hierarchies, International Journal of Forecasting (2021). doi:https://doi.org/10.1016/j.ijforecast.2020.12.003.

[34] T. Van Erven, J. Cugliari, Game-theoretically optimal reconciliation of contemporaneous hierarchical time series forecasts, in: A. Antoniadis, J.-M. Poggi, X. Brossat (Eds.), Modeling and Stochastic Learning for Forecasting in High Dimensions, Springer International Publishing, Cham, 2015, pp. 297–317.

[35] J. Schäfer, K. Strimmer, A shrinkage approach to large-scale covariance matrix estimation and implications for functional genomics, Statistical Applications in Genetics and Molecular Biology 4 (2005) 1–30. doi:10.2202/1544-6115.1175.

[36] R. Hyndman, A. Koehler, Another look at measures of forecast accuracy, International Journal of Forecasting 22 (2006) 679–688. doi:10.1016/j.ijforecast.2006.03.001.

Paper B

USE OF SMART METERS AS FEEDBACK FOR DISTRICT HEATING TEMPERATURE CONTROL

Authors:
Hjörleifur G. Bergsteinsson, Torben Skov Nielsen, Jan Kloppenborg Møller, Sara Ben Amer, Dominik Franjo Dominković and Henrik Madsen.

Published in:
Energy Reports.

17th International Symposium on District Heating and Cooling, DHC2021, 6–9 September 2021, Nottingham, England

Use of Smart Meters as Feedback for District Heating Temperature Control

Hjörleifur G. Bergsteinsson[a,*1], Torben Skov Nielsen[b], Jan Kloppenborg Møller[a], Sara Ben Amer[c], Dominik Franjo Dominković[a], and Henrik Madsen[a]

[a] Technical University of Denmark, Department of Applied Mathematics and Computer Science, Anker Engelunds Vej 1, Building 101A, 2800 Kgs. Lyngby, Denmark
[b] ENFOR, Røjelskær 11, 2840 Holte, Denmark
[c] Technical University of Denmark, Department of Technology, Management and Economics, Produktionstorvet, Building 424, 2800 Kgs. Lyngby, Denmark

Abstract

Smart meters implemented at the end-user in the energy sector create the opportunity to develop data-intelligent methods for district heating systems by using a large amount of fine-granular heat consumption time series from end-users. The current state-of-the-art methods for temperature control in district heating systems rely on predefined critical points in the network and a set reference temperature curve that expresses the minimum forward temperature as a function of the outdoor temperature at the end-users. The critical points are used to ensure that the consumers' supply temperature requirements are met all times. To predefine the critical points at the network, the location of the lowest temperature in the grid needs to be identified at any point in time. Since the lowest temperature often varies over time, one must have a set of critical points in a district heating network. This paper proposes a method to estimate the temperature at an artificial critical point for the network using time-wise quantile estimation using smart meter data at end-users. This novel approach eliminates the need for physical critical points in the net or sensors in wells and creates the possibility of changing the critical point location if needed. The benefits for the provider of using smart meters as feedback, makes the measurement wells redundant and flexibility of the location. The location of low temperature areas in the network can change overtime hence the flexibility of being able to change where the feedback is located. The proposed method to replace the well measurements to provide feedback for temperature control at the production site groups a predefined set of smart meter readings together for each point in time. The grouping is done to have reliable measurements from each smart meter device, excluding some of the meters when a faulty reading occurs. The set of acceptable readings is used to estimate the street pipe's temperature using the estimated quantile of the forward temperature. The approach is tested on two demo cases. The first demo consists of smart meters to estimate the forward temperature of the main street pipe. The second demo uses three smart meters at large apartment buildings as feedback for the control. Initial results show that the estimated temperature of the network can replace the well-measurements which traditionally are used as feedback for temperature control and give a better and more flexible control.

Keywords: District Heating, Temperature Control, Smart Meter Data, Quantile Estimation, Estimated Network Temperature.

1. Introduction

The European Union requires houses connected to a district heating network to be equipped with smart meter devices where feasible [1]. This allows linking the consumption to billing from the district heating and enables the end-user to be more aware of their energy consumption. In Denmark, the digitalization of district heating has been occurring in the past decade by installing smart meters and weather stations in the cities. This digital transformation has given rise in studies relating to district heating and investigating how smart meters data can be used to give valuable insight into the network performance and building energy efficiency, i.e. leakage in the systems or insufficient cooling of the water from inlet to outlet in some buildings. Kristensen and Petersen [2] use smart meter

[1] * Corresponding author. Tel.: +45-45253095
E-mail address: hgbe@dtu.dk

data to derive three heating efficiency indicators of buildings and give an overview of the smart meters system at the district heating utility in Aalborg in Denmark. The same study estimated that 54% of the building area in Denmark is supplied with district heating. Bacher et al. [3] propose a method to separate the total heat load from a single-family residential building into domestic hot water heating and space heating using a non-parametric method to identify the domestic hot water heating. Hence, smart meter data has been studied to improve building performance. However, there have not been any studies on using smart meter data to operate the network and production more efficiently. Smart meter data opens up the possibility to use it as feedback of the network for temperature control at the production site. This is highly valuable for the district heating sector as it is changing from traditional fossil fuels to renewable sources in district heating and is more and more connected to other energy sectors. Therefore, the district heating utilities need to be operated efficiently, utilizing periods when intermittent renewable energy sources (e.g. wind and solar) are available to be used for producing heat, and CO2 reduction. Hence, smart meters can increase the possibility of making the district heating sector more sustainable and flexible.

Operating the district heating network adequately can be done by implementing a control strategy with the objective of reducing heat production and heat losses in the network by minimizing the supply temperature at the production site. Using data-driven temperature optimization in district heating could reach between 240 and 790 million DKK in yearly savings by lowering the supply temperature between 3 and 10 degrees [4]. Madsen et al [5] propose a novel method to introduce a control strategy that minimizes the supply temperature by regulating the flow to match the consumers' heat demand without violating any requirements, i.e. minimum supply temperature to the end-user. The temperature control strategy needs feedback on how the network reacts to changes at the production site to vary the supply temperature adequately; therefore, measurement wells are installed in the network where the network operators believe are the critical points, i.e. where the lowest temperature of the network is. A network usually has only a few of these critical points as they are expensive and need to be maintained regularly. Nielsen and Madsen [6] demonstrate the energy savings at a district heating utility by using data-driven temperature control that is based on Madsen et al [5]. We propose to use smart meters at end-users as feedback of the network's supply temperature for temperature control to get the response characteristics of the network back to production. This novel approach makes the measurement wells redundant and the critical points flexible. For example, as the networks get older or new areas are added to the network, the location of a critical point could change. Therefore, using smart meters to create temperature feedback allows them to change the critical point location when needed.

We used two case studies to demonstrate how smart meters can be used as feedback for temperature control. The first case study used data from a group of single-family houses' smart meters to estimate the network's temperature. This estimated network temperature would then be used as feedback for temperature control. We compared the estimated network temperature to the supply temperature measured at a measurement well that is located before the group of houses that were used to create the estimated temperature. The second case study demonstrates how to use three different smart meters from large apartment buildings as feedback of the network for temperature optimization. A temperature controller uses these three smart meters as input to optimize the supply temperature for the network. We then demonstrate the results of having temperature optimization that uses smart meters as feedback in an online operation. The novelty in this paper is to demonstrate that smart meters can be used as the feedback of the network characteristics either by using a group of meters to estimate the temperature or using them directly.

2. Methodology: Dynamics of Smart Meters in District Heating

The proposed algorithm uses data from a group of smart meters at the end-user to estimate the supply temperature in the main distribution pipe in the street where the end-users are connected to. The mass flow, supply temperature, and timestamp of the readings are the only variables used from the data-set for the algorithm. The flow is used to estimate if the temperature at the smart meter is reliable or not. When there is almost no flow, the water in the pipe leading into the end-user's house becomes still, and the temperature starts to drop due to the heat loss to the surroundings. Therefore, readings with a low flow are removed from the data set and not used to estimate the supply temperature in the main pipe as they do not represent the temperature in the distribution pipe due to the heat loss in the service pipe leading into the end-users. If the data from the smart meters do not send data with fixed time intervals and do not send them at the same time, but rather at random times and different frequencies. The data needs therefore to be aggregated and filled such that they have the same time points and resolution. The readings

from each smart meter are therefore aggregated with a fixed interval by aggregating the readings within the interval to timestamp. The median is computed from those readings to represent the temperature at the timestamp at the given smart meter. For example, for 30 minutes resolution, the rounding will be to 30 minutes and the time is rounded to the nearest hour or half-hour. Thus, timestamps between 45 and 15 are rounded to the hour, and between 15 and 45 are rounded to the half-hour. After rounding the readings for all smart meters, readings with the same timestamps from every smart meter are grouped, and an appropriate quantile of the grouped temperatures is computed. This approach creates an estimation of the distribution pipe temperature using the smart meters at the end-user. Note, it is not recommended to use too high quantile for the estimation. For example, the 90th quantile is more suitable instead of using the 100th quantile, i.e. the maximum temperature of the dwelling for the given timestamp. It is done to remove any measurement noise, or faulty readings thus creating more robust feedback of the street pipe temperature.

For the second case study, a trial was carried out during the heating season 2020/2021 to demonstrate how digitalization can improve the operation of existing district heating networks. The temperature control was conducted using the HeatTO™ software provided by ENFOR that was installed to optimize the supply temperature from the heat exchanger to lower the supply temperature without breaking any restrictions, i.e. minimum supply temperature at the consumer[2]. The methodology behind the temperature controller will not be introduced here as it is quite complex. It can be found in Madsen et al [5] and Madsen et al [7]. However, the temperature controller needs feedback of the network to estimate the network's characteristics, i.e. time-delay in the network and heat losses. Here, data from smart meters are used directly as feedback to the production to optimize the temperature. Hence the data needs to be of high quality and represent the network such that the network characteristics can be estimated and used to optimize the supply temperature and minimize the production cost. To estimate the performance of the new controller, degree days are used to compare supply temperature between two seasons. Degree days are used to compare supply temperature between different heating seasons. The degree days, T^{dd} are computed by estimating the difference between the average ambient temperature, $\overline{T}_a$ over one day, and using 17 °C as the cut-off of heating demand from buildings,

$$T^{dd} = max(0,\ 17\ - \overline{T}_a)$$

The degree days gives the possibility to quantitative the performance between operations as two controllers can not be operated at the same time.

3. Description of Case Studies

3.1. Case study 1: Group of smart meters

In this case study, we apply the smart meter heating data to create an estimated network temperature to be used as temperature feedback of the network to the production. The data is from a district heating network in Brønderslev that is located in the northern part of Jutland in Denmark. The district heating utility, Brønderslev Forsyning has supplied smart meter data from one area in their network. The data set consists of measurements from 15 single-family houses that are also located close to a measurement well. They have also supplied data from the measurement well. Thus, we can compare the result from the proposed algorithm to the measured supply temperature in the network. In order not to violate privacy and comply with GDPR, the smart meter data was anonymized as only Brønderslev Forsyning knows which houses the data belongs to. Hence, the location is unknown and the only information about their location is that each house in the data set is located close to each other. Fig. 1 shows the raw data from the smart meter data set, the plots to the left show the supply temperature and the flow in a winter period while the plots located on the right show the same in a summer period. The bold red line in the temperature plot is the supply netpoint temperature measured at a measurement well that is located close to the houses, or the supply netpoint temperature that is typically used to see how the network is performing. The other lines are the supply temperatures from the smart meters measured at the end-users. The flow is also measured at the

[2] https://enfor.dk/services/heatto/

end-user shown in the lower plots. We can see from the temperature plots that the smart meter temperatures follow the dynamic of the netpoint temperature but there is a temperature level offset. The offset can be explained as the temperature drop due to heat loss to the surroundings in the service pipe from the street to the end-user. The difference between winter and summer can be seen when the flow is compared between the seasons. During the cold months in the winter, houses usually have constant space heating, and thus have a constant flow. However, during summer periods in Denmark, there is no need for space heating. Therefore, the flow during the summer is frequently close to zero except for the peaks as the plot shows, which can be explained by consumers' domestic hot water usage.

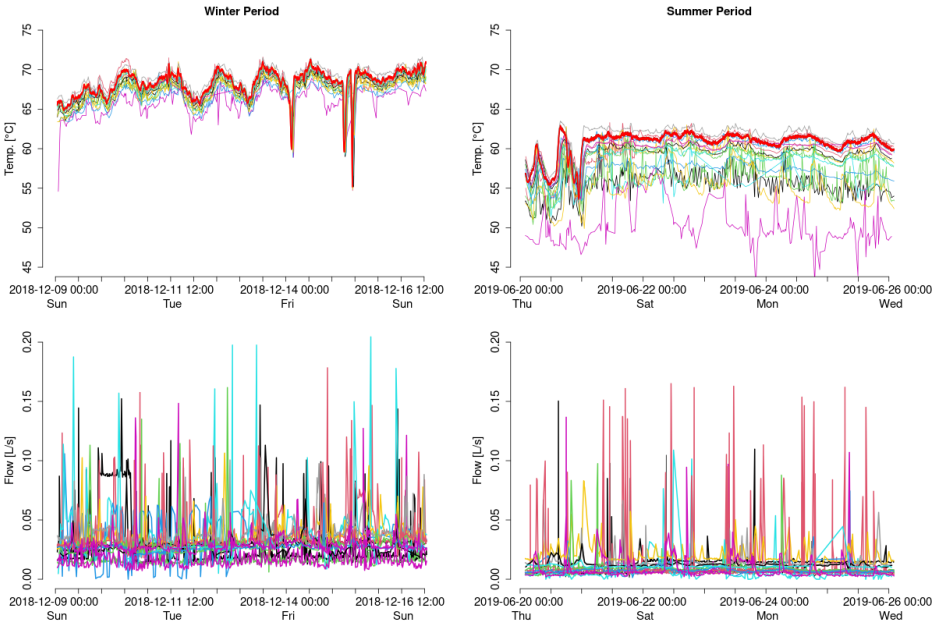

Fig. 1. The temperature and flow of 15 consumers from the district heating network in Brønderslev Forsyning for the winter and summer seasons. The bold red line shows the netpoint temperature measured at a measurement well in the network and used for feedback of the network.

The readings from the smart meters occur with different frequencies and time-interval. This is more visible in Fig. 2 where the plots to the left show temperature and flow readings over two days from two different smart meters. We can see from the plots that readings come at different times, and the frequency is also different. The plot on the right side demonstrates this in more detail, visualising temperature readings from all 15 meters between 17:30 and 18:30. The frequency of the meter readings differs - some give multiple readings in this interval, while some just one, and from one of the smart meters there is no reading at this specific time interval. Hence, it is essential to aggregate the data to a specific resolution before using it to create an estimated supply temperature.

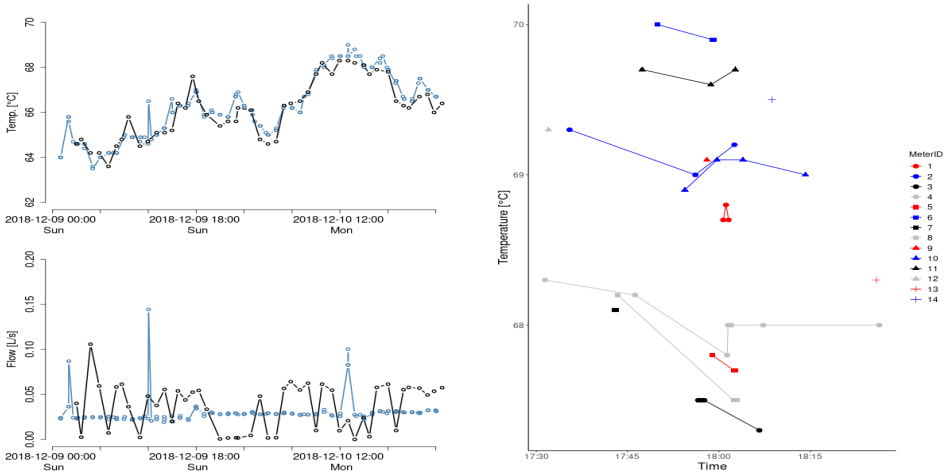

Fig. 2. Differences among the readings (frequency, time of the reading) from smart meters when compared to each other.

3.2. Case study 2: Three large apartment buildings as feedback

This case study is located in the Tingbjerg area in Copenhagen, a small area with large apartment buildings. The network is thought of as an island within HOFORs district heating network as the area is isolated and only supplied with heat from a heat exchanger. HOFOR is Copenhagen's district heating supplier as it both produces heat and operates the distribution network. The network in Tingbjerg is small and has a short time delay, so a measurement well has not been deemed necessary for controlling the supply temperature from the heat exchanger. Consequently, the control of the supply temperature has been operated as an open-loop system, and the temperature is usually determined conservatively to ensure sufficient supply temperature to the consumers. Hydraulic simulation of the system and current ambient temperature were therefore used to optimize the temperature. Hence, no feedback of the network on how it adapts to changes at production and if the requirements are fulfilled. The Tingbjerg area is an ideal case study to demonstrate how to use the digitalization of the network to improve the network's operation as feedback for temperature control. In this study smart meters are used as feedback from large apartment buildings for the temperature control at the heat exchanger. HOFOR supplied smart meter data from 38 apartment buildings inside the Tingbjerg network to be considered as the network's feedback. The data collection started at the beginning of 2020 and data is still being collected as input for the controller that is still in operation. The data in an hourly temporal resolution was supplied once a day, containing readings from each meter since the previous data dump. In November 2020, the data collection changed - now the data is sent every hour with a 15 minutes temporal resolution. Thus, the controller relies on newer and more frequent data. Three smart meters were selected based on the measurements' reliability to give accurate temperature feedback of the network to the controller. The selected meters have relatively constant flow; therefore, the forward temperature signal is of high quality (i.e. the water does not become still in the pipe leading to the house), and measurements from these meters can be used for temperature control. Temperatures from the three meters to be used as feedback are shown in Fig. 3, where it can be seen that they are of good quality despite some peaks in the winter period.

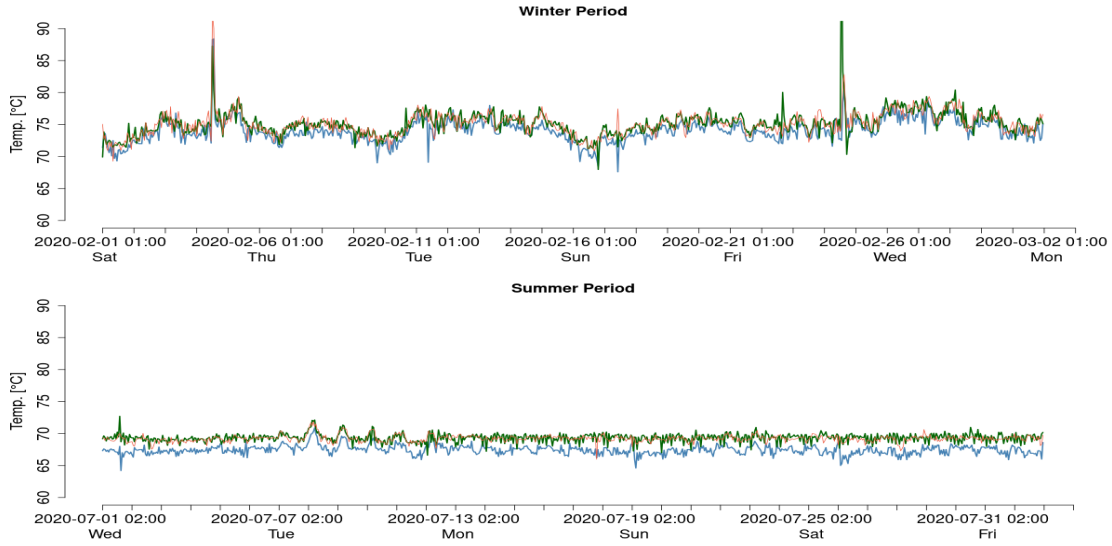

Fig. 3. Supply temperature from three smart meters in Tingbjerg used as feedback for temperature control. The upper and lower plots show the temperature at two different seasons, winter and summer.

4. Discussion and Results

Fig. 4 shows the result from the first case study in Brønderslev where we use a group of smart meters to estimate the netpoint temperature. The grey lines represent the temperature from the end-users. The blue line is the netpoint temperature measured at a measurement well in the network. The well is located before the dwellings that supplied the smart meter data. The red line is the estimated network temperature using the smart meter data using the proposed algorithm from Section 2. We have aggregated the data to 30 minutes resolution and used the 90th quantile to estimate the netpoint temperature from the smart meters. The estimated temperature mimics the measured netpoint temperature adequately over the wintertime. However, it does not perform as well over the summer month. The metric scores, *Mean Absolute Error* (MAE) and *Root Mean Squared Error* (RMSE) are given in the plots to

illustrate the performance of the algorithm and the difference between the two seasons. The accuracy difference between the heating season and the summer months can be due to fewer consumers that use space heating when the air temperature increases - the remaining usage is then domestic hot water, i.e. showering. Thus, the hot water in the pipe into the house is not in constant use; therefore, the water becomes still, and the temperature drops due to the heat loss to the surroundings. This can be seen in Fig. 1 when comparing the flow between the seasons and in Fig. 4 when looking at the different dynamics between the seasons. In the summertime, the temperature tends to drop exponentially and rise quite fast in a short time. The algorithm is however able to capture the street temperature accurately even though it has less information than the winter months due to infrequent usage of space heating. Notice that we are not trying to estimate the supply temperature measured at the well, it is only used to benchmark the performance of the estimated temperature in the street pipe outside the dwellings. Therefore, the algorithm is not optimized to have a perfect representation of the well measurement instead it is used to represent the temperature in the street, which can have different dynamics than the well because of heat loss between the well and dwellings.

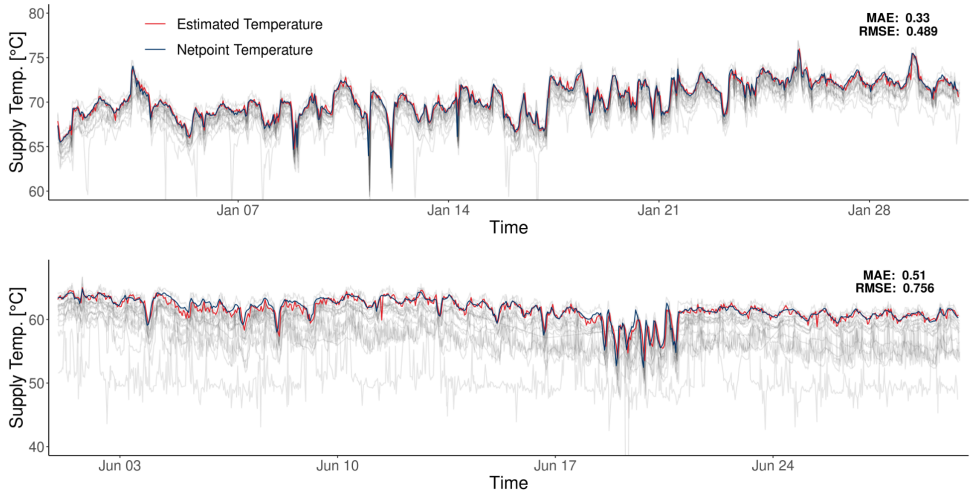

Fig. 4 The estimated network temperature from the proposed method as the red line and the measured netpoint temperature as the blue line. The grey lines are the temperature from the smart meters used in the proposed method.

Fig. 5 shows the result from the second case study in Tingbjerg where it is presented by comparing supply temperature from the previous operation and trial using the new temperature controller and smart meters as feedback of the network. The top-left plot shows the supply and ambient air temperatures for the previous operation when an open-loop controller was in operation, and the middle-left plot shows it for a period when the new controller was used. The bottom plot shows the difference between current and previous supply temperature from the periods in the plots above. The difference plot shows that using the new controller gives a more robust temperature, i.e. the temperature is not changing rapidly as can be seen also when comparing the top and middle-left plots. The variance of the two difference time-series is also given in the plot to illustrate how stable the supply temperature has become using the controller. Hence, having a more stable supply temperature is beneficial for the control of the heating unit at the consumers' and the network's equipment. Large and frequent fluctuations in the supply temperature should be avoided as it increases the maintenance costs compared to stable operations [6].

The plots on the right side of Fig. 5 demonstrate the performance of the control by comparing the supply and return temperature against the degree days. The average supply and return temperature for each day is then computed and plotted against its corresponding degree day as shown in the top and middle-left plots. We can see that the supply temperature is higher when the new controller is in operation as the regression lines also demonstrate. The regression lines (with intercept and slope) are used to highlight the difference between operations. This is a result of the reference curve that was used for the netpoint temperature was transferred from the previous open-loop operation to control the supply temperature to be used when the new controller was operating. The reference curve at the netpoint is to ensure that consumers will receive the required supply temperature as a function of the ambient air temperature.

In the past, the supply temperature control at production did not have any feedback from the network and therefore the reference curve was stated quite conservatively to ensure that the consumers would receive the required supply temperature. Notice that the slope is lower for the new controller hence for colder days the new controller will operate more efficiently even for suboptimal restrictions. The return temperature was also not affected by the new controller as shown in the right bottom plot in Fig. 5. Fig. 6 shows the supply temperature at the three netpoints for both periods. The plots demonstrate that the new controller can maintain the temperature at the netpoints with less spread and a better level of security than what was possible with the previous operation. This can be seen when comparing the prediction intervals (PI) between the two periods in Fig. 6. The intervals were estimated using nonparametric quantile regression using the 10th and 90th quantiles as the lower and upper bounds. The lower bound for the new controller rarely violates the restriction. Hence, it would have been possible to have a lower reference curve during the trial, and still maintain the same level of supply security at the consumer as with previous control. However, this potential for savings was not realized during the trial. The top right plot in Fig. 5 also shows an adjusted supply temperature for the new controller to demonstrate a "what if" situation. The adjusted supply temperatures are lowered by moving the reference curve down such that temperatures from previous operations are above it, we estimated it would by lowering it by 5°C from Fig. 6. The supply temperature for the new control would therefore be 3°C lower compared to the previous operation as can be seen in the regression lines.

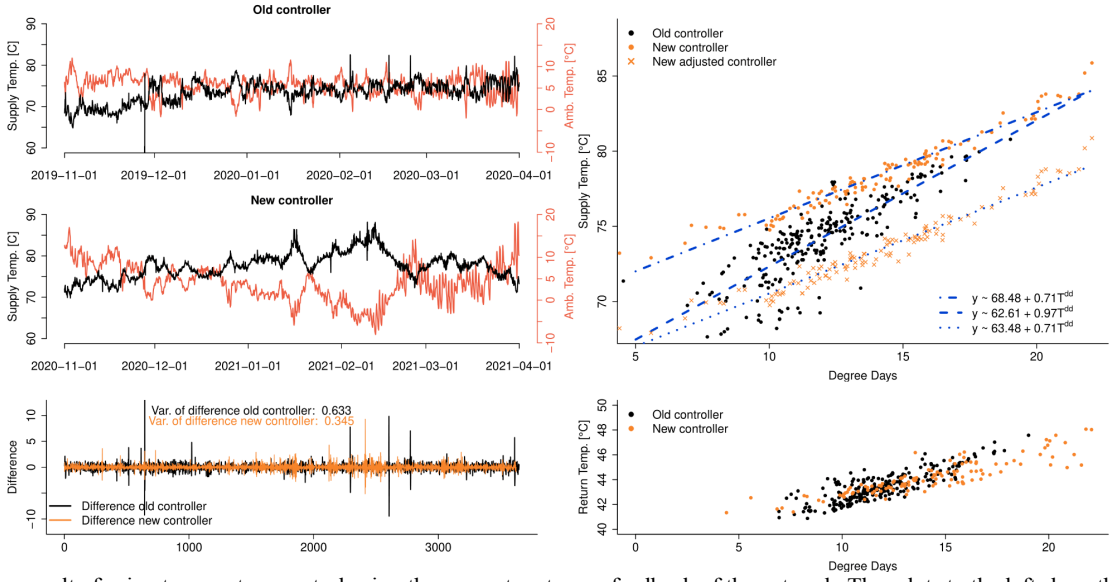

Fig. 5. The result of using temperature control using three smart meters as feedback of the network. The plots to the left show that the supply temperature at the production is more robust with the controller. The plots to the right demonstrate that the supply temperature at the production was higher when the new controller was in operation. The top right plot also shows the supply temperature adjusted by lowering them by 5°C.

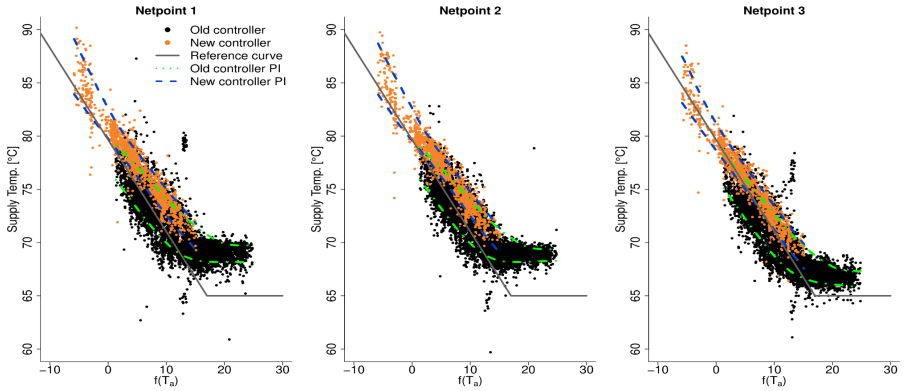

Fig. 6. Supply temperature at the three netpoints (the smart meters) plotted against the rolling average of the ambient air temperature with window length of 24. The reference curve constrains the controller to supply the consumer supply temperature for a given ambient temperature.

5. Conclusion

In this paper, smart meters were used to demonstrate that they can add value to the district heating production and network. A novel algorithm was presented to create feedback of supply temperature from the network using smart meters located at single-family buildings. The estimated network temperature from the algorithm adequately mimics the netpoint temperature measured at measurement well located close to the buildings. Therefore, district heating utilities can reduce their cost by making measurement wells redundant and the corresponding maintenance of the well. The location of the predefined critical points becomes more flexible when using the smart meters as feedback. Using a few smart meters from large apartment buildings that have reliable measurement can be used as direct feedback to be used for temperature control. Results from an online operation of using three smart meters as feedback of the network to temperature control were presented. The results showed that the supply temperature is more robust than without the controller. Unfortunately, restrictions on the supply temperature at the consumer side were set slightly higher than without the new controller and as a result, the supply temperature at the production was higher when the controller was in operation. However, if the restrictions at the consumer were lower during the trial, then the supply temperature at the production side would at least perform with a similar temperature as when the controller was not in operation, and the temperature would also be more robust - with expected savings in network maintenance costs to follow in the longer term. Furthermore, using smart meters as feedback for network operation reduces the need of installing new measurement wells or maintaining the existing ones.

Future research on digitalization in district heating should focus on how to quantify the operational savings for district heating when incorporating digitalization into their operations. Digitalization can have s direct impact on the operational savings for district heating yielding better operation of the system, i.e. temperature optimization in the network, lowering return temperature by identifying bad coolers in the network, improving weather forecast with local climate stations, and leakage in the network.

Acknowledgements

This work is funded by Innovation Fund Denmark through the project Heat 4.0 (8090-00046B) and TOP-UP (9045-00017B), the IDASC (18012745) project funded by Region Hovedstaden, and IEA TS4 annex on Digitalisation of District Heating: Optimised Operation (and Maintenance) of District heating and Cooling Schemes via Digital Processes Management funded by the EUDP Denmark. The authors would like to thank Brønderslev Forsyning and HOFOR for making their data available and for their support. Furthermore, we would like to thank ENFOR, which supplied their software for the trail and their valuable inputs.

References

[1] Directive 2012/27/EU of the European Parliament and of the Council of 25 October 2012 on energy efficiency (EED). (2012). Available online: http://data.europa.eu/eli/dir/2012/27/oj (accessed on 28 March 2021).

[2] Kristensen, Martin Heine, and Steffen Petersen. (2021) "District heating energy efficiency of Danish building typologies." *Energy and Buildings, 231* (2021): 110602.

[3] Bacher, Peder, Philip Anton de Saint-Aubain, Lasse Engbo Christiansen, and Henrik Madsen. (2016) "Non-parametric method for separating domestic hot water heating spikes and space heating." *Energy and Buildings*, 130 (2016): 107-112.

[4] Damvad Analytics. (2019) "Potentialet ved dynamisk datadrevet temperaturregulering i fjernvarmesektoren." February 2019. Available online:https://www.danskfjernvarme.dk/viden/f-u-konto-subsection/rapporter/2018-02-potentialet-ved-dynamisk-datadrevet-temperaturregulering-i-fjernvarmesektoren (accessed on 28 March 2021).

[5] Madsen, Henrik, Ken Sejling, Henning T. Søgaard, and Ólafur P. Pálsson. (1994) "On flow and supply temperature control in district heating system." *Heat Recovery Systems and CHP*, 14 (1994): 613-620

[6] Nielsen, Torben S., and Henrik Madsen (2002) "Control of Supply Temperature in District Heating Systems." In Proceedings of the 8th International Symposium on District heating and Cooling (2002).

[7] Madsen, Henrik, Torben S. Nielsen and Henning T. Søgaard. (1990) "Control of Supply Temperature." *IMM, DTU* (1994).

Paper C

ESTIMATING TEMPERATURES IN A DISTRICT HEATING NETWORK USING SMART METER DATA

Authors:
Hjörleifur G. Bergsteinsson, Phillip B. Vetter, Jan Kloppenborg Møller and Henrik Madsen.

Published in:
Energy Conversion and Management.

Estimating Temperatures In a District Heating Network Using Smart Meter Data

Hjörleifur G. Bergsteinsson[a,*], Phillip B. Vetter[a], Jan K. Møller[a], Henrik Madsen[a,b]

[a]*Technical University of Denmark, Denmark*
[b]*Norwegian University of Science and Technology, Norway*

Abstract

Smart meters at consumers create opportunities to improve operation of the district heating sector using data-driven methods. Information from these meter measurements carries the potential to increase the energy efficiency of both individual houses and the utility network, for example by identifying buildings with too high return temperature, or by detecting leakage in the network. This paper proposes a method for using meter data to estimate network temperatures. Network temperatures can subsequently be used to estimate the network characteristics, namely the nonlinear relationship between network temperature and the plants' temperature and flow. A description of the network characteristics is needed for most temperature-optimisation methods to keep the supply temperature as low as possible without violating the system constraints. Traditionally, measurement wells located in the network have been used. These wells are located at critical points in the network where the largest temperature losses occur. Since the lowest temperature often varies over time, multiple critical points are necessary. The method presented in this paper eliminates the need for these physical critical points in the network. It also makes it possible to change the location of the critical points if needed. The network temperature is estimated using a stochastic state-space model of the heat dynamics from the street level distribution pipe over the service pipe and into individual houses. The parameters in the model are estimated using a maximum likelihood approach, and the Kalman Filter is used to evaluate the likelihood function. The estimation process takes advantage of automatic differentiation using the `R` package Template Model Builder (`TMB`) to reduce the computational workload. The proposed method is validated by comparing the estimated temperature with the temperature measured from a measurement well.

Keywords: Temperature Optimisation, Estimating Network Temperature, Grey-Box Modelling, Kalman Filter, Automatic Differentiation

1. Introduction

Future energy systems need to be flexible because of increasing shares of renewable energy sources that are typically intermittent due to their direct weather dependency. Furthermore, new regulation continues to emerge and awareness of greenhouse gas emissions is increasing, with a consequential transition to more renewable energy. For these reasons, more sophisticated methods are needed to deliver the required energy demand without using fossil fuels. Energy systems for gas, heat and electricity need to be integrated to phase out fossil fuels. Because of their unique capability to store energy, district heating systems play a key role in the transition towards more flexible energy systems [1]. An example is that wind power can be used for heating water to be delivered either directly or stored when the electricity demand is lower than the generated wind power. Hence, in order to fully maximise the flexibility potential of energy systems, they must be integrated. In addition, each energy system needs to perform efficiently to deliver an optimal energy integration. Mathiesen et al. [2] gives an extensive discussion on the integration of all energy systems to increase the flexibility of the system and reach 100 % renewable energy supply.

The work presented here focuses on methods for optimising the temperature control in district heating systems based on information from smart meters installed at consumer level. Today, smart meters that can take very frequent readings are installed in many buildings. This paper suggests methods to take advantage of frequent readings by smart meters.

The role of district heating systems is to meet consumer heat demand while simultaneously minimising both production and operation costs. Optimisation of production planning is concerned with scheduling heating unit operation in order to produce the desired heat demand at the lowest cost [3]. In Denmark, heat is often co-generated (combined heat and power (CHP) production). CHP units typically run during periods with high electricity prices. Any heat that is not used during these periods is stored in the system. During low electricity price periods, district heating operators either use this stored heat or run heat-only units, e.g. gas boilers, heat pumps or solar heat [4]. Hence, the plant's objective is to generate as much power as possible during periods with high electricity prices while satisfying the heating demand. This is achieved by keeping the supply temperature as low as possible. Consequently, lowering the temperature will also reduce heat loss in the network and production cost [5]. Hence, it is crucial to lower the

*Corresponding Author
Email address: hgbe@dtu.dk
Postal Address: Anker Engelunds vej 1, Building 101A, 2800 Kongens Lyngby, Denmark

Nomenclature

Abbreviations

SDE	Stochastic Differential Equation
TMB	Template Model Builder
CHP	Combined Heat and Power
ODE	Ordinary Differential Equation
PDE	Partial Differential Equation
GDPR	General Data Protection Regulation

Mathematical notation

ω	Wiener process
σ	diffusion coefficient
V	observation covariance matrix
e	innovation
S	logistics function
L	likelihood function
$\mathcal{L}$	log-likelihood function
U	random effects
θ	fixed effects
f	drift term
g	diffusion term
h	observation function
x	system states
u	external inputs
$\mathcal{N}$	normal distribution
K	Kalman gain
P	state covariance matrix
Σ	observation covariance matrix

Physical parameters

Q	mass flow rate	[kg/s]
ρ	mass density	[kg/m^3]
c_v	specific heat capacity	[kJ/(kg K)]
A	area	[m^2]
v	flow velocity	[m/s]
p	absolute pressure	[Pa]
f_D	darcy friction coefficient	[-]
S	circumference	[m]
k	thermal conductivity	[W/(m K)]
T	temperature	[°C]
$\dot{q}_\mathrm{e}$	heat loss per meter	[W/m]
C	heat capacity per meter	[J/mK]
λ_p	thermal conductivity	[W/(m K)]
τ	time constant	[s]
d_o	outer diamter	[m]
L	length of pipe	[m]
R	thermal resistance per meter	[(K m)/W]
x_a	insulation thickness	[m]

Subscripts

t	Time
g	Ground

Superscripts

s	Street
(i)	House index

supply temperature in order to optimise the operation of the entire district heating system.

A lower temperature in the network will translate to lower production costs as the operation of most power producing devices can be more efficient. For example, in a CHP plant, lowering the temperature results in increase in ratio of power to heat output, and electricity is more valuable than heat [6]. Also, if the utility uses several energy sources then lowering the temperature will also increase the flexibility of utilising the optimal energy sources at given time-point due to the different limitation of each energy sources. Thus, this makes it possible to utilise heat from new sources like excess heat from comfort cooling. Without a lowering of the temperature these sources which would otherwise have been disregarded due to too low temperatures for entering the network. Likewise, it will be more efficient to operate heat pumps in the network. This increased efficiency will result in better investment feasibility for heat pumps. Heat pumps are likely to play a bigger role in supplying heat to the network as they can utilise heat sources with a low-temperature range, e.g. wastewater, ambient water, industrial heat waste, and solar heat storage [7]. Therefore, for optimal operation of a district heating network with multiple heat sources (e.g. heat pumps and solar heat), temperature optimisation is needed in order to keep the supply temperature as low as possible while increasing both the efficiency and feasibility of heat sources. The complexity of the dynamics of district heating networks implies that data-driven models for temperature optimization are needed to lower the supply temperature in the network.

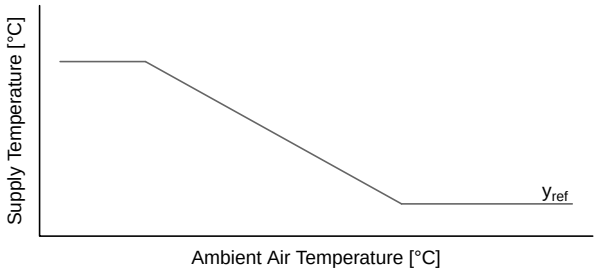

Figure 1: Reference curve, y_ref shows the desired supply temperature of hot water for given ambient air temperature.

Traditionally, the optimisation of supply temperature has been controlled using a reference curve based on the ambient air temperature, as indicated in Figure 1 [8]. The reference curve dictates the minimum allowed supply temperature as a function of the ambient air temperature. The maximum at the

left endpoint of the curve is related to the physical limitation of the system, and the minimum temperature to the right ensures a high enough temperature for domestic hot water usage without risking the formation of bacteria, e.g. legionella. However, this control scheme is a conservative estimate since it only considers one variable to ensure that the heat supply is sufficient at any time, and this often leads to an unnecessarily high supply temperature. Obviously, heat demand is not only dependent on the ambient air temperature but also by other climate variables, e.g. solar radiation and wind. The weather effect on heat demand is also not instantaneous, as buildings are known to have slow thermal reactions. The slow thermal reaction in a single building is demonstrated in Madsen and Holst [9]. In Nielsen and Madsen [10] it is shown how the physical knowledge about heat consumption in buildings can be used to establish a heat consumption model for the district heating network. Preferably, the controller should include heat demand forecasts and the network characteristics of the system (transportation time, heat loss).

Madsen et al. [5] propose a control strategy for the supply temperature at the plant by utilising heat demand forecasts and network characteristics. The strategy estimates future set-points of the supply temperature at the production plant by using a model for the network characteristics determined from receiving feedback of the network temperatures, and the approach allows the system to adapt to changes in the system. This stochastic time-varying system is both nonlinear and non-stationary, and consequently, a nonlinear and time-varying transfer function was proposed in order to model the relationship between the supply temperature and flow at the plant and the network temperature from a critical point in the network. This strategy has been demonstrated to lower the supply temperature at the district heating plant and thereby reduce production costs and the heat losses in the system [11].

This control strategy requires feedback from the system. Usually, this information is measured at a selected number of so-called critical points in the network at street level using sensors in measurement wells. A critical point in the network is typically located close to the end-users showing the largest temperature losses in transportation from the plant. A single critical point is sub-optimal, since the location of the critical points in the network can change over time as a consequence of the diurnal pattern and consequently several (e.g. five) critical points are normally used. However, networks change, e.g. pipes get older, pipes are replaced, new areas are connected to the network, buildings are refurbished and so forth, and hence it would be advantageous if system feedback could be received at any particular location in the network. This is possible by using data from smart meters installed at consumers, to establish temperature feedback. This approach furthermore eliminates the need to install temperature sensors in measurement wells. For a number of years, smart meters have been installed, allowing consumers to link their consumption to bills from the district heating utility. There is a requirement from the European Union that all buildings have individual energy meters where feasible, including heat meters for houses connected to district heating networks [12].

Smart meters create new opportunities to develop data-intelligent methods for district heating operations. This digital transformation has fostered research related to district heating at consumers. Data from smart meters can be used to give valuable insight into network performance and building energy efficiency through investigations of possible leakage in the network or insufficient cooling at consumers. For example, Kristensen and Petersen [13] use smart meter data to derive three heating efficiency indicators of buildings to compare the energy performance of the buildings, and also give insights on district heating and smart meters in Denmark. However, the feedback of the network temperature needs to be robust, and using measurements from smart meters from single-family houses without a quality check could give a wrong representation of the network characteristics. Bergsteinsson et al. [14] propose a simple method to estimate the network temperature by resampling and aggregating data from a group of smart meters. Based on this data, an artificial network temperature is estimated using time-wise quantile estimation at each time step. However, the method is quite naive and is not robust when the measurements are either of bad quality or they lack extended periods of time. A more advanced model is therefore required to give more accurate and reliable continuous feedback of the network. Hence, a model derived from physical knowledge of the system is needed which uses meter measurements to estimate the parameters of the network, i.e. grey-box modelling.

Hence, by establishing a robust method to estimate the network temperature at multiple points in the network by using smart meter data will enable feedback options which again makes it possible for utilities to implement controllers for temperature optimization. Data-driven temperature optimization lowers supply temperature in the network, thus decreasing the production cost and reducing heat losses in the network.

1.1. Pipe dynamics and grey-box modelling in district heating networks

Thermodynamics modelling of hot water pipes in district heating has been studied extensively, both to provide a deeper understanding of the dynamics and to obtain information that can be used to minimise cost: e.g. by selecting the optimal size of pipes or by reducing heat losses in the network. For planned new district heating systems, models of the network can be used to simulate scenarios for the design of new pipelines and to analyse the hydraulic and thermal behaviour for the purpose of minimising the costs of establishing the new networks and to decide on their operation. For already established systems, the methods can be used to simulate different scenarios, e.g. peak loads or to identify locations where maintenance is needed. Hence, adequate physical models of district heating networks are important for efficient operation. There are studies that have proposed different methods for modelling the thermodynamics in hot water pipes for operational purposes. The most widespread approach is a finite element method in which the pipe is divided into infinitesimal segments in order to solve a governing partial differential equation for computations of the temperature difference and heat loss in the pipes.

This can be solved using finite volume schemes, e.g. in Vandermeulen [15] the first-order upwind finite volume model is used to study district heating network flexibility by storing heat in the network by altering the supply temperature. Other studies have proposed using a finite element method to simulate the thermodynamics of having the temperature distribution from plant to end-user for operation during changes in the system [16]. Benonysson et al. [17] propose a node method and derive a mathematical model of the pipes using heat transfer equations to compute the outlet temperature of the pipe when the inlet temperature is measured. Søgaard [18] propose modelling the network as a dynamic input-output system, describing the network response characteristics between measurements from the plant and a point in the network. Hence, different possibilities have been proposed in the literature to describe the network dynamics, from white-box methods entirely based on physics to black-box models based solely on measurements. There is no globally optimal method to describe the dynamics, as each method has its pros and cons. White-box modelling is suitable for network design and flexibility/peak-shaving simulations. White-box methods are usually computationally heavy and require human maintenance to validate and select appropriated values, while black-box methods are fast and usually need no maintenance. Therefore, black-box methods are better suited for control applications where the computation time needs to be low such that the control can respond quickly to changes, as proposed in Madsen et al. [19]. An extensive summary of hot water pipe dynamics is given in van der Heijde [20] and Vandermeulen [15].

The method proposed in this paper, will take advantage of this physical knowledge of pipe thermodynamics to estimate street-level network temperature (critical point) using smart-meter measurements. This will be done using statistical methods to estimate the parameters of equations that are derived from physics, i.e. by using a grey-box modelling approach. Grey-box models bridge the gap between physical and statistical modelling and are frequently studied in literature and have shown promising results for parameter estimation and control purposes. Madsen and Holst [9] demonstrates how grey-box models can be used to describe the dynamics of the indoor air temperature in buildings and its dependency on weather and heat input. The forecasting ability of the method has proved to have high accuracy. Bacher and Madsen [21] describe an approach for using data for optimal model selection of grey-box models and for estimating model parameters for a particular building. They also discuss grey-box model applications for validating the energy performance of buildings, energy consumption forecasting and indoor climate control. Thilker et al. [22] propose a grey-box model to describe the heating dynamics of an old school building, and Thilker et al. [23] demonstrates the potential of using the grey-box model to control the return temperature to lower the operational cost of the building.

1.2. Contribution of the paper and overview

The main contribution of this work is the formulation and application of a set of partially observed stochastic differential equations to be used for inference on the temperature of the distribution network at street level based on smart-meter readings from individual consumers. The system of stochastic differential equations SDEs is derived from partial differential equation PDE) that describe the heat transfer dynamics from the distribution pipe into single-family houses over a service pipe. The estimated network temperature can be used to gain information about the network response characteristics; i.e. how the network reacts to changes in temperature or flow at the plant. This information can be used as feedback for temperature control, for example. Thus, this result makes physical measurement wells redundant, as a group of houses with smart meters can be used to estimate the network temperature.

The paper is organized as follows. Section 2 presents the reasoning behind the chosen SDE formulation and explains the chosen estimation method used for both parameter and state estimation. Section 3 presents the result from applying the methods to the presented data, including parameter interpretation. Finally, Section 4 and Section 5 discuss the presented results and draw some general conclusions.

2. Methods

In this section, a model will be established which describes the heat transfer between the distribution pipe and houses through a service pipe. The model is derived from physics and takes the form of a partial differential equation which is subsequently approximated by a stochastic differential equation for continuous temperature estimates in order to incorporate the information from smart meters. The heat dynamics is affected by the thermodynamic properties of the pipes and their relationship to the surroundings.

2.1. Stochastic differential equations

The models consider in this paper, will be on the following continuous-discrete time stochastic state space form [24],

$$\mathrm{d}x_\mathrm{t} = f(x_\mathrm{t}, u_\mathrm{t}, \theta)\,\mathrm{d}t + g(\theta)\,\mathrm{d}\omega_\mathrm{t} \tag{1}$$

$$y_{\mathrm{t_k}} = h(x_{\mathrm{t_k}}) + e_{\mathrm{t_k}} \tag{2}$$

where the system state x_t evolves in continuous time as determined by the drift f and the diffusion g. The observation function h relates measurements to system states. The drift function depends on the state itself, external inputs u_t and the system parameters θ. The drift term, accounts for most of the known phenomena of the system and draws on physical knowledge, while the diffusion relates to unaccounted for and unknown system drivers as well as noise. The aim is to describe these effects by the random perturbations imposed by the Wiener process ω_t whose non-overlapping increments are independent and Gaussian-distributed i.e. $\omega_\mathrm{t} - \omega_\mathrm{s} \sim N(0, t-s)$. The (hidden) state x_t is observed (indirectly) through the measurements $y_{\mathrm{t_k}}$ which become available at certain discrete times $t = t_\mathrm{k}$. In this paper, it will be assumed that a subset of the states (the individual smart meters) are directly observed, although contaminated by Gaussian noise $e_{\mathrm{t_k}}$.

2.2. Pipe modelling

The thermodynamics modelling of the processes inside a pipe is rather complicated, and consequently a number of assumptions are introduced to approach the problem and make it feasible. The water in the pipes will be assumed to be incompressible, the pipe is a grounded, insulated single pipe (which implies no influence from the return pipe), and that the system is in a steady-state. It will also be assumed that the ground temperature is constant throughout each month using the information from Grunnet Wang et al. [25].

van der Heijde et al. [26] uses a dynamic thermo-hydraulic pipe model for district energy systems for the purpose of creating dynamic simulations of the temperature in district heating networks and cooling pipe systems. The pipe model is a partial differential equation describing the temporal evolution (t) of the energy across the axial dimension (x) in the pipe. The PDE describes the heat transfer and the associated heat losses to the surroundings through a combination of the energy and the continuity equation. The equation is

$$\underbrace{\frac{\partial(\rho c_{\mathrm{v}} T A)}{\partial t}}_{\substack{\text{time}\\ \text{derivative}}} + \underbrace{\frac{\partial(\rho v(c_{\mathrm{v}} T + p/\rho)A)}{\partial x}}_{\substack{\text{spatial}\\ \text{derivative}}} = \underbrace{vA\frac{\partial p}{\partial x}}_{\substack{\text{pressure}\\ \text{difference}\\ \text{energy}}} + \underbrace{\frac{1}{2}\rho v^2 |v| f_{\mathrm{D}} S}_{\substack{\text{wall friction}\\ \text{dissipation}}} + \underbrace{\frac{\partial}{\partial x}(kA\frac{\partial T}{\partial x})}_{\substack{\text{axial heat}\\ \text{diffusion}}} - \underbrace{\dot{q}_e}_{\substack{\text{heat}\\ \text{loss}}}, \tag{3}$$

where ρ $[kg/m^3]$ is the mass density of the fluid in the pipe, c_{v} $[kJ/(kgK)]$ is the specific heat of the fluid in the pipe, A $[m^2]$ is the cross sectional area of the pipe, v $[m/s]$ is the flow velocity, p $[kg/m^3]$ is the absolute pressure, f_{D} $[-]$ is the Darcy friction coefficient, S $[m]$ is the pipe circumference, k $[W/(mK)]$ is the thermal conductivity, T $[^\circ C]$ is the temperature inside the pipe, and $\dot{q}_{\mathrm{e}}$ $[W/m]$ is the heat loss per unit length.

van der Heijde et al. [27] and Vandermeulen [15] argue that most of the terms in Eq. (3) can be assumed to be negligible, and with the additional assumption that the water is incompressible the equation simplifies to

$$\frac{\partial(\rho c_{\mathrm{v}} A T)}{\partial t} + \frac{\partial(\rho c_{\mathrm{v}} A v T)}{\partial x} = -\dot{q}_e. \tag{4}$$

The equation remains a PDE which describes the heat transfer through a pipe in the form of an advection equation with the loss term $-\dot{q}_{\mathrm{e}}$. The equation can be solved analytically when assuming steady-state operation, but it is a challenging task. Hence, it is usually solved instead by using finite volume methods, where the pipe is split into multiple smaller sections. The solution is then obtained by integrating across each of these sections. For further details, see van der Heijde et al. [26], Dénarié et al. [28], and Grosswindhager et al. [29].

Eq. (4) will be used to estimate the supply temperature in the distribution pipe in the street before entering the service pipe into the house. This is illustrated in Figure 2, where the street hot water $T_t^{(\mathrm{s})}$ enters the service pipe and the temperature, $T_t^{(\mathrm{i}),\mathrm{obs}}$ is measured by the smart meters after travelling over the service pipe. The temperature loss in the system is assumed to be caused by heat loss to the surroundings. The flow is assumed to be constant through the service pipe. The only information that is known is the measurements from the smart meters and the assumed constant ground temperature for each month using the average temperature given by Grunnet Wang et al. [25]. The goal is to use only the smart meter data to estimate the street temperature. The data-driven model is a reformulation of Eq. (4) using a stochastic differential equation as shown in Section 2.1.

The first step is to transform Eq. (4) into a standard Resistance-Capacitance (RC) form. First, the heat capacity over the service pipe per unit length of the water is defined as $C = Ac_{\mathrm{v}}\rho$ $[J/mK]$. Next, it will be assumed that the mass flow is constant through the service pipe, $Q = \rho A v$ $[kg/s]$. The final assumption for this model derivation is that the heat loss is proportional to the temperature difference between the water and ground with the proportionality constant being the inverse thermal resistance (between pipe and ground) as shown in Wallentén [30]. Hence, Eq. (4) becomes,

$$C\frac{\partial T}{\partial t} + \frac{\partial(c_{\mathrm{v}} Q T)}{\partial x} = \frac{T_g - T^{(i)}}{R}. \tag{5}$$

The next step is to discretise the equation in space

$$\frac{\partial(c_{\mathrm{v}} Q T)}{\partial x} \underbrace{\longrightarrow}_{\text{discretisation}} \frac{c_{\mathrm{v}} Q \Delta T}{\Delta x}. \tag{6}$$

Thus, the temperature difference is that between the street and house $\Delta T = T^{(\mathrm{s})} - T^{(\mathrm{i})}$ over the service pipe length Δx. The length of the pipe is then multiplied through the equation, which redefines the heat capacity as $C = C\Delta x$ and the thermal resistance as $R = R/\Delta x$. The system now becomes

$$\frac{\partial T_{\mathrm{t}}}{\partial t} = C^{-1}\left(c_{\mathrm{v}} Q_t (T_t^{(\mathrm{s})} - T_t) - R^{-1}(T_t - T_{\mathrm{g}})\right). \tag{7}$$

This differential equation describes how the house temperatures change with time. It is now straight-forward to formulate the proposed stochastic differential equation (for one house)

$$\mathrm{d}T_{\mathrm{t}} = C^{-1}\left(c_{\mathrm{v}} Q_t (T_t^{(\mathrm{s})} - T_t) - R^{-1}(T_t - T_{\mathrm{g}})\right)\mathrm{d}t + \sigma\,\mathrm{d}\omega_t, \tag{8}$$

where σ is the diffusion coefficient and ω_t is a standard Wiener process, i.e. the source of the noise in the system. Hence, we have established a time-dependent stochastic state space model to describe the energy exchange between the street and individual houses. Jointly with the observation equation, this type of model is also referred to as a grey-box model in the literature, see for instance [9].

As already mentioned, the only information available is the smart-meter data, and the objective of this study is then to estimate the supply temperature at the distribution pipe in the street before the water enters the houses. To clarify; the information available is from multiple smart-meter measurements installed

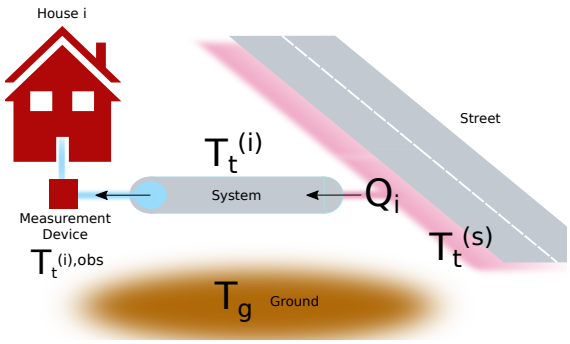

Figure 2: A drawing that demonstrates the interaction between street, house, and ground with labels for states, observations, and inputs. The modelled system is that of the proposed model and it is centred around the pipe. The temperature of the water travelling through the pipe is the result of mixing the already present water of temperature $T_t^{(i)}$ with that of the street of temperature $T_t^{(s)}$, while heat is transferred through the pipe to the surrounding colder ground proportional to the temperature difference $T_t^{(i)} - T_g$.

in houses connected to the same street, but there is no information on the street temperature. The street temperature must in principle be greater or equal to the highest observed house temperature, although estimated values of it can be lower when accounting for measurement uncertainty. The actual behaviour of the street temperature is determined by the plant output and the network characteristics. However, in order to keep the specifications simple, these detailed descriptions are considered to be outside the scope of the current model, and, in the absence of any "real" drivers for the street temperature, its dynamics are simply modelled by a random walk. The model for the forwarded water temperature at the individual houses $T_t^{(i)}$ is that presented in Eq. (4). These equations are combined to obtain the considered system of equations

$$\mathrm{d}T_t^{(\mathrm{i})} = C_i^{-1}\left(c_\mathrm{v} Q_t^{(\mathrm{i})}(T_t^{(\mathrm{s})} - T_t^{(\mathrm{i})}) - R_i^{-1}(T_t^{(\mathrm{i})} - T_t^{(\mathrm{g})})\right)\mathrm{d}t + \sigma_\mathrm{i}\,\mathrm{d}\omega_t^{(\mathrm{i})}, \tag{9a}$$

$$\mathrm{d}T_t^{(\mathrm{s})} = \sigma_\mathrm{s}\,\mathrm{d}\omega_t^{(\mathrm{s})}, \tag{9b}$$

which is illustrated in Figure 2. The model takes as inputs the mass flow rate $Q_t^{(\mathrm{i})}$ [kg/s], and the ground temperature T_g [°C]. The constant $c_\mathrm{v} \approx 4.186$ [kJ/(kgK)] is the specific heat capacity of water. In this formulation, three parameters are associated with each house - namely the thermal capacity C_i [J/K], the thermal resistance R_i [K/W], and the diffusion scaling σ_i [°C/$\sqrt{\mathrm{s}}$]. The individual house temperatures $T_t^{(\mathrm{i})}$ are observed directly by the measurement device $T_t^{(\mathrm{i,obs})}$, but it is assumed that the uncertainty of these measurements can be approximated by Gaussian noise, which gives rise to the following observation equation

$$T_t^{(\mathrm{i,obs})} = T_t^{(\mathrm{i})} + e_t^{(\mathrm{i})}, \quad e_t^{(\mathrm{i})} \sim N(0, V_t^i), \tag{10}$$

In order to establish a flow-dependent variance construction, the logistics function S is used

$$V_t^i = \mathbb{V}\left[e_t^{(\mathrm{i})}\right] = \sigma_{\mathrm{obs}}^2 + S(-Q_\mathrm{i}(t) + b) = \sigma_{\mathrm{obs}}^2 + \frac{K}{1 + \mathrm{e}^{a(Q_\mathrm{i}(t)-b)}}, \tag{11}$$

which acts as a smooth approximation to the Heaviside step function $H(Q) = \mathbf{1}_{Q<b}$, as illustrated in Figure 3. The purpose of this variance construction is to decrease the weight of observations that are gathered under low flow conditions, where the observations obviously contain very limited information about the street temperature - as is illustrated in Figure 6 and Figure 7. This construction was seen to be necessary in order to prevent the street temperature estimates from decreasing according to the decreasing house temperatures during such periods of low flow, i.e. no energy consumption. The parameter b in Eq. (11) was chosen based on inspection of flow rate histograms for the buildings in that neighbourhood. An individual neighbourhood assessment must be performed, since the lower flow rate threshold at which rapid temperature decreases occur varies. A typical bi-modal flow histogram from one particular house is shown in Figure 4. The red dashed line marks the threshold value of $b = 15$ below which a given observation whose flow satisfies $Q \leq b$ will have its variance increased. The magnitude of this increase lies in the range $k \in [K/2, K)$. Figure 5 demonstrates the correlation between temperature and flow rate, by marking with black dots the observations with flow rates below the set threshold. In this study, the parameters of the logistics function will be set to; $K = 10000$ (maximum value), $a = 0.25$ (steepness of the curve), and $b = 15$ (midpoint of the function).

The house-specific thermal parameters C and R summarise the thermal properties of the transportation system between street and house. The *system* here should be interpreted as

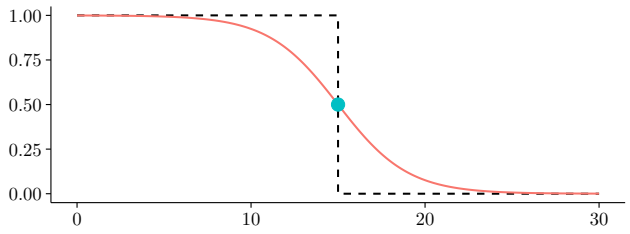

Figure 3: A plot of the logistics function in (11) with parameters $K = 1$, $a = 0.5$ and $b = 15$ (as used in this work) which approximates to a step-function at b. The variance increases towards K for $Q \to 0$ and vanishes for $Q \to \infty$.

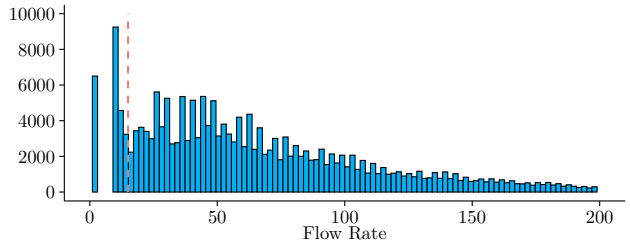

Figure 4: A histogram of the flow rates for a specific house which illustrates the location of the cut-off level as determined by the value of b in (11).

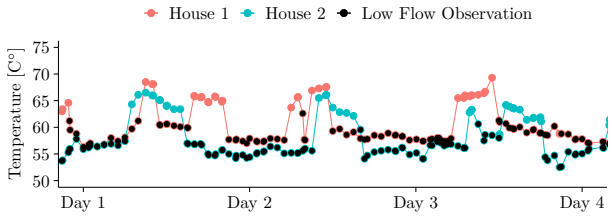

Figure 5: Observations of temperatures for two houses with indication (black dots) of the observations where the flow rate is less than $Q \leq 20L/h$ in correspondence with the cut-off level seen in Figure 4 and the associated curve in Figure 3.

an idealisation of the entire pipe installations and the different environments that those are embedded in. The underlying assumption that the potentially huge variety of such installations can be described by the simple model in Eq. (9a)-(9b) enables a crude way of analysing houses for potential issues, like detecting poorly insulated pipes or potential leakages.

2.3. Parameter estimation

The parameter estimation is carried out using the maximum likelihood approach. In this paper, two slightly different computational methods are applied, both of which rely on the R-package Template Model Builder (TMB) [31]. The first method applies the intended methodology of TMB which is that of a generalized mixed-effects model, while the latter method employs the discrete Kalman filter. The filter is also implemented with TMB simply to draw on the package's employment of automatic differentiation which yields the gradient (and also the hessian with respect to the parameters) of the likelihood function which drastically reduces computation times. A comparison between these methods will be presented here in terms of speed, parameter estimates and uncertainty.

In the mixed effects formulation, TMB employs the Laplace approximation to integrate out the states Eq. (9a)-(9b) which are considered random effects in this framework. The TMB procedure; the marginal likelihood $L(\theta)$ of the fixed effects θ by integrating out the random effects U of the joint likelihood $L(U, \theta)$ using the Laplace approximation. The marginal likelihood is

$$L(\theta) = \int L(U, \theta)\, \mathrm{d}U, \tag{12}$$

and the log-likelihood $\mathcal{L}$ is therefore

$$\mathcal{L}(\theta) = \log \int \exp \mathcal{L}(U, \theta)\, \mathrm{d}U. \tag{13}$$

A second-order Taylor expansion around the random effects maximum $\hat{U}$

$$\hat{U} = \arg\max_{U} \mathcal{L}(U, \theta), \tag{14}$$

of the joint log-likelihood $\mathcal{L}$ is

$$\mathcal{L}(U, \theta) \approx \mathcal{L}(\hat{U}, \theta) + \frac{1}{2}(U - \hat{U})H(\hat{U}, \theta)(U - \hat{U})^T. \tag{15}$$

where $H(\hat{U}, \theta)$ is the hessian. The integral in Eq. (13) can then be approximated by inserting Eq. (15), which evaluates to unity (after having corrected for the missing factor of the normal distribution), since the integrand is then a multivariate normal density with mean $\hat{U}$ and covariance matrix H. The approximation is only exact if the joint distribution is also Gaussian. The marginal likelihood can now be computed as

$$\mathcal{L}(\theta) \approx \mathcal{L}(\theta, \hat{U}) + \frac{(M+1)N}{2} \log 2\pi - \frac{1}{2} \log \det H(\hat{U}, \theta), \tag{16}$$

where M is the number of smart meters and N is the number of observations. For additional information about the Laplace approximation and its properties, see e.g. Madsen and Thyregod [32].

In practice, when employing the TMB framework, the user writes the negative log-likelihood as a C++ file which is then compiled and imported as a function into R. In the present case, the likelihood contributions from a stochastic differential equation system come from 1) the (hidden) state transitions and 2) the state observations. The former contribution is given by

$$U_\mathrm{t} \sim \mathcal{N}(\mathcal{M}_\mathrm{t}, P_\mathrm{t}), \tag{17}$$

where $U_\mathrm{t} = \begin{bmatrix} T_t^{(1)} & T_t^{(2)} & \dots & T_t^{(\mathrm{M})} & T_t^{(\mathrm{s})} \end{bmatrix}^T$ is the random effects state vector with mean $\mathcal{M}_\mathrm{t}$ and covariance P_t (see Eq. (30)) which directly depend on $U_{\mathrm{t}-1}$ through the one-step transition density. The contribution from the observations is given by

$$Y_\mathrm{t} \sim \mathcal{N}(U_\mathrm{t}, V_t), \tag{18}$$

where V_t is given by (11).

In the case of the Kalman filter there is only one contribution to the likelihood, due to the state updating scheme. The posterior state $X_{\mathrm{t}|\mathrm{t}}$ and covariance $P_{\mathrm{t}|\mathrm{t}}$ estimates are obtained from updating the prior estimates $X_{\mathrm{t}|\mathrm{t}-1}$ and $P_{\mathrm{t}|\mathrm{t}-1}$ once new information Y_t becomes available. The updating scheme is

$$X_{\mathrm{t}|\mathrm{t}} = X_{\mathrm{t}|\mathrm{t}-1} + K_\mathrm{t} e_\mathrm{t}, \tag{19}$$

$$P_{\mathrm{t}|\mathrm{t}} = (I - K_\mathrm{t} H) P_{\mathrm{t}|\mathrm{t}-1} (I - K_\mathrm{t} H)^T + K_\mathrm{t} V_t K_t^\mathrm{T}, \tag{20}$$

where K_t is the associated Kalman gain given by

$$K_\mathrm{t} = P_{\mathrm{t}|\mathrm{t}-1} H^T {\Sigma_{\mathrm{Y,t}}}^{-1}, \tag{21}$$

$$\Sigma_{\mathrm{Y,t}} = H P_{\mathrm{t}|\mathrm{t}-1} H^T + V_t, \tag{22}$$

with innovation

$$e_\mathrm{t} = Y_\mathrm{k} - h(X_{\mathrm{k}|\mathrm{k}-1}), \tag{23}$$

and with $H = \dfrac{\mathrm{d}h(u)}{\mathrm{d}u}$. The likelihood contribution arises from the innovation e_t and the covariance matrix $\Sigma_{\mathrm{Y,t}}$ as

$$\mathcal{L}(\theta)_\mathrm{t} = \frac{1}{2}\left[\log \det \Sigma_{\mathrm{Y,t}} + d_\mathrm{t} \log 2\pi + e_\mathrm{t} \Sigma_{\mathrm{Y,t}} e_\mathrm{t}^T\right], \tag{24}$$

where d_t is the number of available observations at time t.

2.4. Theoretical Resistance-Capacitance Parameters in Pipe

The theoretical values of the Resistance-Capacitance (RC) model of a pipe can be computed based on physical knowledge of the pipe. The thermal capacitance of water in the service pipe is computed as

$$C = Ac_{\mathrm{v}}\rho. \tag{25}$$

The pipes from the distribution pipe to the substation are usually DN25 pipes. If it is assumed that the inner diameter of the pipe is 0.0273 $[m]$ then $C = 2.45$ $[kJ/mK]$.

The resistance of the pipe is

$$R = \frac{1}{2\pi\lambda_{\mathrm{p}}} \ln \frac{d_{\mathrm{o}} + 2x_{\mathrm{a}}}{d_{\mathrm{i}}}, \tag{26}$$

with the thermal conductivity of insulation of the pipe as, $\lambda_{\mathrm{p}} = 0.028$ $[W/mK]$, outer diameter $d_{\mathrm{o}} = 0.0337$ $[m]$, and insulation thickness $x_{\mathrm{a}} = 0.0182$ $[m]$, which gives a resistance of $R = 4.16$ $[mK/W]$.

The theoretical estimates of the RC parameters will be used to compare them to the estimated RC parameters.

2.5. First and second-order moments

The likelihood contributions of the stochastic differential equation require computing the expectation and variance of the one-step predictions regardless of the use of the two methods. In the case of `TMB`, this amounts to computing $\mathcal{M}_{\mathrm{t}}$ and P_{t} and similarly for the Kalman filter $U_{\mathrm{t|t-1}}$ and $P_{\mathrm{t|t-1}}$. This will generally require integrating the first and second-order moments of the SDE forward in time, and while that is possible using standard ordinary differential equation (ODE) solvers, such an approach introduces both the choice of integral method and a time-step Δt. The accuracy of the integration will depend on both, but since integration methods are standard (e.g. a 4^{th} order Runge-Kutta method), the time-step will introduce a trade-off between computing time and integration accuracy. However, this can be avoided altogether if the moment equations can be solved analytically, and that is possible for the presented model. In particular, the system in Eq. (9a)-(9b) is linear but parameters are time-dependent due to the flow rate, and an analytical solution is therefore not tractable. A common solution is to impose a zero-order hold (ZOH) condition on the time-dependence (i.e. $Q_i(t) = Q_{i,k}$ for $t \in [t_k\,,\, t_{k+1}]$) such that the system becomes piece-wise linear and time invariant. Imposing this on the system in Eq. (9a)-(9b) for the M houses yields the matrix-vector form

$$\mathrm{d}\begin{bmatrix} T_t \\ T_t^{(s)} \end{bmatrix} = \left(\underbrace{\begin{bmatrix} A_1 & S \\ 0 & 0 \end{bmatrix}}_{A} \begin{bmatrix} T_t \\ T_t^{(s)} \end{bmatrix} + \underbrace{\begin{bmatrix} B_1 \\ 0 \end{bmatrix}}_{B} T_g \right) \mathrm{d}t + \underbrace{\begin{bmatrix} \sigma & 0 \\ 0 & \sigma_s \end{bmatrix}}_{G} \begin{bmatrix} \mathrm{d}\omega_t \\ \mathrm{d}\omega_t^{(s)} \end{bmatrix} \tag{27}$$

$$= \left(A \begin{bmatrix} T_t \\ T_t^{(s)} \end{bmatrix} + B \right) \mathrm{d}t + G \begin{bmatrix} \mathrm{d}\omega_t \\ \mathrm{d}\omega_t^{(s)} \end{bmatrix}, \tag{28}$$

where the auxiliary matrices and vectors have been introduced,

$$A_1 = \begin{bmatrix} a_1 & 0 & 0 & 0 & 0 \\ 0 & a_2 & 0 & 0 & 0 \\ 0 & 0 & \ddots & 0 & 0 \\ 0 & 0 & 0 & a_{M-1} & 0 \\ 0 & 0 & 0 & 0 & a_M \end{bmatrix}, \qquad S = \begin{bmatrix} s_1 \\ s_2 \\ \vdots \\ s_{M-1} \\ s_M \end{bmatrix},$$

$$B_1 = \begin{bmatrix} b_1 \\ b_2 \\ \vdots \\ b_{M-1} \\ b_M \end{bmatrix}, \qquad \sigma = \begin{bmatrix} \sigma_1 & 0 & 0 & 0 & 0 \\ 0 & \sigma_2 & 0 & 0 & 0 \\ 0 & 0 & \ddots & 0 & 0 \\ 0 & 0 & 0 & \sigma_{M-1} & 0 \\ 0 & 0 & 0 & 0 & \sigma_M \end{bmatrix},$$

with the lowercase variables given as

$$a_{\mathrm{i}} = -(s_{\mathrm{i}} + b_{\mathrm{i}}), \qquad s_{\mathrm{i}} = C_{\mathrm{i}}^{-1} Q_{\mathrm{i,k}} c_{\mathrm{p}}, \qquad b_{\mathrm{i}} = (C_{\mathrm{i}} R_{\mathrm{i}})^{-1},$$

and where

$$T_t = \left[T_t^{(1)}, T_t^{(2)}, \ldots, T_t^{(M)} \right]^T, \quad \mathrm{d}\omega_t = \left[\mathrm{d}\omega_t^{(1)}, \mathrm{d}\omega_t^{(2)}, \ldots, \omega_t^{(M)} \right].$$

The solution to this $(M+1)$-dimensional system of stochastic differential equations from time $t = t_{\mathrm{k}}$ to $t = t_{\mathrm{k+1}}$ can be written using the Itô interpretation of the stochastic differential equation as

$$x_{\mathrm{k+1}} = \mathrm{e}^{A\Delta t} x_{\mathrm{k}} + \int_{t_{\mathrm{k}}}^{t_{\mathrm{k+1}}} \mathrm{e}^{A\Delta t_{\mathrm{s}}} B \,\mathrm{d}s + \int_{t_{\mathrm{k}}}^{t_{\mathrm{k+1}}} \mathrm{e}^{A\Delta t_{\mathrm{s}}} G \,\mathrm{d}\omega_{\mathrm{s}}, \tag{29}$$

using the time differences $\Delta t = t_{\mathrm{k+1}} - t_{\mathrm{k}}$ and $\Delta t_{\mathrm{s}} = t_{\mathrm{k+1}} - s$ and using the zero-order hold assumption across these time intervals. Only the transitional mean and covariance must be computed and these remain Gaussian by the linearity of the system. They are given by

$$\mathbb{E}[x_{\mathrm{k+1}}] = \hat{A}\,\mathbb{E}[x_{\mathrm{k}}] + \hat{B}, \tag{30a}$$

$$\mathbb{V}[x_{\mathrm{k+1}}] = \hat{A}\,\mathbb{V}[x_{\mathrm{k}}]\,\hat{A}^T + \hat{Q}, \tag{30b}$$

where the involved matrices are

$$\hat{A} = \mathrm{e}^{A\Delta t}, \tag{31}$$

$$\hat{B} = \int_{t_k}^{t_{k+1}} \mathrm{e}^{A\Delta t_s} B \,\mathrm{d}s, \tag{32}$$

$$\hat{Q} = \int_{t_k}^{t_{k+1}} \mathrm{e}^{A\Delta t_s} G G^T \mathrm{e}^{A^T \Delta t_s} \,\mathrm{d}s, \tag{33}$$

which can be calculated by computing the matrix exponential of the augmented matrices

$$\exp\left(\begin{bmatrix} A & B \\ 0 & 0 \end{bmatrix} \Delta t \right) = \begin{bmatrix} \hat{A} & \hat{B} \\ 0 & I \end{bmatrix}, \tag{34}$$

$$\exp\left(\begin{bmatrix} -A & GG^T \\ 0 & A^T \end{bmatrix} \Delta t \right) = \begin{bmatrix} V_{11} & V_{12} \\ 0 & V_{22} \end{bmatrix}, \tag{35}$$

and subsequently extracting the variance as $\hat{Q} = V_{22}^T V_{12}$ [33]. The elements of these matrices are

$$\hat{Q}_{i,i} = \frac{s_i^2}{2a_i^3}\sigma_s^2\left[e^{2a_i\Delta t} - 4e^{a_i\Delta t} + 2a_i\Delta t + 3\right] + \frac{1}{2a_i}\sigma_i^2\left[e^{2a_i\Delta t} - 1\right], \tag{36}$$

$$\hat{Q}_{i,j} = \frac{s_i s_j}{(a_i + a_j)a_i^2 a_j^2}\sigma_s^2 \cdot \Big(e^{(a_i+a_j)\Delta t}a_i a_j - e^{a_i\Delta t}a_j(a_i + a_j) - e^{a_j\Delta t}a_i(a_i + a_j) + a_i^2(a_j\Delta t + 1) + (a_j^2\Delta t + a_j)a_i + a_j^2\Big), \tag{37}$$

$$\hat{Q}_{M+1,i} = \frac{s_i}{a_i^2}\sigma_s^2\left[e^{a_i\Delta t} - a_i\Delta t - 1\right], \tag{38}$$

for $i = 1, 2, ..., M$ and $j = i + 1, i + 2, ..., M$. The last diagonal element is $\hat{Q}_{M+1,M+1} = \sigma_s^2\Delta t$. The non-zero elements of $\hat{A}$ lie on the diagonal and the last column i.e.

$$\hat{A}_{i,i} = e^{a_i\Delta t}, \tag{39}$$

$$\hat{A}_{M+1,i} = \frac{s_i}{a_i}(e^{a_i\Delta t} - 1), \tag{40}$$

for $i = 1, 2, ..., M$. The last diagonal element is $\hat{A}_{M+1,M+1} = 1$. The first M entries of $\hat{B}$ are

$$\hat{B}_i = \frac{b_i}{a_i}(e^{a_i\Delta t} - 1)T_g, \tag{41}$$

for $i = 1, 2, ..., M$, and $B_{M+1} = 0$.

The likelihood computations can therefore be carried out without having to invoke costly integration techniques, simply by computing the one-step moments in Eq. (30) by calculating these elements necessary to directly construct $\hat{A}$, $\hat{B}$ and $\hat{Q}$. The source code for the implemented likelihood functions for the two methods is available through a GitHub repository[1].

3. Results

This section presents the results in terms of parameter estimation using the two proposed approaches, the `TMB` and Kalman filter to estimate the street temperature from smart meter data. Computation time and difference between the parameter estimations will be compared between the two approaches to investigate if one method is more advantageous. The results between the two different areas are also compared to validate the model's generalizability as the dynamics between the areas are different. The performance of the model in different seasons is also investigated. Finally, the challenges of the model are discussed. The data used in this paper is presented before the results from the proposed approaches are discussed.

[1] `https://lab.compute.dtu.dk/hgbe/smartmeters-kalmanfilter-tmb`

3.1. Data

The data used in this study was provided by the district heating utility in Brønderslev, *Brønderslev Forsyning*. The data consists of measurements from smart meters in individual buildings and measurement wells from two critical areas inside the Brønderslev district heating network. Only a subset of the smart meters is needed for the proposed method. The smart-meter measurements used in this study are from 30 different single-family houses: 15 from each area. They consist of time series of *supply temperature [°C]*, *return temperature [°C]*, *flow [L/h]*, *energy [kWh]*, *volume [m^3]* with associated *timestamps*. The measurement wells are located in the distribution network at street level before the houses and they measure the *forward temperature [°C]* and the *return temperature [°C]*. In order not to violate privacy and comply with the General Data Protection Regulation (GDPR), Brønderslev Forsyning made the smart-meter data anonymous before making it available for the study by not disclosing the location of the houses. Hence, the only information given in the data is which of the two areas the house belongs to, and that houses are grouped closely together for each area. Only metering data has been used to estimate the distribution network temperature at street level, whereas the temperature measurement from the wells has only been used to validate the results.

3.2. Measurements from smart meters and measurement Wells

In this paper, only temperature and flow measurements from the smart meters are used. The objective of this study is to estimate the supply temperature in the distribution pipe in the street that feeds into the houses where the smart meters are installed, without using the nearby measurement well. The temperature measurements from the well will only be used for model validation. The measured variables are denoted as follows:

$$\text{Network Temp. (Street)} : T_t^{\text{s}}, \quad t = 1, \ldots, N, \tag{42}$$

$$\text{Temp. at houses} : T_t^{(\text{i})}, \; t = 1, \ldots, N, \; i = 1, \ldots, M, \tag{43}$$

$$\text{Flow at house} : Q_t^{(\text{i})}, \; t = 1, \ldots, N, \; i = 1, \ldots, M, \tag{44}$$

where the subscript t is the time index (N number of observations), and the superscript i is a label of the smart meter (or house) number and M is the number of smart meters. The measurements from the smart meters in Area 1 and Area 2 were obtained from 1 July 2018 to 1 July 2019, and 1 January 2018 to 25 September 2020, respectively. The resolution of the smart-meter data is not fixed and the number of readings each day changes over time. The measurements from the wells that represents the temperatures in the streets from Area 1 and Area 2 are on two minutes resolution, and the periods are from 4 December 2018 to 8 January 2020 and 12 September 2019 to 12 September 2020.

Smart meters are usually located close to the substation inside the houses where they measure district-heating information. The hot water is delivered to the substation by a service pipe that is connected to the district heating distribution network at street level as shown in Figure 2. Measurement wells are located at the critical areas in the distribution network and

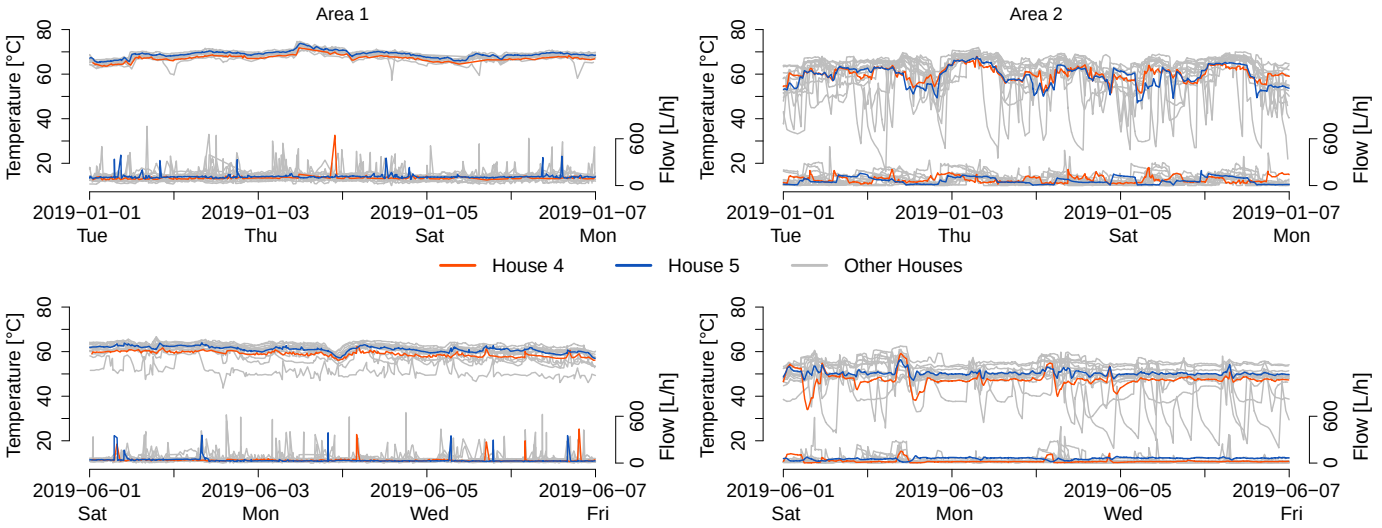

Figure 6: The plot demonstrates for an interval of one hour, a different number of readings at different timestamps when comparing smart meter data from each house. Two houses are shown with colours while the other houses are illustrated as grey. The two houses are highlighted with colour to illustrate the difference in the dynamics between the two areas.

are usually placed before the hot water enters the first house in the area at street level. Unfortunately, the measurements from the two wells do not overlap during a cold period (January) or warm period (summer) as Figure A.15 shows. Also, there are some errors with the measurements over a longer period for both areas.

Time series plots of the supply temperature and flow from the smart meters at the houses in both areas for a subset of the data are shown in Figure 6. Two random houses from both areas are visualised with coloured lines, while the rest are presented as grey. These two houses are selected to highlight the different dynamics between the areas and seasons. The dynamics are very different between the areas, while houses in the same area tend to show similar behaviour, i.e. similar houses with similar control strategies. There are also seasonal patterns, e.g. usually there is no need for space heating during summer so the temperature is lower during summer as the heat demand is significantly lower compared to cold periods, e.g. January. This can also be seen in Table 1, which lists the quantiles of the temperature and flow measurements in January and June. As expected, the temperatures are higher during January. Due to the significant difference between the areas, each area will be discussed separately in the following.

Area 1: demonstrates rather constant temperature, especially during cold periods, as shown in both Figure 6 and Table 1. Hence, there is always some consumption of heat, as the flow is almost never zero (see Q_5 for flow in Table 1). However, during a warmer period, less heat consumption is needed and therefore there is frequently no heat consumption.

Area 2: demonstrates a more dynamic operation where the flow is frequently turned on and off (see Figure 7 for an example). The hot water in the service pipe frequently loses heat to the surroundings, as it becomes still in the pipe for long periods when the flow is low. Hence, there is frequent temperature

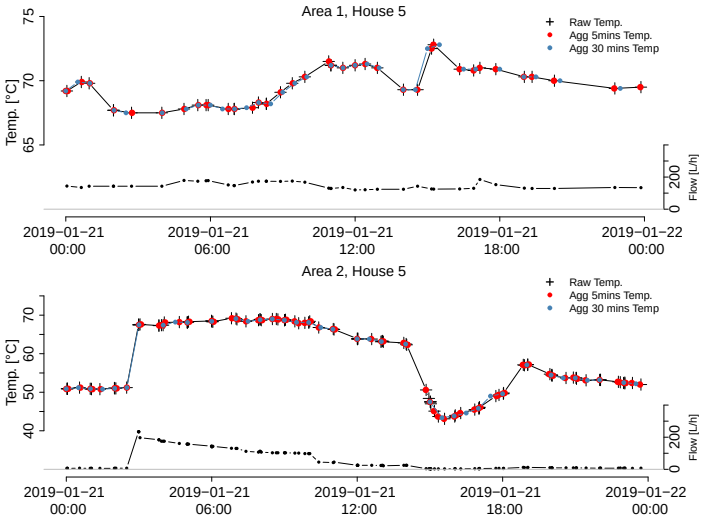

Figure 7: The plots visualise how the resolution of the aggregation can result in less information on the temperature dynamics of the pipe. The upper and lower plots show data from Area 1 and Area 2. The raw supply temperature and flow with two different aggregation resolutions, 5 minutes and 30 minutes are visualised together. The flow demonstrates the heating dynamics in more detail, i.e. how the temperature increases and decreases with the flow behaviour.

drops in Figure 6 and a large gap in the quantiles in Table 1, most notably during January.

Area 1 has a more constant behaviour, as both the temperature and flow do not change rapidly. Area 2 shows more of an on/off behaviour, i.e. turning on and off the heating over time. There is a significantly larger variability in Area 2 in the temperature readings. Area 2 might therefore be low-energy buildings where the substation is controlled intelligently compared to the traditional operation of substations where only the outside temperature is used. It can be assumed that Area 1 uses traditional operations and the buildings are likely to be older.

Table 1: The table presents flow and temperature measurement quantiles from both areas for the periods January and June in 2019. The differences in both flow rate and temperature between June (summer) and January (winter) are very noticeable for all shown quantiles.

		Q_5	Q_{25}	Q_{50}	Q_{75}	Q_{95}
		January				
Area 1	Flow [L/h]	53.36	87.67	107.5	133.09	182.5
	Temp [°C]	66.10	68.30	69.7	71.30	72.9
Area 2	Flow [L/h]	3.59	12.73	40.75	94.0	172.30
	Temp [°C]	43.34	55.70	61.10	65.1	68.70
		June				
Area 1	Flow [L/h]	13.00	21.00	30.0	46.09	83.0
	Temp [°C]	52.50	56.80	59.5	61.20	63.3
Area 2	Flow [L/h]	7.00	13.00	17.33	22.0	56.37
	Temp [°C]	42.80	48.90	50.30	51.8	56.60

Therefore, they usually have lower energy efficiency and need more heat to keep the indoor climate comfortable. The on/off control strategy impairs the quality of the temperature measurement from the smart meters. The water in the service pipe becomes still and therefore it does not give an accurate representation of the temperature in the distribution pipe. Consequently, this needs to be considered when estimating the temperature, to avoid including these periods where the signal becomes unreliable. The seasonal behaviour demonstrated in both areas occurs when the ambient air temperature has increased above a certain cutoff temperature where no space heating is required to feel comfortable inside, as space heating is highly correlated with the ambient air temperature. Hence, the quality of the temperature signal is also seasonal-dependent, as the amount of flow is affected by the desired space heating.

Bergsteinsson et al. [14] describe that each smart meter is unique, as the quality of the signal depends on the smart meter and the quality of the temperature signal is flow-dependent. Smart meters send instantaneous values with different resolutions, as is highlighted in Figure 8. The plot visualises the temperature readings from all houses in Area 1 over a one-hour period. Note that one house did not send any readings during this period. Some meters only send one value, while others send multiple. Therefore, the data from the smart meters needs to be resampled, as required by the method used to estimate the supply temperature.

In this study, the discrete-time Kalman filter is used to estimate the house and street temperatures. It is therefore convenient to have the smart-meter data resampled at the same time points. As demonstrated in Figure 8, the data does not have a fixed resolution. The purpose of this method is to estimate the temperature in the street distribution pipe from historical data, but it is not used to predict future temperature values. The smart-meter data is therefore aggregated with the desired resolution by rounding to the nearest time point. An appropriate resolution needs to be used such that as little as possible of the information vital for the analysis is lost. Figure 7 shows the time series of the raw supply temperature data and aggregates

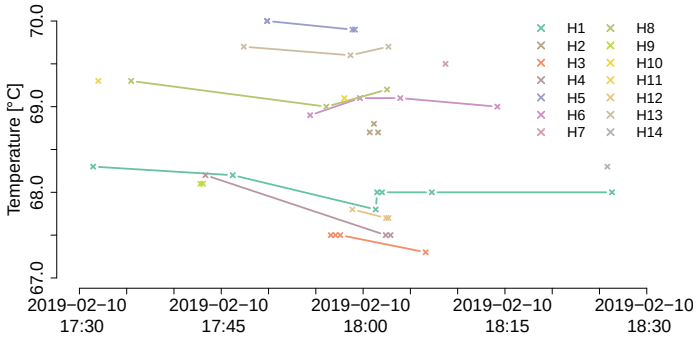

Figure 8: The plot demonstrates an interval of one hour. There is a different number of readings with different timestamps when comparing smart-meter data from each house. There is also one house missing in this period.

by rounding to 5 minutes resolution and 30 minutes. The flow is also plotted to highlight the heating dynamics of the houses and between the areas. Because of the on/off dynamics of the heating for House 5 in Area 2, the importance of the resolution can be seen more clearly than in Area 1. When the flow is shut off, the temperature starts to decrease and the rapid drop happens quite fast. Therefore, having all information on how the temperature decreases is important. Hence, using the correct aggregation resolution becomes essential to capture the heat dynamics of the service pipe. Investigating the difference between 5 minutes and 30 minutes in the drop, the 30 minutes resolution results in less information, as expected.

The smart-meter data in this study will be resampled to the same 5 minutes resolution and using one month from cold and warm periods to demonstrate the performance of the proposed method of estimating the distribution network temperature at street level. In Area 1, measurements from January 2019 and June 2019 will be used to estimate the parameters of the model. The estimated temperature will be validated by comparing it to the measured temperature from the well. The validation will be done in the same period as the parameters are estimated, as the well temperatures are not used in the model estimation and this method's purpose will only be used to estimate the temperature from the previous day when the smart-meter data arrives. Area 2 will use January 2020 and June 2020 to estimate the model parameters. However, during January it is not possible to validate the estimated temperature, as the measured temperature at the well for this period is wrong, as seen in Figure A.15.

3.3. Settings and procedure

The results presented in the following section were obtained using the Kalman Filter method, which proved the fastest and most robust. The parameter estimation is carried out using the Kalman Filter and the smoothed estimate of the states and uncertainties are then found with the mixed-effects method by supplying the estimated parameters. The parameter estimation is roughly ten times faster using the Kalman filter method, as shown in Table A.2. A linear fit to this data in the log-log domain shows that the computation time for the Kalman filter and mixed-effects method grows by $M^{2.6}$ and $M^{2.2}$ respectively. The faster computation time using the Kalman Filter is explained by

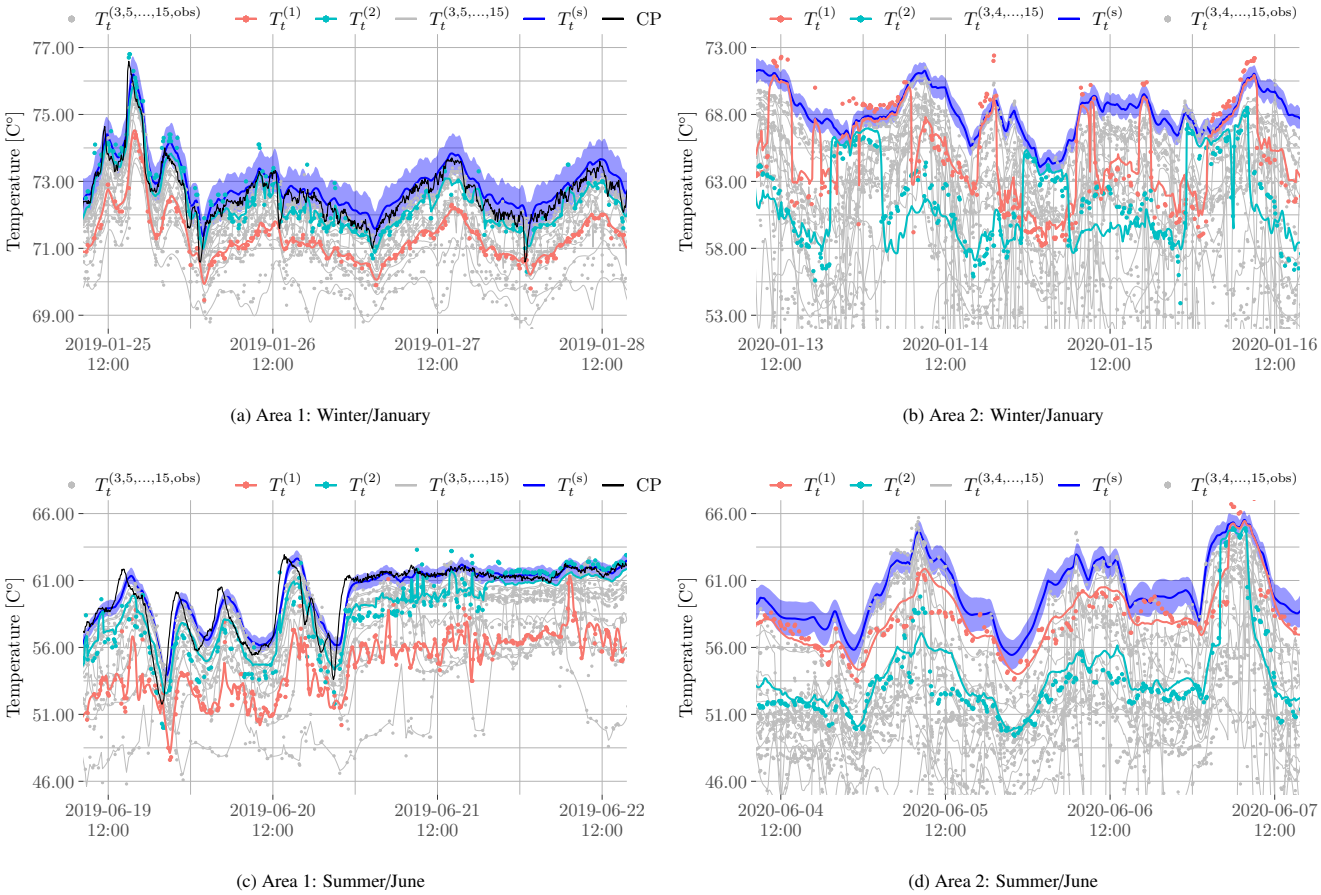

(a) Area 1: Winter/January

(b) Area 2: Winter/January

(c) Area 1: Summer/June

(d) Area 2: Summer/June

Figure 9: Plots showing the estimated street temperature and the measured house temperatures. The top and bottom plots show the estimation from winter and summer respectively, and plots on the left and right belong to Area 1 and Area 2 respectively. The measured street temperature that is the critical point (CP) temperature is plotted for Area 1 to demonstrate the performance of the proposed method. These measurements were not available for Area 2.

a much lower intercept. There was no drawback to the faster parameter estimation, since results were consistent regardless of the method used (Table A.3). A total of fifteen houses were used, and data from a one-month period containing 15×8353 observations of the house temperatures ($\approx 80\%$ are NA-values) and the same amount of house flow rate observations.

The following upper and lower parameter bounds were used

$$\theta^{(\mathrm{i})}_{\mathrm{upper}} = \left[C^{(\mathrm{i})}_{\mathrm{upper}}, R^{(\mathrm{i})}_{\mathrm{upper}}, \sigma^{(\mathrm{i})}_{\mathrm{upper}}\right] = [500, 1500, 2]\,, \tag{45}$$

$$\theta^{(\mathrm{i})}_{\mathrm{lower}} = \left[C^{(\mathrm{i})}_{\mathrm{lower}}, R^{(\mathrm{i})}_{\mathrm{lower}}, \sigma^{(\mathrm{i})}_{\mathrm{lower}}\right] = \left[1, 1, 10^{-5}\right], \tag{46}$$

for each house to ensure realistic estimates (see Eq. (25) and Eq. (26)). The parameter bounds were imposed naturally by introducing the inverse logit domain transformation

$$\theta = \left(\theta_{\mathrm{upper}} - \theta_{\mathrm{lower}}\right) S\left(\theta_{\mathrm{logit}}\right) + \theta_{\mathrm{lower}}, \tag{47}$$

where $S(x) = (\exp(-x) + 1)^{-1}$. The diffusion parameters σ_i were generally very uncertain, thus difficult to estimate. It was necessary to introduce a lower bound constraint to avoid numerical instability as a consequence of diffusion parameters tending towards zero. The observation variance behaved in a similar fashion and was fixed at $\sigma^2_{obs} = 1$ (which amounts to allowing temperature fluctuations on the order of ±2 degrees). The uncertainty is a rough approximation of the uncertainties informed by a specific smart meter manufacturer [34], and is reported to be 3-5 % (depending on the flow), which amounts to 1.5 - 2.5 °C assuming a temperature of 50 °C.

3.4. Empirical results

The smoothed-state reconstructions for the two areas are presented along the rows, with an estimate for a winter month (January) and a summer month (June) along the columns, in Figure 9. A pronounced difference can be seen between the dynamics in the two areas; Area 1 displays rather stable house temperatures relative to Area 2 where they surge up and down. This surging is driven by similar behaviour in the flow, as discussed in Section 3.1. Hence, the necessity of having the flow-dependent observation noise as shown in Eq. (11) is clear. One advantageous effect of the flow-dependent implementation is that the street temperature estimate is less prone to rapid changes during periods where all flows simultaneously decrease below the set threshold. This occurs because the estimated street temperature attempts to follow the fast dynamics of the houses, but such behaviour is contrary to its slower temperature dynamics. The implementation was therefore necessary in

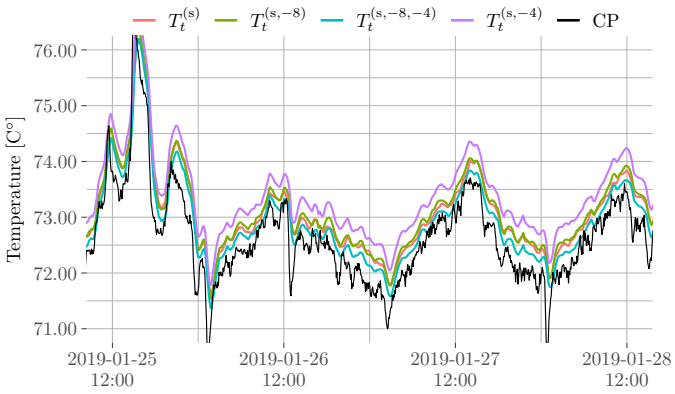

Figure 10: The estimated street temperature using all houses and after removing the two houses whose thermal resistance estimates hit the boundary $R_i = 1500$. The estimate is not significantly improved until the second house is removed.

order to achieve more realistic and accurate results. It can be argued that, even though increasing the number of houses would decrease the probability that all flows decrease below the set threshold simultaneously, the implementation remains valuable in order to appropriately penalize the impact of low-flow observations.

The model performance can be validated by comparing the estimated street temperature with the critical point measurements in Area 1. The street temperature predictions seem to have captured the overall trends and oscillations quite well, although a systematic bias is evident during summer, and in winter a slight temporal delay is seen. Comparing the two areas, the effect of having accurate information (due to the high flow) is evidently a smoother temperature curve for all houses and the street. The estimation in Area 2 displays faster variations as a result of the rapid variance changes and differences in the number of houses that provide information. In particular, inspecting the two highlighted temperatures ($T_t^{(1)}$ and $T_t^{(2)}$) emphasises that whenever information becomes available (and the flow is high) the individual houses surge upwards to some limit from where they determine the street temperature estimate. A primary challenge with the estimation in Area 2 is to prevent the slower dynamics of the street from being controlled by the much faster dynamics of the houses. This is a trade-off between measurement variance and obtaining what appears to be adequately slow changes in the street temperature estimates. The effect of the variance increase cannot be assessed due to the lack of critical point measurements in this area.

It was discovered that the estimation accuracy is sensitive to temperature observations that are much higher than expected, e.g. from poorly calibrated sensors. The challenge of detecting such problematic houses is difficult to solve in general. These houses are identified by inspection of the thermal resistances R_i because they tend to hit the upper bound. The modeller should be aware of the possibility that such houses have a controlling influence on the street temperature. In this particular case, the thermal resistances of both House 8 and House 4 were converging to the boundary, although only the former had observations that were much above the critical point measurements. The effect of omitting these two houses from the estimation was investigated after omitting none, one or both. The estimated street temperatures are compared to the critical point temperatures in Figure 10, and the associated mean average errors (MAEs) are

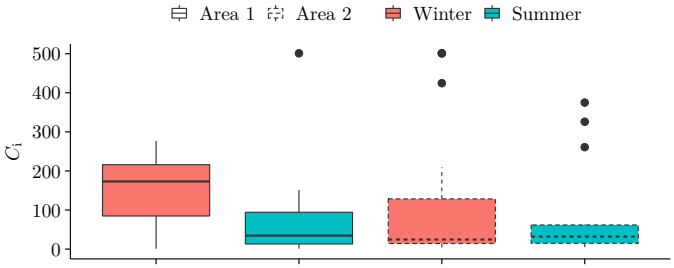

Figure 11: Boxplots of estimated thermal capacities C for the two areas and two months. The median values are clearly largest in Area 1 during the winter month but a few outliers are seen in Area 2 regardless of season.

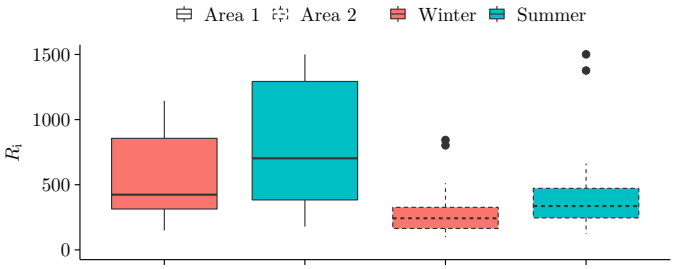

Figure 12: Boxplots of the logarithmic value of the thermal resistance R. There is a clear difference between the two areas, and generally speaking the interquartile range is larger during the summer.

$$\begin{aligned} \text{None}: &\quad \text{MAE} = 0.46 \\ \text{House 4}: &\quad \text{MAE} = 0.76 \\ \text{House 8}: &\quad \text{MAE} = 0.50 \\ \text{House 4 \& 8}: &\quad \text{MAE} = 0.32 \end{aligned}$$

This shows that the MAE increases (0.3/65%) after removing House 4, increases slightly (0.04/8%) after removing House 8 and decreases (0.14/30%) after removing both House 4 and House 8. The surprising outcome in the former case was further investigated and apparently it occurs because only a single house will act as a controller of the street temperature. In this case, House 4 dominates the street temperature estimation, so if removed House 8 will start to dominate the temperature estimation which has significantly higher observations and thus pulls the estimated street temperature up towards itself. In contrast, there is no effect from omitting House 8, because House 4 remains in control.

3.5. Thermal parameters and house dynamics

A statistical overview of the obtained values of C_i and R_i is provided in Figure 11 and Figure 12, respectively. The former figure shows that the estimated winter capacities from Area 1 stand out from the others by having a significantly higher median value, and the parameters are generally more dispersed in

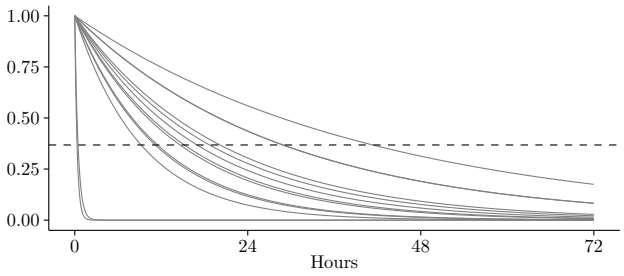

(a) Area 1 - January

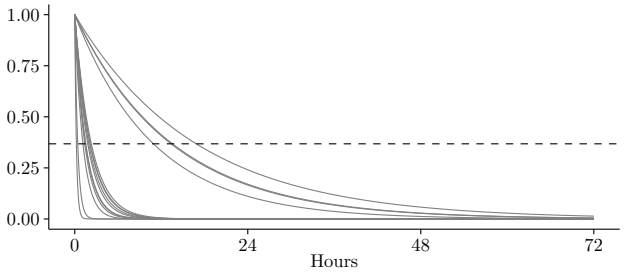

(b) Area 2 - January

Figure 13: The decaying exponential functions $\exp(-t/\tau)$ which depict how fast the temperature decreases from some initial set-point temperature down to the ground temperature T_g in the individual house systems. The dashed line shows when it has decayed to 37% of its initial value.

Area 1 regardless of period (as evident from the interquartile ranges of 131, 81, 114 and 47, respectively). Comparing areas shows that the dispersion is high during winter and low during summer. This conclusion is reversed when inspecting the estimated resistances, but could not be explained by strong correlations between R_i and C_i. The estimates and the estimated confidence intervals for winter and summer estimation are shown in Table A.4.

The stochastic differential equations describing each house give rise to *time constants* (τ) which can be interpreted as the duration after which the temperature has decreased by roughly 36%. This yields a straight-forward way to compare and characterise the thermal properties of the service pipe at each house in a single quantity, and this enables quick identification of e.g. bad insulation, leakage or out-dated installations. The time constants are the eigenvalues of the system matrix A in Eq. (27)

$$\tau_\mathrm{i} = \left(c_\mathrm{p} \frac{Q_\mathrm{i}}{C_\mathrm{i}} + \frac{1}{R_\mathrm{i} C_\mathrm{i}} \right)^{-1} \approx R_\mathrm{i} C_\mathrm{i}. \tag{48}$$

The approximation is only valid when the flow is low (approaching zero), since the former term then becomes very small. The time constants are illustrated as decaying exponential functions in Figure 13 for the two areas in the winter month. Evidently, certain houses can be identified as cooling down much more rapidly than the others; the houses where temperature converges to zero instantly. In Area 2, there seems to be two groups of houses: one which slowly converge to zero and another which cool down faster. There also seem to be two houses that cool down faster than other houses in the same area. This information could be used to investigate whether there is a problem with the service pipe.

3.6. Model validation

Model validation is always important, but in the presented case there are a few challenges of applying e.g. the autocorrelation function (ACF). One is that the data used in this study does not have a fixed resolution (Figure 8), and the more difficult one is that the model does not try to capture the behaviour at all times. When the flow is low, the dynamics of the temperature is not captured well by the model, as the dynamics will be more influenced by e.g. the placement of the smart meter inside the house and higher heat loss in the service pipe due to the slow transportation time. In the model, this is handled by increasing the observation variance. However, the 1-step residuals will show a very large autocorrelation in these situations (Figure A.16 and Figure A.17) leading to a very large lag-one autocorrelation for the overall data. Disregarding data when the flow is high results in smaller autocorrelation (Figure 14). Even with these reservations, it seems clear that, at least for some houses, the lag-one autocorrelation is high, indicating that some model deficiencies could be addressed in future work.

Also, as an approximation, the temperature distribution over the service pipe is is assumed to be uniform; i.e. the discretisation which results in Eq. (7). This assumption is only valid when flow is high. Hence, the temperature distribution is not uniform during situations when flow is low. Also, the transportation time of the water moving will therefore take longer. It can be calculated from

$$Q = \rho V = \frac{\rho A L}{t} \rightarrow t = \frac{\rho A L}{Q}, \tag{49}$$

where Q $[kg/s]$ is the mass flow rate, ρ $[kg/m^3]$ is the density of the water, V $[m^3/s]$ is the volumetric flow rate, A $[m^2]$ is the cross sectional area of the pipe, L $[L]$ is the length of the pipe, and t $[s]$ is the time. Some of the pipe properties needs to be assumed, for instance that the pipe is a DN25 with inner diameter of 0.0273 $[m]$, the length of the pipe is 10 $[m]$ and the flow is 50 $[kg/h]$, then the transportation time can be computed,

$$t = \frac{997[\mathrm{kg/m3}](\pi(0.0273[\mathrm{m}])^2/4)10[\mathrm{m}]}{50[\mathrm{kg/h}]/60[\mathrm{h/min}]} \approx 7 \text{ min}. \tag{50}$$

In this example, the transportation time is longer than the resolution time of the data used in this study, which indicates that the uniform temperature distribution is not valid and could lead to higher autocorrelation in the errors.

4. Discussion

This paper modelled the thermodynamics of a service pipe that delivers heat from a distribution pipe in a district heating network to consumers' substations. A partial differential equation of the thermodynamics in the service pipe was presented, and then approximated by a stochastic differential equation. Using this set of equations, an approach was established to obtain temperature feedback at arbitrary points in the district heating network by using measurements from a group of smart meters located inside single-family houses close to the chosen point. The network temperature is the temperature of the hot water

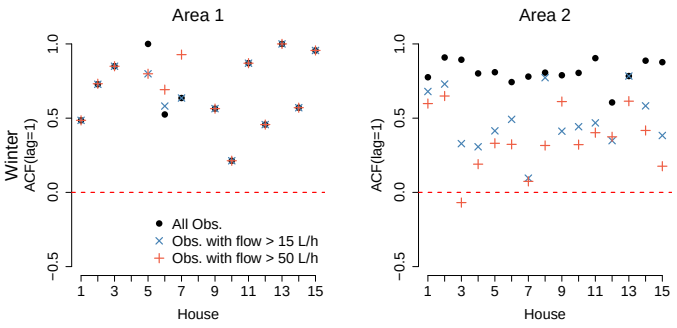

Figure 14: The figure shows the one lag value of the ACF from the prediction errors for each house for the two areas in separate plots. It highlights that, when the flow is low, the model does not capture the dynamics of the temperature adequately.

inside the distribution pipe before it delivers hot water to the consumers' service pipe. The street temperature was modelled as a random walk. A combined SDE model composed of the street temperature connected to all of the houses' service pipes through identical Resistance-Capacity SDEs was constructed. This automatically restricts the street pipe temperature to at least as high as the temperature measurements from the smart meters. It was shown that the presented methodology produced accurate results when compared to the measured network temperatures in a single area with fifteen houses. In another area of fifteen houses, it was shown how to deal with rapidly decreasing temperatures when the flow rates become too low, by increasing the uncertainty of these observations.

The parameter estimation was carried out using two different methods, namely a discrete Kalman filter and a mixed-effects method, both implemented as `C++` files and used with the R package `TMB`. The latter was the intended formulation using `TMB` while the former only took advantage of the automatic differentiation provided by `TMB`. The speed-up gained using `C++` together with automatic differentiation is substantial, and results showed that the computation time for the Kalman filter method was approximately ten times faster than that of the mixed-effect formulation. For this reason, parameter estimation was carried out using the Kalman filter, while smoothed-state estimates and variances were calculated using the mixed-effects formulation.

The smoothed-state estimates are presented here since these are more appropriate than one-step prediction estimates due to the in-sample use case. Thus, because the model is intended to assess the feedback from the network for computing a transfer function for the network characteristics, e.g. time delay and temperature loss in the system, the posterior state estimates that use all available information are appropriate. In a concrete scenario, one could imagine that the previous 24 hours of data are sent from the smart meters once per day. This information can then be used for temperature optimisation of the supply temperature at the plant with the aim of reducing the heating cost and heat losses in the system by lowering the supply temperature at the plant. The ability to establish temperature feedback in district heating networks opens many possibilities for utility companies to improve their operation. Most importantly, it gives the ability to use controllers for data-driven temperature optimization of the network. Temperature optimization reduces the heat loss and lower the needed network supply temperature, and thereby the operational cost of the heat production is reduced. It furthermore makes physical measurement wells in the network redundant, reducing costs, planning time, installations and maintenance. The feedback becomes flexible since any group of houses can be selected to establish a new network temperature. This is highly beneficial for the utility company, since the location where the highest temperature loss in the network will vary across time, due to deterioration of pipes, replacement of older pipes with newer ones, leakage and so forth. Finally, using smart meters can make multi-temperature zones inside the network more feasible, where more local heat sources can be added to the network. Thus, lowering the operating temperature in the grid using temperature optimisation and having more detailed information from the grid will give rise to more decentralised heat sources. For instance, heat sources such as heat pumps, waste heat from industries, and solar thermal collectors (with thermal storage systems) can be included to provide heat to consumers. This will also increase the efficiency of sector coupling with the electricity sector, where the renewable energy systems are better utilised.

The authors argue that the presented model is advantageous due to 1) its relative simplicity, 2) its ability to handle scenarios with a lot of information, and 3) its ability to assess thermal properties and the outlier detection possibilities that this enables. In particular, it was proposed that houses with resistance lying on the upper boundary should be omitted from the analysis. A few issues are, however, worth mentioning: Firstly, the choice of a random walk to estimate the street temperature generally creates poor conditions for long-term forecasting because its variance increases linearly with time. While this is of no concern here due to the in-sample model purposes, should one use the model for predictions it is crucial that the uncertainty assessment is corrected. Therefore, in order to improve the forecastability of the proposed method, it would be necessary to replace the random walk using some model. For instance, the Ornstein-Uhlenbeck process could be used, although that would require some assumptions or knowledge about the parameters that will have to be derived from the forwarded temperature at the district heating facility. A second issue is the assumption that the temperature distribution over the service pipe is uniform. however, this assumption is not valid when the flow is low as the flow would also not be constant through the service pipe. It also depends on the properties of the pipe, e.g. the diameter and length. One suggestion for model extension would be to increase the model order by dividing the pipe into multiple segments to model the temperature distribution over the pipe. This extension could lower the autocorrelation of the prediction errors.

5. Conclusion

This paper has demonstrated how smart-meter data can be used to improve the operation of a district heating network by

establishing temperature feedback of the network. Temperature feedback is highly valuable for temperature control of the network. Simplified descriptions using stochastic differential equations are formulated from considering partial differential equations that describe the thermodynamics of the hot water in the service pipe from the distribution pipe to consumer substations. A random walk is used to model the temperature variations in the distribution pipe that is connected to all houses. The street temperature and thermal parameters of the model are estimated by minimizing the negative log-likelihood function using a discrete Kalman filter. Smoothed state estimates are subsequently computed using the mixed-effects formulation in `TMB`. The results show that the proposed method can mimic the measured street temperature accurately when compared to observations. It is important that houses whose thermal parameters hit the upper boundary are removed to reduce the estimation bias in the street temperature. The estimation procedure entails that smart-meter data arrives daily and contains measurements with same resolution from the past 24 hours. The procedure uses this information to estimate the network temperature for the past 24 hours. The proposed method further allows for potential identification of houses with e.g., bad insulation or leakages. They can be identified by inspecting whether the estimated thermal resistance, capacity and resulting time constant of a particular house, lie outside of the expected range.

The possibilities for future work for using smart meters to enhance the operation of district heating networks are endless. Hence, the utilities have become very data-rich and can use this information to learn how the network is performing. Utilising this opportunity will help district heating to become more energy efficient in the transformation to RES and increase its flexibility. The extension of the proposed method in this paper would be to improve the forecasting ability and validate its potential for on-line control of supply temperature at the plant. In future work, the authors also aim to investigate model extensions e.g., by dividing the service pipe into multiple segments. Furthermore, a simulation study could be beneficial to investigate the estimation of R and C parameters of the service pipe.

6. Acknowledgement

This work is funded by Innovation Fund Denmark through the projects HEAT 4.0 (8090-00046B), CITIES (1305-00027B), Flexible Energy Denmark (8090-00069B), TOP-UP (9045-00017B), and Region H through the IDASC project (18012745). Finally, the project is partly funded under the Norwegian FME-ZEN project financed by the ZEN partners and the Research Council of Norway (257660).

We thank the district heating utility *Brønderslev Forsyning* for their support and for supplying data & information for this study. We would also like to thank Torben Skov Nielsen and ENFOR for providing their support and knowledge about the temperature optimisation of a district heating network.

Appendix A.

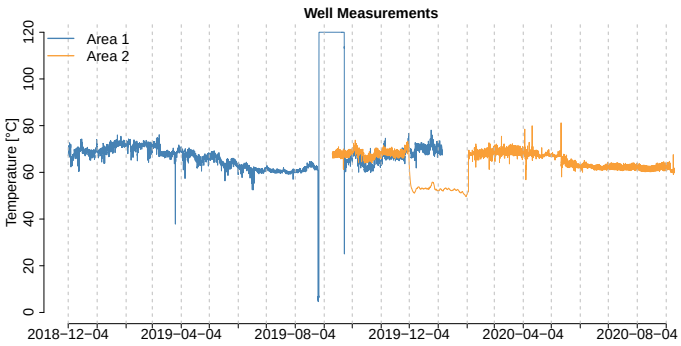

Figure A.15: The figure shows the time series of the available temperature measured at the wells for both areas. The wells are located before the houses. Notice, that there are some problems with the measurements, e.g. during October 2019 for Area 1 and during December 2019 and January 2020 for Area 2.

Table A.2: The approximate computation times (seconds) for parameter estimation using either the Kalman Filter or the mixed-effects method as a function of the number of houses M. The dimensions of the system are $M + 1$ and the number of parameters are $3M + 1$. The estimation is based on one month of data with a sampling time of 5 minutes, which corresponds to 8350 temperature observations for each house, although roughly $\approx$ 80% are missing values (NA-values).

Number of Houses [M]	Kalman	Mixed-Effects	Ratio
2	1	29	1:22
3	4	59	1:15
4	6	101	1:18
5	11	140	1:13
6	16	206	1:13
7	26	282	1:11
8	33	416	1:13
9	54	518	1:10
10	69	685	1:10
11	86	829	1:10
12	118	1157	1:10
13	143	2112	1:15
14	168	2273	1:14
15	198	2266	1:11

Table A.3: The thermal parameter estimates for 10 houses using the Kalman or the mixed-effects model formulations. The estimates are seen to be identical for all practical purposes.

Parameter	Kalman	TMB	Difference
C_1	153	155	2
C_2	11	12	1
C_3	254	258	3
C_4	29	30	1
C_5	355	358	3
C_6	472	478	6
C_7	5	5	0
C_8	8	14	5
C_9	335	341	6
C_{10}	238	240	2
R_1	245	245	0
R_2	648	648	0
R_3	219	219	0
R_4	1085	1084	1
R_5	185	185	0
R_6	125	125	0
R_7	619	619	0
R_8	1501	1501	0
R_9	266	266	0
R_{10}	450	450	0

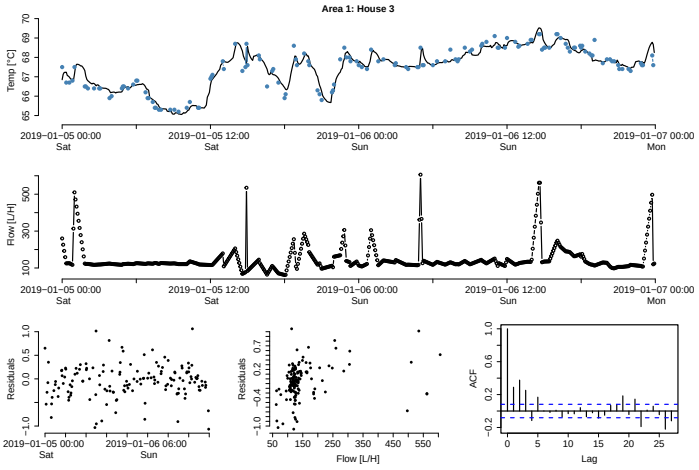

Figure A.16: The figure displays difficulties in validating the model by doing residuals analysis on the one-step predictions in Area 1. The top plot shows the measured temperature (black solid line) and the one-step prediction (blue points). The flow for the same period is illustrated in the second top plot. In the bottom are three plots that show the residuals analysis of the one-step prediction errors. The first plot on the bottom shows the residuals over time, the next plot shows the residuals versus the flow, and the last plot shows the autocorrelation function (ACF) of the residuals.

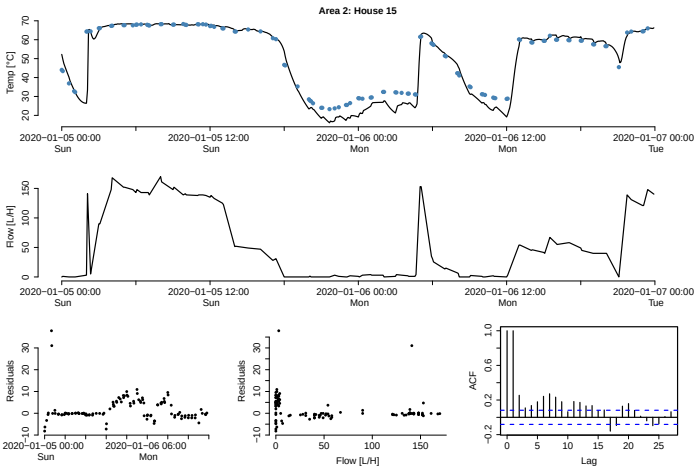

Figure A.17: The figure displays difficulties in validating the model by doing residuals analysis on the one-step predictions in Area 2. The top plot shows the measured temperature (black solid line) and the one-step prediction (blue points). The flow for the same period is illustrated in the second top plot. In the bottom are three plots that show the residuals analysis of the one-step prediction errors. The first plot on the bottom shows the residuals over time, the next plot shows the residuals versus the flow, and the last plot shows the autocorrelation function (ACF) of the residuals.

Table A.4: Parameters and their confidence interval for the winter and summer estimation for both areas. For Area 1, there is no CI as the hessian could not be computed for the summer period.

	Area1		Area2	
Parameter	Winter Estimate (CI)	Summer Estimate (CI)	Winter Estimate (CI)	Summer Estimate (CI)
C_1	101.9 (95.56, 108.55)	23.18 (,)	1 (1, 501)	63.31 (36.46, 105.9)
C_2	1 (1, 501)	1 (,)	14.92 (13.23, 16.83)	28.31 (19.14, 41.72)
C_3	173.1 (165.92, 180.43)	34.48 (,)	33.3 (33.16, 33.45)	38.28 (37.87, 38.69)
C_4	(,)	(,)	209.22 (201.04, 217.5)	44.98 (44.24, 45.73)
C_5	245.16 (232.6, 257.76)	94.28 (,)	13.69 (13.56, 13.83)	3 (1.01, 215.3)
C_6	276.88 (256.16, 297.25)	501 (,)	6.9 (6.37, 7.47)	260.84 (238.7, 282.82)
C_7	1.91 (1, 482.67)	1 (,)	26.77 (26.65, 26.89)	32.02 (31.8, 32.25)
C_8	(,)	(,)	21.09 (20.58, 21.6)	11.68 (6.21, 22.63)
C_9	245.27 (234.49, 256.07)	92.26 (,)	6.53 (6.12, 6.97)	16.62 (7.25, 38.95)
C_{10}	188.13 (175.59, 201.01)	103.34 (,)	424.55 (386.15, 451.75)	325.95 (316.36, 335.29)
C_{11}	108.01 (103.37, 112.79)	61.95 (,)	501 (1, 501)	374.99 (356.71, 391.66)
C_{12}	215.99 (207.73, 224.32)	12.88 (,)	501 (1, 501)	60.27 (58.5, 62.09)
C_{13}	55.7 (47.37, 65.31)	13.16 (,)	23.83 (22.2, 25.58)	10.44 (8.32, 13.17)
C_{14}	84.73 (74.02, 96.66)	31.55 (,)	24.55 (24.32, 24.79)	25.76 (23.5, 28.24)
C_{15}	214.41 (210.22, 218.61)	151.78 (,)	47.72 (47.37, 48.06)	13.17 (9.27, 18.85)
R_1	325.13 (324.43, 325.82)	376.67 (,)	802.19 (796.28, 808.1)	1501 (1, 1501)
R_2	1144.12 (1101.73, 1183.31)	1501 (,)	334.75 (334.42, 335.08)	505.99 (504.8, 507.19)
R_3	316.06 (315.02, 317.1)	413.17 (,)	248.26 (248.06, 248.47)	369.16 (367.44, 370.89)
R_4	(,)	(,)	187.79 (187.53, 188.05)	295.92 (295.72, 296.12)
R_5	249.1 (248.48, 249.72)	383.63 (,)	511.81 (511.11, 512.52)	662.55 (659.75, 665.35)
R_6	149.37 (149.27, 149.46)	178.89 (,)	214.6 (214.54, 214.66)	336.28 (335.31, 337.24)
R_7	863.29 (858.63, 867.94)	1344.51 (,)	242.79 (242.52, 243.05)	315.6 (315.07, 316.13)
R_8	(,)	(,)	255.48 (255.43, 255.54)	359.45 (358.9, 359.99)
R_9	423.67 (420.63, 426.72)	702.91 (,)	843.58 (840.51, 846.64)	1376.71 (982.53, 1478.17)
R_{10}	790.97 (768.8, 813.06)	1314.43 (,)	113.28 (113.12, 113.43)	121.07 (121.03, 121.12)
R_{11}	962.96 (921.01, 1003.51)	911.85 (,)	121.09 (120.94, 121.25)	126.53 (126.49, 126.58)
R_{12}	243.45 (243.17, 243.73)	402.5 (,)	97.4 (97.26, 97.53)	228.06 (227.93, 228.19)
R_{13}	722.03 (711.75, 732.32)	917.09 (,)	318.39 (318.29, 318.49)	438.79 (437.89, 439.69)
R_{14}	856.07 (822.25, 889.46)	1292.8 (,)	166.48 (166.43, 166.52)	236.79 (236.63, 236.96)
R_{15}	313.96 (313.48, 314.44)	371.96 (,)	162.23 (162.13, 162.33)	253.68 (253.45, 253.91)
σ_1	1.005e-05 (1.000e-05, 2.000e+00)	2.908e-02 (,)	2.456e-01 (2.363e-01, 2.553e-01)	1.014e-05 (1.000e-05, 2.000e+00)
σ_2	1.123e-05 (1.000e-05, 2.000e+00)	1.000e-05 (,)	2.406e-02 (1.911e-02, 3.027e-02)	1.202e-05 (1.000e-05, 2.000e+00)
σ_3	1.001e-05 (1.000e-05, 2.000e+00)	4.019e-02 (,)	8.904e-02 (8.872e-02, 8.937e-02)	1.083e-01 (1.076e-01, 1.090e-01)
σ_4	(,)	(,)	4.869e-02 (4.851e-02, 4.886e-02)	1.151e-05 (1.000e-05, 2.000e+00)
σ_5	6.549e-03 (5.815e-03, 7.376e-03)	2.245e-02 (,)	4.023e-02 (3.938e-02, 4.109e-02)	8.395e-05 (1.000e-05, 2.000e+00)
σ_6	1.180e-05 (1.000e-05, 2.000e+00)	4.628e-02 (,)	7.117e-02 (6.957e-02, 7.279e-02)	1.981e-02 (1.950e-02, 2.012e-02)
σ_7	1.026e-05 (1.000e-05, 2.000e+00)	1.445e-05 (,)	1.117e-01 (1.113e-01, 1.120e-01)	6.177e-02 (6.088e-02, 6.267e-02)
σ_8	(,)	(,)	1.001e-05 (1.000e-05, 2.000e+00)	1.232e-05 (1.000e-05, 2.000e+00)
σ_9	1.002e-05 (1.000e-05, 2.000e+00)	1.013e-05 (,)	1.006e-05 (1.000e-05, 2.000e+00)	1.158e-05 (1.000e-05, 2.000e+00)
σ_{10}	1.006e-05 (1.000e-05, 2.000e+00)	1.014e-05 (,)	4.317e-02 (4.309e-02, 4.325e-02)	1.545e-02 (1.539e-02, 1.551e-02)
σ_{11}	1.002e-05 (1.000e-05, 2.000e+00)	1.010e-05 (,)	3.288e-02 (3.283e-02, 3.294e-02)	1.149e-02 (1.139e-02, 1.159e-02)
σ_{12}	1.002e-05 (1.000e-05, 2.000e+00)	2.857e-02 (,)	4.341e-02 (4.336e-02, 4.345e-02)	1.170e-05 (1.000e-05, 2.000e+00)
σ_{13}	1.001e-05 (1.000e-05, 2.000e+00)	1.025e-05 (,)	1.000e-05 (1.000e-05, 2.000e+00)	1.409e-05 (1.000e-05, 2.000e+00)
σ_{14}	1.004e-05 (1.000e-05, 2.000e+00)	1.050e-05 (,)	6.046e-02 (5.997e-02, 6.096e-02)	1.149e-05 (1.000e-05, 2.000e+00)
σ_{15}	1.000e-05 (1.000e-05, 2.000e+00)	1.690e-02 (,)	9.479e-02 (9.456e-02, 9.501e-02)	1.208e-02 (1.759e-05, 1.813e+00)
σ_{16}	0.01 (0.01, 0.01)	0.01 (,)	1.542e-02 (1.540e-02, 1.544e-02)	1.479e-02 (1.471e-02, 1.487e-02)

References

[1] E. Guelpa, V. Verda, Thermal energy storage in district heating and cooling systems: A review, Applied Energy 252 (2019) 113474. doi:https://doi.org/10.1016/j.apenergy.2019.113474.

[2] B. Mathiesen, H. Lund, D. Connolly, H. Wenzel, P. Ostergaard, B. Möller, S. Nielsen, I. Ridjan, P. KarnOe, K. Sperling, F. Hvelplund, Smart energy systems for coherent 100% renewable energy and transport solutions, Applied Energy 145 (2015) 139–154. doi:10.1016/j.apenergy.2015.01.075.

[3] I. Blanco, D. Guericke, A. Andersen, H. Madsen, Operational planning and bidding for district heating systems with uncertain renewable energy production, Energies 11 (2018). doi:10.3390/en11123310.

[4] I. Blanco, A. Andersen, D. Guericke, H. Madsen, A novel bidding method for combined heat and power units in district heating systems, Energy Systems 11 (2019) 1137–1156. doi:10.1007/s12667-019-00352-0.

[5] H. Madsen, K. Sejling, H. T. Søgaard, O. P. Palsson, On flow and supply temperature control in district heating systems, Heat Recovery Systems and CHP 14 (1994) 613–620. doi:https://doi.org/10.1016/0890-4332(94)90031-0.

[6] L. Arvastson, Stochastic Modeling and Operational Optimization in District Heating Systems, Ph.D. thesis, Centre for Mathematical Sciences, Lund University, 2001.

[7] A. David, B. V. Mathiesen, H. Averfalk, S. Werner, H. Lund, Heat roadmap europe: Large-scale electric heat pumps in district heating systems, Energies 10 (2017). doi:10.3390/en10040578.

[8] I. Oliker, Steam Turbines for Cogeneration Power Plants, Journal of Engineering for Power 102 (1980) 482–485. doi:10.1115/1.3230281.

[9] H. Madsen, J. Holst, Estimation of continuous-time models for the heat dynamics of a building, Energy and Buildings 22 (1995) 67–79. doi:https://doi.org/10.1016/0378-7788(94)00904-X.

[10] H. A. Nielsen, H. Madsen, Modelling the heat consumption in district heating systems using a grey-box approach, Energy and Buildings 38 (2006) 63–71. doi:https://doi.org/10.1016/j.enbuild.2005.05.002.

[11] T. Nielsen, H. Madsen, Control of supply temperature in district heating systems, in: Proceedings of the 8th International Symposium on District heating and Cooling, 2002.

[12] European Parliament and Council of the European Union, Directive 2012/27/EU, 2012. Accessed on 28 March 2021.

[13] M. H. Kristensen, S. Petersen, District heating energy efficiency of danish building typologies, Energy and Buildings 231 (2021) 110602. doi:https://doi.org/10.1016/j.enbuild.2020.110602.

[14] H. G. Bergsteinsson, T. S. Nielsen, J. K. Møller, S. B. Amer, D. F. Dominković, H. Madsen, Use of smart meters as feedback for district heating temperature control, Energy Reports 7 (2021) 213–221. doi:https://doi.org/10.1016/j.egyr.2021.08.153, the 17th International Symposium on District Heating and Cooling.

[15] A. Vandermeulen, Quantification and Optimal Control of District Heating Network Flexibility, Ph.D. thesis, KU Leuven, 2020.

[16] V. D. Stevanovic, B. Zivkovic, S. Prica, B. Maslovaric, V. Karamarkovic, V. Trkulja, Prediction of thermal transients in district heating systems, Energy Conversion and Management 50 (2009) 2167–2173. doi:https://doi.org/10.1016/j.enconman.2009.04.034.

[17] A. Benonysson, B. Bøhm, H. F. Ravn, Operational optimization in a district heating system, Energy Conversion and Management 36 (1995) 297–314. doi:https://doi.org/10.1016/0196-8904(95)98895-T.

[18] H. T. Søgaard, Stochastic systems with embedded parameter variations - applications to district heating, Ph.D. thesis, Technical University of Denmark, Department of Applied Mathematics and Computer Science, 1993.

[19] H. Madsen, T. Nielsen, H. Søgaard, Control of Supply Temperature: EFP 1323/93-07, Informatics and Mathematical Modelling, Technical University of Denmark, 1996.

[20] B. van der Heijde, Optimal Integration of Thermal Energy Storage and Conversion in Fourth Generation Thermal Networks, Ph.D. thesis, KU Leuven, 2019.

[21] P. Bacher, H. Madsen, Identifying suitable models for the heat dynamics of buildings, Energy and Buildings 43 (2011) 1511–1522. doi:https://doi.org/10.1016/j.enbuild.2011.02.005.

[22] C. A. Thilker, P. Bacher, H. G. Bergsteinsson, R. G. Junker, D. Cali, H. Madsen, Non-linear grey-box modelling for heat dynamics of buildings, Energy and Buildings 252 (2021) 111457. doi:https://doi.org/10.1016/j.enbuild.2021.111457.

[23] C. A. Thilker, H. G. Bergsteinsson, P. Bacher, H. Madsen, D. Calì, R. G. Junker, Non-linear model predictive control for smart heating of buildings, in: E3S Web of Conferences, volume 246, EDP Sciences, 2021, p. 09005.

[24] A. Jazwinski, Stochastic Processes and Filtering Theory, Academic Press, San Diego, CA, USA, 1970.

[25] P. Grunnet Wang, M. Scharling, K. Pagh Nielsen, C. Kern-Hansen, K. Wittchen, 2001 – 2010 Danish Design Reference Year: Reference Climate Dataset for Technical Dimensioning in Building, Construction and other Sectors, number 13-19 in DMI Technical Report, Danmarks Meteorologiske Institut, 2013.

[26] B. van der Heijde, M. Fuchs, C. Ribas Tugores, G. Schweiger, K. Sartor, D. Basciotti, D. Müller, C. Nytsch-Geusen, M. Wetter, L. Helsen, Dynamic equation-based thermo-hydraulic pipe model for district heating and cooling systems, Energy Conversion and Management 151 (2017) 158–169. doi:https://doi.org/10.1016/j.enconman.2017.08.072.

[27] B. van der Heijde, A. Aertgeerts, L. Helsen, Modelling steady-state thermal behaviour of double thermal network pipes, International Journal of Thermal Sciences 117 (2017) 316–327. doi:https://doi.org/10.1016/j.ijthermalsci.2017.03.026.

[28] A. Dénarié, M. Aprile, M. Motta, Heat transmission over long pipes: New model for fast and accurate district heating simulations, Energy 166 (2019) 267–276. doi:https://doi.org/10.1016/j.energy.2018.09.186.

[29] S. Grosswindhager, A. Voigt, M. Kozek, Linear finite-difference schemes for energy transport in district heating networks, in: Proceedings of the 2nd international conference on computer modelling and simulation, 2011, pp. 5–7.

[30] P. Wallentén, Steady-state heat loss from insulated pipes, Ph.D. thesis, Division of Building Physics, Lund University, 1991.

[31] K. Kristensen, A. Nielsen, C. W. Berg, H. Skaug, B. M. Bell, Tmb: Automatic differentiation and laplace approximation, Journal of Statistical Software 70 (2016) 1–21. doi:10.18637/jss.v070.i05.

[32] H. Madsen, P. Thyregod, Introduction to General and Generalized Linear Models, Texts in Statistical Science Series, CRC Press, 2010.

[33] C. Van Loan, Computing integrals involving the matrix exponential, IEEE Transactions on Automatic Control 23 (1978) 395–404. doi:10.1109/TAC.1978.1101743.

[34] 5810827_Z2_GB_06.pdf, Kamstrup, 2021. URL: https://koce1-kamstrup.ocecdn.oraclecloud.com/content/published/api/v1.1/assets/CONT4AA1D9C46651406CA6AC17F2FAAFCC99/native/5810827_Z2_GB.pdf?channelToken=ed241bbb18f444908a8fc9ed97ca5d5b, accessed: 2021-12-20.

Paper D

NON-LINEAR GREY-BOX MODELLING FOR HEAT DYNAMICS OF BUILDINGS

Authors:
Christian Ankerstjerne Thilker, Peder Bacher, Hjörleifur G. Bergsteinsson, Rune Grønborg Junker, Davide Cali and Henrik Madsen.

Published in:
Energy and Buildings.

Non-Linear Grey-Box Modelling for Heat Dynamics of Buildings

Christian Ankerstjerne Thilker, Peder Bacher, Hjörleifur G. Bergsteinsson, Rune Grønborg Junker, Davide Cali, Henrik Madsen

Technical University of Denmark, Department of Applied Mathematics and Computer Science,
Asmussens Allé, Building 303B, DK-2800 Kgs. Lyngby, Denmark

Abstract

This paper introduces a non-linear grey-box (GB) model based on stochastic differential equations that describes the heat dynamics of a school building in Denmark, equipped with a water-based heating system. The building is connected to a local district heating network through a heat exchanger. The heat is delivered to the rooms mainly through radiators and partially through a ventilation system. A monitoring system based on IoT sensors provides data on indoor climate in the rooms and on the heat load of the building. Using this data, we estimate unknown states and parameters of a model of the building's heating system using the maximum likelihood method. Important novelties of this paper include models of the water flow in the circuit and the state of the valves in the radiator thermostats. The non-linear model accurately predicts the indoor air temperature, return water temperature and heat load. The ideas behind the model lay a foundation for GB models of buildings that use different kinds of water-based heating systems such as air-to-water/water-to-water heat pumps. Such GB models enable model predictive control to control e.g. the indoor air climate or provide flexibility services.

Keywords: Grey-box models, Stochastic differential equations, Non-linear models, District heating, Smart energy systems

1. Introduction

The use of fossil-based energy sources does not belong in a sustainable future [1]. Society must shift to energy sources where CO_2-emissions lie within the planetary boundaries; i.e. we need to use resources that are renewable [2]. This future low-carbon society calls for fundamental changes of the energy system. Today the systems are operated such that the production follows the demand. However, an efficient implementation of a low-carbon society calls for a system where the demand follows the weather-driven energy production. Most importantly we need methods for unlocking the flexibility at all levels of the society; examples being buildings, supermarkets, wastewater treatment plants, industrial process facilities, districts, municipalities and cities. A lot of recent work, therefore, centres around the concept known as energy flexibility [3, 4]. The core idea is to control the energy consumption to align it with energy production. For this purpose, model-based predictive control is a very promising control framework [5]. This paper introduces a novel grey-box (GB) model based on stochastic differential equations (SDEs) that is designed for controller based optimisation of the heat load of buildings. The ultimate purpose of developing such a GB model is to intelligently control buildings in order to minimise the CO_2-emissions and unlock the flexibility. A reliable model (together with weather forecasts) is essential for a good performance of model predictive control (MPC) for buildings [6].

MPC for control of buildings' indoor climate requires reliable building models that describe the heat dynamics. Complex building energy performance models based exclusively on physical equations, known as white–box models, are often used for providing simulations. Occasionally, in white box building models, stochastic models are used to simulate occupants behaviour, as in [7, 8]. However, they are demanding to build, computationally heavy, and difficult or impossible to tune to real-world data, which makes them infeasible for control. Especially for the existing building stock. On the contrary, black–box models can be fast in terms of simulation time. But they do not include laws of physics, and thus may be hard to interpret and lack the ability to extrapolate and generalise beyond training data. GB models bridge the gap between white- and black-box models by leveraging both physical and statistical properties [9]. They are based on *simple* physical principles and considerations of the system, which make them computationally light and ideal for parameter calibration using available data. Linear GB models for buildings are widely seen in the literature [10]. Wang and Xu [11] use a genetic algorithm to estimate a linear heat dynamics model that describes the thermal conditions in the wall envelopes and internal mass for an office building. The goal is to predict the heat load and the indoor air temperature. Massano et al. [12] uses an unscented Kalman filter to estimate parameters in a linear RC-inspired model to predict the indoor air temperature. Bacher and Madsen [13] outlines a model development procedure for SDE-based GB models. However, it is a well known fact that *non-linear systems* exhibit vast richness in the solution structure, far beyond what is seen in linear systems [14]. For instance, non-linear models are necessary to sufficiently describe the heat dynamics of build-

Email addresses: chant@dtu.dk (Christian Ankerstjerne Thilker), pbac@dtu.dk (Peder Bacher), hgbe@dtu.dk (Hjörleifur G. Bergsteinsson), rung@dtu.dk (Rune Grønborg Junker), dcal@dtu.dk (Davide Cali), hmad@dtu.dk (Henrik Madsen)

ing integrated photo-voltaic modules [15, 16]. Non-linear GB models can also be found within industrial robotics [17] and in aquatic ecosystems modelling [18], just to mention a few areas. To the knowledge of the authors, the literature on non-linear GB models for radiator-based heating systems is scarce.

We propose SDEs as the modelling framework for the building model [19]. This has many advantages: First, SDEs provide a natural method to model physical phenomena as they are formulated in continuous-time. Second, they include probabilistic uncertainty that accounts for modelling approximations, unrecognised exogenous variables, and uncertainty related to the provided input variables. Last, they lay a solid foundation providing predictions of the system behaviour and for model-based optimal control, to predict system behaviour. It is well-know that solutions to Ordinary Differential Equations (ODEs) are functions of time, and this implies that an ODE modelling framework assumes that we are able to predict the exact evolution in time of the states. Solutions to SDEs are stochastic processes, which are characterised by the family of finite-dimensional densities, and this implies that the future evolution of the states is encumbered with uncertainty, and this uncertainty can be quantified. Optimal control theory based on SDEs is well-established in the literature with numerous examples of applications, e.g. for control of glucose concentration in humans [20], building thermal control [21], and operation of waste-water treatment plants [22].

1.1. Main contributions

The existing literature contains various examples of linear GB models of the heat dynamics of buildings. However, the literature seems to contain limited work on SDE-based non-linear GB models for water-based heating systems, especially related to district heating (DH). This paper presents and analyses the development of a non-linear GB model for a school building in Denmark with water-based heating. We base the analysis and estimation on a single week of data using meteorological weather observations as inputs, and we will demonstrate that one week of data is sufficient for identifying a good model. Due to the generality of the model, it is argued that the model is applicable to a wide range of buildings with water-based heating systems and different heat sources (including heat pumps).

An important contribution of this paper is the model of the thermostatic valves of the radiators. The radiator valves are mechanically adjusted by the thermostats that are configured with a set-point. The valves open and close proportionally to the difference between the set-point and actual air temperature. The valves naturally do not behave discontinuously when heat is or is not needed. Models for thermostatic valves exist in applications of white-box models [23]. Most are modelled as P, PI, or PID-based controllers for white-box building models [24]. Hansen [25] suggested detailed physical models of radiators and thermostats. However, the models end up being too large and detailed for grey-box purposes. To the knowledge of the authors, the literature contains no examples of models for thermostat valves formulated as GB-models. This paper presents a sigmoid-function to describe the continuous sensitivity of the valves due to changes in the indoor air temperature according to the set-point. Another important contribution is a model of the water flow in the building heating system.

1.2. Structure and outline of the paper

This paper has the following structure. Section 2 introduces the building and its engineering systems, together with the overall experiment. Here, we also describe the data and how it was gathered. Section 3 describes the model development process and the ideas behind the suggested model. Next, we present and discuss the results; the parameter estimates, a simulation of the variables compared to data, and a 1-step residual analysis. Lastly, Section 6 sums up the essential findings of the paper.

2. The building and the experimental setup

This section introduces the building and describes the experimental data and the generation process.

2.1. The building

The building, a school with three floors and a basement, is located in Høje Taastrup, Denmark. The uppermost floor is a part-refurbished roof attic. Bruun [26] provides all technical information about the building.

Being built in 1929, the building is not insulated according to modern standards. Figure 1 shows a digital reconstruction and a photo of the building. It includes 10 classrooms that are ventilated by mechanical ventilation using an air handling unit (AHU) for air circulation. The facade and internal walls consist of solid bricks (300 mm and 180 mm thickness, respectively). The windows have wooden frames and double-paned low-E glazings. Floors are made from wood joists and the roof is partly uninsulated and partly insulated slate roof. The building is connected to the local electricity and heat grid, where the ladder is a DH system. The building uses district heating for domestic hot water (DHW), the AHU, and space heating. The latter term governs the heating (and cooling) system of the indoor air. For this building, the space heating is a separate water-based circuit with dedicated pumps. Radiators of different types (cast-iron, panel convectors, plane conductors) with individual thermostats establish the space heating system in the individual rooms of the building. Individual thermostatic valves automatically regulate the water flow into the radiator units as to maintain a pre-defined *set-point*. The space heating system is separated from the DH system by a plate heat exchanger. Independent PI-controllers regulate the water flow on both the district heating and the building side of the heat exchanger.

2.2. The experiment

The experiment carried out was planned in advance and designed to generate data suitable for system identification purposes. The main focus was to change the control input, the thermostat set point, such that information about the essential dynamics of the system can be estimated. A sequence of the set point was designed with four different parts. First part contains a few long steps with set points set to a minimum (10 °C) and back again to a base level (21 °C) to get information about

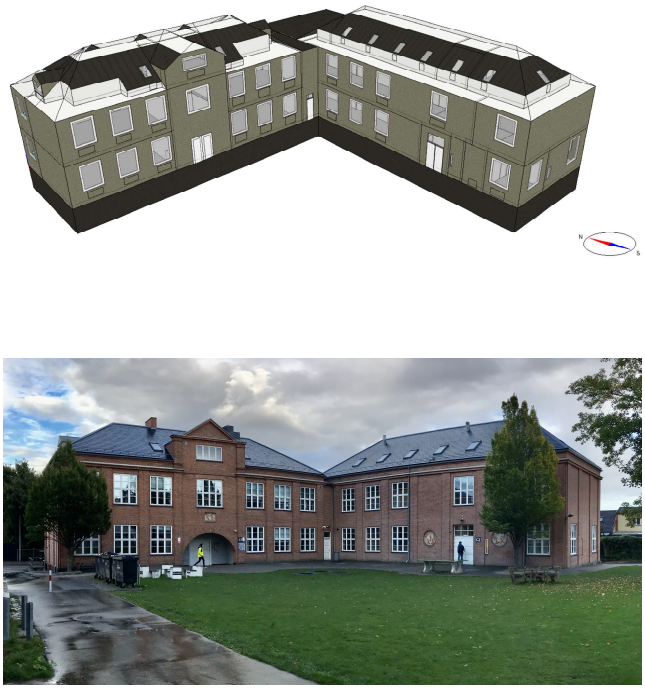

Figure 1: Visual illustrations of the building site. The upper digital reconstruction is supplied by [26]

the dynamics governing the system. Second part is a multi-level signal, where the extremes (14 and 27 °C) are kept for the longest time and then shorter periods are kept for relatively shorter time. Third part contains short periods with drops to a minimum from the base temperature. Finally, a step sequence where the set point is stepped from 23 °C in two hours steps down to 17 °C and up again. The forward temperature of the space heating water is set constant to 55 °C at all times. The entire sequence was slightly shorter than 7 days and was executed during the Christmas vacation, where the building was unoccupied.

2.3. The data

Table 1 lists all the variables of the data. Figure 2 shows the experimental data in the period December 21 through December 27. The upper graph displays the heat load of the building. It seems to be characterised by a large peak whenever the heat turns on, before reverting to a lower and steady level. The second graph shows the forward- and return water temperature, which go to and from the space heating system. The forward temperature fluctuates a lot when the thermostat set point is set very low – because the thermostat valves are closed and thus the flow in the radiator circuit is nearly stopped, which results in poor control of the forward temperature since the control was not designed for this situation. The return temperature quickly becomes large when the space heating is turned on. In absence of heat load, the return temperature quickly decreases. But, the reversion and behaviour in absence of heat load seem to be rather inconsistent. The third graph shows the indoor temperature of each room (in grey) and the mean of all rooms (in black). Lastly, the bottom graph shows the exterior weather conditions, i.e. the ambient air temperature and the global solar radiation. The latter is relatively small throughout the period, which complicates the estimation of the solar radiation gain for the model. We return to this matter later in Section 5.

The variables of the building we wish to be able to predict are the following

- The mean indoor air temperature, T_t^i.
- The heat load of the building delivered by the DH system, ϕ_t^h.
- The temperature of the returning water in the SH system of the building, T_t^{ret}.

The subscript t indicates the dependence on time. These variables are of special interest when it comes to optimal control of the indoor climate. In Denmark, building operators pay for the amount of heat they consume. Additionally, the operators pay fees for too high return temperatures since it is a source of poor energy efficiency in the DH network. First, the DH operators have to increase the mass flow rate of the water, if the users do not cool the return water. Secondly, if the return water to the DH facilities is too hot, the efficiency of the central heat production plant decreases. This payment scheme makes it economically advantageous for the building operators to use heat when it is cheap and minimise the return temperature.

Table 1: Data interpretation.

Name	Quantity	Unit
$T_t^{(i,j)}$	Indoor air temperature in room j	[°C]
T_t^i	Mean indoor air temperature in the building	[°C]
T_t^{for}	Forward temperature	[°C]
T_t^{ret}	Return temperature	[°C]
T_t^{set}	Temperature set-point	[°C]
ϕ_t^h	Delivered heat to the building	[kW]
T_t^a	Ambient air temperature	[°C]
ϕ_t^s	Solar radiation on a horizontal surface	[kW/m^2]

3. Model development

This section describes the model development process.

Due to the large parameter and state space, it is advantageous to perform the modelling in small steps. The main idea is to split up the modelling processes into two parts. To identify the steps, we need to realise that the building heat dynamics consist of two parts (as a first simple assumption). The delivered heat from the water in the radiator system operates independently and only interacts with the indoor air temperature of the building by the radiators themselves. This interaction involves only two parameters. Therefore, we split the modelling part into the following three steps:

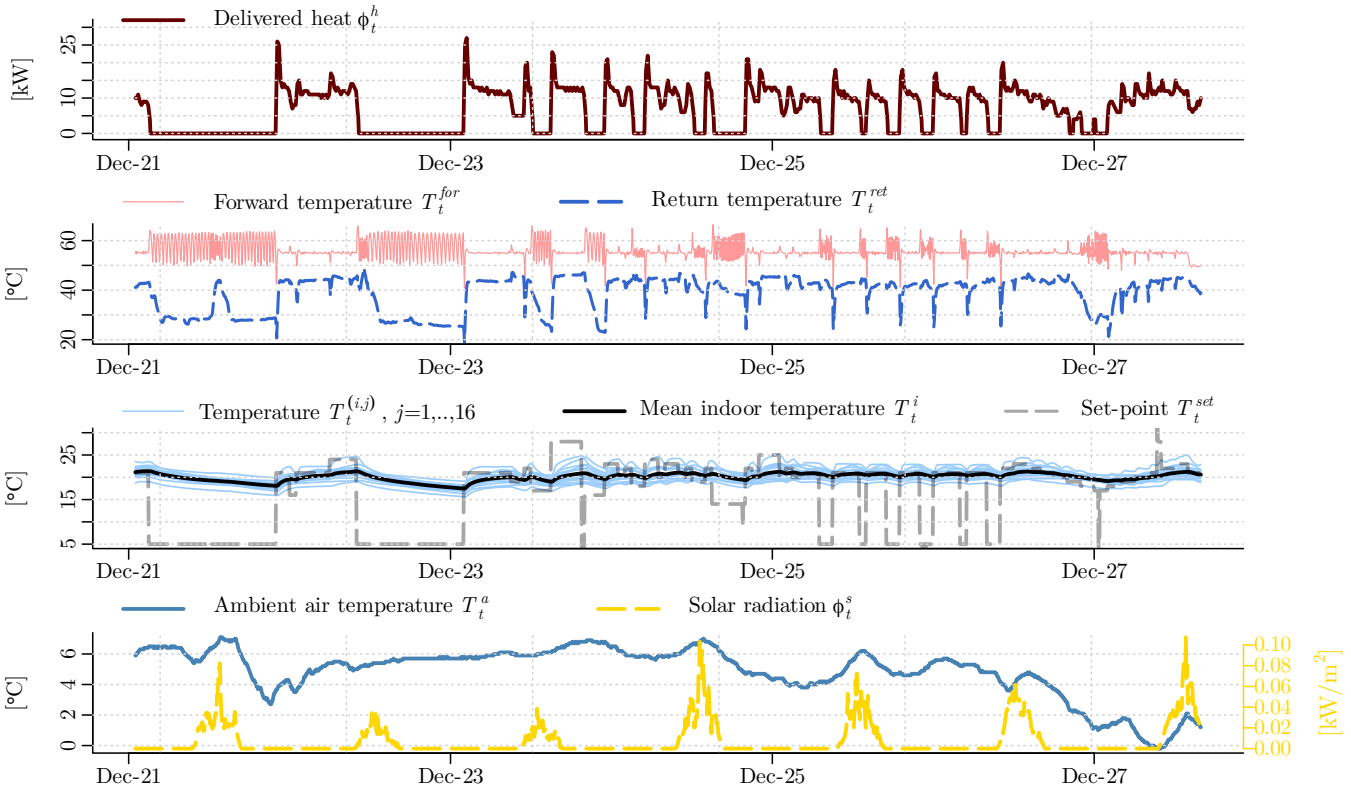

Figure 2: The data from the experiment performed in December 2019.

1. Given the observed time series of the delivered heat from the district heating system, we develop a model that predicts only the indoor air temperature of the building.
2. Given the observed time series of the indoor air temperature and set point, we develop a model that predicts only the heat load from the district heating system to the building. That is while keeping the parameters fixed, that concerns the indoor air temperature model obtained in step 1.
3. We combine the two models and start the parameter optimisation from the results of the two independently sufficient models to obtain a combined model structure.

By developing the two system models individually at first, it also becomes much easier to identify the necessary dynamical features that govern the systems.

3.1. Stochastic differential equations

The model will be formulated using SDEs. A SDE typically has the following form

$$\mathrm{d}\boldsymbol{X}_t = f(\boldsymbol{X}_t, t)\mathrm{d}t + g(\boldsymbol{X}_t, t)\mathrm{d}\boldsymbol{\omega}_t \tag{1}$$

where f and g are the drift and diffusion terms, respectively, and the subscript t denotes the dependence on time. The diffusion term makes a SDE differ from an ordinary differential equation. ω_t is known as Brownian motion and is a fundamental process for stochastic calculus. It is governed by independent Gaussian increments; $\omega_t - \omega_s \sim N(0, t-s)$, for $s \leq t$. This has remarkable consequences and relates it to the physical diffusion equation. The purpose of the diffusion term is to describe chaotic phenomena that are too complex to include in the drift part of the model structure.

3.2. The building heat dynamics model

The literature contains numerous examples of developing heat dynamic models for buildings using continuous-time GB models, see e.g. [27–31]. We do not give the model identification steps explicitly for our case though but simply report the final result.

Figure 2 gives insights into what elements the building model should include. Inspections of the two long periods, where the heat is turned off, show that the mean indoor temperature seems to drop fast at first and then flatten to a certain decay rate. This indicates that we should include two time constants; one for the fast and initial drop and one for the slow long-term decay. We may interpret these fast and slow dynamics as the temperature

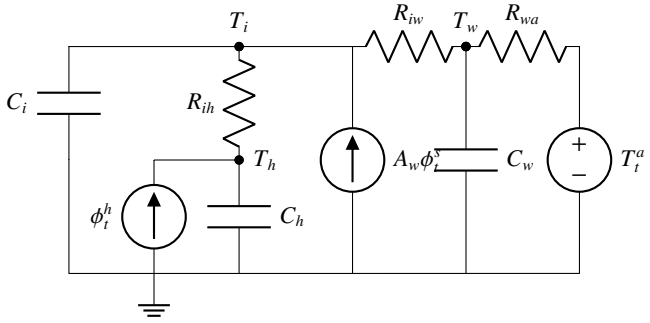

Figure 3: The model structure of the building heat dynamics. The model predicts the indoor air temperature, $\hat{T}_i$, given the delivered power ϕ_h.

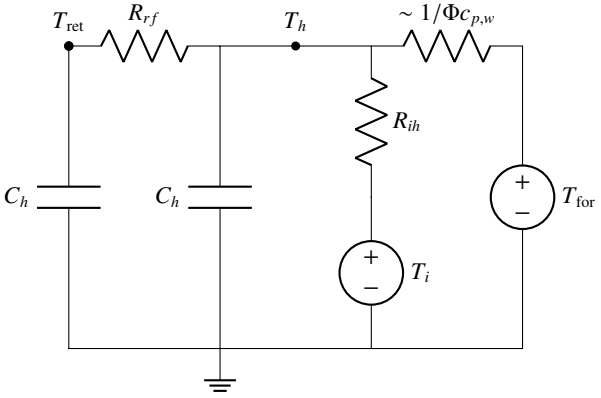

Figure 4: The radiator circuit dynamics model visualised as an RC-diagram. The flow, Φ_t, acts as a heat transfer coefficient in the radiator temperature state. We simply write the resistance as the flow state Φ_t.

of the indoor air and the temperature of the building walls. For this reason, the wall temperature state exchanges heat with the ambient air temperature and acts as a low-pass filter between the interior and exterior. We also choose to model the radiators as an accumulating medium where the heat input enters directly. The solar radiation gain enters the room air directly through windows.

Figure 3 shows the heat dynamics structure for the building as an RC-diagram. The equivalent SDE model has the following form

$$\mathrm{d}T_t^i = \frac{1}{C_i}\left(\frac{1}{R_{ih}}\left(T_t^h - T_t^i\right) + \frac{1}{R_{iw}}\left(T_t^w - T_t^i\right) + A_w\phi_t^s\right)\mathrm{d}t + \sigma_1\mathrm{d}\omega_t^1\,, \tag{2a}$$

$$\mathrm{d}T_t^w = \frac{1}{C_w}\left(\frac{1}{R_{iw}}\left(T_t^i - T_t^w\right) + \frac{1}{R_{wa}}\left(T_t^a - T_t^w\right)\right)\mathrm{d}t + \sigma_2\mathrm{d}\omega_t^2\,, \tag{2b}$$

$$\mathrm{d}T_t^h = \frac{1}{C_h}\left(\frac{1}{R_{ih}}\left(T_t^i - T_t^h\right) + \phi_t^h\right)\mathrm{d}t + \sigma_3\mathrm{d}\omega_t^3\,. \tag{2c}$$

(2d)

3.3. The radiator circuit dynamics model

The thermostatic valves regulate the water flow through the radiators. An important novelty of this paper is to model the thermostatic valves using the non-linear sigmoid function. The idea is that the vales open when it is too cold and close when it is too warm. Assuming that the valves react continuously to the indoor air temperature, the sigmoid function corresponds to some kind of proportional control (0 being closed and 1 being open).

3.3.1. The thermostatic valve function

To describe the thermostatic control, i.e. the amount of heat that the heat exchanger delivers, we use a sigmoid function. To be specific, it describes the sensitivity of the heating system to deviations in the mean indoor temperature. We use the following formulation

$$f_t^{\text{valve}} = \frac{1}{1 + \exp\left(-\alpha(T_t^{\text{set}} + T_{\text{offset}} - T_t^i)\right)}\,. \tag{3}$$

α is the slope of the sigmoid function and determines how quickly the heating system turns on and off. T_{offset} acts as an offset: since the observations of the indoor temperature typically is taken some place in the rooms (probably not right next to the radiators), the thermostats may respond to a temperature that is warmer or colder than the observed one.

The sigmoid function has the disadvantage that it cannot reach 1 nor 0. For the purpose of this paper, it means that even though the set-point is, say, 18 and the observed temperature is 20, the model predicts that the radiators still deliver some heat (depending on the slope and offset). Depending on the specific thermostats and the valves in the radiators, this prediction may be wrong. We address this problem further in Section 5.

3.3.2. Derivation of the space heating model structure

The space heating system proved hard to model. It is difficult to describe all necessary dynamics in a simple manner. However, we found it fruitful to model the *water flow* in the radiator circuit as a dynamical equation governed by a time-delay. The governing physical equation of the net energy transferred to the radiator, Q_t^h, from the water is [24]

$$\mathrm{d}Q_t^h = \Phi_t c_{p,w}(T_t^{\text{for}} - T_t^{\text{ret}})\mathrm{d}t \tag{4}$$

where Φ_t is the flow of the water in the SH system, $c_{p,w}$ is the specific heat capacity of water.

As Section 2.1 describes, pumps that are controlled by some PI controllers regulate the water flow in the SH system to maintain a certain pressure. Therefore, there is a delay from when the valves open until the pumps increase the water accordingly. For this reason, it seems reasonable to model the water flow as a differential equation itself.

To model the return temperature, a few observations are important. It is obvious that when the space heating system delivers heat, the hot/cold water has been round in the radiator circuit and returns to the heat exchanger in a colder state. This is clearly visible in Figure 2 that the return temperature varies between 40-50 °C when heat is delivered. However, when the heat load equals zero, the return temperature acts rather inconsistently. At these times, the return temperature mostly responds

with a quick decay to below 30 °C. But as Figure 2 shows, it sometimes drifts indescribably. Most times, when the water flow stops, the water in the return pipes quickly delivers its heat to the surroundings and arrives at some equilibrium. But when no heat is delivered, the building operators do not pay for larger return temperatures and gives no additional insights into the system. Therefore, we have no means to model the return temperature, when the heat load is zero. We thus disregard the return temperature observations when the flow is zero for simplicity. Section 5 explains how we implement this in the parameter estimation.

To recap the above thoughts; when heat is needed, the water flows into the radiators to deliver heat and afterward returns to the heat exchanger to be heated again. By combining (4) with a state for the flow and the return temperature, we arrive at the following model

$$\mathrm{d}\Phi_t = \frac{1}{C_f}\left(\Phi_{\max} f_{\text{valve}} - \Phi_t\right)\mathrm{d}t + \sigma_3 \mathrm{d}\omega_t^3\,, \tag{5a}$$

$$\mathrm{d}T_t^h = \frac{1}{C_h}\left(\Phi_t c_{p,w}\left(T_t^{\text{for}} - T_t^h\right) + \frac{1}{R_{ih}}\left(T_t^i - T_t^h\right)\right)\mathrm{d}t + \sigma_4 \mathrm{d}\omega_t^4\,, \tag{5b}$$

$$\mathrm{d}T_t^{\text{ret}} = \frac{1}{C_h}\left(\frac{1}{R_{fr}}\left(T_t^h - T_t^{\text{ret}}\right)\right)\mathrm{d}t + \sigma_5 \mathrm{d}\omega_t^5\,. \tag{5c}$$

where $\Phi_{\max}$ the maximum attainable flow of the radiator circuit water. Figure 4 depicts the structure as a RC-diagram.

3.4. Heat load estimation equation

From (4) we estimate the heat load, i.e. the power from the DH to the indoor air, as

$$\phi_t^h = \Phi_t c_{p,w}\left(T_t^{\text{for}} - T_t^{\text{ret}}\right)\,. \tag{6}$$

It should be natural to assume that the heat difference between the forward and return water is due only to the delivered heat by the space heating system. The temperature difference multiplied by the flow and the specific heat capacity of water is thus an estimate of the heat load. The flow state creates a time delay on the heat load. The term $\left(T_t^{\text{for}} - T_t^{\text{ret}}\right)$ is almost always large when the heat is turned off. Had there been no delay, e.g. for the equation $C_1 f_{\text{valve}}\left(T_t^{\text{for}} - T_t^{\text{ret}}\right)$, where C_1 is an arbitrary parameter, the heat load would immediately spike when the valves open. However, the heat load data is governed by delay which suggests that such time delay is needed.

3.5. The combined model

The combined model has the form

$$\mathrm{d}T_t^i = \frac{1}{C_i}\left(\frac{1}{R_{ih}}\left(T_t^h - T_t^i\right) + \frac{1}{R_{iw}}\left(T_t^w - T_t^i\right) + A_w \phi_t^s\right)\mathrm{d}t + \sigma_1 \mathrm{d}\omega_t^1\,, \tag{7a}$$

$$\mathrm{d}T_t^w = \frac{1}{C_w}\left(\frac{1}{R_{iw}}\left(T_t^i - T_t^w\right) + \frac{1}{R_{wa}}\left(T_t^a - T_t^w\right)\right)\mathrm{d}t + \sigma_2 \mathrm{d}\omega_2\,, \tag{7b}$$

$$\mathrm{d}\Phi_t = \frac{1}{C_f}\left(\Phi_{\max} f_{\text{valve}} - \Phi_t\right)\mathrm{d}t + \sigma_3 \mathrm{d}\omega_t^3\,, \tag{7c}$$

$$\mathrm{d}T_t^h = \frac{1}{C_h}\left(\Phi_t c_{p,w}\left(T_t^{\text{for}} - T_t^h\right) + \frac{1}{R_{ih}}\left(T_t^i - T_t^h\right)\right)\mathrm{d}t + \sigma_4 \mathrm{d}\omega_t^4\,, \tag{7d}$$

$$\mathrm{d}T_t^{\text{ret}} = \frac{1}{C_h}\left(\frac{1}{R_{fr}}\left(T_t^h - T_t^{\text{ret}}\right)\right)\mathrm{d}t + \sigma_t \mathrm{d}\omega_t^5\,. \tag{7e}$$

where $\Phi_{\max}$ are the flow speed of the water on the building site. The observation equations are

$$y_k^i = T_{t_k}^i + v_1\,, \qquad v_1 \sim N_{iid}(0, R_1)\,, \tag{8a}$$

$$y_k^h = \Phi_t c_{p,w}\left(T_t^{\text{for}} - T_t^{\text{ret}}\right) + v_2\,, \qquad v_2 \sim N_{iid}(0, R_2)\,, \tag{8b}$$

$$y_k^{\text{ret}} = T_{t_k}^{\text{ret}} + v_3\,, \qquad v_3 \sim N_{iid}(0, R_3)\,. \tag{8c}$$

4. Model identification and estimation

This section describes the identification method and the details governing the parameter estimation process. This paper proposes maximum likelihood inference for parameter estimation in stochastic differential due to its ability to estimate noise parameters. See e.g. Madsen [32] or Pawitan [33] for an introduction to maximum likelihood methods.

4.1. The maximum likelihood principle

Given the sequence of observations $\mathcal{Y}_N = \{\boldsymbol{Y}_i\}_{i=1}^N$, $\boldsymbol{Y}_k = [y_k^i, y_k^h, y_k^{\text{ret}}]^\intercal \in \mathbb{R}^{n_y}$, and set-points $\mathcal{U}_{N-1} = \{T_{\text{set},i}\}_{i=0}^{N-1}$, define the likelihood function as the product of the one-step ahead conditional densities:

$$\mathcal{L}(\boldsymbol{\theta}|\mathcal{Y}_N, \mathcal{U}_{N-1}) = p(\boldsymbol{X}_0)\prod_{k=1}^{N} p(\boldsymbol{Y}_k|\mathcal{Y}_{k-1}, \mathcal{U}_{k-1}, \boldsymbol{\theta})\,. \tag{9}$$

Here, p is the probability of observing $\boldsymbol{Y}_k$ given the previous observations, set-points, and parameters $\boldsymbol{\theta}$. For linear stochastic differential equations, where the noise is state-independent and driven by Brownian motion, the conditional densities are also Gaussian. For non-linear systems though, this is not the case and the analytical density is in general hard (or impossible) to find. But when the time between observations are small, the Gaussian density approximates the analytical (unknown) density well. This motivates our choice of using the Gaussian density in the likelihood function. The Gaussian density is completely characterised by its conditional mean and variance; by introducing the one-step prediction error

$$\boldsymbol{\epsilon}_k = \boldsymbol{Y}_k - \hat{\boldsymbol{Y}}_{k|k-1}\,, \tag{10}$$

where $\hat{\boldsymbol{Y}}_{k|k-1} = \mathrm{E}[\boldsymbol{Y}_k|\mathcal{Y}_{k-1}, \mathcal{U}_{k-1}, \boldsymbol{\theta}]$, and the associated covariance $\boldsymbol{R}_{k|k-1} = \mathrm{Var}[\boldsymbol{Y}_k|\mathcal{Y}_{k-1}, \mathcal{U}_{k-1}, \boldsymbol{\theta}]$, we can write the likelihood function as

$$\mathcal{L}(\boldsymbol{\theta}|\mathcal{Y}_N, \mathcal{U}_{N-1}) = p(\boldsymbol{X}_0) \prod_{k=1}^{N} \frac{\exp\left(-\frac{1}{2}\boldsymbol{\epsilon}_{k|k-1}^{\mathsf{T}} \boldsymbol{R}_{k|k-1}^{-1} \boldsymbol{\epsilon}_{k|k-1}\right)}{\sqrt{\det(\boldsymbol{R}_{k|k-1})(2\pi)^{n_y}}} . \quad (11)$$

Taking the logarithm on both sides, we obtain the log-likelihood function

$$\begin{aligned} \ell(\boldsymbol{\theta}|\mathcal{Y}_N, \mathcal{U}_{N-1}) = \log\left(p(\boldsymbol{X}_0|\boldsymbol{\theta})\right) - \frac{1}{2}\sum_{k-1}^{N} \boldsymbol{\epsilon}_{k|k-1}^{\mathsf{T}} \boldsymbol{R}_{k|k-1}^{-1} \boldsymbol{\epsilon}_{k|k-1} \\ + \log\left(\det(\boldsymbol{R}_{k|k-1})(2\pi)^{\frac{n_y}{2}}\right) \end{aligned} \quad (12)$$

The log-likelihood has some attractive advantages over the ordinary likelihood when it comes to numerical properties, which is why it is often preferred. The parameter estimates $\hat{\boldsymbol{\theta}}$ is found by maximising the log-likelihood function

$$\hat{\boldsymbol{\theta}} = \arg\max_{\boldsymbol{\theta}} \ell(\boldsymbol{\theta}|\mathcal{Y}_N, \mathcal{U}_{N-1}) \quad (13)$$

To evaluate the log-likelihood function, we need to compute the one-step prediction errors, $\boldsymbol{\epsilon}_k$, and the associated covariance $\boldsymbol{R}_{k|k-1}$ (due to our assumption of Gaussian densities). The continuous-discrete extended Kalman filter supplies exactly these.

4.2. The continuous-discrete extended Kalman filter

The continuous-discrete extended Kalman filter (CDEKF) is a variant of the celebrated Kalman filter [34]. It considers system models governed by continuous-time dynamics where the observer observers parts of the system at discrete times. In short, the CDEKF consists of a *prediction* step and an *update* step. The extended Kalman filter relies on a linearisation of the non-linear system (1), which causes troubles if g is state-dependent. In such cases, the Lamperti-transformation is an important tool to transform (1) into a state-independent SDE [35, 36]. The literature contains many introductions and applications to Kalman filtering and the CDEKF, see e.g. [19, 37].

4.2.1. The prediction scheme

In this step, the CDEKF predicts the state of the building $\hat{\boldsymbol{X}}_{k|k-1} = \mathrm{E}[\boldsymbol{X}_k|\mathcal{Y}_{k-1}, \mathcal{U}_{k-1}, \hat{\boldsymbol{\theta}}]$, $\boldsymbol{X}_k = \left[T_{t_k}^i, T_{t_k}^w, \Phi_{t_k}, T_{t_k}^h, T_{t_k}^{\mathrm{ret}}\right]^{\mathsf{T}}$, together with the state covariance $\hat{P}_{k|k-1} = \mathrm{Var}[\boldsymbol{X}_k|\mathcal{Y}_{k-1}, \mathcal{U}_{k-1}, \hat{\boldsymbol{\theta}}]$ at the next time step t_k given the estimated state at time t_{k-1}, $\hat{\boldsymbol{X}}_{k|k-1}$. This involves solving a set of coupled ordinary differential equations (ODEs). Any ODE-solver is sufficient for this task.

4.2.2. The update scheme

The updating scheme is about estimating the underlying state and its covariance, denoted $\hat{\boldsymbol{X}}_{k|k}$ and $\hat{P}_{k|k}$, at the next time instance t_k, given our predictions, $\hat{\boldsymbol{X}}_{k|k-1}$ and $\hat{P}_{k|k-1}$, and an observation $\boldsymbol{Y}_k$. Informally speaking, the updating scheme finds a weight K, typically called the *Kalman gain*, which "measures" how much weight the observation should have on the estimate $\hat{\boldsymbol{X}}_{k|k}$. Consider the update equations for the state estimate

$$\hat{\boldsymbol{X}}_{k|k} = \hat{\boldsymbol{X}}_{k|k-1} + K\boldsymbol{\epsilon}_k . \quad (14)$$

If K is small, the prediction weights more compared to the observation in the estimate of the state. The covariance of the one-step prediction error, $\boldsymbol{R}_{k|k-1}$, is usually calculated in the updating scheme as well. With $\boldsymbol{\epsilon}_k$ and $\boldsymbol{R}_{k|k-1}$ at hand, we can evaluate the conditional density associated with the k'th observation. This recursion is applied to all observations in $\mathcal{Y}_N$, and with a given initial condition $\boldsymbol{X}_0$, the log-likelihood in (12) can be computed.

4.3. Details in the parameter estimation

As previously described, the return temperature exhibits inconsistent behaviour when the heat load is zero. Also, at these times, the return temperature is not of interest for control purposes. For these reasons, we choose to disregard the return temperature in the parameter estimation at times where the heat load is zero. That is, we need to ensure that the return temperature for these times does not affect the likelihood function. We thus add a very large constant (say 10^{20}) to the observation variance in the Kalman filter when the heat load is close to zero (say < 0.01 [kW]). As a result, the *observed* return temperature has negligible effect on the likelihood estimates during these times. Such actions are crucial to implement for applications in general, e.g. MPC, where indescribable dynamics occur or observations are not of interest and a Kalman filter is applied for state estimation. The larger variance on the observed return temperature ensures that it contributes very little to the state estimate at that point in time.

To evaluate the log-likelihood in Eq. (12), this paper uses the software CTSM-R [38]. To maximise Eq. (13), we use the `NLopt` optimisation library in `R` [39].

5. Results and discussion

This section presents the results in terms of parameter estimates, simulation of the model, and residual analysis. We compare a simulation of the model with the experiment data to see the model's performance over the entire data set given only the initial conditions. For the simulations, we use the same weather observations and set-points as inputs. Finally, we discuss the capabilities and strengths/weaknesses of the model.

5.1. Simulation results

Table 2 displays the parameter estimates for the model presented in Section 3. All parameters are strongly significant, except the solar radiation gain A_s. The explanation is likely that the data contains no significant solar radiation. The parameter thus becomes hard to determine without large uncertainty. But the solar radiation gain is an important disturbing factor for building climate control [40]. We thus intend to describe the solar radiation gain better in the future, when more experiments/data are available. The literature contains interesting approaches to model this, such as using B-splines to describe the varying solar gain during the day [41].

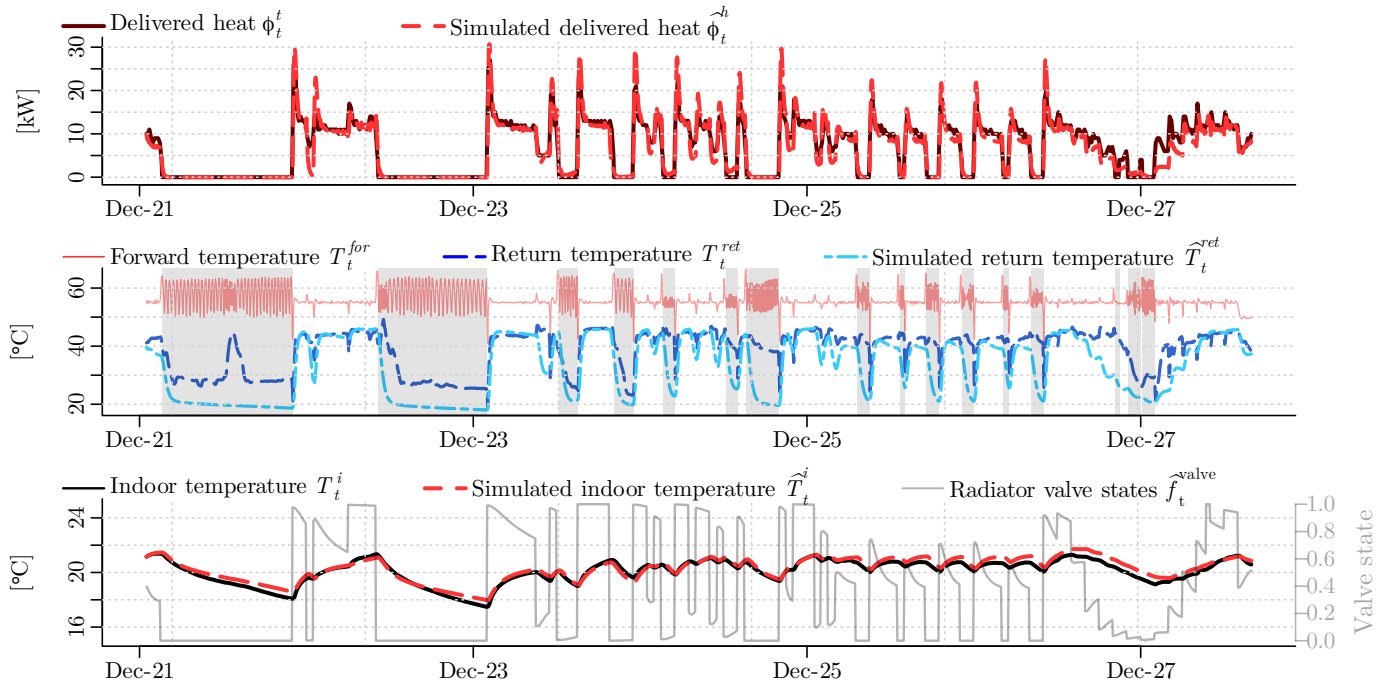

Figure 5: A simulation of the resulting model: given initial conditions, the model predicts the entire week. The model performs well given the long prediction horizon. The grey periods in the second plot depicts the periods where the observation variance on the return temperature is large.

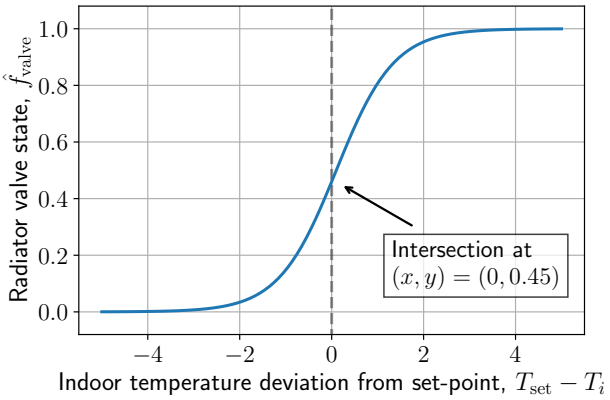

Figure 6: The estimated valve function

Figure 5 shows a simulation of the experiment given only the initial states. The model does a good job for all three variables. It predicts long into the future and still shows good accuracy without large drifts. That is, the model performs very well on long prediction horizons. This is crucial for the performance of MPC. This simulation, however, uses the same weather disturbances as the data. For practical purposes, weather forecasts are not perfect, which affects the prediction performance [42, 43].

The predicted heat was a challenge to model, but the simulation suggests that the model captures the most crucial dynamics. However, the simulation also indicates that the model is not able to "turn off" the delivered heat fast enough compared to data, as it seems to go slower to zero. This flaw comes from the fact that we model the flow as an SDE itself, Φ_t. Thus, the flow goes exponentially towards the term $\Phi_{\text{max}} f_t^{\text{valve}}$ (which never equals zero due to the sigmoid curve).

We found that the dynamics of the return temperature were hard to mimic and capture, especially when the heat is turned off. Figure 5 confirms that we are somewhat capable of predicting the return temperature whenever the heating system is turned on.

The indoor air temperature in Figure 5 seems to catch the overall dynamics of the data. The building model, however, does seem to be a bit too well insulated by the looks of the long periods where no heat is delivered. The simulated temperature decreases slower compared to data. Also, from around December 25th and onwards, the simulated indoor air temperature seems to drift a bit upwards compared to data. In this period, the estimated valve states are never fully opened, indicating that the set-point and observed temperatures are very close. Thus, the upwards drift of the simulation could come from the sigmoid curve of the valve function, since it never fully closes and is still open even when the observed temperature is above the set-point.

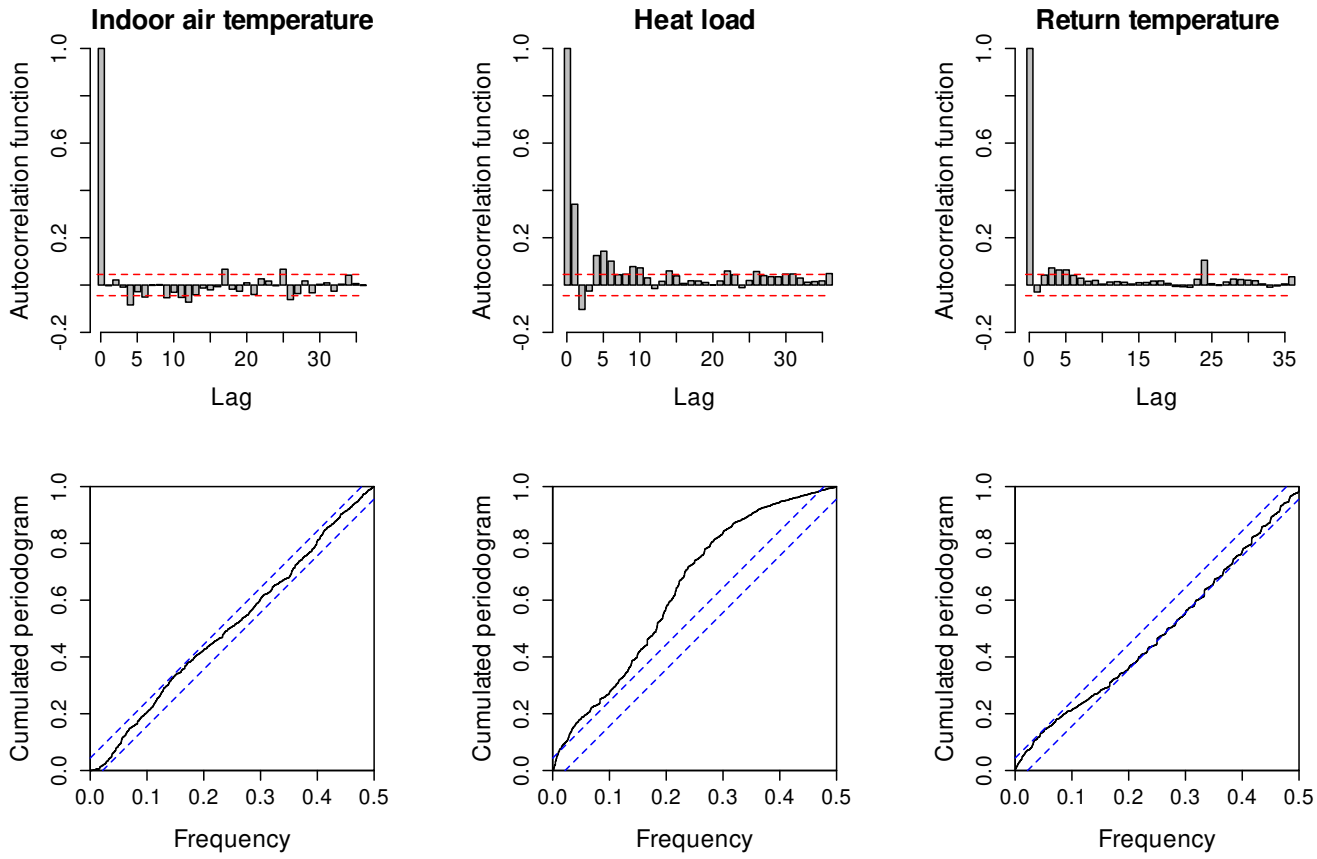

Figure 7: The estimated autocorrelation function and the cumulated periodogram of the 1-step prediction errors. Each column corresponds to a variable.

5.2. *Residual analysis*

Figure 7 shows the estimated autocorrelation function and the cumulated periodogram of the 1-step prediction errors for each of the variables. Both the autocorrelation function and the cumulated periodogram indicate that the residuals of the indoor air temperature and the return temperature can be classified as white noise. However, the heat load residuals are governed by some minor autocorrelation in the first few lags. Inspection of the spectral density and the residual plots confirms that non-uniformity of the spectrum primarily comes from the heater state's exponential decay towards zero when the heat is turned off. However, this is a minor autocorrelation that is not going to impact MPC performance significantly.

5.3. *Future work*

Since the experiment took place during the Christmas holiday, the building was not occupied at any time. However, human occupancy/behaviour is important to model and include in MPC [44]. Also, due to the lag of occupancy, we do not know how open windows affect the indoor air temperature. These are important topics to investigate further to accurately model the thermal dynamics of the building [45].

6. Conclusion

This paper introduced a physically inspired SDE-based non-linear model to describe the complex heat dynamics of a school building with water-based heating. The purpose of the model is to predict the indoor air temperature, the heat load, and the return temperature of the water in the space heating (SH) system. We model the thermostats in the radiators using a Sigmoid function to describe the level of water flow through the radiators. We fitted the parameters in the model from time-series data using maximum likelihood estimation. To validate the estimated model, we compared a simulation of the model, only given the initial conditions and disturbances, to data. This showed great accuracy over an entire week. The residual analysis indicated that the model lacks some dynamical descriptions of the heat load. We believe the reason might be that the model does not shut down the heat load fast enough. Beside this, the model looks promising for enabling MPC and e.g. embedded forecasts.

7. Acknowledgement

The authors received funding from the following projects; *Sustainable plus energy neighbourhoods (syn.ikia)* (H2020 No. 869918), *Centre for IT-Intelligent Energy Systems (CITIES)*

Table 2: The parameter estimates together with their statistical properties

Parameter	Estimate	95% confidence interval	Unit
T_{offset}	-0.101	[-0.081, -0.121]	[°C]
C_h	0.134	[0.128, 0.140]	[kJ/°C]
C_f	0.198	[0.194, 0.202]	
R_{fr}	2.030	[1.898, 2.162]	[°C h / kJ]
C_i	9.57	[9.40, 9.742]	[kJ/°C]
C_w	45.36	[42.80, 47.92]	[kJ/°C]
R_{ih}	2.151	[2.121, 2.181]	[°C h / kJ]
R_{iw}	0.199	[0.195, 0.203]	[°C h / kJ]
R_{wa}	2.251	[1.775, 2.727]	[°C h / kJ]
A_s	7.600	[-1.443, 16.64]	[m^2]
σ_1	8.6e-4	[9.7e-5, 0.008]	[°C]
σ_2	0.429	[0.419, 0.439]	[°C]
σ_3	111.6	[107.6, 118.0]	[kg/h]
σ_4	1.647	[1.144, 2.370]	[°C]
σ_5	6.469	[6.327, 6.612]	[°C]
R_1	9.6e-7	[1.1e-7, 8.5e-6]	[°C]
R_2	2.7e-4	[5.2e-6, 0.014]	[kW]
R_3	5.4e-3	[1.4e-3, 0.021]	[°C]
$\Phi_{\max}$	1145.3	[1133.5, 1157.1]	[kg/h]
α	1.592	[1.550, 1.634]	[1/°C]

(DSF 1305-00027B), Top-Up (Innovation Fund Denmark 9045-00017B), SCA+ (Interreg Öresund-Kattegat-Skagerrak), *Research Centre on Zero Emission Neighbourhoods in Smart Cities (FME-ZEN)* (Research Council of Norway, No. 257660), and *Flexibile Energy Denmark (FED)* (IFD 8090-00069B).

References

[1] M. Z. Hauschild, S. Kara, I. Røpke, Absolute sustainability: Challenges to life cycle engineering, CIRP Annals 69 (2020) 533 – 553. doi:https://doi.org/10.1016/j.cirp.2020.05.004.

[2] D. W. O'Neill, A. L. Fanning, W. F. Lamb, J. K. Steinberger, A good life for all within planetary boundaries, Nature Sustainability 1 (2018) 88–95. doi:10.1038/s41893-018-0021-4.

[3] N. O'Connell, P. Pinson, H. Madsen, M. O'Malley, Benefits and challenges of electrical demand response: A critical review, Renewable and Sustainable Energy Reviews 39 (2014) 686–699. doi:10.1016/j.rser.2014.07.098.

[4] S. Østergaard Jensen, J. Parker, P. Engelmann, A. Joanna, Examples of Energy Flexibility in Buildings, Technical Report, 2019.

[5] G. Serale, M. Fiorentini, A. Capozzoli, D. Bernardini, A. Bemporad, Model predictive control (mpc) for enhancing building and hvac system energy efficiency: Problem formulation, applications and opportunities, Energies 11 (2018) 631. doi:10.3390/en11030631.

[6] J. Drgoňa, J. Arroyo, I. Cupeiro Figueroa, D. Blum, K. Arendt, D. Kim, E. P. Ollé, J. Oravec, M. Wetter, D. L. Vrabie, L. Helsen, All you need to know about model predictive control for buildings, Annual Reviews in Control (2020).

[7] D. Cali, D. Müller, H. Madsen, Benefits of the inclusion of occupant behaviour profiles in the simulation of the energy performance of buildings, in: Proceedings of 16th IBPSA International Conference & Exhibition Building Simulation 2019, 2019. URL: http://buildingsimulation2019.org, building Simulation 2019, BS 2019 ; Conference date: 02-09-2019 Through 04-09-2019.

[8] S. Wolf, D. Cali, M. Alonso, R. Li, R. Andersen, J. Krogstie, H. Madsen, Room-level occupancy simulation model for private households, volume 1343, IOP Publishing, 2019. URL: https://cisbat.epfl.ch/, cISBAT 2019 : Climate Resilient Cities - Energy Efficiency & Renewables in the Digital Era, CISBAT 2019 ; Conference date: 04-09-2019 Through 06-09-2019.

[9] A. Duun-Henriksen, S. Schmidt, R. Røge, J. Møller, K. Nørgaard, J. Jørgensen, H. Madsen, Model identification using stochastic differential equation grey-box models in diabetes, Journal of diabetes science and technology 7 (2013) 431–440. doi:10.1177/193229681300700220.

[10] A. Afram, F. Janabi-Sharifi, Gray-box modeling and validation of residential HVAC system for control system design, Applied Energy 137 (2015) 134 – 150. doi:https://doi.org/10.1016/j.apenergy.2014.10.026.

[11] S. Wang, X. Xu, Simplified building model for transient thermal performance estimation using ga-based parameter identification, International Journal of Thermal Sciences 45 (2006) 419 – 432. doi:https://doi.org/10.1016/j.ijthermalsci.2005.06.009.

[12] M. Massano, E. Macii, E. Patti, A. Acquaviva, L. Bottaccioli, A grey-box model based on unscented kalman filter to estimate thermal dynamics in buildings, in: 2019 IEEE International Conference on Environment and Electrical Engineering and 2019 IEEE Industrial and Commercial Power Systems Europe (EEEIC / I CPS Europe), 2019, pp. 1–6.

[13] P. Bacher, H. Madsen, Identifying suitable models for the heat dynamics of buildings, Energy and Buildings 43 (2011) 1511 – 1522. doi:https://doi.org/10.1016/j.enbuild.2011.02.005.

[14] A. Scott, The Nonlinear Universe, Springer, Berlin, Heidelberg, 2007.

[15] M. Jiménez, H. Madsen, J. Bloem, B. Dammann, Estimation of non-linear continuous time models for the heat exchange dynamics of building integrated photovoltaic modules, Energy and Buildings 40 (2008) 157 – 167. doi:https://doi.org/10.1016/j.enbuild.2007.02.026.

[16] C. Lodi, P. Bacher, J. Cipriano, H. Madsen, Modelling the heat dynamics of a monitored test reference environment for building integrated photovoltaic systems using stochastic differential equations, Energy and Buildings 50 (2012) 273 – 281. doi:https://doi.org/10.1016/j.enbuild.2012.03.046.

[17] E. Wernholt, S. Gunnarsson, Nonlinear grey-box identification of industrial robots containing flexibilities, IFAC Proceedings Volumes 38 (2005) 356 – 361. doi:https://doi.org/10.3182/20050703-6-CZ-1902.00060, 16th IFAC World Congress.

[18] J. K. Møller, H. Madsen, J. Carstensen, Parameter estimation in a simple stochastic differential equation for phytoplankton modelling, Ecological Modelling 222 (2011) 1793 – 1799. doi:https://doi.org/10.1016/j.ecolmodel.2011.03.025.

[19] B. Øksendal, Stochastic Differential Equations (3rd Ed.): An Introduction with Applications, Springer-Verlag, Berlin, Heidelberg, 1992.

[20] D. Boiroux, J. B. Jørgensen, A nonlinear model predictive control strategy for glucose control in people with type 1 diabetes, IFAC-PapersOnLine 51 (2018) 192 – 197.

[21] C. A. Thilker, R. G. Junker, P. Bacher, J. B. Jørgensen, H. Madsen, Model predictive control based on stochastic differential equations, in: C. Ghiaus, M. Amayri, S. Ploix (Eds.), Towards Energy Smart Homes: Algorithms, technologies, and applications, Springer, 2021. Publication expected in 2021.

[22] P. A. Stentoft, T. Munk-Nielsen, L. Vezzaro, H. Madsen, P. S. Mikkelsen, J. K. Møller, Towards model predictive control: online predictions of ammonium and nitrate removal by using a stochastic ASM, Water Science and Technology 79 (2018) 51–62.

[23] B. Xu, L. Fu, H. Di, Dynamic simulation of space heating systems with radiators controlled by trvs in buildings, Energy and Buildings 40 (2008) 1755 – 1764. doi:https://doi.org/10.1016/j.enbuild.2008.03.004.

[24] M. S. Mohseni, D. Gotthardsson, R. Hållbus, H. Vallmark, Analysis of the thermal performance of hydronic radiators and building envelop: Developing experimental (step response) and theoretical models and using simulink to investigate different control strategies, Civil Engineering Research Journal 2 (2017). doi:10.19080/CERJ.2017.02.555595.

[25] L. H. Hansen, Stochastic modelling of central heating systems, Ph.D. thesis, Technical University of Denmark, 1997. URL: http://www2.imm.dtu.dk/pubdb/pubs/2460-full.html.

[26] C. G. Bruun, Optimization of Building Operation Using High-resolution Sensor Data, Master's thesis, Technical University of Denmark, Department of Civil Engineering, 2019. Can be found at https://www.findit.dtu.dk.

[27] H. Madsen, J. Holst, Estimation of continuous-time models for the heat dynamics of a building, Energy and Building 22 (1995) 67–79.

[28] K. K. Andersen, H. Madsen, L. H. Hansen, Modelling the heat dynamics of a building using stochastic differential equations, Energy and Buildings 31 (2000) 13 – 24.

[29] B. Nielsen, H. Madsen, Identification of a linear continuous time stochastic model of the heat dynamics of a greenhouse, Journal of Agricultural Engineering Research 71 (1998) 249 – 256. doi:https://doi.org/10.1006/jaer.1998.0322.

[30] O. Brastein, D. Perera, C. Pfeifer, N.-O. Skeie, Parameter estimation for grey-box models of building thermal behaviour, Energy and Buildings 169 (2018) 58 – 68. doi:https://doi.org/10.1016/j.enbuild.2018.03.057.

[31] H. Harb, N. Boyanov, L. Hernandez, R. Streblow, D. Müller, Development and validation of grey-box models for forecasting the thermal response of occupied buildings, Energy and Buildings 117 (2016) 199 – 207. doi:https://doi.org/10.1016/j.enbuild.2016.02.021.

[32] H. Madsen, Time series analysis, Chapman & Hall, 2007.

[33] Y. Pawitan, In All Likelihood: Statistical Modelling and Inference Using Likelihood, Oxford: Clarendon Press, 2006.

[34] R. E. Kalman, A New Approach to Linear Filtering and Prediction Problems, Journal of Basic Engineering 82 (1960) 35–45.

[35] J. K. Møller, H. Madsen, From State Dependent Diffusion to Constant Diffusion in Stochastic Differential Equations by the Lamperti Transform, IMM-Technical Report-2010-16, Technical University of Denmark, DTU Informatics, Building 321, 2010.

[36] C. A. Thilker, Optimization for Smart Energy Systems, Master's thesis, Technical University of Denmark, Department of Applied Mathematics and Computer Science, 2020.

[37] P. Frogerais, J.-J. Bellanger, L. Senhadji, Various ways to compute the continuous-discrete extended kalman filter, IEEE Transactions on Automatic Control - IEEE TRANS AUTOMAT CONTR 57 (2012) 1000–1004. doi:10.1109/TAC.2011.2168129.

[38] R. Juhl, N. R. Kristensen, P. Bacher, J. Kloppenborg, H. Madsen, CTSM-R User Guide, 2013. Technical University of Denmark.

[39] S. G. Johnson, The nlopt nonlinear-optimization package, 2020. URL: http://github.com/stevengj/nlopt.

[40] T. Mingfang, Solar control for buildings, Building and Environment 37 (2002) 659 – 664. doi:https://doi.org/10.1016/S0360-1323(01)00063-4.

[41] C. Rasmussen, L. Frölke, P. Bacher, H. Madsen, C. Rode, Semi-parametric modelling of sun position dependent solar gain using b-splines in grey-box models, Solar Energy 195 (2020) 249 – 258.

[42] C. Thilker, H. Madsen, J. Jørgensen, Advanced forecasting and disturbance modelling for model predictive control of smart energy systems, Applied Energy 292 (2021). doi:10.1016/j.apenergy.2021.116889.

[43] F. Oldewurtel, A. Parisio, C. N. Jones, D. Gyalistras, M. Gwerder, V. Stauch, B. Lehmann, M. Morari, Use of model predictive control and weather forecasts for energy efficient building climate control, Energy and Buildings 45 (2012) 15 – 27. doi:https://doi.org/10.1016/j.enbuild.2011.09.022.

[44] F. Oldewurtel, D. Sturzenegger, M. Morari, Importance of occupancy information for building climate control, Applied Energy 101 (2013) 521 – 532. doi:https://doi.org/10.1016/j.apenergy.2012.06.014, sustainable Development of Energy, Water and Environment Systems.

[45] H. B. Rijal, P. Tuohy, F. Nicol, M. A. Humphreys, A. Samuel, J. Clarke, Development of an adaptive window-opening algorithm to predict the thermal comfort, energy use and overheating in buildings, Journal of Building Performance Simulation 1 (2008) 17–30. doi:10.1080/19401490701868448.

Paper E

NON-LINEAR MODEL PREDICTIVE CONTROL FOR SMART HEATING OF BUILDINGS

Authors:
Christian Ankerstjerne Thilker, Hjörleifur G. Bergsteinsson, Peder Bacher, Henrik Madsen, Davide Calì and Rune G. Junker.

Published in:
E3S Web Conf.

Non-linear Model Predictive Control for a Smart Heating System of old Danish School Building

Technical University of Denmark, Department of Applied Mathematics and Computer Science, Asmussens Allé, Building 303B, DK-2800 Kgs. Lyngby, Denmark

Abstract

Intelligent and flexible operation of components in district heating systems can play a crucial role in integrating larger shares of renewable energy sources in energy systems. Buildings are one of the crucial components that will enable flexibility in the district heating by using intelligent operation. Recent work suggests that such improved operation at the same time can increase thermal comfort and lower economic costs. We have digitalised the heating system in a Danish school by adding IoT devices, such as thermostats and temperature sensors to demonstrate the possibilities of making buildings intelligent. Based on experimental data, this paper introduces a non-linear grey-box model of the thermal dynamics of the building based on results from Thilker et al. [1]. A non-linear model predictive control method is presented for the thermostatic set-point control of the building's radiators. Based on the building model and the control algorithm, simulation studies are carried out to show the flexibility potential of the building – when used for lowering the return temperature the results suggest that economic costs can be lowered by around 10% using predictive control.

Keywords: Grey-box models, Stochastic differential equations, Non-linear model predictive control, District heating, Smart energy systems

1. Introduction

Digitisation of heating systems, i.e., smart thermostats and indoor climate sensors creates the possibility of making buildings smart by having heating data. This, however, does not yet make the building (or the heating system) smart as it does not yet use the data to make the system efficient or flexible. Without intelligence, the system is just data-rich. The system becomes intelligent when it uses the data to e.g. lower some cost functions, i.e. lower the heating costs without violating thermal comfort or reduce heat consumption during peak hours (known as peak shaving). Model Predictive Control (MPC) is a method that enables a system to become intelligent [2]. It is a control method that minimises some predefined cost function while satisfying a set of constraints. MPC has become very popular for the HVAC sector in the past years as it makes the system intelligent by making it either efficient or flexible [3–5]. The advantage of the MPC over other control methods is its ability to predict the future behaviour of the system. Thereby, the MPC can take weather predictions and future activities into account when optimising the manipulated variables (e.g., desire temperature in a room) of the system [6]. MPC setups usually run in a closed-loop where the controller gets feedback on how the system reacted to the latest input or disturbance. The MPC is based on a model (e.g. a set of differential equations) that describes the behaviour of the system and generates predictions of the system's future behaviour. This paper considers the heating system of an old Danish school building that has been digitalised with the objective to make it intelligent using MPC. For this building, a non-linear grey-box model is presented based on physics and data to describe the behaviour of the building's heat dynamics. Grey-box modelling is a well-known procedure used for system identification and modelling the dynamics of buildings [7, 8]. The parameters of the building model are estimated using the CTSMR software [9]. A non-linear MPC uses the grey-box model to control the heating system according to some thermal comfort constraints. The MPC utilises weather predictions of the solar irradiance and ambient air temperature to compute the optimal radiator set-points to obtain a desired indoor air temperature. The objective of the controller presented in this paper is to lower the heating cost of the building. This enables us to use various price signals that reflect how expensive heat is at any given time. Consequently, the controller minimises the heating costs by heating when the price is low. This methodology has shown to be fruitful for lowering the electricity consumption of a smart solar tank modelled by a grey-box model, where the MPC takes advantage of future disturbances and its flexibility [10]. Another example is electrical heating of a residential house equipped with a heat pump where the goal is to lower the electricity costs with varying electricity prices [11].

The main contribution of this work is to demonstrate how to use a non-linear grey box model for model predictive control. We present a multiple shooting method to solve the optimal control problem related to the MPC and incorporate numerical weather forecasts as future inputs. Another contribution is to illustrate the effects of two different price signals for the MPC in simulation studies. The first shows how to make the building flexible by utilising the right price signal. The second shows

Table 1: The parameter estimates together with their statistical properties

Parameter	Estimate	Unit
T_{offset}	-0.101	[°C]
C_h	0.134	[kJ/°C]
C_f	0.198	
R_{fr}	2.030	[°C h / kJ]
C_i	9.57	[kJ/°C]
C_w	45.36	[kJ/°C]
R_{ih}	2.151	[°C h / kJ]
R_{iw}	0.199	[°C h / kJ]
R_{wa}	2.251	[°C h / kJ]
A_s	7.600	[m^2]
σ_1	8.6e-4	[°C]
σ_2	0.429	[°C]
σ_3	111.6	[kg/h]
σ_4	1.647	[°C]
σ_5	6.469	[°C]
R_1	9.6e-7	[°C]
R_2	2.7e-4	[kW]
R_3	5.4e-3	[°C]
Φ_{max}	1145.3	[kg/h]
α	1.592	[1/°C]

how to optimally control the building in order to minimise the economic costs associated with heating a Danish building in a District heating network. The result of the MPC is compared to a simple control strategy which is among the current standards in buildings and the influence of different price signals is investigated.

1.1. Structure and outline of the paper

The remainder of this article is organized as follows. Section 2 presents the building and modelling scheme along with the parameter identification method and its results. In Section, 3 introduces the non-linear MPC method that is used to control the building. The simulation results are presented and discussed in Section 4. The paper is then concluded in Section 5.

2. Non-linear building model

This section introduces the non-linear building model used in the present paper. The model is thoroughly introduced and discussed in Thilker et al. [1] and it is recommended that the reader reads it for further details on the building and model. The building acts as a school and has 17 rooms in total divided over three floors. Fig. 1 displays a picture and an electronic reconstruction of the building. The building was built in 1929, and the insulation has not been renovated to meet today's standards. The building is supplied by water-based heating from the local district heating service. To distribute heat to the rooms, radiators are placed in the rooms that are connected to the building heat exchanger. Analyses of the building performance indicate that the heating power of the radiators are under-dimensioned in some rooms, meaning that they do not get a sufficient amount of heat to maintain a comfortable temperature [12].

To make the building intelligent and enable real-time control, sensors are needed to measure various attributes of the building. Accordingly, temperature sensors have been installed in each room (which are also able to measure CO_2-levels and humidity), equipped each radiator with an electronic thermostat, and installed heat load-meters in the BMS-system. Furthermore, the temperature of the water is measured both when it returns to- and leaves the heat exchanger, which is connected to the district heating. All sensors are connected to the cloud and the system receives data frequently. This setup allows the building operator to monitor and control all of the components related to the heating system in the building.

2.1. Building model

We consider a non-linear model on the form

$$\mathrm{d}\boldsymbol{x}(t) = f(\boldsymbol{x}(t), u(t), \boldsymbol{d}(t))\mathrm{d}t + g(\boldsymbol{x}(t), u(t), \boldsymbol{d}(t))\mathrm{d}\boldsymbol{\omega}(t)\,, \quad (1a)$$

$$\boldsymbol{y}_k = h(\boldsymbol{x}(t_k)) + \boldsymbol{w}_k, \qquad \boldsymbol{w}_k \sim N(0, \boldsymbol{R})\,, \quad (1b)$$

where $\boldsymbol{x}(t)$ is the state vector, $u(t)$ is the control input, $\boldsymbol{d}(t)$ is the disturbances, and $\boldsymbol{R}$ is the observation covariance. $\boldsymbol{\omega}(t)$ is Brownian motion and reflects the uncertainty of the model. Eq. (1) is structurally similar to ordinary differential equations except for the diffusion term. It has the advantages that it describes random effects that are too complex to model deterministically and it predicts uncertainty as well, e.g. the variance of the estimates Øksendal [13].

We simplify the thermal control problem by considering the *average* of the measured temperature in all 17 rooms

$$T_i = \frac{1}{n}\sum_{k=1}^{n} T_k. \quad (2)$$

Consequently, some rooms are going to be warmer or colder, however the problem is simplified significantly in terms of dimensionality. It is important for real-time MPC that the model is small enough to compute the control input without too much delay. In the following, we consider a system with the states $\boldsymbol{x}(t) = [T_i(t), T_w(t), \Phi(t), T_h(t), T_{\text{ret}}(t)]$, where T_i is the average indoor air temperature, T_w is the temperature of the building wall, Φ is the flow of the water in the radiator circuit, T_h is the temperature of the radiators, and T_{ret} is the temperature of the returning water (going to the heat exchanger of the building). The control input to the model, $u(t)$, is the set-points of the radiator thermostats. To estimate the valve-opening state of the thermostats, the following sigmoid function is used:

$$f_{\text{valve}}(t) = \frac{1}{1 + e^{-\alpha(u(t) - T_i(t) + T_{\text{offset}}(t))}}\,, \quad (3)$$

where u is the thermostat set-point, α determines the slope of the sigmoid function, and $T_{\text{offset}}(t)$ is an offset that models the physical distance between the temperature sensors in the room and the thermostats of the radiators. f_{valve} therefore estimates how open the radiator valves are (1 being fully open and 0 being fully closed), i.e. how much water flows through the radiators.

Figure 1: The building in question

The disturbances include the ambient air temperature and solar irradiance $\boldsymbol{d}(t) = [T_a(t), \phi_s(t)]^\intercal$.

The building dynamics model is

$$f(\boldsymbol{x}(t), u(t), \boldsymbol{d}(t)) = \begin{bmatrix} \frac{1}{C_i}\left(\frac{1}{R_{ih}}\left(T_t^h - T_t^i\right) + \frac{1}{R_{iw}}\left(T_t^w - T_t^i\right) + A_w\phi_t^s\right) \\ \frac{1}{C_w}\left(\frac{1}{R_{iw}}\left(T_t^i - T_t^w\right) + \frac{1}{R_{wa}}\left(T_t^a - T_t^w\right)\right) \\ \frac{1}{C_f}\left(\Phi_{\max} f_{\text{valve}} - \Phi_t\right) \\ \frac{1}{C_h}\left(\Phi_t c_{p,w}\left(T_t^{\text{for}} - T_t^h\right) + \frac{1}{R_{ih}}\left(T_t^i - T_t^h\right)\right) \\ \frac{1}{C_h}\left(\frac{1}{R_{fr}}\left(T_t^h - T_t^{\text{ret}}\right)\right) \end{bmatrix}. \tag{4}$$

In Eq. (4), A_w is the effective area of the solar radiation gain, $c_{p,w}$ is the specific heat capacity of water, and $\Phi_{\max}$ is the maximum water flow in the radiator circuit. T_{for} is the temperature of the forward going water on the building side of the heat exchanger and is kept constant at 55 °C at all times. The diffusion term in (1), g, has the simple form

$$g(\boldsymbol{x}(t), u(t), \boldsymbol{d}(t)) = \text{diag}(\sigma_1, \sigma_2, \sigma_3, \sigma_4, \sigma_5)\,. \tag{5}$$

Naturally, not all states of the building are observed. Instead, we are limited to the information available in the non-linear observation equation (1b):

$$\boldsymbol{y}_k = h(\boldsymbol{x}(t_k)) = \left[T_i(t_k),\ \phi_h(t),\ T_{\text{ret}}(t_k)\right]^{\text{T}}\,. \tag{6}$$

That is, we observe the average indoor air temperature $Ti(t_k)$, the heat load $\phi_h(t_k) = \Phi(t_k)(T_{\text{for}} - T_{\text{ret}}(t_k))$, and the return temperature $T_{\text{ret}}(t_k)$. Recall that the forward temperature is known and is $T_{\text{for}} = 55$ °C.

2.2. Model parameter estimation

We use CTSMR [9] to estimate the parameters in the continuous-time stochastic model. The parameter estimation is based on the maximum likelihood principle [14]. That is, we maximise the likelihood function, which is a function of the parameters

$$\mathcal{L}(\boldsymbol{\theta}) = p(\boldsymbol{x}_0)\prod_{k=1}^{N} p(\boldsymbol{y}_k|\boldsymbol{\theta}) \tag{7}$$

where p is the probability of observing $\boldsymbol{Y}_k$ with the model in Eq. (4) and Eq. (5) given the parameters $\boldsymbol{\theta}$. Given the model structure in Eq. (4) and Eq. (5), as well as appropriate informative data, any unknown parameters can be estimated.

Table 1 lists the parameter estimates from the estimation procedure. Fig. 2 compares the fit of the resulting model to the data and indicates a good match. It shall be noted, that the return temperature measurements are not representative when the heat load is zero and the water flow in the building is zero. We thus put very low weight on the return temperature observations in the estimation procedure in these time intervals (indicated by the greyed-out periods in the figure).

3. Non-linear Model Predictive Control: a multiple shooting method

This section introduces a direct multiple-shooting method for solving the particular non-linear MPC problem. It also discusses a method to discretise the optimisation problem to make it numerically tractable. The optimisation problem lies the basis for computing the set-points for the radiators. However, solving the optimisation problem requires us to know the entire state of the system, $\boldsymbol{x}$. Reconstructing the system states based on observation, $\boldsymbol{y}$, the continuous-discrete extended Kalman filter is used [15].

This paper considers an optimal control problem on the following form

$$\min_{\boldsymbol{x},\boldsymbol{u}} \quad \varphi = \int_{t_k}^{t_k+T} \ell(\boldsymbol{x}(t), u(t), \boldsymbol{d}(t))\,, \tag{8a}$$

$$s.t. \quad \boldsymbol{x}(t_k) = \boldsymbol{x}(0)\,, \tag{8b}$$

$$\dot{\boldsymbol{x}}(t) = f(\boldsymbol{x}(t), u(t), \boldsymbol{d}(t))\,, \tag{8c}$$

$$u_{\min}(t) \le u(t) \le u_{\max}(t)\,, \tag{8d}$$

$$T_{\min}(t) \le T_i(t) \le T_{\max}(t)\,, \tag{8e}$$

where T is the prediction and control horizon, ℓ is cost function, and $f(\boldsymbol{x}(t), u(t), \boldsymbol{d}(t))$ is the model equations in (4).

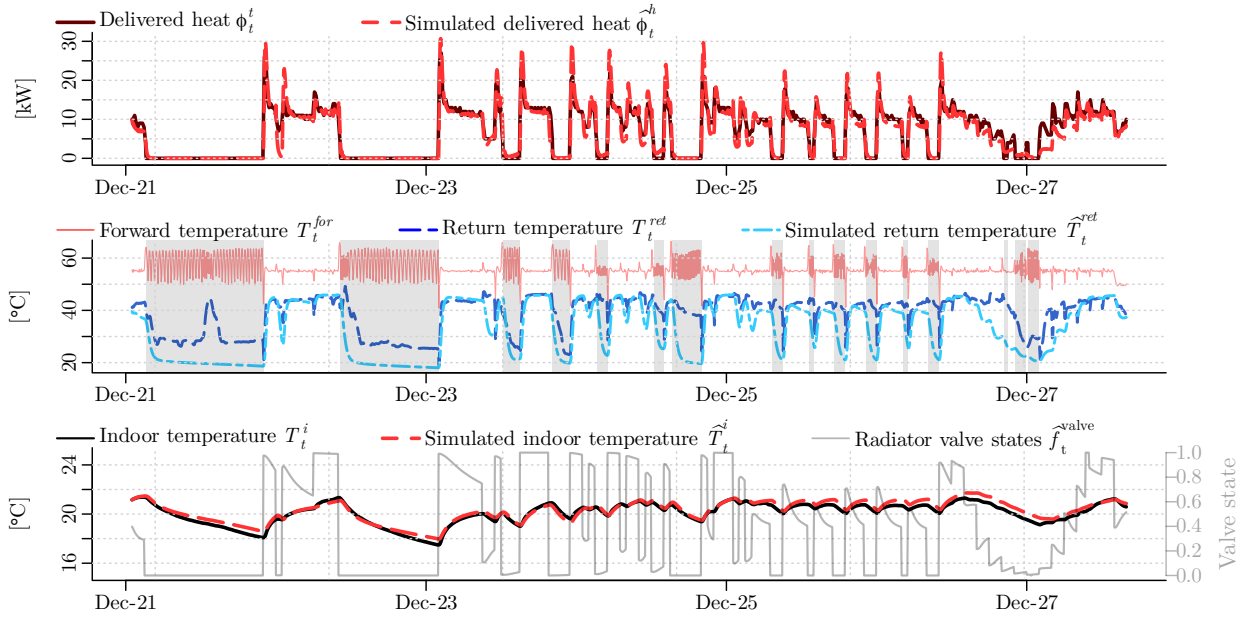

Figure 2: Experimental data together with the estimated heat load, air temperature, and return water temperature by the model. The greyed-out periods in the second graph indicates periods where the return temperature is disregarded, because the observations do not represent the actual return temperature.

3.1. *Discrete-time approximation of the optimal control problem*

To make the optimal control problem in Eq. (8) numerically tractable, we propose a multiple shooting method to discretise the problem. Multiple shooting is a *simultaneous* method in the sense that the state variables also are a part of the optimisation problem.

The problem is discretised in the sense that the system consider $\boldsymbol{x}$ at discrete time points $t_k, t_{k+1}, \ldots, t_{k+N}$ starting from the initial time t_k till $t_k + T$. Now, define a function $\phi(\boldsymbol{x}(t), u(t), \boldsymbol{d}(t))$ that computes the solution to the following initial value problem

$$\dot{\boldsymbol{x}}(t) = f(\boldsymbol{x}(t), u(t), \boldsymbol{d}(t)) \tag{9a}$$
$$\boldsymbol{x}(t_k) = \boldsymbol{x}_k\,, \qquad \text{(initial condition)} \tag{9b}$$

at time t_{k+1}. Hence, $\phi(\boldsymbol{x}(t_k), u(t), \boldsymbol{d}(t)) = \boldsymbol{x}(t_{k+1})$ is a function that integrates the system forward to the next time instance given the input and disturbances in the time interval $[t_{k+i}, t_{k+1+i}[, \forall i$. To simplify the optimisation problem, we assume that the set-points, $u(t)$, and the disturbances, $\boldsymbol{d}(t)$, are piece-wise constant in each time interval $[t_k, t_{k+1}[$

$$u(t) = u_k, \qquad t \in [t_k, t_{k+1}[\,, \tag{10}$$
$$\boldsymbol{d}(t) = \boldsymbol{d}_k, \qquad t \in [t_k, t_{k+1}[\,. \tag{11}$$

The optimal control problem therefore simplifies to

$$\min_{\{\boldsymbol{x}_k\}_{i=0}^{N},\, \{u_k\}_{i=0}^{N-1}} \quad \varphi = \sum_{i=0}^{N-1} L_k(\boldsymbol{x}_k, u_k, \boldsymbol{d}_k)\,, \tag{12a}$$
$$s.t. \quad \boldsymbol{x}_{k+1} = \phi(\boldsymbol{x}_k, u_k, \boldsymbol{d}_k)\,, \tag{12b}$$
$$\boldsymbol{u}_{\min,k} \le u_k \le \boldsymbol{u}_{\max,k}\,, \tag{12c}$$
$$T_{\min,k} \le T_{i,k} \le T_{\max,k}\,. \tag{12d}$$

In the above,

$$L_k = \{ \int_{t_k}^{t_{k+1}} \ell(\boldsymbol{x}(t), u_k, \boldsymbol{d}_k)\mathrm{d}t \;:\; \boldsymbol{x}_{k+1} = \phi(\boldsymbol{x}_k, u_k, \boldsymbol{d}_k),\, \boldsymbol{x}(t_k) = \boldsymbol{x}_k \} \tag{13}$$

is the quadrature of $\boldsymbol{x}(t)$ w.r.t ℓ in the time interval $[t_k, t_{k+1}]$.

For numerical computation of the minimisation problem in (12), we use CasADi [16], which offers easy numerical implementation and automatic differentiation for optimal control problems.

4. Simulation results

This section presents the results of two simulation studies. The first simulation investigates the flexibility of the building. The second investigates the ability to minimise the economic operational costs of heating the building using NMPC. We use the Euler-Maruyama simulation scheme to simulate from the SDE-model and the continuous-discrete extended Kalman filter to reconstruct the system state.

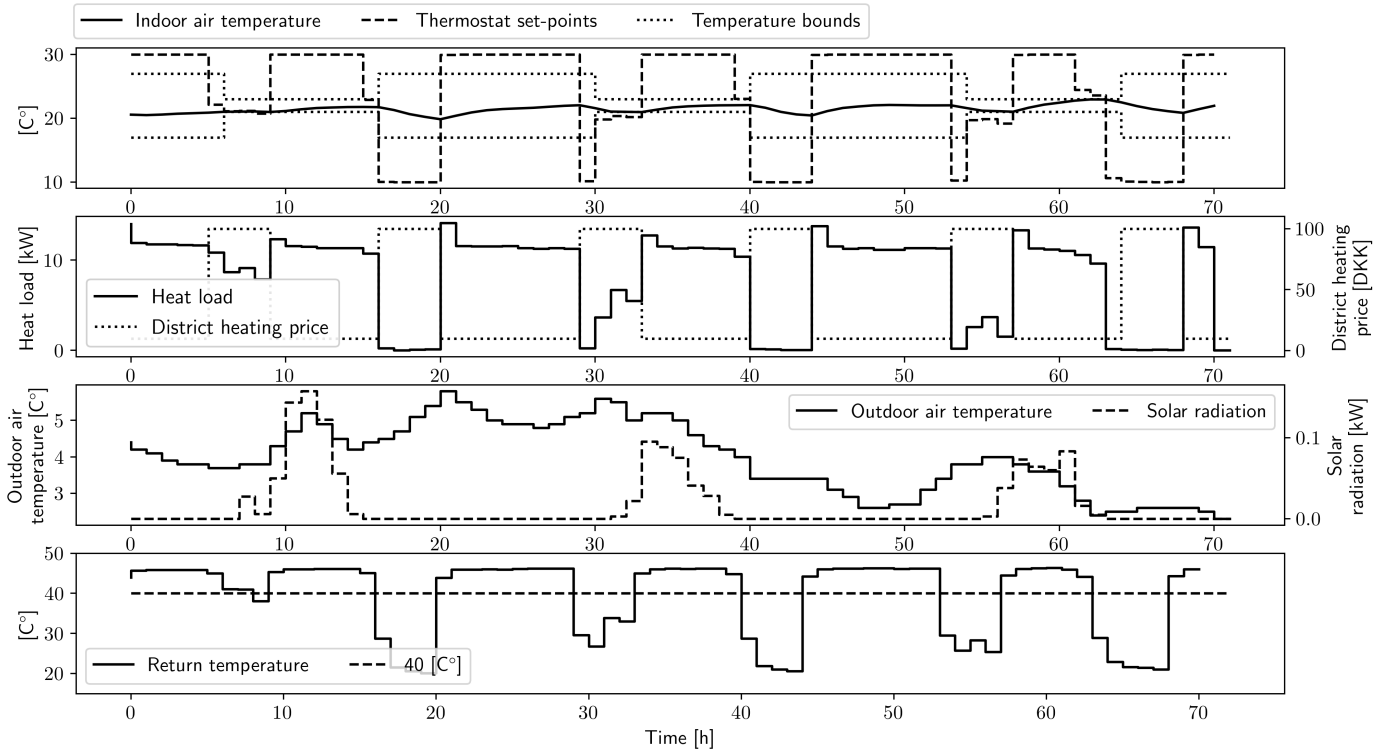

Figure 3: A small simulation of thermostatic set-point control of the building using a price signal that reflects peak hours and displays flexibility. The controller keeps the heat usage to a minimum during peak hours where the heat is expensive.

4.1. Simulation: Flexibility of the building

To investigate the flexibility of the building in a smart energy grid, we use a cost function in the MPC that takes a price signal. In a flexibility setting, the price signal reflects how "expensive" it is to heat the building at any given time. The cost function is defined as

$$\ell_1(\boldsymbol{x}(t), u(t), \boldsymbol{d}(t), s(t)) = c(t)\Phi(t)(T_{\text{for}} - T_{\text{ret}}(t)) + \rho s(t) \quad (14)$$

where c is the price signal, s is a slack variable to soften the indoor air temperature constraints and ρ is the slack penalty. Fig. 3 presents a simulation of the building model in Eq. (4) using the optimal control problem in Eq. (14) in a closed-loop setting with the time between control inputs and the prediction horizon equal to one hour and 24 hours, respectively. Furthermore, the controller has access to the future weather disturbances. In the simulation, the heating price is simply designed in order to see the effect of the MPC: it is most expensive at 100 DKK per kWh during peak hours in the mornings and evenings. The heat price is otherwise low at 10 DKK per kWh. As a result, the controller mainly heats outside peak hours and only does so if the temperature gets too low. Due to the under-dimensionality of the heating system and the building's poor insulation level, the controller still needs to supply some heat during the peak hours to maintain the desired temperature. The results suggest that the building can supply some flexibility under these circumstances. However, considering that the outdoor temperature in Denmark can become even lower than in the present simulation, the building will have less flexibility in such situations.

4.2. Simulation: Minimisation of operational costs by lowering return temperature

As a building owner in the danish district heating, one pays an additional fee if the return temperature is high since it is a source of heat loss in the grid. The pricing scheme is very different between district heating areas, both the price of heat and the penalty for not cooling the return adequately. In the present analysis, we set it quite progressively, namely as follows: if the return temperature is above 40 °C, the heat price increases 2% per extra degree Celsius of the return temperature. The cost-function where this is accounted for is

$$\ell_2(\boldsymbol{x}(t), u(t), \boldsymbol{d}(t), s(t), v(t)) = c(t)\Phi(t)(T_{\text{for}} - T_{\text{ret}}(t))\,(1 + 0.02v(t)) + \rho s(t) \quad (15)$$

where v is a slack variable that softens the upper constraint at 40 °C on the return temperature and 0.02 is the percent-wise increase in heat cost.

Fig. 4 displays a simulation study of the building model in (4) using the optimal control problem Eq. (15). The figure also depicts a baseline, which uses a simple set-point control that turns down the temperature during the night and back on during the day. The baseline represents the current practice in most buildings using rule-based control: a fixed set-point pattern used every day. This experiment reflects the actual economic costs of

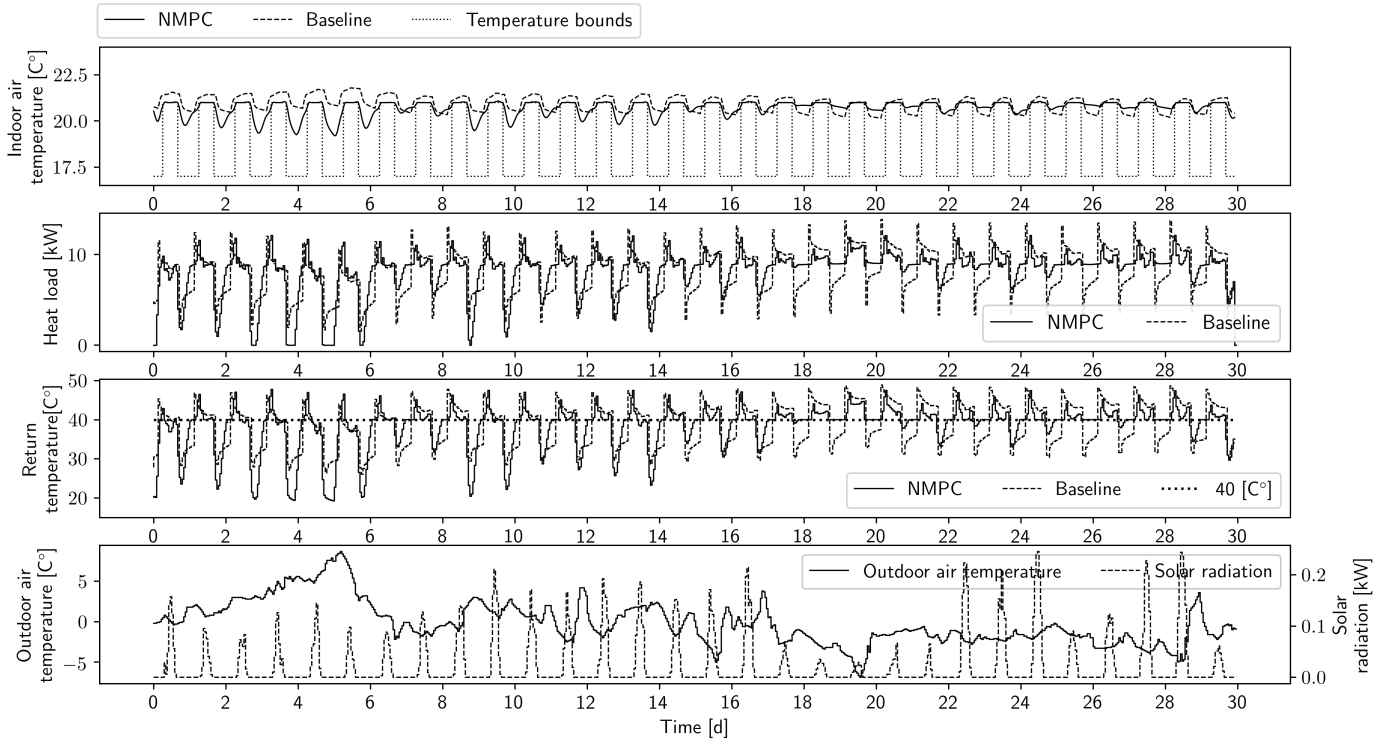

Figure 4: A simulation study that compares a current standard set-point control in today's buildings (Baseline) and the NMPC presented in this paper. The heat costs are constant at 0.1 EUR/kWh plus a penalty of 2% for each °C the return temperature is above 40 °C. Results suggest an economic reduction by around 10%.

operating the building together with the extra fee when the return temperature is too high. The results demonstrate the emphasis the controller puts on keeping the return temperature below 40 °C while supplying enough heat to comply with the constraints. The actual economic costs associated with each control strategy during the one simulated month are 4522.9 DKK and 4066.6 DKK for the baseline and MPC, respectively. This points toward economic savings of around 10% by using the proposed control strategy. Much of this reduction is explained by the ability of the controller to lower the return temperature and avoid extra penalties, which account for 382.2 DKK and 89.5 DKK, respectively. Especially during the cold periods, where extra heat is needed, the economic savings are high during those periods. Heat load is reduced from and 5891.4 kWh to 5742.5 (around 2.5 %) kWh by the MPC, which comes from the ability of the MPC to lower the temperature closer to the constraints. This optimisation not only benefits the building operators, but also benefits the district heating operators by significantly decreasing the amount of heat loss in the district heating system.

It should be stressed that these results apply only to the current settings and may vary according to different district heating areas and pricing schemes. Also, in a realistic setup with meteorological weather forecasts, building occupants, etc., the control performance may also be affected.

5. Conclusion and future work

This article introduced a non-linear grey-box model describing the heat dynamics of a large building. This model enabled us to predict and control the future evolution of temperatures and heating in a building. We presented a non-linear model predictive control method and used it in a simulation study to show it behaves. The results suggest that intelligent control of the heat supply unlocks the building's flexibility and supplies economic savings under a particular, but realistic, pricing scheme of up to 10% are available. The specific savings may vary depending on the district heating area since pricing schemes vary. Also, the controller had access to future weather, which in a realistic setting must be replaced with weather forecasts potentially decreasing the savings.

6. Acknowledgements

The research was supported through the project "Smart Cities Accelerator 2016–2020" funded by the EU program Interreg Öresund-Kattegat-Skagerrak, the European Regional Development Fond, as well as the CITIES project (DSF1305-00027B) and the FED project (Innovation fund Denmark ref. 8090-00069B).

References

[1] C. A. Thilker, H. G. Bergsteinsson, R. G. Junker, H. Madsen, D. Cali, P. Bacher, Non-linear grey-box modelling for the heat dynamics of buildings (2021). To be submitted to Energy and Buildings.

[2] J. Drgoňa, J. Arroyo, I. Cupeiro Figueroa, D. Blum, K. Arendt, D. Kim, E. P. Ollé, J. Oravec, M. Wetter, D. L. Vrabie, L. Helsen, All you need to know about model predictive control for buildings, Annual Reviews in Control 50 (2020) 190 – 232. URL: http://www.sciencedirect.com/science/article/pii/S1367578820300584. doi:https://doi.org/10.1016/j.arcontrol.2020.09.001.

[3] A. Q. Santos, B. N. Jørgensen, H. Kazmi, R. Ru-usu, A. Hasan, T. Péan, Y. Zhou, S. Cao, J. C. Young Jae Yu, R. G. Junker, C. Finck, A. S. Derakhtenjani, J. A. Candanedo, A. Athienitis, D. Christantoni, A. Kathirgamanathan, D. P. Finn, K. Zhang, Control strategies and algorithms for obtaining energy flexibility in buildings".EBC Annex 67: EnergyFlexible Buildings, Technical Report, 2019.

[4] H. Madsen, J. Parvizi, R. F. Halvgaard, L. E. Sokoler, J. B. Jørgensen, L. H. Hansen, K. B. Hilger, Control of electricity loads in future electric energy systems, Handbook of Clean Energy Systems (2015).

[5] C. A. Thilker, R. G. Junker, P. Bacher, J. B. Jørgensen, H. Madsen, Model predictive control based on stochastic differential equations, in: C. Ghiaus, M. Amayri, S. Ploix (Eds.), Towards Energy Smart Homes: Algorithms, technologies, and applications, Springer, 2021. Publication expected in first half of 2021.

[6] F. Oldewurtel, A. Parisio, C. N. Jones, D. Gyalistras, M. Gwerder, V. Stauch, B. Lehmann, M. Morari, Use of model predictive control and weather forecasts for energy efficient building climate control, Energy and Buildings 45 (2012) 15 – 27. URL: http://www.sciencedirect.com/science/article/pii/S0378778811004105. doi:https://doi.org/10.1016/j.enbuild.2011.09.022.

[7] P. Bacher, H. Madsen, Identifying suitable models for the heat dynamics of buildings, Energy and Buildings 43 (2011) 1511 – 1522. doi:https://doi.org/10.1016/j.enbuild.2011.02.005.

[8] K. K. Andersen, H. Madsen, L. H. Hansen, Modelling the heat dynamics of a building using stochastic differential equations, Energy and Buildings 31 (2000) 13 – 24.

[9] R. Juhl, N. R. Kristensen, P. Bacher, J. Kloppenborg, H. Madsen, CTSM-R User Guide, 2013. Technical University of Denmark.

[10] R. Halvgaard, P. Bacher, B. Perers, E. Andersen, S. Furbo, J. Jørgensen, N. Poulsen, H. Madsen, Model predictive control for a smart solar tank based on weather and consumption forecasts, volume 30, 2012, pp. 270–278. doi:10.1016/j.egypro.2012.11.032.

[11] H. A. Schlüter, D. Boiroux, N. K. Poulsen, H. Madsen, J. B. Jørgensen, Economic model predictive control for energy systems in smart homes, in: 2019 IEEE Conference on Control Technology and Applications (CCTA), 2019, pp. 598–604. doi:10.1109/CCTA.2019.8920663.

[12] B. C. G., Optimization of Building Operation Using High-resolution Sensor Data, Master's thesis, Technical University of Denmark, Department of Civil Engineering, 2019. Can be found at https://www.findit.dtu.dk.

[13] B. Øksendal, Stochastic Differential Equations (3rd Ed.): An Introduction with Applications, Springer-Verlag, Berlin, Heidelberg, 1992.

[14] Y. Pawitan, In All Likelihood: Statistical Modelling and Inference Using Likelihood, Oxford: Clarendon Press, 2006.

[15] P. Frogerais, J.-J. Bellanger, L. Senhadji, Various ways to compute the continuous-discrete extended kalman filter, IEEE Transactions on Automatic Control 57 (2012) 1000–1004.

[16] J. A. E. Andersson, J. Gillis, G. Horn, J. B. Rawlings, M. Diehl, CasADi – A software framework for nonlinear optimization and optimal control, Mathematical Programming Computation (In Press, 2018).

Paper F

DATA-DRIVEN METHODS FOR EFFICIENT OPERATION OF DISTRICT HEATING SYSTEMS

Authors:
Hjörleifur G. Bergsteinsson, Jan Kloppenborg Møller, Christian Ankerstjerne Thilker, Daniela Guericke, Alfred Heller, Torben Skov Nielsen, and Henrik Madsen.

Published in:
Handbook of Low Temperature District Heating.

Data-Driven Methods for Efficient Operation of District Heating Systems

Hjörleifur G. Bergsteinsson[1][0000−0001−8808−5528], Jan Kloppenborg Møller[1][0000−0002−6100−043X], Christian Ankerstjerne Thilker[1][0000−0002−5886−9652], Daniela Guericke[1][0000−0002−0668−9094], Alfred Heller[1][0000−0001−6062−3494], Torben Skov Nielsen[2], and Henrik Madsen[1][0000−0003−0690−3713]

[1] Department of Applied Mathematics and Computer Science, Technical University of Denmark, 2800 Kongens Lyngby, Denmark
[2] ENFOR, Holte 2840, Denmark

Abstract. In this chapter, data-driven methods for the efficient operation of DHSs are described. DHSs are inherently non-linear and time-varying systems as the heating demand is highly influenced by non-linear dependencies on the weather conditions as well as the occupancy behaviour. Furthermore, the dependency on flow and temperature in delivering the needed heat demand using the district heating network gives a non-linear dependency on these two signals. This chapter presents several data-driven models to handle the non-linear and time-varying phenomena in order to ensure an efficient operation. First, we introduce forecasts that are used to reach an optimal operation as forecasts are needed for both control and production planning, e.g heat demand and electricity price forecasts. Second, temperature control of a DHN will be introduced with a focus on how the physical characteristics of the network can be incorporated into a control scheme. A special focus will be on how to ensure that the temperatures in the network are high enough to ensure the needed heat supply for the attached buildings in the entire district heating network is met. We shall also briefly look at the role of smart buildings integrated into a DHN that can be used to enhance the efficiency and flexibility of a DHS.

Keywords: Data-Driven Methods · Adaptive Control · Heat Load Forecasting · Temperature optimization · Sector Coupling · Renewable Energy · Flexibility.

1 Introduction

The transition to a low-carbon society calls for fundamental changes in the energy system. Today energy systems are operated and planned such that the production follows the demand. However, an efficient implementation of a low-carbon society calls for the exact opposite approach, namely systems where demand follows production which will be dominated by renewables. This highlights

a need for new methods for planning and operation of energy systems. Most importantly the flexibility at virtually all types of end-users on all aggregation levels has to be unlocked. The typical variation in time of the energy produced by wind and solar power implies a need for flexibility that can be offered by well designed district heating and district cooling systems, and with the increased sizing of the district heating systems even seasonal energy storage solutions often become feasible.

The examples used in this chapter originate from Denmark. The history of Danish district heating is more than 100 years old and today about 65% of the households in Denmark are supplied with district heating. What started as a way of getting rid of waste in an efficient way is today a billion-dollar business and a cornerstone in the Danish energy system. In Denmark in 2020 more than 50% of the electricity load was covered by wind power, and this implies a further need for flexibility which effectively can be offered by the district heating system. The methods and results described in this chapter originates from a large number of national research and innovation projects like HEAT 4.0, CITIES, FED and IDASC[3]. We refer to the homepage of these projects for further information.

In the first generation of energy systems, central power plants were established to deliver electricity, leading to huge energy losses in the form of waste heat. In the next step, this 'waste heat' was utilised in the form of district heating by introducing combined heat and power plants (CHPs), an environmental leap step towards a sustainable future. District heating developed since to a 4^{th} generation, where intermediate renewable energy sources and waste heat from industrial and other sources are utilised. To accommodate the mentioned sources efficiently, the district heating operation has to be adjusted to *low-temperature* operation, improving the efficiency of the whole system.

Diversity in production, large district heating networks, advanced distribution and time-varying demand side characteristics, result in increased complexity. Digitalisation is proposed to facilitate the transformation from a rather simple to a highly complex system. Digitalisation introduces new possibilities and hereby complexities, naming wireless monitoring with IoT, increased connectivity and communication, AI and big-data analytics, systems-of-systems, and distributed system layouts with cloud-, fog- and edge computing. Digitalisation and the increased possibility of getting frequent sensor and meter data open up for the next generation of data-driven methods for the operation of district heating systems.

Common for the data-driven methods is, that the methods involve dynamical modelling based on grey-box and data-driven digital twin techniques, which again leads to new methods for real-time forecasting, control and optimization. The methods are adaptive, i.e. that they are automatically adapted to observed changes which can be deduced based on the received data from the system.

[3] CITIES: https://smart-cities-centre.org/, HEAT 4.0 will get a subpage on CITIES homepage, FED: https://www.flexibleenergydenmark.com/ and IDASC: Pamphlet at https://orbit.dtu.dk/en/publications/digitalisering-af-fjernvarme-erfaringer-derluner.

Adaptive methods are crucial in order to handle the inherently non-linear and time-varying characteristics of heat loads and the district heating system.

State-of-the-art methods for forecasting are crucial for efficient operation of district heating systems. This is partly due to the fact that often the heat has to be produced several hours (or days) before the heat is delivered in the houses. Methods for forecasting are outlined in Section 2.

Obviously, methods for heat load forecasting have to be based on methods for weather forecasting supplied by meteorological weather forecast services. Standard MET forecasts are targeting rural areas whereas district heating systems most often are seen in highly populated areas, and in Denmark, all the largest cities are supplied with district heating. Since the city weather often is rather different from the weather in nearby rural areas methods for forecasting city weather have to be considered, and such methods are briefly discussed in Section 3.

The heating is supplied by controlling the supply temperature at all plants and by controlling the flow in the network. In general, efficient operations of district heating networks call for data-driven methods for keeping the network temperature as low as possible, yet the temperature must be high enough to ensure that the heating systems of the supplied buildings are able to ensure a reasonable indoor temperature and a minimum temperature of the hot water. The temperature level is decisive for the efficient utilisation of renewable energy sources via e.g. heat pumps and for effective use of waste heat from low-temperature sources, e.g. from the cooling of data centres. Methods for temperature optimization are described in Section 4.

Section 5 briefly presents methods for the operation of buildings connected to a district heating network. The term *Smart buildings* is applied to highlight that it is necessary to be able to control the buildings if they have to play an active role in the optimisation and control of the overall district heating or power system. For smart buildings, the heat load can be partly controlled by the operator of the district heating network, and hereby we will be able to avoid or reduce peak demands and solve issues with bottlenecks in the district heating or the local power network.

Most of the examples in this chapter originates from studies conducted in HEAT 4.0 using data from the district heating system in Brønderslev, Denmark. This heating system is servicing approx. 5,000 customers in 3 subnets. The heat production is characterised by several different units such as CHP units, gas boiler, electric boiler, ORC, heat pumps, and concentrating, tracking solar thermal units and a large thermal storage capacity.

Some of the methods presented in this chapter links to Chapter 8 which considers a method for optimal production planning and power market bidding.

2 Forecasting for DHS

Forecasts are needed for optimal operation of DHS, e.g. for preparing the system for the future load by investigating future scenarios and selecting the appropriate

strategy, e.g. what heating units to produce the heat, and what supply temperature to push into the DHN. Three forecasts usually are needed: heat demand, electricity price and weather. Heat demand forecasts play a crucial role in enhancing the operation of DHS as it is needed in both production and network. Weather forecasts are used as inputs for the heat demand model therefore accuracy of them in predicting the future weather inside cities is crucial. Electricity price forecasts are used for bidding strategy into the electricity market, hence an important task when scheduling the production and creating bids for the day-ahead and balancing electricity markets. The focus of this section will be on methods for heat load forecasting, while electricity price forecasting will be briefly considered.

Section 2.1 will introduce what drives heat consumption. This will be used to identify a heat load model and to introduce a suitable framework for producing online heat load forecasts. Section 2.1.1 introduces the forecasting toolbox and framework of producing an online forecast for heat demand using a linear regression model that adapts over time to handle the non-stationary. Section 2.1.2 describes the identification of important model elements using physical knowledge of the system, and it describes how this knowledge can be used in grey-box modelling procedure for identifying an optimal model for heat load forecasting. Finally, an example of a heat demand forecast is introduced in Section 2.1.3 to demonstrate the benefits of understanding the underlying physics of the dynamics.

District heating is mostly applied in cities where the climate is quite different to rural areas. A procedure to localize NWP to the climate in cities is proposed in a subsequent section. Section 2.2 introduces briefly the need for them and a short introduction to the paper for further reading. In addition, Section 2.3 discusses other frameworks to produce heat demand forecasts.

2.1 Heat Load Forecast

Heat load in district heating consists of two components at the demand side; space heating and domestic hot water, plus related heat losses in the DHS. Due to the nature of the heat load, it varies during the day (the diurnal profile), week, and year creating a time-varying process. [20] splits the heat load into two categories, physical heat load and social heat load. The physical heat load depends on the climate and heat losses in the system, while the social heat load depends on the occupancy behaviour of the consumers and includes the usage of domestic hot water. Furthermore, the variation in heat load is described in [20] as follows:

- *Seasonal Heat Load Variation:* The physical heat load follows the increase and decrease in ambient air temperature throughout the year with a negative correlation. The social heat load also has a seasonal component, since, e.g., people tend to stay outside or away for longer periods during summer while people stay inside and consume more hot water during winter.

– *Daily Heat Load Variation:* The magnitude of the heat load during the day follows the ambient air temperature however the shape of it depends on the social behaviour. Social heat demand variations are explained by individual and collective social heat behaviour, e.g., harmonised working hours. The daily variation usually has two peaks, in the morning and the late afternoon.

Regional climate affects the individual district heating case, and the heat load related social behaviour can vary between systems. For example, a coastal locality is influenced by the damping thermal inertia of the sea. Taking into account the climate is very important to achieve efficient operation of DHS. Since DHSs are mostly used in urban areas, we will introduce the climatic characteristics inside cities and highlight the effects of climate variables on heat consumption in Section 3.

Thus, when developing a model to describe the dynamics of a system, it is important to understand the underlying process that drives the system. For DHS systems it is the physical description of how the weather influences the heat consumption through the buildings properties, e.g. walls and windows. Also, social heat consumption usually drives the peak consumption, e.g. in the mornings and evenings. A clear understanding of what drives the consumption from physics and social components will improve the models' performance in forecasting the future load of the system.

We will focus on utilizing the grey-box methodology which is using the physical understanding of the system, along with statistical methods to identify and estimate parameters of the model to establish an adequate forecasting model to predict the heating load. First, the heat load forecast framework for producing the predictions needs to be robust and simple to update the model and the predictions when new information is available. It needs to be stable and handle the non-stationary variation in the heat demand due to the weather-driven consumption of space heating. Also, being able to translate the nonlinear relationship between the effects of the weather on the consumption. This is addressed in Section 2.1.1 where the first part introduces a linear regression model where the parameters of the model are updated recursively and past information is exponentially down-weighted, i.e. addresses the non-stationarity. The second part proposes using a two-stage modelling procedure to be able to transform inputs that have a non-linearity relationship to the heat demand such that the coefficients of the linear model can be estimated. The second is the identification of the optimal model to be used in the forecasting framework. In Section 2.1.2, we introduce a model with a full physical representation of the heat consumption in a DHS. However, the full physical model is too computational heavy and would probably not be able to produce the forecasts before they become invalid. Therefore, an adequate heat load forecast model that uses parts of the physical insights of the consumption and what variables are known to influence the demand (e.g. ambient air temperature, wind, and solar radiation). Then combine this knowledge with statistical methods to identify the optimal forecasting model. Hence, it is preferable to use a model, which gives physically inspired descriptions of known relationships between climate variables and historical data to forecast the

heat load. The coefficients are then estimated using data-driven methods based on a time series of observations from the system where the model will be used. However, in order to formulate such a data-driven digital twin model, where the parameters can be assimilated based on observed time series, the model cannot be too complex. Basically, the model has to be structurally identifiable.

2.1.1 Framework for Heat Load Forecasts

We will present here the framework of establishing heat load forecasts that can handle the non-stationary and non-linearity that comes with the heat load. This type of forecasting framework has been proven accurate and able to perform in online operations [39] [41]. The framework allows the forecasting model to adapt to the changes in the parameters of the model for heat load forecasting, like the transitioning from cold (winter) to warm periods (summer). This adaptive method has shown adequate results in multiple research fields, especially in energy applications that are greatly dependent on physical behaviour. See e.g. [41] [5] [9] for heat load forecasting, [4] for solar power forecasting, and [64] [46] for wind power forecasting.

The model used in the framework is a linear regression model that can be recursively updated when new information arrives and exponentially down-weights past information. The model is also able to handle dynamical and non-linear effects, e.g. the ambient air temperature influence on heat load. There will also be an emphasis on creating individual models for each k-step prediction horizon since the lagged values in the model have to be tailored to the horizon, i.e. a unique model with its own coefficients is created for each k-step forecast.

Linear regression models are well known and frequently used for statistics and forecasting. Regression models as defined in [35] are used to describe a static relationship between a dependent variable Y_t and p independent variables, $\boldsymbol{x}_t = \{x_{1,t}, \ldots, x_{p,t}\}$. Therefore, an index t is introduced to denote the variable at time t. A regression model can be written as

$$Y_t = f(\boldsymbol{x}_t, t; \boldsymbol{\theta}) + \epsilon_t \tag{1}$$

where $f(\boldsymbol{x}_t, t; \boldsymbol{\theta})$ is a known mathematical function of the $p + 1$ independent variables and t, but with unknown parameters $\boldsymbol{\theta} = (\theta_1, \ldots, \theta_m)^T$. ϵ_t is a random variable with mean $\mathbf{E}[\epsilon_t] = 0$ and variance $\mathrm{Var}[\epsilon_t] = \sigma_t^2$. In many applications the function f is linear, and the models describe how the parameters $\boldsymbol{\theta}$, are defining a linear combination of the independent variables, $\boldsymbol{x}_{\mathrm{t}}$. Given N observations of the dependent variable and the p explanatory variable, we are able to estimate the model parameters.

Estimation of the parameters of a linear regression model is usually done with either Least Squares (LS) or Weighted Least Squares (WLS) methods. Given the estimated parameters $\hat{\boldsymbol{\theta}}$, we are able to predict a future value $Y_{\mathrm{t+k|t}}$ given the independent variable $\boldsymbol{x}_{\mathrm{t+k}}$. However, this type of forecasting model will not perform well for heat loads as it lacks adaptive properties to update when the system changes. This is necessary for heat load forecast due to the time-varying

and non-linear characteristics of the heat load. Therefore, the model needs to be extended allowing for time-varying parameters that can adapt to changes and transformations of the input variables.

Hence, an ordinary linear regression model will not perform well for heat load forecast as it lacks adaptation as the heat load changes over time. [33] propose an adaptive method that uses exponential weights using a forgetting factor, λ to discount old information. For auto-regressive models with exogenous input, this is called Recursive Least Squares (RLS) method with exponential weighting. It allows the model parameters to adapt over time when new information becomes available, therefore making it feasible to handle time-varying phenomena. The algorithm updates the parameters every time new information becomes available and uses the forgetting factor to discount old information and thereby increase the relative importance of the most recent observations. That is, the coefficients are recursively updated by an LS estimation with the weights exponentially decaying over time. The rate of decay is determined by the forgetting factor. The forgetting factor can change over time, but usually, it is set constant. The forgetting factor is between 0 and 1 and controls the level of adaptivity, where values close to 1 implies more weight equal weight on both new and older observations. While values of the forgetting factor close to 0 emphasise more on newer observations in the estimation. The optimal λ is found by minimizing the Root Mean Square Error (RMSE) as shown in [41]. This proposed method of allowing the parameters to update as new information becomes available and discount old observation is desired for online operations, as recursive methods offer relatively simple and few computations to estimate the parameters for every iteration.

Apart from the time-varying dynamics of the heat load, the non-linear dependency between the heat load and suggested input variables need to be considered and combined with the recursive update. For this, a two-stage modelling procedure can be used as proposed in [44] and [53]. Before estimating the model, the observed input variables are used either by mapping input variables by using some function (e.g. splines) or by using them directly (instant effect of the independent variables) in a so-called *transformation stage*. After the transformation, it is possible to create a linear model of the transformed data to predict the dependent variable by using the RLS scheme to recursively estimate the parameters in the so-called *regression stage*.

In the *transformation stage*, transfer functions (e.g., low-pass filter), basis splines, and Fourier harmonic series can be used to transform the non-linear relationship between the independent and dependent variables, to a linear relationship. Other functions can also be used to create the regression vector, e.g., kernels. The low-pass filter has proven a useful tool for explaining the effect of climate variables on the heat dynamics of buildings. As buildings are insulated and have thermal mass, they do not react instantaneously to changes in ambient air temperature. Hence, low-pass filtering of the ambient air temperature leads to a better description of the temperature effect on the heat consumption. Climate variables are typically transformed using rational transfer functions to model their effect on heat demand adequately [5] [41]. Low-pass filters can be

created from rational transfer functions with a stationary gain equal to one. For example, the simple first-order transfer function,

$$H_{a_{\mathrm{T}}}(q) = \frac{1-a}{1-aq^{-1}}. \tag{2}$$

where q^{-1} is the backward shift operator, i.e. $q^{-1}x_t = x_{t-1}$, and $a \in [0, 1]$ is the time constant, can be used to describe the dynamics between the dependent variable and independent variable that is being filtered. For instance, in a model describing heat transfer between indoor air temperature in a building and ambient air temperature, a high time constant would mean that the building has a high thermal mass and good insulation.

2.1.2 Model Identifications Procedure

Model identification of a system is a tedious process, as many things need to be considered. For instance, for model identification of heat consumption, questions arise on how to model the diurnal variation, is there a weekly tendency that needs to be modelled specifically, which climate variables are significant, how to translate their relationship to the heat consumption, and how to consider the physical representation of heat consumption to keep the desired indoor temperature in buildings. Here, we will first investigate the heat consumption from a physical point of view of the system, then we will demonstrate the use of statistical methods to enhance the model.

During winter in cold regions, the heat load is dominated by the demand for space heating, i.e. keeping the indoor temperature at the desired level to satisfy the thermal comfort of the consumers. Therefore, the physical model of heat demand can be viewed according to the heat loss characteristics of a building, the passive heat loss through the construction and the active contributions due to ventilation. These characteristics can then be used to suggest a forecasting model of a total heat load for a district heating network where other heat losses are added, e.g., losses related to the transportation of heat from production to consumer. This is the first step of establishing a grey-box model, identifying a model from physical knowledge. For instance, [34] demonstrate a grey-box modelling approach for the heat dynamics of a building using a continuous-time model based on stochastic differential equations. The final model is validated by both simulation and forecasting of the indoor temperature. In [42] this analysis is extended by considering the physical knowledge not just on a single building level but also the thermal characteristic of the entire DHS to create a model of the total consumption in the DHS. Climate measurements of ambient air temperature, wind speed, and global radiation are used to create an appropriate model. We will present the results from [42] on creating a heat load model of a DHS, first from a purely physical derived model then introduce their proposal of a model where the parameters of the model can be estimated by statistical methods. Thus, utilizing the Grey-box methodology of using physical knowledge and statistical information embedded in data to reduce the model-space of a purely physical model to describe the dynamics of the system.

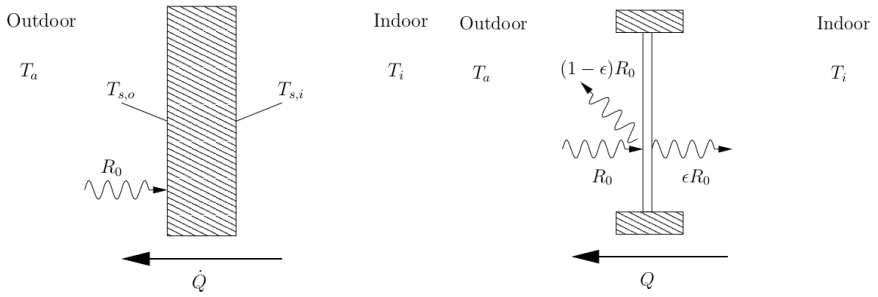

(a) Heat transfer through a wall. (b) Solar Radiation through a window.

Fig. 1: The figures demonstrate the stationary heat transfer through wall and window of a building. Missing is the effect of the ventilation. The source of these figures are found in [41].

[41] uses Figure 1 to illustrate the heat dynamics of houses by describing the heat transfer through a wall and a window. The figure demonstrates the stationary heat transfer of a building without considering ventilation. Plot (a) shows the energy exchange through a wall as the heat flux, $\dot{Q}$ which is found from the stationary relation between the outdoor, wall and indoor environment,

$$\text{Outdoor: } \dot{Q} = h_{\mathrm{o}}(T_{\mathrm{s,o}} - T_{\mathrm{a}}) - \epsilon R_0, \tag{3}$$

$$\text{Wall: } \dot{Q} = U_{\mathrm{w}}(T_{\mathrm{s,i}} - T_{\mathrm{s,o}}), \tag{4}$$

$$\text{Indoor: } \dot{Q} = h_{\mathrm{i}}(T_{\mathrm{i}} - T_{\mathrm{s,i}}), \tag{5}$$

$$\text{Overall: } \dot{Q} = U(T_{\mathrm{i}} - T_{\mathrm{a}}) - U\frac{\epsilon}{h_{\mathrm{o}}}R_0, \tag{6}$$

where T_{a} is the ambient air temperature, T_{i} is the indoor temperature, $T_{\mathrm{s,o}}$ is the temperature of the outdoor surface of the wall, $T_{\mathrm{s,i}}$ is the temperature of the indoor surface of the wall, R_0 is the solar radiation orthogonal to the wall, and $U = (1/U_{\mathrm{w}} + 1/h_{\mathrm{i}} + 1/h_0)^{-1}$. The heat is transferred as convection, thus from warmer to colder areas and the h coefficients are the convection heat coefficient of the outside of the wall and inside the wall. They describe the impact of a boundary layers on the outer and inner wall surface where the former is influenced by wind speed and wind direction while the latter can be assumed constant due to the constant indoor environment. U_{w} is the thermal conductivity of the wall divided by the wall thickness. Hence, the overall stationary heat flux through the wall is described from the stationary relation between the ambient air temperature, indoor temperature, the solar radiation orthogonal to the wall by convection.

The energy transfer through the window (Figure 1 (b)) consists of both convection and conduction. Thus,

$$\dot{Q} = -\epsilon R_0 + U(T_{\mathrm{i}} - T_{\mathrm{a}}) \tag{7}$$

where ϵ is the fraction of solar radiation orthogonal to the window (R_0) is not reflected by the window and the $U(T_\mathrm{i} - T_\mathrm{a})$ describes the energy conduction through the window between the ambient air temperature and indoor temperature, and $U = (1/U_\mathrm{win} + 1/h_\mathrm{i} + 1/h_0)^{-1}$. U_win is the thermal conductivity of the window divided by the window thickness.

Ventilation, where the warm air from the buildings is replaced by the cold air gradually, also needs to be considered. The heat flux of the ventilation is

$$\dot{Q} = \mathcal{C}\dot{V}(T_\mathrm{i} - T_\mathrm{a}) \tag{8}$$

where $\mathcal{C}$ is the product of the specific heat capacity of the air and the mass density of the air and $\dot{V}$ is the flow of air through the building.

Heat load of an area can therefore be assumed to be the heat loss of the heat transfer through walls, and windows and by ventilation, thus $Q_{L,t} = Q_{Wall,t} + Q_{Window,t} + Q_{Ventilation,t}$ plus the energy needed for domestic hot water usage, $Q_{W,t}$ for all buildings in the area. There is also "free" heat $Q_{F,t}$ that contributes to the indoor temperature coming from e.g. electrical equipment but also humans. Thus, energy needed for space heating $Q_{H,t}$ can be expressed as,

$$Q_{H,t} = [Q_{L,t} - Q_{F,t}]_+. \tag{9}$$

The truncation of negative values is used since when the quantity inside the squared brackets gets negative the indoor temperature will increase or ventilation will be used to prevent this. Considering the total heat consumption as

$$Q_t = \delta_{c,t}[Q_{L,t} - \delta_{p,t}Q_{F,t}]_+ + \delta_{p,t}Q_{W,t} \tag{10}$$

where $\delta_{c,t}$ is the fraction of the consumer heat load reacting to the climate. $\delta_{p,t}$ is the fraction of the potential consumption active at time t. This means $\delta_{p,t}$ accounts for holidays and $\delta_{c,t}$ accounts for the fact that during the summer almost no consumers react to the climate and that they do not all start/stop reacting on the climate at the same time of year. The dependence on holidays might be negligible since it will to a large extent only affect the demand for domestic hot water.

In [42] it is argued that these quantities of the model above can not be estimated using available measurement, i.e., the total heat consumption, ambient air temperature, wind speed and global radiation. Therefore, it is suggested to create a model structure, inspired by the physical quantities of the heat transfer qualities, that can be estimated only from data. Detail explanation of how to go from the physical orientated equation in Eq. (10) to the model in the equation below is found in [41],

$$\begin{aligned} Q_t = {} & \mu(h_t^{24}, \Upsilon_t) + a_{20}H_2(q)R_t \\ & + a_{111}H_1(q)W_t + a_{120}H_1(q)T_{a,t} + a_{121}H_1(q)W_tH_1(q)T_{a,t} \\ & + a_{100}H_1(q)R_t + a_{101}H_1(q)W_tH_1(q)R_t \\ & + a_{211,0}W_t + a_{211,1}W_{t-1} + a_{220,0}T_{a,t} + a_{220,1}T_{a,t-1} + e_t, \end{aligned} \tag{11}$$

the filters $H_1(q)$ and $H_2(q)$ are found to be

$$H_1(q) = \frac{0.066}{1 - 0.934q^{-1}}, \tag{12}$$

$$H_2(q) = \frac{0.350 + 0.612q^{-1} - 0.226q^{-2}}{1 - 1.703q^{-1} + 0.739q^{-2}}. \tag{13}$$

The function $\mu(\cdot, \cdot)$ models the diurnal variation, R_t is the solar radiation on a square pillar (see [41]), W_t is the wind speed, T_a is the ambient air temperature, e_t is the model error (iid and $\mathbf{N}(0, \sigma^2)$), $a_\cdot$ are the coefficients of the model. The filters $H_1(q)$ and $H_2(q)$ are rational transfer functions and are proposed to filter the climate variables to model their effect on them to heat load. [41] argues for instance that constant indoor temperature will effectively eliminate the heat storage capacity of the floor and internal walls, it seems therefore reasonable to use a rational transfer function when filtering the climate variables for heat demand modelling. Also, low-pass filtering is ideal as the dynamics at the boundary layers of the wall can be neglected however the conduction through the wall needs to be modelled. The inertia due to thermal mass indicates a slow response which again suggests that low-pass filtering of the climate variables is preferable.

This model is proposed to be used to predict the future heat consumption and is demonstrated to perform adequately, however it was shown that there was auto-correlation in the errors. A noise model was included in the model to remove the correlation using an auto-regressive model on the errors using lagged error on the past; 1, 2, 3, 23, and 24 lagged errors. The model including auto-regression on the noise showed that the auto-correlation of residuals was removed. When comparing the result between the model in Eq. (11) and the same model using the noise model, it showed that for online applications the model including auto-regressive terms performed significantly better for short-term horizons while the model without auto-regressive terms performed better for longer horizons. Hence, using information from the prediction errors of the model can improve the predictions on shorter horizons. The model in Eq. (11) was used for demonstrating online prediction and adaptive RLS update on parameters, it showed high accuracy for forecasting the heat consumption for a large area where the different methods of modelling diurnal profile and different forgetting factors are compared.

Model identification of a heat load model can also be done using traditional statistical approaches, thus finding the optimal structure of the forecasting model by finding which independent variables that describe the heat load adequately. Also, how those independent variables enter the model, i.e. how they are mapped in the *transformation stage*, here the physical understanding of the system will be helpful. Model building procedure is described in detail in [35, Ch. 6] along with the estimation of parameters and validating the model (model checking). Model building is an iterative process where the model is changed until it is adequate for its purpose, e.g. forecasting. The initial step of the identification procedure for heat load is to find a suitable initial model describing the response variable, the heat load. For instance, a model that includes an intercept and some

factor that models the social effect of the heat load, i.e. the diurnal profile. It is recommended not to include climate variables in the initial model so that one can use the initial model to identify which climate variables are significant. Modelling the social effect in the initial model is necessary as it is common knowledge that the heat demand has a diurnal profile. The profile can be modelled using Fourier harmonic functions [41] [5], splines or even using the hour of the day as input [17]. After modelling the social component, it is necessary to identify which climate variables influence the heat load and how to include them. The cross-correlation function can be used to identify what climate variable to include by computing the cross-correlation between the prediction residuals (e.g. one-step-ahead prediction) from the initial model and the climate variable time series. The climate variable with the most significant correlation is then added first to the initial model. Selecting what transfer function to use can be done by computing an error metric score that is suitable, for instance, the RMSE score. The model extensions are then compared to find which function (or the instant effect, i.e. no mapping function) describes the dynamics between the climate variable and heat load adequately. After finding the optimal function, these steps are iteratively repeated until the final model accuracy can not be improved. This model identification process is demonstrated in [5] where the climate variables are used to find the ideal model of the heat load of single-family houses. This procedure is data-driven as the variables are added or transformed and then the results are investigated to see if an adequate model has been found, this is repeated until the result is satisfactory. Here, the understanding of the physical nature of the system dynamics is an advantage, as it gives insight on what variables to be included and how they affect the system, e.g. low-pass filter of the ambient temperature to describe the heat consumption. We will present an example of forecasting performance by including a low-pass filter of the ambient temperature in Section 2.1.3. Additionally, the performance of a state-of-the-art forecasting model is also illustrated.

2.1.3 Demo: Brønderslev

In this section, we will demonstrate the performance of a forecasting model that is created based on the methodologies discussed earlier in a real district heating case. We will highlight how the model can be improved by implementing physical knowledge of the relation between heat consumption and weather. A state-of-the-art forecasting model supplied by HeatFor™[4] will be compared with the simple model proposed here.

A first attempt to forecast the heat load in Brønderslev would be to include an intercept, Fourier harmonics series to describe the diurnal variations and NWP of the ambient air temperature. Hence, an initial model is

$$\hat{Y}_{t+k|t} = \theta_{0,k} + \mu(t, n_{\text{har}}, \alpha_{\text{diu}}) + \theta_{1,k} T^{\text{a,NWP}}_{t+k|t}, \tag{14}$$

[4] https://enfor.dk/services/heatfor/

where the subscripts t is the time and k is the prediction horizon. This model will be referenced as *Model 1*. The estimation and forecast of this model is done using the `R` package, `onlineforecast` [3]. The package gives the opportunity to create forecasting models with estimated coefficients for each prediction horizon and estimates the coefficients adaptively using recursive least squares with an exponential forgetting factor as described in Section 2.1.1. Model 1 will be extended to model the physical knowledge on how changes in ambient air temperature affect the heat consumption's by using the low-pass filter $H_a(q)$,

$$\hat{Y}_{t+k|t} = \theta_{0,k} + \mu(t, n_{\text{har}}, \alpha_{\text{diu}}) + \theta_{1,k} H_a(q) T^{\text{a,NWP}}_{t+k|t}, \tag{15}$$

and will be referenced as *Model 2*. The results are shown in Figure 2 where on the left plot, the RMSE for prediction horizons from one to 72 hours ahead is illustrated. The improvement of including the low-pass filter is significant as the RMSE for all horizons has decreased. The right plot in Figure 2 shows one realizations of the forecasts into the future made at 2021-01-10 12:00 (vertical grey dashed line). This is only one realization and therefore should not be used to compare the accuracy between the models (use the RMSE plot). These plots clearly show that the improved model represents the dynamic changes in more detail. Thus, exploiting the physical dynamics will improve the forecast accuracy. The state-of-the-art performance is also visualized and a quite significant improvement in accuracy is clear when compared to the simple models where the most significant difference is in the short term forecast. High accuracy in the short term forecast is crucial for enhancing the temperature optimization of the network, i.e. lowering production cost and reducing heat losses by keeping the supply temperature low.

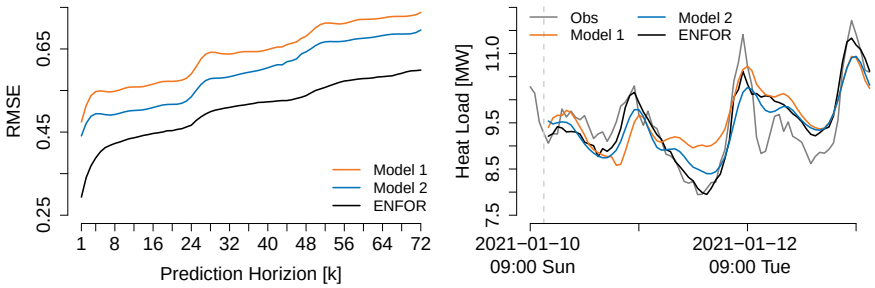

Fig. 2: Figure demonstrates performance of three heat demand forecast models. The left plot shows the RMSE for different k-step prediction horizons while on the right plot shows realizations of the forecast from the models.

2.2 Electricity Price Forecast

Here, we will briefly describe the forecast of electricity prices to optimize the economic performance of district heating with respect to the production balancing of heat and electricity, see further information in Chapter 8. In the case of electricity markets dependent units such as CHP or heat pumps, successful trading on the electricity market is important for the district heating operator. Most of the electricity is usually traded on the day-head market or wholesale markets and prices are often quite volatile and not determined by physical laws. The volatility and complexity increases also with the increasing share of decentralized RES. Therefore, electricity price forecasts are generally only based on data-driven models. [28] propose a two-step method to forecast the electricity price and its dependency on forecasts of load and wind power production. Uncertainty of the price forecast is not given in [28], however information on the uncertainty is often beneficial and needed e.g. for including the forecast into stochastic optimization methods [11] [56] [19]. There is a trend towards research on electricity price forecasting including uncertainties and probabilistic forecasts are gaining interest [47]. For an overview and more detailed information, we refer the reader to the review papers [31], [47], [65] or [1].

2.3 Further Heat Load Forecast Models

We have proposed grey-box modelling in order to formulate a model for heat load forecasting and a framework to deliver online predictions of the heat load in this section. However, black-box models that are purely driven by data are often also suggested for heat load forecasting. For example, linear models where coefficients are estimated using ordinary least squares [14] [17]. Autoregressive Integrated Moving Average (ARIMA) models are also a common method, often with a seasonal component. In [23] it is proposed to use seasonal ARIMA (SARIMA) models using a state-space framework. Using the state-space formulation implies that they are able to produce online predictions using the Kalman filter to update the forecasts and generate new forecasts when new information is available. These methods are easy to understand and computationally fast however they fail to describe the nonlinear dependency without applying transformations of the input. More advanced black-box models, e.g. machine learning methods, have demonstrated that they can capture the complex non-linear dependency that can be difficult to model using linear models.

In [57] it is proposed to use a feedforward neural network using two layers of neurons and using a sigmoid as the activation function. The neural network is applied without using any physical knowledge about the DHS. The results from the neural network are compared to a grey-box model where the nonlinear relationship between demand and weather variables are taken into account by transfer functions to create a linear regression model. Similar inputs are used for both models however the number of coefficients is significantly higher for the neural network with two layers of five and two neurons. In this case, the grey-box model has a better performance in the out-of-sample comparison.

[15] compares a linear model, a feed-forward neural network, and Support Vector Regression (SVR) where the input variables are also investigated. They find that the SVR has the highest accuracy. Additionally, the results show that including holidays and calendar data as inputs improved all models. However, this method has problems with the time-varying dynamics of the heat load. Recurrent Neural Network (RNN) have been proposed to solve the time-varying issue as demonstrated in [29] where it is shown that RNN handles trends in heat demand data while a feedforward neural network cannot.

3 City Weather Forecasting

Both heat load forecast and temperature control are heavily dependent on the most recent weather observations, in particular, the ambient temperature is important for the operation. The climate where the district heating system is located needs to be analyzed to operate the network in an optimal setting as mentioned at the beginning of this section. [41] and [39] suggest that the climate variables: ambient temperature, solar radiation, and wind speed (including direction) have the largest effect on the heating demand (in order of decreasing importance). [41] gives a detailed description of how these climate variables influence heat consumption based on physical and stationary considerations:

- *Ambient Temperature:* The ambient temperature affects the indoor climate through heat conduction in the outer walls and windows, but also through ventilation. It is shown that the outdoor temperature affects the indoor temperature through a low-pass filtered signal, and a simple transfer function model is suggested to describe how the variations in the outdoor temperature affect variations in the indoor temperature.
- *Solar Radiation:* The solar radiation affects the indoor climate based on the angle of beams hitting the building, where the orientation of the beams through the windows and the window area are most important. Basis functions are used to translate the non-linear dynamics of the solar radiation to its contribution to heating consumption.
- *Wind Speed:* Wind speed and wind direction affect the indoor climate as natural ventilation, the effect is depending on the tightness of the building. The wind speed also affects the convection heat coefficient on the outside of the buildings. It is therefore modelled using a low-pass filter for describing the contribution to the consumption.

Hence, in order to use these climate variables for enhancing the operational performance of the district heating system, both a clear understanding of how these variables affect the heat consumption and a forecast of them is needed. However, weather forecasts models are usually tuned for rural areas, not urban areas where district heating is applied. The difference between the climate in rural and urban areas is quite significant. Historically climate variables have been measured in rural areas. For instance, airports usually have climate stations and they are usually located outside the cities in an open area, where the only

impact is from the natural environment including lack of woody vegetation and directly exposed to natural rain, sun, and wind. However, climate inside cities is different from rural areas and therefore variables like the air temperature measured at the airports may deviate from the temperature inside cities, where the air temperature is exposed to human activities and the built environment. Research shows that the ambient air temperature typically is higher in urban areas than in rural areas [59]. The effect is termed urban heat island (UHI). A UHI is an urban area that is warmer than its surrounding rural areas due to human activities or build human infrastructures. The variation of outdoor air temperature data is both spatial and temporal. Several studies point towards a typical average difference in urban and rural temperatures of 2-3 K. For instance, [8] investigates the UHI phenomenon in Copenhagen and demonstrate it by using three climate stations that are located inside the city, in the outskirts and in a rural area outside of the city. The hourly temperature average from the climates stations is computed and illustrates that the temperature measured at each station is different. The climate station inside the city always has the highest temperature, whereas the station in the rural area has the lowest. It is also shown that the spatial temperature has time-varying characteristics, both a diurnal and an annual variation is seen. For example, comparing the result during summer shows that the average temperatures in the mornings are very similar. However, later in the day, it differs, with higher temperature difference in the city, and during the night it gets colder at the rural side. Thus, the city does not lose heat as fast as the rural part. While during winter, there is quite a constant offset between the stations.

3.1 Localizing Numerical Weather Predictions

NWPs are obtained by a physical model of the atmosphere and ocean to predict climate variables. They are computed over a grid of the earth and are then interpolated to a specific location where weather predictions are needed. However, NWPs are designed for rural weather forecasting, and they often have problems adjusting to the local climate in cities due to the local climate phenomenon. The models seem to have trouble adjusting to local heat contributions, e.g., solar heat in the street, heat from buildings, etc. DHS relies heavily on NWP to operate their system efficiently, therefore it is important to correct the NWPs before using them as input e.g. models for heat load forecasting. Especially, for temperature control of the district heating system as it is done on a short-term horizon (between one and 24 hours) and is heavily dependent on the current local climate. Using a local climate station to localize the NWP, corrects the short term NWP by adapting them to the climate using real-time climate measurement [41]. Hence, this yields an optimal weather forecast for a certain area that can be used to operate the temperature control in the most optimal setting. In [22] it is proposed to use Model Output Statistics (MOS) to bind NWP to local climate stations observations, e.g., localize the forecasts. The MOS is a simple technique that uses linear regression where the observed climate variable is the response variable and predictors are the NWP variables which therefore

bind the NWP to the local climate. It is a simple and frequently used method that will reduce systematic bias in the NWP if there is any. [13] propose using an adaptive method to reduce the systematic bias and lower the RMSE of the NWP.

4 Temperature Optimization

Efficient operation of district heating networks implies an objective to minimize production cost of the heat production and reduce heat waste without compromising the consumer comfort. For most district heating systems, this is achievable by minimizing the supply temperature at the production site while providing the total desired heat and fulfilling the requirement of minimum supply temperature at any time point at all points in the district heat network. Decreasing the supply temperature is also highly valuable for heat production sites where electricity is also produced (e.g. with CHP units) since it implies an increase in the ratio of power to heat output, and, as electricity is more valuable than heat, a more profitable operation is achieved.

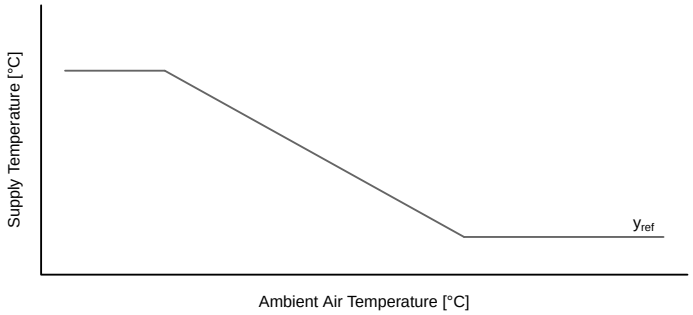

Fig. 3: Example of an reference curve for controlling the supply temperature at the production site.

Temperature Optimization aims at optimizing the supply temperature. Traditionally, the supply temperature has been either controlled using a reference curve schema or experience of the network operators [48]. A typical reference curve is illustrated in Figure 3 where the supply temperature is determined as a function of the current ambient air temperature. The supply temperature is kept constant during high ambient temperature for providing solely heat for domestic hot water during non-heating periods. The temperature is set high enough to reduce any bacteria risk. As the ambient temperature decreases the supply temperature is increased until the maximum supply temperature of the system is reached. This control schema aims at ensuring that the consumers receive

sufficient heat as it models the relationship between supply temperature and ambient air temperature as a worst-case scenario, i.e. the temperature should not be lower than the reference curve. The curve also considers that the supply temperature for the given ambient temperature will have sufficiently large safety margin for the flow to be adjust to satisfy the heat demand before it reaches the maximum limit of the system.

The use of a single reference curve at the plant usually results in a high supply temperature which is high enough to account for all other variables not taken into account. This results in higher costs and more heat losses in the system. Thus, the curve is not an optimal strategy as it does not consider other climate variables that are known to influence heat consumption like wind speed, wind direction and solar radiation. Furthermore, it ignores the social behaviour and the time-varying relationship to demand. One of the important time-varying relationships is the time-varying time delay between time of production and the time when the water reaches the end-users. It is also not a predictive controller, i.e. it does not look into the future while selecting the supply temperature. Even though this method is naive, it is more advantageous than operating the supply temperature at the maximum limit and allowing the flow to vary. However, the supply temperature is kept unnecessarily high when discarding other factors. When including all of the above factors into a control schema using predictive methods, new possibilities of lowering the supply temperature without violating any restrictions are possible. It is also important to point out that the hydraulics limitation of the DHN implies some restrictions on the minimal supply temperature in order to ensure that the flow is below the maximum limit with some safety margins. Thus, then the supply temperature should be increased to meet the demand thus a predictive method is needed to optimal select supply temperature before the flow reaches the maximum.

With the increase in both computational power and research in adaptive control during the 1980s, more sophisticated methods for DH have been developed. Adaptive controllers are able to operate time-varying and non-linear systems since they can adapt to changes based on the output feedback of a system (see, e.g., [2] for further reading). For DH, this improvements in research related to control theory was necessary as the system is inherently non-linear and non-stationary.

In this chapter, we will focus on temperature optimization based on adaptive control using feedback from critical points in the DHN. Our description is based on the research published in amongst others [36], [39], [40], [43] and [52]. The described method also resulted in commercial software that is used for temperature optimization in DHN in Denmark[5]. The method has proven to be able to provide significant reductions.

Section 4.1 will introduce and describe the characteristics of a DHN. Section 4.2 will build on the knowledge of the network characteristics to introduce temperature optimization and control. A demo case will be used in Section 4.3

[5] https://enfor.dk/services/heatto/

to demonstrate the performance and savings. Finally, we give an overview of additional work on temperature optimization in Section. 4.4.

4.1 DHN Characteristics for Control

The dynamics in a DHN are driven by the physical dynamics and the consumer consumption, i.e., the heat load. As discussed in Section 2, consumption is driven by intra-day variations, climate, and local social behaviour. [7] propose an extensive physical-simulation model of a DHN and describe the important physical factors to consider. The physical factors are *time delay*, *heat loss*, *pressure*, and *friction loss*. These factors are very important for accurate simulation and understanding of the DHN to maintain acceptable temperature and differential pressure in the network. Thus, they are used to operating the network efficiently by delivering the desired heat to the consumer while minimizing the operation cost. Heat loss from pipe to surroundings is determined by the time of being transported in the DHN from production to consumer (time delay), along with the temperature of the hot water and the resistance of the pipe insulation. The time delay or the transport time of sending hot water to a point in the network is determined by the flow that is controlled by pumps based on the differential pressure applied in the system. The pressure in the system controls how fast the hot water travels in the network. Hence, the temperature loss between consumers and the production plant depends on these variables. Pressure loss in the system depends on the fiction which is a function of flow rate and other pipe properties. The network is designed to maintain fixed differential pressure in the network which the pump control tries to maintain by adjusting the flow in the system. If the flow becomes too high then it becomes impossible to maintain the needed differential pressure due to limitation of the pump. In these cases, the differential pressure over the substations at the consumers will start to decrease, and the consumer will then not receive the desired heat. Thus, these four factors contribute to the two main components of delivering heat to the consumer; temperature and flow. They are jointly linked together in a nonlinear relationship. For instance, heat losses in the system are based on the temperature and the flow, hence, the temperatures in the network are determined by the temperature and the flow in the system.

The DHS usually also has restrictions based on these factors on the operation of the network due to physical limitations and additional constraints made by the utility. [43] describes the following usual restrictions:

- *A maximum allowable flow rate in the system:* The restrictions in the flow rate are due to the (always) limited pumping capacity, the risk of cavitation in heat exchangers and difficulties maintaining a sufficiently high differential pressure in the remote parts of the network during periods with high flow rates.
- *A minimum guaranteed inlet temperature at the consumers:* This restriction is due to limitations in the consumer installations as well as minimum domestic water usage temperature requirements imposed by hygienic concerns.

- *A maximum allowable supply temperature:* This restriction ensures not to damage pipelines and consumer installations.
- *Limited short term variation in the supply temperature:* The stresses inflicted on the network by large and frequent fluctuations in the supply temperature dictate that the short term variations in supply temperature should be limited.
- *Maximum allowable diurnal variations of the supply temperature:* In some systems the size of the expansion tanks may impose limitations on the allowable diurnal variation of the supply temperature.

Hence, the framework of the temperature control has many constraints and physical limitations, thus operating a network consequently needs to consider multiple aspects to result in an optimal operation. Optimal operation of the network is achieved by minimizing the production cost without; compromising the safe operation of the system, adversely affecting the maintenance cost of the system, or sacrificing consumer satisfaction. Note that the physical description and system restriction listed here are not valid for every system. Each network is unique with different physical limitations or restrictions. However, they need to be considered before implementing the temperature control to reduce the risk of failure and achieve optimal operation of the network.

4.2 Controlling the Supply Temperature

In this section, we will introduce a control schema for the optimal operation of the DHN. To simplify the schema, we will only present a DHS that is supplied from a single station. We also assume that the return temperature and the diurnal peak load are not affected by the optimization. The operation costs can then be minimized by minimizing the supply temperature without violating the requirements discussed in previous sections. The methodology introduced here is a statistical approach of estimating the supply temperature by using adaptive estimation and transfer functions to model the network as initially proposed by [58], [39] and [40] and extended in [43] and [52]. Hence, we do not use a detailed physical model of the network, but statistical methods along with measurement data to describe the dynamics.

Since the adaptive controller incorporates feedback information, the DHN needs to have some measurement wells in the network so that the controller gets feedback from the network. The placement of the wells should be located at points with the lowest temperature in the network, i.e. largest temperature loss. These wells are usually decided at the time of the design of the network or added later in the network when needed. Hence, if the supply temperature requirement is satisfied at those points in the network, we expect that the requirements of all consumers in these areas are satisfied. These measurement points will be referred to as *critical netpoints*. A more recent study outlines how to use frequent meter readings at the consumer as a cheaper and flexible alternative to measurement wells in the network [10].

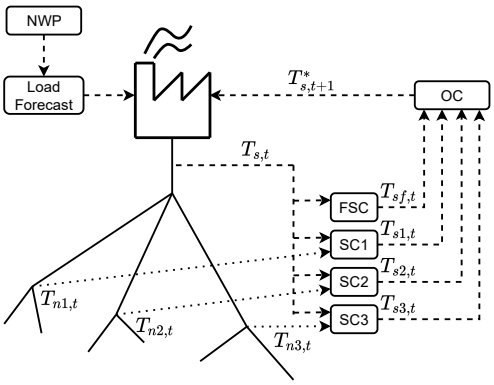

Fig. 4: Schematic view of the proposed controller: The supply temperature sub-controller (SC) which models the relationship between supply temperature, $T_{\mathrm{s,t}}$ at the production and netpoint temperature, $T_{\mathrm{n,t}}$. The flow sub-controller (FSC) that uses heat load and flow to estimate the supply temperature. Finally, the overall controller (OC) that select the highest supply temperature, $T^*_{\mathrm{s,t+1}}$ from the sub-controllers.

The main concept of the proposed control strategy is illustrated schematically in Figure 4. The overall controller (OC) selects the highest supply temperature from the flow sub-controller (FSC) and the supply temperature sub-controllers (SC) at the critical netpoints to be used in the next time step. The SCs estimate the lowest possible supply temperature from the plant using statistical identified transfer function without violating any restrictions at the consumers. The FSC computes the supply temperature without violating the maximum flow limit. The main principle of the proposed control schema is to keep the supply temperature as low as possible. The heat demand can be satisfied, by varying the mass flow through the network and by varying the supply temperature,

$$p_{\mathrm{t}} = q_{\mathrm{t}} c_{\mathrm{w}} (T_{\mathrm{s,t}} - T_{\mathrm{r,t}}), \tag{16}$$

where p_{t} is the heat load, q_{t} is the flow, c_{w} is the specific heat constant of the water in $[\mathrm{Jkg}^{-1^\circ}C^{-1}]$, $T_{\mathrm{s,t}}$ is the supply temperature and $T_{\mathrm{r,t}}$ is the return temperature.

The proposed control strategy here ensures the heat demand is met by varying the mass flow before increasing the supply temperature. Thus, the supply temperature is kept as low as possible while the flow is varied, the temperature is only increased when the flow is at the maximum value. The methodology behind the controllers and the transfer function between the production and network to establish a temperature controller is introduced briefly below. More detailed explanation can be found in [38], [39], [40], [36], [43], [49], [57], [45], and [52].

Flow Sub-Controller (FSC): The FSC is to perform an online control of the supply temperature in order to ensure that the flow rate q is kept close

but below a maximum value $q_{\max}$. The controller uses prediction models of the heat load to find the optimal supply temperature while keeping the flow as high as possible. The future control signal, $T_{\mathrm{s,t+1}}$ is found by solving for the supply temperature in Eq. (16),

$$T_{\mathrm{s,t+1}} = \hat{T}_{\mathrm{r,t+1|t}} + \frac{\hat{p}_{\mathrm{t+1|t}}}{c_{\mathrm{w}} q^{ref}_{\mathrm{t+1|t}}} \tag{17}$$

where $\hat{T}_{\mathrm{r,t+1|t}}$ is the forecasted return temperature, $\hat{p}_{\mathrm{t+1|t}}$ is the forecasted heat load, $q^{ref}_{\mathrm{t+1|t}}$ is the mass flow from a reference that is a prediction of the flow and it is below the maximum flow, $q^{\max}$ with a large probability, e.g. 99%. This simple controller only considers one-step ahead into the future however the change of supply temperature will affect the consumers after different time delays. [38] propose to use weights on the j-step predictors to estimate the desired supply temperature, $T_{\mathrm{s,t+1}}$. More advanced flow controllers are proposed in [43] and [36] where the uncertainty from prediction of heat load and return temperature are considered.

The Transfer Function Model: An important aspect of the supply temperature sub-controllers based on a statistical approach is to identify the dynamic relationship between the temperature at the critical netpoints in the network and the supply temperature and flow rate at the plant for the supply temperature sub-controller. An accurate estimate of the network characteristics, including the time-delay is important for enhancing the controller. Hence, a model of the district heating network is needed for optimal control. A pure physical model of the network (white-box) is almost impossible to establish and would also be likely to be too complex for control purposes. The white-box models can be adjusted in operational settings when calibrating the models based on previous operational data however it can take too long due to the complexity of the model or the assumptions being wrong.

We therefore propose, a model of the network that is simple and is easily updated when new information or data is available about the network characteristics. The model is found by identifying a statistical transfer function of the network. It can be modelled using a single-input single-output AutoRegressive-eXtraneous (ARX) structure using a fixed time delay. An example of a model for a DHN is given here

$$T_{\mathrm{np,t}} = a_1 T_{\mathrm{np,t-1}} + b_{0,\mathrm{t}-\tau} T_{\mathrm{s,t}-\tau} + b_{1,\mathrm{t}-\tau-1} T_{\mathrm{s,t}-\tau-1} + b_{2,\mathrm{t}-\tau-2} T_{\mathrm{s,t}-\tau-2} + \epsilon_{\mathrm{t}}, \tag{18}$$

where T_{np} is the netpoint temperature, T_s is supply temperature, ϵ is the noise, and τ is the time-delay between the two temperatures (at the plant and the critical point). The coefficients; b_0, b_1, and b_2 describe the diurnal variation in the system, see [58] for further details. The model coefficients can be constants, but since DHS are non-stationary it is better to estimate them recursively. The model is formulated as a 1-step prediction model. The k-step predictions are then obtained by recursive use of the 1-step prediction model. [58] propose different advanced methods to estimate the time-delay recursively, however [39]

and [43] propose a simple scheme to estimate the time-delay for the transfer function. The simple method estimate the time delay as the lag with the largest numerical value of the cross-covariance function between the time series of the supply temperature and the netpoint temperature. [52] suggest more advanced methods to forecast the netpoint temperature and estimate the time-delay using conditional Finite Impulse Response (cFIR) as a transfer function of the supply and netpoint temperature. This allows for a nonlinear estimation of the time delay depending on the mass flow in the system.

Supply Temperature Sub-Controller (SC): The SC controller focus on creating a control schema to vary the supply temperature at the production without violating the requirements at the critical net points. It utilizes the transfer function model of the supply temperature and netpoint temperature, then it is possible to create a control scheme based on the network characteristics. In [50] and [51] an eXtended Generalized Predictive Controller (XGPC) is proposed to be used to control the supply temperature based on the transfer function between the supply and netpoint temperature. The Generalized Predictive Controller (GPC) is modified to handle non-stationary systems as it assumes that the predicted output can be expressed as a linear combination of present and future controls. Traditionally, this is obtained by solving the Diophantine equations as proposed in [12] however these equations are formulated for time invariant systems, and hence modifications are needed. [50] propose a modification of the GPC where the optimal prediction is the general conditional expectation of the system output. The GPC is only reasonable if the underlying process has a slow time-variation, while the XGPC can handle time-varying processes. This is important since the transfer functions that describes the network characteristics are inherently non-stationary due to the flow which results in a time-varying time-delay and different heat losses as a function of the hot water temperature and the flow.

The XGPC is based on an ARMAX model with time-varying parameters, and the transfer function of the network is given by

$$A_{\mathrm{t}}(q^{-1})y_{\mathrm{t}} = B_{\mathrm{t}}(q^{-1})u_{\mathrm{t}} + C_{\mathrm{t}}(q^{-1})e_{\mathrm{t}} \tag{19}$$

where A_{t}, B_{t} and C_{t} are time-varying coefficients polynomials and q^{-1} is a back shift operator. The coefficients can be estimated adaptively using recursive methods. As the model is based on a time-varying process, the j-step predictor $\hat{y}_{\mathrm{t+j|t}}$ is described by the conditional expectation of $y_{\mathrm{t+j}}$ given observations of the output up to time t

$$\hat{\boldsymbol{y}}_{\mathrm{t}} = \boldsymbol{H}_{\mathrm{t}}\boldsymbol{u}_{\mathrm{t}} + \boldsymbol{v}_{\mathrm{t}}, \tag{20}$$

where $\hat{\boldsymbol{y}}$ is the vector of predicted reference netpoint temperature, $\boldsymbol{u}$ is the vector of the future control signal (supply temperatures), $\boldsymbol{v}$ is the vector containing the expected response from the input free system, and $\boldsymbol{H}_t$ is a matrix of size $N \times N$ containing the time-varying impulse response of the system. The N is the maximum prediction horizon of the controller, $N \geq 1$. In relation to an ARMAX model and Eq. (20), it can be seen that h is the weight of the time-varying

impulse function on future control values and v is the expected response from the past, see [50] and [43] for detailed explanation. It is possible to use individual models for each j-step prediction. Then each row in $\boldsymbol{H}_t$ and $\boldsymbol{v}$ corresponds to the j-step prediction model.

The object of the controller is to minimize the difference between measured netpoint temperature and the reference temperature without violating the restriction with a certain probability. Additionally, the controllers object function contains a term with the purpose of minimizing the supply temperature fluctuations. A derivation of the optimization problem using the XGPC control law and model of the network as shown in Eq. (20) in detailed is provided in [43]. The cost function is optimized over multi-step ahead horizons due to fact that the time-delay in the system varies, and hence a minimum and maximum time period is considered for the cost calculation. It is then shown that solving for the input vector, $\boldsymbol{u}_{\mathrm{t}}$ (supply temperature) results in

$$u_{\mathrm{t}} = -[\boldsymbol{H}_{\mathrm{t}}^{T}\boldsymbol{H}_{\mathrm{t}} + \boldsymbol{F}^{\mathrm{T}}\boldsymbol{\Lambda}_{\mathrm{t}}\boldsymbol{F}]^{-1}[\boldsymbol{H}_{t}^{\mathrm{T}}(\boldsymbol{v}_{\mathrm{t}} - \boldsymbol{y}_{t}^{\mathrm{ref}}) + \boldsymbol{F}^{\mathrm{T}}\boldsymbol{\Lambda}_{\mathrm{t}}\boldsymbol{g}_{\mathrm{t}}]. \tag{21}$$

At each timestep, the control signal is estimated, therefore only the first element of the control vector is implemented, i.e.,

$$u_t = [1, 0, \ldots, 0]\boldsymbol{u}_t \tag{22}$$

In [36] extended formulations of the XGPC controller are propose. This includes for instance equality constraints.

Overall Controller (OC): The OC selects the highest supply temperature, $T^{*}_{s,t+1}$ from the required supply temperatures computed from the FSC and the SC sub-controllers. This temperature is used as the supply temperature for the plant in the following hour; the model parameters, the predictions and the controller signals are updated each hour.

4.3 Temperature Control: Demo Case

Supply temperature at Brønderslev is now being operated by data-driven temperature optimization in an on-line operation. It has been operating since beginning of 2020. Before the data-driven operation, the network was operated by a simple algorithm that did not receive online feedback from the network. The data-driven temperature optimization was done with the `HeatTO`™[6]. Figure 5 illustrate the results from the previous operation and the current data-driven operation. The operations are compared using three months (February, March and April) in 2019 and 2020 for the old and new controller. The plot compares the performance of the controllers versus the degree days. Degree days are used to compare supply temperature between heating seasons when comparing different operations. Degree days, T^{dd}, are computed as the positive difference between the average ambient temperature ($\bar{T}_a$) over one day, and a cut-off of heating demand from buildings (we use 17°C here), i.e.

$$T^{dd} = \max(0, 17 - \bar{T}_a). \tag{23}$$

[6] https://enfor.dk/services/heatto/

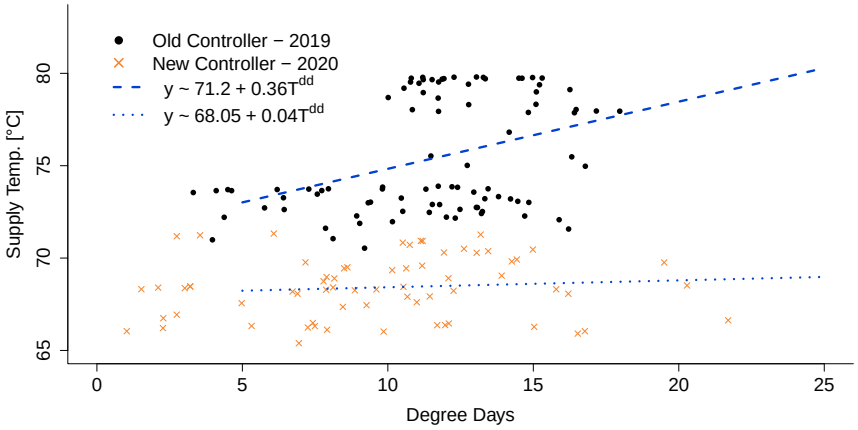

Fig. 5: The figure demonstrates the difference in performance using simple operations of controlling the supply temperature and using advanced data-driven method.

The average supply temperature for the given day is then computed and plotted versus the degree day as seen in Figure 5. To compare these operations, a regression model using Ordinary Least Squares to estimate the parameters of a model with an intercept and slope have been fitted to each operation as shown in Figure 5 and the result is,

$$\text{New controller: } T_{supply} = 68.05 + 0.04T^{dd} \tag{24}$$

$$\text{Old controller: } T_{supply} = 71.2 + 0.36T^{dd} \tag{25}$$

We see that the supply temperature regression line is approximately 3°C lower at y-intercept of for the new controller when $T^{dd} = 0$, and further the slope is lower for the new controller resulting in larger differences for higher T^{dd}. To compare the difference between the two methods, for instance, if we investigate a degree day of 10 then the difference is around 7°C.

Decreasing the supply temperature for operation leads to an increase in savings for the utility. [40] suggest a rule of thumb for savings resulting from lowering the supply temperature in CHP plant; For each degree lowered, the savings for the heat loss in the network is 0.5 % and the savings from more efficient production is 1 %, thus the savings can be compute as

$$\text{Savings } = (\text{Cost}_{\text{before}} * x\ [^{\circ}C] * 0.5\%) + \text{Shares}_{Production}(\text{Cost}_{\text{before}} * x\ [^{\circ}C] * 1\%) \tag{26}$$

where x is the supply temperature difference between operations to evaluate the savings.

Thus, the estimated savings from using data-driven temperature optimization would be roughly 7°C and the savings would be around 10.5%. The equations are

just a rule of thumb to demonstrate potential savings when sufficient production data is not available.

4.4 Further Temperature Optimization Methodologies

[24] propose similar temperature optimization approaches where statistical methods are used as describe above. They also take advantage of adaptive methods to model the time-varying behaviour. However, instead of the XGPC and cFIR, they use fuzzy modeling to describe the response characteristics of the network when supply temperature and flow are varied. They model the relationship between the supply temperature and network temperature using multiple local linear models that are only valid in certain regions, i.e., a fuzzy modeling to handle the non-linear dynamics of the system.

All above mentioned methods utilize statistical descriptions of the network to estimate the transfer function between two points (plant and netpoint temperatures) and derive controllers from there. The goal was keeping the supply temperature as low as possible to minimize the operational production cost for the system. However, other approached have been suggested in literature on how to estimate the network characteristics to ensure consumer are receiving the required temperature. For example, models derived from a purely physical representation of the network by modelling heat and mass transfer of the hot water. [7] propose a physical model of the network to describe the dynamics of it. This is done by predicting the network temperature and flows at nodes in the network, referred to as the node method. In this, the DHN is represented by a set of nodes by their physical description, e.g. heat capacities and pipe diameters. Therefore, it is possible to simulate the network characteristics and estimate suitable supply temperatures. This model gives promising results in offline settings where it optimize both the production and the network however for online purposes, the optimization is too computational heavy to deliver on time. [32] suggest a simplified physical representation of the network by reducing the number of pipes into a tree structure without a circular loop. They demonstrate this in a case study where pipe system is reduced from 1079 nodes to approximately 10 nodes without sacrificing any accuracy and reducing the computation time by 99 %. In [54] and [55] also physical models are derived to operate a DHS using predictive control where the time-delay and temperature are estimated from a physical simulation model. Additional research where both production and network or only the network are modelled using physical derivation of the network has been carried out (e.g. frequently, the node-method proposed by [7] is used). Further Mixed Linear Integer Programming (MILP) is utilized to optimize the production and network to minimize the cost [6] [21] [54]. [63] gives an extensive review on DH for further reading.

5 Smart Buildings in DH Network

Buildings are traditionally controlled to achieve a good indoor environment. This is done independently from the surrounding energy systems. When introducing the dependency of building control on its surroundings, a decision has to be made, whether the building is in charge to meet the demand of e.g. the energy system (indirect control), or the external system can control the building system by direct control. The indirect method can be achieved by broadcasting prices that change according to the energy system demands as described below. In both schemas, direct and indirect control, the ability to control the many heating-, cooling- and ventilation devices has to be in place in the involved buildings. The below introduced MPC demonstrates such an enabling technology for simple buildings. In general, smart control of buildings require a central monitoring and control system.

In this section, we describe the interaction of smart buildings with DHN as an appetiser to this extensive topic. First, we describe the motivations and values of smart buildings in DHNs. Next, we introduce MPC as a technology to deliver smart control of thermal dynamics for buildings. We round up the section by introducing a hierarchical control setup, where buildings work together in the DHN to deliver flexibility.

5.1 Value propositions by Smart Building to district heating

There are numerous reasons for controlling buildings in smart ways; many of them targets the efficiency of the building itself, others target the interaction with the surrounding neighbourhood and infrastructures. In this section, we give a single example of a control technology, MPC, that enables both, internal and external smart solutions.

A traditional, non-smart building can be controlled by different means, manual control by the occupants, individual control on each heating and cooling device (radiators and floor heating) and also individual apartment control in multi-family buildings. Advanced buildings, such as office buildings and schools, are controlled by central installations that connect all the controllers and collect data centrally. Such systems are often controlled by skilled professionals and perform better compared to individual control. However, literature shows clearly that there are huge improvements to the performance of the latter type of buildings.

In recent years, the objectives for control are increasingly complex. On top of a request for the primary goal of optimal indoor climate conditions, energy efficiency with related CO2-footprints and economical efficiency are the most dominating of many target objectives of control. Later years, flexibility is added to the list of objectives for the control. Control strategies that can handle these diverse demands call for rather advanced control units in buildings, smart buildings, that are able to e.g. shift heat demand in time and power, compute advanced demand patterns that compensate for temporal variations and other things.

Many methods and solutions have been proposed in the literature. One of these methods is Model Predictive Control (MPC), which is a well-established and developed method for building climate control [18]. Its popularity is due to its simplicity and natural way of incorporating constraints and accounting for disturbances (such as the weather) in the optimal control problem. And not at least, the methods can be implemented in almost any hardware with computational abilities that are necessary to do essential predictive computations. Basically, these computations are very similar to the modelling and forecasting tasks applied to whole district heating systems and discribed in this book, because buildings are similarly affected by climate conditions [62]. However, the number of controllable units may be extremely high compared to district heating systems and so is the number of sensors involved.

Below, a simple building system is described as an example of the MPC methodology.

5.1.1 Building Control by MPC-Methodology

We give a brief example of how to model and represent a building in a DH network to describe the necessary states and properties.

Consider the following system of stochastic differential equations with observations taken at times t_k

$$\mathrm{d}\boldsymbol{x}(t) = f(\boldsymbol{x}(t), u(t), \boldsymbol{d}(t))\mathrm{d}t + g(\boldsymbol{x}(t))\mathrm{d}\boldsymbol{\omega}(t)\,, \tag{27a}$$

$$\boldsymbol{y}_k = h(\boldsymbol{x}(t_k)) + \boldsymbol{v}_k, \quad \boldsymbol{v}_k \sim N(0, R) \tag{27b}$$

where $\boldsymbol{x}$, u, and $\boldsymbol{d}$ are the system states, controllable input, and non-controllable input respectively. f is the deterministic dynamics, g is the diffusion function, $\boldsymbol{\omega}$ is a standard Brownian motion, and $\boldsymbol{v}_k$ is the observation noise. The above model formulation is an example of a grey-box model that includes physical dynamics in f and describes stochastic elements by the Brownian motion that are too complex to otherwise model. [60] modelled the thermal dynamics of a Danish school building in a DH network using a hydraulic heating system with thermostatic controlled radiators. Such a model enables an MPC to control the thermal dynamics and perform tasks such as peak shaving or load balancing. They find the following states useful to include in a grey-box model

- The indoor air temperature, $T_i(t)$. This is typically the variable that is important to maintain a comfortable indoor climate.
- The temperature of the building envelope, $T_w(t)$. This contains important information about the insulation-level and how much heat is stored in the walls.
- The flow of the water in the space heating system, $\Phi(t)$; this varies as the thermostats in the radiators open and closes.
- The temperature of the radiators in the building, $T_h(t)$, is the component that delivers the heat in the rooms.
- The return water temperature, $T_{\mathrm{ret}}(t)$, is important since it determines the amount of heat the building uses.

Together, the above states form a model that is sufficient in describing the important thermal dynamics in a large building in a DHN. The controllable input to the system, $u(t)$, is the set-point to the radiator thermostats in the building. The map between the difference in set-point and room air temperature and how open the valves are was modelled using a sigmoid function. Using a system with the above states, it was possible to optimise the operations of the building while shifting the loads to desired times [61].

5.2 Implementation of MPC in a Smart Buildings

On a building level, the MPC has the objective to satisfy certain constraints (e.g. a comfortable indoor air climate at all times) while at the same time minimising some objective. For economic MPC [30] [25], the objective typically is to minimise the cost of the heat consumption while only adjusting the set-point of the radiator valves (in the case of a building as described above). Such an optimisation problem can be formulated as

$$J(\hat{\boldsymbol{x}}_{k|k},\hat{\boldsymbol{d}}(t)) = \min_{u_k} \sum_{i=0}^{N-1} \boldsymbol{c}_{k+i} u_{k+i} \tag{28}$$

$$\begin{aligned}
s.t. \quad & \hat{\boldsymbol{x}}(t_0) = \hat{\boldsymbol{x}}_{k|k} \ , && && (29)\\
& \hat{\boldsymbol{x}}(t) = f(\boldsymbol{x}(t), u(t), \boldsymbol{d}(t)) \ , && t \in [t_i, t_{i+N}[\ , && (30)\\
& \hat{\boldsymbol{y}}_{k+i} = h(\hat{\boldsymbol{x}}_{k+i}) \ , && i = 1, \ldots, N \ , && (31)\\
& \hat{\boldsymbol{y}}_{k+i} \leq \boldsymbol{y}_{\max,k+i} \ , && i = 1, \ldots, N \ , && (32)\\
& \hat{\boldsymbol{y}}_{\min,k+i} \leq \hat{\boldsymbol{y}}_{k+i} \ , && i = 1, \ldots, N \ , && (33)\\
& \Delta u_{\min,k+i} \leq \Delta u_{k+i} \leq \Delta u_{\max,k+i} \ , && i = 0, \ldots, N-1 \ , && (34)\\
& u_{\min,k+i} \leq u_{k+i} \leq u_{\max,k+i} \ , && i = 0, \ldots, N-1 \ , && (35)\\
& u(t) = u_i \, , && t \in [t_i, t_{i+1}[\, . && (36)
\end{aligned}$$

In the above, N is the prediction horizon, $\{\boldsymbol{c}_{k+i}\}_{i=0}^{N-1}$ is the price signal that reflects the price of the heat. The price indicates the degree to which it is acceptable to heat at any time. This could be with respect to e.g. a CO2-signal that carries information about the CO2-density of the heat. The objective of the MPC is then to minimise this cost while satisfying constraints on the building states and the controllable input.

A well designed MPC is capable of optimising both the thermal comfort of the indoor environment and the energy usage [16]. However, the above mentioned MPC solves only the problem of optimal control on a local building level. The building does not interact with other nearby buildings and does not necessarily align their collective heat load. The heat load of an entire district may consequently be unfavourable for the district heating company. Hence, an obvious question to pose is: How should the heat load of the building stock be shifted to be optimal for both the building and the network? Somehow, we want the entire

heat load of a larger district to follow some reference heat load. For instance to follow the power production curve of renewable energy sources.

In the next section, we briefly describe an overall methodology that enables the application of the presented MPC of buildings to cooperate with the surrounding smart energy networks, amongst these district heating.

5.3 Hierarchical Control

We explained earlier in this book, the complexity of the energy systems are increasing and the demand for adaptability and flexibility has to follow to ensure robust user services at any time. In this example, we focus on 'flexibility' which is the ability to shift demands in time and power, according to the request by the surrounding energy systems, i.e., we might be requested to shift the heat loads of entire building stocks to match some reference heat load on a district level.

[37] has proposed a hierarchical control approach, where the overall optimisation is distributed between involved 'layers'. One layer could be the electrical net, another the district heating system, and at the lower end the individual buildings.

The price signal is determined by a reference type of controller such that a reference heat load can be followed. This gives possibilities for peak shaving or for controlling the load such that it matches e.g. the local renewable energy production.

The goal of a hierarchical control setup is to send out a price signal to the considered building stock that gives the buildings information about the energy prices (which could reflect e.g. the CO2-density of the available energy). A flexibility function is a model that determines the heat demand of a building (or any dynamical system) given a price signal. Figure 6 illustrates this dependence. In [26], a linear model is suggested, but in [27] it is shown that a non-linear grey-box model using stochastic differential equations is more appropriate. By sending out a price signal to the building stock and receiving back the entire heat demand, the upper-level controller is able to shape the price signal such that the heat demand of the entire building stock follows for instance a reference heat load.

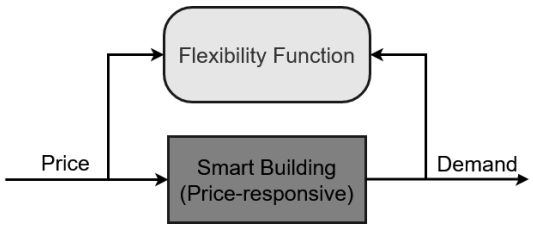

Fig. 6: The heat demand of a building can be predicted from the price signal.

The flexibility function can be determined also for a larger DH network from which we are able to construct a two-layer control setup. The upper layer solves the problem of making the heat load of the building stock follow a reference power signal given by the district heating operators to minimise e.g. the CO2-emissions, while the lower level consists of individual controllers for each building. The latter controllers solve the problem of keeping each building comfortably heated, i.e. maintaining the indoor temperature within the comfort constraint.

References

1. Aggarwal, S.K., Saini, L.M., Kumar, A.: Electricity price forecasting in deregulated markets: A review and evaluation. International Journal of Electrical Power & Energy Systems **31**(1), 13–22 (2009). https://doi.org/10.1016/j.ijepes.2008.09.003
2. Åström, K., Wittenmark, B.: Adaptive Control (2 rev. Dover ed.). Dover Publications (2008)
3. Bacher, P., Bergsteinsson, H.G., Frölke, L., Sørensen, M.L., Lemos-Vinasco, J., Liisberg, J., Møller, J.K., Nielsen, H.A., Madsen, H.: onlineforecast: An r package for adaptive and recursive forecasting. arXiv preprint arXiv:2109.12915 (2021)
4. Bacher, P., Madsen, H., Nielsen, H.A.: Online short-term solar power forecasting. Solar Energy **83**(10), 1772–1783 (2009). https://doi.org/10.1016/j.solener.2009.05.016
5. Bacher, P., Madsen, H., Nielsen, H., Perers, B.: Short-term heat load forecasting for single family houses. Energy and Buildings **65**, 101–112 (2013). https://doi.org/10.1016/j.enbuild.2013.04.022
6. Bavière, R., Vallée, M.: Optimal temperature control of large scale district heating networks. Energy Procedia **149**, 69–78 (2018). https://doi.org/10.1016/j.egypro.2018.08.170
7. Benonysson, A., Bøhm, B., Ravn, H.F.: Operational optimization in a district heating system. Energy Conversion and Management **36**(5), 297–314 (1995). https://doi.org/10.1016/0196-8904(95)98895-T
8. Bergsteinsson, H., Ben Amer, S., Nielsen, P., Madsen, H.: Digitalization of District Heating. Technical University of Denmark (2021), https://orbit.dtu.dk/en/publications/digitalization-of-district-heating
9. Bergsteinsson, H.G., Møller, J.K., Nystrup, P., Ólafur Pétur Pálsson, Guericke, D., Madsen, H.: Heat load forecasting using adaptive temporal hierarchies. Applied Energy **292**, 116872 (2021). https://doi.org/10.1016/j.apenergy.2021.116872
10. Bergsteinsson, H.G., Nielsen, T.S., Møller, J.K., Amer, S.B., Dominković, D.F., Madsen, H.: Use of smart meters as feedback for district heating temperature control. Energy Reports **7**, 213–221 (2021). https://doi.org/10.1016/j.egyr.2021.08.153
11. Blanco, I., Guericke, D., Andersen, A., Madsen, H.: Operational planning and bidding for district heating systems with uncertain renewable energy production. Energies **11**(3310) (2018). https://doi.org/10.3390/en11123310
12. Clarke, D., Mohtadi, C., Tuffs, P.: Generalized predictive control—part i. the basic algorithm. Automatica **23**(2), 137–148 (1987). https://doi.org/10.1016/0005-1098(87)90087-2
13. Crochet, P.: Adaptive kalman filtering of 2-metre temperature and 10-metre windspeed forecasts in iceland. Meteorological Applications **11**(2), 173–187 (2004). https://doi.org/10.1017/S1350482704001252

14. Dahl, M., Brun, A., Andresen, G.B.: Using ensemble weather predictions in district heating operation and load forecasting. Applied Energy **193**, 455–465 (2017). https://doi.org/10.1016/j.apenergy.2017.02.066
15. Dahl, M., Brun, A., Kirsebom, O.S., Andresen, G.B.: Improving short-term heat load forecasts with calendar and holiday data. Energies **11**(7) (2018). https://doi.org/10.3390/en11071678
16. De Coninck, R., Helsen, L.: Practical implementation and evaluation of model predictive control for an office building in brussels. Energy and Buildings **111**, 290–298 (2016). https://doi.org/10.1016/j.enbuild.2015.11.014
17. Dotzauer, E.: Simple model for prediction of loads in district-heating systems. Applied Energy **73**(3), 277–284 (2002). https://doi.org/10.1016/S0306-2619(02)00078-8
18. Drgoňa, J., Arroyo, J., Cupeiro Figueroa, I., Blum, D., Arendt, K., Kim, D., Ollé, E.P., Oravec, J., Wetter, M., Vrabie, D.L., Helsen, L.: All you need to know about model predictive control for buildings. Annual Reviews in Control **50**, 190–232 (2020). https://doi.org/10.1016/j.arcontrol.2020.09.001
19. Fang, T., Lahdelma, R.: Optimization of combined heat and power production with heat storage based on sliding time window method. Applied Energy **162**, 723–732 (2016). https://doi.org/10.1016/j.apenergy.2015.10.135
20. Gadd, H., Werner, S.: Daily heat load variations in swedish district heating systems. Applied Energy **106**, 47–55 (2013). https://doi.org/10.1016/j.apenergy.2013.01.030
21. Giraud, L., Merabet, M., Bavière, R., Vallée, M.: Optimal control of district heating systems using dynamic simulation and mixed integer linear programming. In: Proceedings of the 12th International Modelica Conference (2017). https://doi.org/10.3384/ecp17132141
22. Glahn, H.R., Lowry, D.A.: The use of model output statistics (MOS) in objective weather forecasting. Journal of Applied Meteorology (1962-1982) **11**(8), 1203–1211 (1972)
23. Grosswindhager, S., Roither-Voigt, A., Kozek, M.: Online short-term forecast of system heat load in district heating networks. In: Proceedings of the 31st international symposium on forecasting, Prag, Czech Republic (2011)
24. Grosswindhager, S., Voigt, A., Kozek, M.: Predictive control of district heating network using fuzzy dmc. In: 2012 Proceedings of International Conference on Modelling, Identification and Control. pp. 241–246 (2012)
25. Halvgaard, R., Poulsen, N.K., Madsen, H., Jørgensen, J.B.: Economic model predictive control for building climate control in a smart grid. In: 2012 IEEE PES Innovative Smart Grid Technologies (ISGT). pp. 1–6 (2012). https://doi.org/10.1109/ISGT.2012.6175631
26. Junker, R.G., Azar, A.G., Lopes, R.A., Lindberg, K.B., Reynders, G., Relan, R., Madsen, H.: Characterizing the energy flexibility of buildings and districts. Applied Energy **225**, 175–182 (2018). https://doi.org/10.1016/j.apenergy.2018.05.037
27. Junker, R.G., Kallesøe, C.S., Real, J.P., Howard, B., Lopes, R.A., Madsen, H.: Stochastic nonlinear modelling and application of price-based energy flexibility. Applied Energy **275**, 115096 (2020). https://doi.org/10.1016/j.apenergy.2020.115096
28. Jónsson, T., Pinson, P., Nielsen, H., Madsen, H., Nielsen, T.: Forecasting electricity spot prices accounting for wind power predictions. IEEE Transactions on Sustainable Energy **4**(1), 210–218 (2013). https://doi.org/10.1109/TSTE.2012.2212731
29. Kato, K., Sakawa, M., Ishimaru, K., Ushiro, S., Shibano, T.: Heat load prediction through recurrent neural network in district heating and cooling systems. In: 2008

IEEE International Conference on Systems, Man and Cybernetics. pp. 1401–1406 (2008). https://doi.org/10.1109/ICSMC.2008.4811482
30. Kuboth, S., Heberle, F., König-Haagen, A., Brüggemann, D.: Economic model predictive control of combined thermal and electric residential building energy systems. Applied Energy **240**, 372–385 (2019). https://doi.org/10.1016/j.apenergy.2019.01.097
31. Lago, J., Marcjasz, G., De Schutter, B., Weron, R.: Forecasting day-ahead electricity prices: A review of state-of-the-art algorithms, best practices and an open-access benchmark. Applied Energy **293**, 116983 (2021). https://doi.org/10.1016/j.apenergy.2021.116983
32. Larsen, H., Palsson, H., Bøhm, B., Ravn, H.: Aggregated dynamic simulation model of district heating networks. Energy Conv. Manag. **43**, 995–1019 (2002). https://doi.org/https://doi.org/10.1016/S0196-8904(01)00093-0
33. Ljung, L., Söderström, T.: Theory and practice of recursive identification, The MIT press series in signal processing, optimization, and control, vol. 4. MIT Press (1983)
34. Madsen, H., Holst, J.: Estimation of continuous-time models for the heat dynamics of a building. Energy and Buildings **22**(1), 67–79 (1995). https://doi.org/10.1016/0378-7788(94)00904-X
35. Madsen, H.: Time series analysis. Chapman & Hall (2007). https://doi.org/10.1201/9781420059687
36. Madsen, H., Nielsen, T., Søgaard, H.: Control of Supply Temperature: EFP 1323/93-07. Informatics and Mathematical Modelling, Technical University of Denmark (1996)
37. Madsen, H., Parvizi, J., Halvgaard, R., Sokoler, L.E., Jørgensen, J.B., Hansen, L.H., Hilger, K.B.: Control of electricity loads in future electric energy systems. Handbook of Clean Energy Systems pp. 1–26 (2015)
38. Madsen, H., Sejling, K., Søgaard, H.T., Palsson, O.P.: On flow and supply temperature control in district heating systems. Heat Recovery Systems and CHP **14**(6), 613–620 (1994). https://doi.org/10.1016/0890-4332(94)90031-0
39. Madsen, H., Søgaard, H., Sejling, K., Palsson, O.: Models and Methods for Optimization of District Heating Systems.: Part I: Models and Identification Methods. Informatics and Mathematical Modelling, Technical University of Denmark (1990)
40. Madsen, H., Søgaard, H., Sejling, K., Palsson, O.: Models and Methods for Optimization of District Heating Systems.: Part II: Models and Control Methods. Informatics and Mathematical Modelling, Technical University of Denmark (1992)
41. Nielsen, H., Madsen, H.: Predicting the Heat Consumption in District Heating Systems using Meteorological Forecasts. Informatics and Mathematical Modelling, Technical University of Denmark (2000)
42. Nielsen, H.A., Madsen, H.: Modelling the heat consumption in district heating systems using a grey-box approach. Energy and Buildings **38**(1), 63–71 (2006). https://doi.org/10.1016/j.enbuild.2005.05.002
43. Nielsen, T.S.: Online prediction and control in nonlinear stochastic systems. Ph.D. thesis, Technical University of Denmark, Department of Applied Mathematics and Computer Science (2002)
44. Nielsen, T.S., Madsen, H., Nielsen, H.A., Pinson, P., Kariniotakis, G., Siebert, N., Marti, I., Lange, M., Focken, U., V, L., Hal, H.I.: Short-term wind power forecasting using advanced statistical methods. In: Proceedings of The European Wind Energy Conference, EWEC (2006)

45. Nielsen, T., Madsen, H., Holst, J., Søgaard, H.: Predictive control of supply temperature in district heating systems. Informatics and Mathematical Modelling, Technical University of Denmark (2002)
46. Nielsen, T., Madsen, H., Nielsen, H., Giebel, G., Landberg, L.: Prediction of regional wind power. In: Proceedings of the 2002 Global Windpower Conference, Paris, France (2002)
47. Nowotarski, J., Weron, R.: Recent advances in electricity price forecasting: A review of probabilistic forecasting. Renewable and Sustainable Energy Reviews **81**, 1548–1568 (2018). https://doi.org/10.1016/j.rser.2017.05.234
48. Oliker, I.: Steam Turbines for Cogeneration Power Plants. Journal of Engineering for Power **102**(2), 482–485 (04 1980). https://doi.org/10.1115/1.3230281
49. Palsson, O.P.: Stochastic modeling, control and optimization of district heating systems. Ph.D. thesis, Technical University of Denmark, Department of Applied Mathematics and Computer Science (1993)
50. Palsson, O.P., Madsen, H., Søgaard, H.T.: Generalized predictive control for non-stationary systems. Automatica **30**(12), 1991–1997 (1994). https://doi.org/10.1016/0005-1098(94)90061-2
51. Palsson, O., Madsen, H., Søgaard, H.: Application of predictive control in district heating systems. Proceedings of the Institution of Mechanical Engineers, Part A: Journal of Power and Energy **207**(3), 157–163 (1993). https://doi.org/10.1243/PIME_PROC_1993_207_029_02
52. Pinson, P., Nielsen, T., Nielsen, H., Poulsen, N., Madsen, H.: Temperature prediction at critical points in district heating systems. European Journal of Operational Research **194**(1), 163–176 (2009). https://doi.org/10.1016/j.ejor.2007.11.065
53. Rasmussen, L.B., Bacher, P., Madsen, H., Nielsen, H.A., Heerup, C., Green, T.: Load forecasting of supermarket refrigeration. Applied Energy **163**, 32–40 (2016). https://doi.org/10.1016/j.apenergy.2015.10.046
54. Sandou, G., Font, S., Tebbani, S., Hiret, A., Mondon, C., Tebbani, S., Hiret, A., Mondon, C.: Predictive control of a complex district heating network. In: Proceedings of the 44th IEEE Conference on Decision and Control. pp. 7372–7377 (2005). https://doi.org/10.1109/CDC.2005.1583351
55. Sandou, G., Font, S., Tebbani, S., Hiret, A., Mondon, C.: Global modelling and simulation of a district heating network. In: Proc. of the 8th International Symposium on District Heating and Cooling, Espoo, Finland (2004)
56. Schledorn, A., Guericke, D., Andersen, A., Madsen, H.: Optimising block bids of district heating operators to the day-ahead electricity market using stochastic programming. Smart Energy **1** (2021). https://doi.org/10.1016/j.segy.2021.100004
57. Sejling, K.: Modelling and prediction of load in district heating systems. Ph.D. thesis, Technical University of Denmark, Department of Applied Mathematics and Computer Science (1993)
58. Søgaard, H.T.: Stochastic systems with embedded parameter variations - applications to district heating. Ph.D. thesis, Technical University of Denmark, Department of Applied Mathematics and Computer Science (1993)
59. Steeneveld, G., Koopmans, S., Heusinkveld, B., Van Hove, L., Holtslag, A.: Quantifying urban heat island effects and human comfort for cities of variable size and urban morphology in the netherlands. Journal of Geophysical Research Atmospheres **116**(20) (2011). https://doi.org/10.1029/2011JD015988
60. Thilker, C.A., Bacher, P., Bergsteinsson, H.G., Junker, R.G., Cali, D., Madsen, H.: Non-linear grey-box modelling for heat dynamics of buildings. Energy and Buildings (2021). https://doi.org/10.1016/j.enbuild.2021.111457

61. Thilker, C., Bergsteinsson, H., Bacher, P., Madsen, H., Cali, D., Junker, R.: Nonlinear model predictive control for smart heating of buildings. In: Proceedings of Cold Climate HVAC & Energy 2021 (2021)
62. Thilker, C.A., Madsen, H., Jørgensen, J.B.: Advanced forecasting and disturbance modelling for model predictive control of smart energy systems. Applied Energy **292**, 116889 (2021). https://doi.org/10.1016/j.apenergy.2021.116889
63. Vandermeulen, A., van der Heijde, B., Helsen, L.: Controlling district heating and cooling networks to unlock flexibility: A review. Energy **151**, 103–115 (2018). https://doi.org/10.1016/j.energy.2018.03.034
64. Vlasova, J., Kotwa, E., Nielsen, H.A., Madsen, H.: Spatio-temporal modelling of short-term wind power prediction errors. Informatics and Mathematical Modelling, Technical University of Denmark (2007)
65. Weron, R.: Electricity price forecasting: A review of the state-of-the-art with a look into the future. International Journal of Forecasting **30**(4), 1030–1081 (2014). https://doi.org/10.1016/j.ijforecast.2014.08.008

Paper G

DIGITALIZATION OF DISTRICT HEATING

Authors:
Hjörleifur G. Bergsteinsson, Sara Ben Amer, Per Sieverts Nielsen and Henrik Madsen.

Published in:
Technical University of Denmark.

DTU

Technical
University
of Denmark

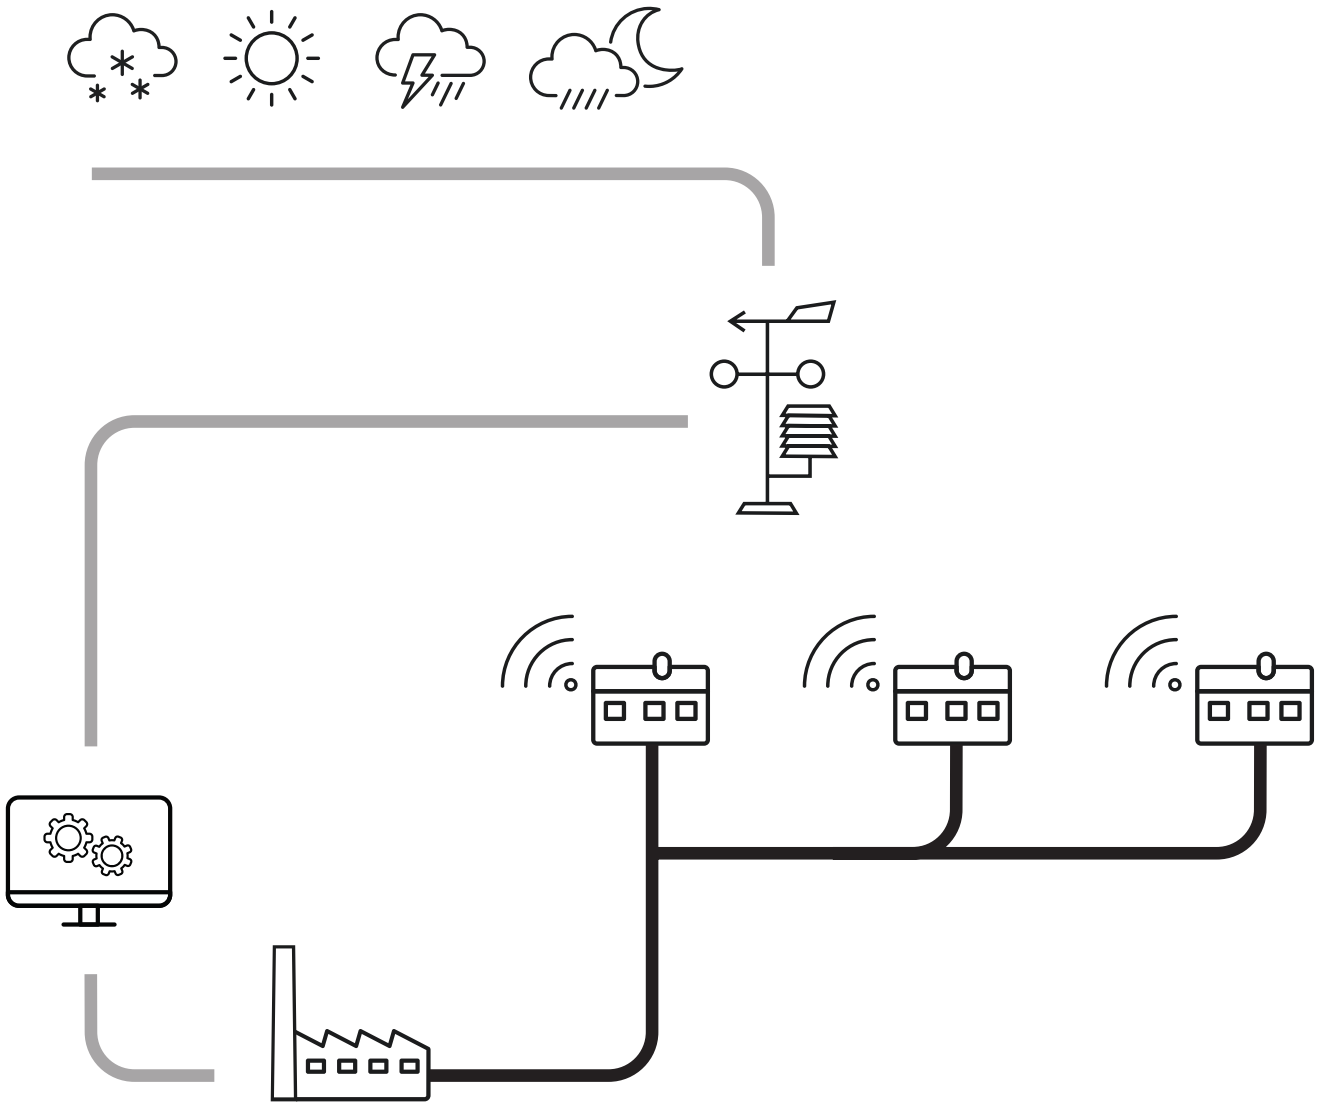

Digitalization of District Heating

Digitalization of District Heating

Experience of using Data-Driven Methods in District Heating: Digitalized Operation in Tingbjerg, Copenhagen

June, 2021

Hjörleifur G. Bergsteinsson[1], Sara Ben Amer[2], Per Sieverts Nielsen[2], and Henrik Madsen[1]

[1] DTU Compute
[2] DTU Management

Cover photo: Bjarne Erick, 2021

Published by: DTU Compute, Department of Applied Mathematics and Computer Science, Richard Petersens Plads, Building 324, 2800 Kgs. Lyngby Denmark
https://www.compute.dtu.dk/

Summary

Operating district heating network using additional data with the traditional data extracted from the SCADA system at the production is discussed. The benefits of including new data in daily operation for the utility is demonstrated. For the past decade, more and more data is becoming available for district heating utilities with the smart meters being installed in every home connected to the district heating network. More local climate stations inside the city are also being installed, and made accessible for everyone. In this report, data from smart meters are presented and how they can be used to operate the network more efficiently. Also, weather forecast in cities is discussed and how they can be improved by localizing them to the local climate using climate stations.

The case study in this report is a on-line operation of temperature control in Tingbjerg which is a small area that is operated by HOFOR. HOFOR is a utility company in Copenhagen which handles for example the district heating, and waste water. They also produce energy for the Copenhagen area. The case study demonstrate how to localize heat demand forecast and operate closed-loop temperature control for a small area. The result for the operation is compared to the previous operation where it was done using open-loop temperature control, i.e. no feedback of the system. The report emphasizes how current state-of-the-art methods can be improved by using newly available data (e.g. smart meters as feedback of the network) and thereby enhancing the efficiency of the operation.

The present report is followed by two other deliverables: the report "Energy data: mapping, barriers and value creation" and the report in Danish, entitled "Digitalisering af fjernvarmen - erfaringer der luner" ("Digitalization of district heating").

Acknowledgements

First of all the authors wish to thank the Capital Region of Denmark for financial support of this project under the project name IDASC (Intelligent Data Anvendelse i Smart Cities - Intelligent data use in smart cities). We also wish to thank Charles Lainez and Kim Mygind from HOFOR for giving us the opportunity of doing an on-line temperature optimization trial in Tingbjerg and for providing the data and information about their system. Many thanks to Tina Hjøllund, Copenhagen Municipality, and Karolina Huss, Gate 21 for fruitful discussions about how to approach and engage people in a technical discussion about district heating.

Special thanks to Torben S. Nielsen at ENFOR for providing us detailed information about the district heating system and insights on how data-driven methods can be applied. Furthermore, we would like to acknowledge ENFOR for allowing us to use their state-of-the-art software to demonstrate how digitalization can improve state-of-the-art methods.

Contents

Chapter 1

Introduction

The IDASC, *Intelligent Data-Anvendelse i Smart Cities* (Intelligent Data Use in Smart Cities) project's goal is to investigate potentials from several different data sources that are now available because of digitalization in district heating systems and generally in cities. The overall purpose is to consider and combine all relevant data from meteorological services, city weather data, production data, SCADA data, and end-user smart meter data to enhance the operation of a network. We will evaluate the advantages of combining different data sources and analyze the improvements compared to the typical situation today using only data from the SCADA system. Data from the SCADA system are usually measurements measured at the production site, e.g. supply temperature, return temperature, flow, and ambient air temperature. The focus is to demonstrate the potential of how the new data sources can improve the operating district heating network, i.e. delivering heat from production to consumers in a more optimal setting. In this study, we do not directly consider how digitalization can lead to more optimal production of the heat; however, operating the network efficiently will obviously also have a positive influence on the possibilities for optimized production planning, hopefully lowering the production cost. Therefore, in this report, we only consider heat demand forecast and temperature optimization to increase the operation of the district heating network. The additional data used in this report are the smart meters at consumers in the area where the heat is delivered and a local climate station that is located close to the area. In theory, using this additional data will enhance the operation as it gives more detailed information on the response characteristics of the network and the local climate in the area. We will demonstrate this by using state-of-the-art and off the shelves algorithms for forecast and control provided by ENFOR, an energy forecasting company. ENFOR is a spin-off company from DTU where the initial ideas of these algorithms were established. We will show how to include the proposed additional data in the algorithms and demonstrate the improvements in the operation of the district heating network. In the case study, we will apply these digitalized methods in an on-line operation trial to analyze the gain of using smart meters when used as feedback to obtain a closed-loop temperature control of the network, and the importance of binding numerical weather prediction to a local climate using a local weather station. The trial was conducted in Tingbjerg (Copenhagen), an area which is operated by the district heating utility, HOFOR.

1.1 District Heating (DH) Operation

During the past decades, the district heating sector has been transformed from using primarily traditional fossil fuels to using renewable heat sources and biomass. During the same period, the DH systems have become more digitalized, e.g. with sensors in the district heating network and smart meters at the consumers. At the same time, district heating is becoming a crucial part of the overall integrated energy system because of its flexibility potential, e.g. by storing excess energy as thermal energy [1, 2]. Hence, optimal operation of district heating is crucial.

The inherent flexibility of the DH system is highly valuable for the future integrated and low-carbon energy system. An important aspect is a fact that DH systems can store energy when there is

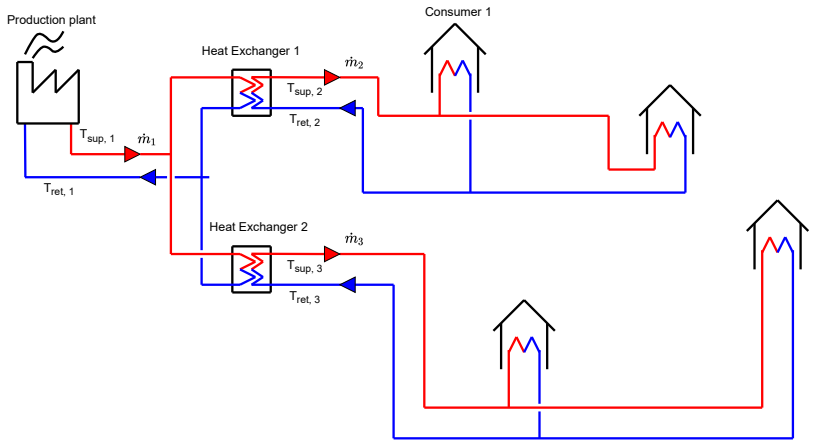

Figure 1.1: Simple schematic view of district heating; The heat production, the transmission lines to heat exchangers then distribution lines that deliver the hot water to the consumer substations.

a surplus of energy from intermittent renewable energy sources (e.g. wind and solar). However, to maximize the flexibility potential of district heating, they would need to operate efficiently by optimizing the production and the temperatures of the network. This report will focus on how to improve the efficiency of operating an existing district heating network by applying data-driven methods using additional data.

District heating consists of heat production, a network of pipes (transmission and distribution) where the hot water is either delivered to substations (heat exchangers) at the consumers and returned to the production facility, and the final component is the consumers. This is illustrated in Figure 1.1. The supply temperature is generated at the production by heating the water, for example, at a Combined Heat and Power (CHP) plant where the temperature is increased by cooling the steam after it has generated electricity in the turbine. The mass flow in the pipes is then controlled using pumps at the production plant. Frequently, additional pumps are required in the network to maintain the desired pressure in the system. First of all, an optimal operation of DH systems implies that the supply temperature and the network temperature should be kept as low as possible without violating any requirements, e.g. supply temperature at a given outdoor temperature. Lowering the supply temperature will reduce the heat loss in the network, and improve the efficiency of the electricity production at CHP plants [3, 4]. Furthermore, a lower temperature implies also more optimal use of, for instance, heat pumps.

Delivering the heat demand is controlled by varying the supply temperature [5]. Controlling the operation of a district heating network rely on either an open-loop or a closed-loop controller to estimate how the heat should be delivered, by regulating the supply temperature where the flow is indirectly varied to meet the demand of the consumers. The open-loop controllers use either a white-box simulation of the system to operate the supply temperature in the network or a simple algorithm based on the knowledge of the system to regulate the temperature. Hence, the open-loop operation does not have any feedback from sensors and data in the network and therefore such controllers can not adapt to any disturbance in the system or changes to the network characteristics. Thereby, they do not use the information from the network to adapt to achieve more optimal operation.

There have been proposed control schemes that operate the supply temperature in a closed-loop [5, 6]. Such a system typically uses a few measurement points located in the network. These points are usually located in the network where the operator believes that the lowest (critical) temperature is, i.e. where the largest temperature loss occurs. Therefore, the supply temperature at

the production site is controlled to satisfy the requirements at these critical points. The controllers also control the flow in the system to match the heat demand of the system, and therefore the optimal operation is implemented with a sequence of controllers trying to deliver the heat while keeping the supply temperature at a minimum.

The approach for temperature control in this report is found by lowering the supply temperature while keeping the flow close to the operation limit of the system; see [7]. At the same time, a lower supply temperature will enhance the efficiency of power generation at the CHP plant. Heat demand forecasts are needed for production optimization, and these forecasts are also used for finding the optimal supply temperature [8]. The heat demand is highly correlated with ambient air temperature and therefore usually the forecast model uses numerical weather prediction (NWP) of the ambient air temperature as input. Accurate NWP will be beneficial to the district heating operation as they improve the heat demand forecast accuracy. In addition, both smart meter and local climate station data can help to improve the closed-loop control of the system. The models used are self-calibrating, and consequently, the models automatically adapt to network characteristics as well as local climate conditions.

1.2 Structure of this Report

In Section 1.3, we will describe the case study used in this report. In Section 2, we will discuss how to use data from smart meters and how it can lead to additional cost savings related to district heating operations. Section 3 describes how Numerical Weather Predictions (NWPs) and local climate data can be beneficial for the district heating operations. Heat demand forecasting is characterised in Section 4, while control of district heating network temperatures is described in Section 5. The report finally concludes in Section 6.

1.3 Case study: Tingbjerg district heating area

The case study used in this report is the district heating network located in Tingbjerg, which is a small area with large apartment buildings located in the northwestern part of Copenhagen. The area is supplied by heat from a heat exchanger that connects the central Copenhagen transmission system operated by CTR to the distribution network in Tingbjerg operated by HOFOR. The transmission system operated by CTR supplies heat to approximately 250.000 buildings in central Copenhagen. HOFOR is the distribution network operator, and consequently, HOFOR is the district heating supplier to buildings in central Copenhagen. There are 45 buildings connected to the network inside the Tingbjerg area and 39 of them are equipped with a smart meter. Figure 1.2 shows the layout of the network in Tingbjerg.

Previously, HOFOR has operated Tingbjerg as an open-loop system using the TERMIS tool to simulate and adjust the supply temperature and flow from the heat exchanger. Thus, they had no knowledge of what temperature consumers were receiving, i.e. how the system was working except when consumers complained because of too low temperatures. However, each apartment building in Tingbjerg has a smart meter that is connected to the district heating side. These smart meters can provide the forward temperature, return temperature, flow, and energy consumption for each building. HOFOR also needs a heat demand forecast to regulate the temperature for the heat exchanger, which they get by scaling demand forecasts for a larger area that contains Tingbjerg heating demand. The scaling factor is the ratio between the larger area's historical demand and the historical demand from Tingbjerg.

Therefore, Tingbjerg is an ideal case for demonstrating how the operation of an existing network can be improved with data-driven methods and digitalization. In this study, the smart meters will be used to provide feedback for closed-loop control using the data from the smart meters to increase the efficiency of the network operations and to show potential savings by lowering the supply temperature at the heat exchanger while satisfying all requirements. Heat demand forecast

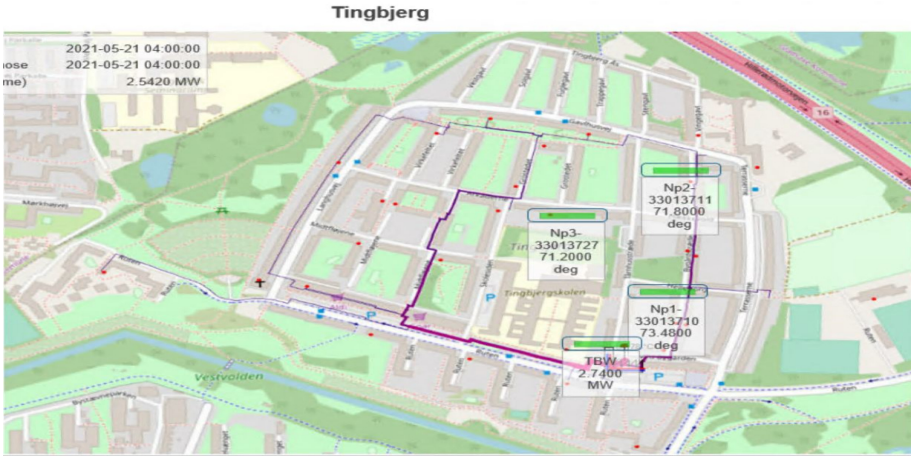

Figure 1.2: Layout view of the district heating network in Tingbjerg. The heat exchanger is located where the box with the TBW text is located. The other three shows the status of the smart meters used to give feedback.

for the Tingbjerg area will be created based on the historical demand and NWPs that will be localized to the area by using a local climate station that is close by. We will then demonstrate the benefits of using automated feedback techniques from an on-line operation use of our setup. The period from the start of the on-line operation trial from 1st of November 2020 until 1st of April 2021 will be used to compare the new data-driven approaches with the methods used previously.

Chapter 2

Smart Meters in DH systems

This chapter introduces smart meters in district heating and their role in transforming district heating systems into the digital age. We will also discuss how smart meters can give value for both consumers, where they are installed, as well as for the district heating utility by enhancing the network performance. This report will focus on using smart meter data as feedback to the production and use the response of the network to improve the performance, as the objective of this report is to increase operational savings for the utility using additional data.

Smart meters are and have been installed at district heating consumers in Europe for the past decade because of requirements from the European Union. The requirement is that consumers that are connected to district heating networks have to be equipped with smart meter devices where feasible [9]. This enables the consumer to be more aware of their current energy consumption and allows linking the consumption to the billing from the utility. They can now see their consumption on higher resolution and even in some cases, they have it available on-line. This has prompted a different payment schema from the district heating operator. For example, consumers are penalized with a fine if their daily average return temperature in a period is higher than a certain limit because it is costly to the system. District heating utilities attempt to recover these costs by penalizing consumers that have a bad cooling effect in their buildings and hence a higher return temperature. A higher return temperature implies extra heat loss in the return pipes, higher pumping costs, and less efficient production at the plant [10, 5]. Hence, the large amount of data that is now available due to the smart meters, that can be used to identify the energy performance of buildings and the network. This can be used to give valuable insights into the network performance and building energy efficiency, i.e. leakage in the system or insufficient cooling of the water from inlet to outlet in some buildings.

Current studies that use smart meter data often only have the building's energy performance improvements as the center of attention. Kristensen and Petersen [10] use smart meter data to derive three heating efficiency indicators of buildings and give an overview of the smart meters system at the district heating utility in Aalborg in Denmark. The three heating efficiency indicators are annual heating energy use intensity, daily heat load variation, and cooling efficiency. Thilker et al. [11] demonstrate that it is possible to lower the operation cost by 10% by a date-driven control of the heating system of a Danish school building, and they use model predictive control to estimate future set-points of the thermostats of the radiators to lower the return temperature of the district heating to the school. Bacher et al. [12] suggest a method to separating the heat load used to hot tap water from the load needed for space heating. Most of the research is focused on individual building energy performance. Improving the aggregated performance of a larger network using smart meters is usually not investigated except for identifying bad coolers in the network, i.e. high return temperature from substations. Lowering the return temperature from substations in the building is immensely important for district heating, especially networks that have combined heat and power production (CHP). As also stated in Arvastson [4] a higher return temperature

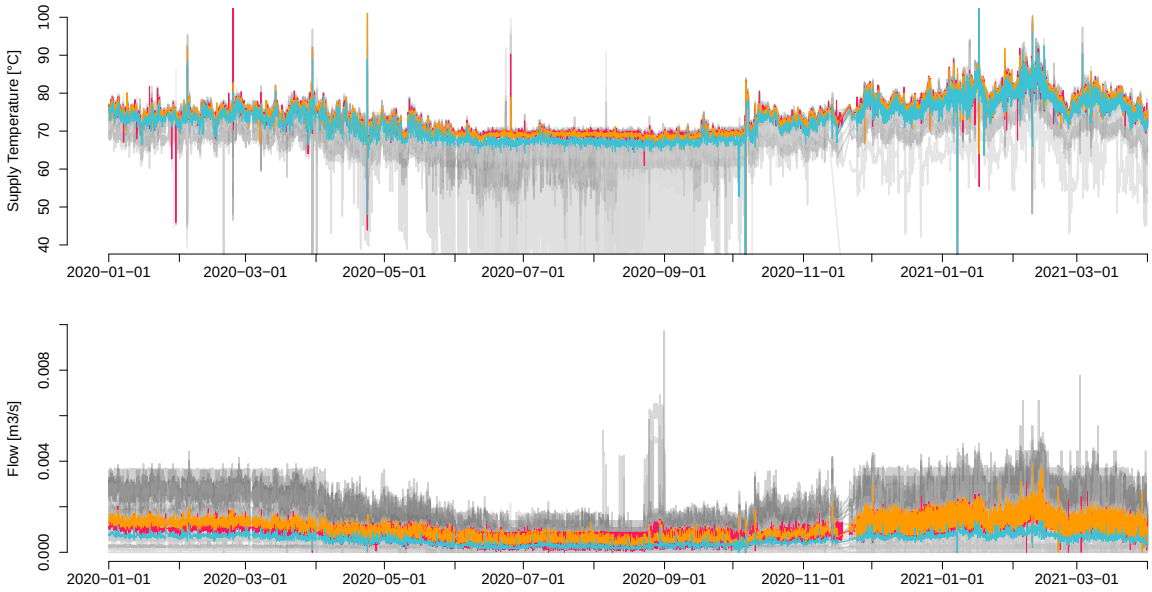

Figure 2.1: Supply temperature and flow from smart meters at Tingbjerg visualized. The coloured lines are the meters that were selected to be used as feedback for the temperature optimization.

from the network to the condenser will decrease the efficiency of the operation of the CHP plant. Lower return temperature will also lower the heat loss in the return pipes. It also implies a lower necessary flow rate or supply temperature for the given energy.

Data from smart meters have not previously been used to enhance the operation of the network, i.e. used as feedback of the temperature of the network to the temperature optimization at production. Until now, for closed-loop temperature control, only data from a few measurement points in wells have been used. These measurement points are called critical points and selected where the operators believe to have the highest temperature loss. The temperature control uses the feedback of the system as input to estimate the model parameters and time delay of the system to control the supply temperature and flow with the objective of keeping the temperature as low as possible [5]. Hence, smart meters can be used as a feedback signal to the controller either by estimating the supply temperature in the street pipe using groups of smart meters or by using a single smart meter from a large apartment building where the heat loss in the service pipe to the building is negligible because of the high flow [13].

The smart meter data used in this report will be introduced in Section 2.1 and a more detailed description of using the smart meter as feedback to the control will be given in Section 2.2.

2.1 Smart meter data

The smart meter data comes from the case study area, Tingbjerg where HOFOR provided access to on-line smart meter data from 39 meters located in large apartment buildings from January 2020. In the beginning, the data was sent only at 09:00 each morning where the data had an hourly resolution. These readings contained data from each smart meter for the past 24 hours. However, by the end of November 2020, the resolution was updated to 15 minutes and the data was sent each hour containing the past four data readings. Table 2.1 shows the variables that are logged by the smart meter. Usually, the utility companies use this information to bill the consumers based on their energy consumption. Moreover, if the return temperature is too high then there could be a penalty payment scheme in place as discussed before. Otherwise, the data is often not analyzed further.

Variables	Units
Time	Date, Time
Cumulative Energy	MWh
Cumulative Volume	m^3
Supply Temperature	°C
Return Temperature	°C

Table 2.1: Variables from the smart meters.

Figure 2.1 shows the supply temperature and flow from the smart meters in Tingbjerg. The flow rate is computed from the cumulative volume by taking the difference between the volume at each time-step and divide with the corresponding time in seconds between, resulting in a flow rate in cubic meters per second. Notice, the difference between the readings in the summer and winter periods. The winter period consists of stable temperatures, it has a quite constant variation and does not drift off towards zero. In the summer period, the temperature is noisier and fluctuates more. For some of the meters, the temperature seems to drift off towards zero. The reason for this is the low heat demand during the summer when there is almost no need for space heating in Denmark because the ambient air temperature is around 20 °C. Only domestic hot water usage is needed during warm periods. This seasonal variation of the heating demand can be seen in the flow plot as the flow decreases over the summer periods. Therefore, when there is no heating consumption, the water in the service pipe to the building becomes still and the temperature starts to drop because of the heat loss to the surroundings [13]. Thus, the readings from smart meters over the summer period are more unreliable compared to the winter periods as they do not give an accurate representation of the temperature of the hot water in the distribution pipes. When selecting smart meters to be used as feedback to the controller, this needs to be considered. The selected smart meters need to have a very stable and constant flow during the summer period, or create an algorithm that addresses the temperature drop.

During a short period at the end of November 2020, readings from the smart meters are missing. This is happening when HOFOR increased the resolution of the readings from hourly to 15 minutes and updated the frequency of the readings. Therefore, the period from 2020-11-15 to 2020-11-22 has almost no information. We also see frequent spikes in the data for both high and low temperatures. These can be faulty readings in the meters as the figures present the raw data (instantaneous values), i.e. no quality check of the values has been conducted. There is also a significant peak period in the flow, just before September 1st, 2020. The peak could be a consequence of the fact that the ambient temperature dropped rather quickly and the heating demand therefore increased while the supply temperature at the production has not increased. Notice also that after the increase in resolution the data seems to be more volatile compared to the hourly resolution period. Higher resolution leads to an increased risk of outliers and also that we are able to see more dynamics in the heating demand than before.

2.2 Smart meter data used in the feedback loop

The main objective of this report is to demonstrate the value of smart meter data for operating the district heating network. More specifically, we will use smart meter data as feedback to temperature optimization. The feedback will be used to give the controller signal on how the network is reacting to changes in the supply temperature, and hence creating a closed-loop controller. Previous closed-loop controllers used measurement wells in the network as feedback. This replacement of a well measurement with the use of meter data constitutes a digital transformation of the closed-loop control as demonstrated in Figure 2.2. This digital transformation reduces the need for measurement wells in the network and reduces the maintenance effort. We think that this gives significant savings potential for the district heating operator as it reduces the cost of having feedback control for operators without the need of installing measurement wells in the system.

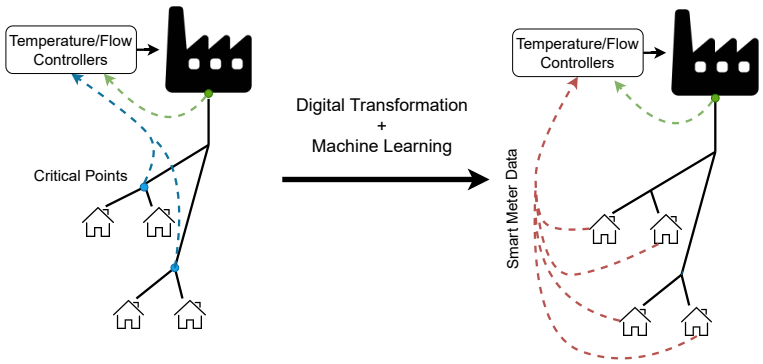

Figure 2.2: A sketch of DH network demonstrating that the production site uses feedback from smart meters for the temperature control, instead of the measurement wells (critical points).

In Tingbjerg, the operation in the past did not allow for a closed-loop installation as the area does not have any measurement wells in the network to send feedback. In this study, Tingbjerg will be operated using a closed-loop controller and demonstrate the benefits of having a closed-loop by comparing the result to the open-loop operation. The controller will use feedback from the network using data from smart meters at large apartment buildings. The methodology behind the temperature control is described in more detail in Section 5. However, in short, district heating systems are complex non-stationary and time-varying systems therefore methods for tracking the time-varying parameters are needed as a part of the modeling process. In practice, the parameters are updated based on different input data, and here the feedback from the network data is essential due to how the system reacts to different flows, e.g. how the time delay varies. Thus, we have to select smart meters that send reliable signals to guarantee suitable feedback for real-time estimation of the parameters of the model. As mentioned before, three smart meters were selected as feedback. More meters could have been used however the Tingbjerg district heating network is rather small and three feedback or critical points are considered to be sufficient.

Finding suitable meters to be used as feedback is critical. An obvious first task is to identify a set of meters that we consider ideal for representing the entire district heating area. Ideal meters are meters that during the summer period, they have a stable supply temperature which implies that the flow is usually high during longer periods. From this group of ideal meters, three meters were selected based on a few numbers of missing values and seemed ideal to be used as feedback as historically they have sent reliable data and usually constant flow, i.e. heating of the house was not stopped by closing the flow of the water. Figure 2.1 shows data from all of the meters and the three meters are highlighted with bolder colored lines. Other meters from the ideal group could also have been used. In the future, other meters can replace the current feedback meters if deemed necessary for being able to include the lowest temperature in the network at all points in time. The three selected meters are shown in color in Figure 2.1 to visualize their reliable signal. The meters were selected without knowing their exact location in Tingbjerg. Knowing the location of the meters gives additional useful information, as selecting meters that are placed the furthest away from the production could be ideal for the feedback loop. Usually, the consumers with the highest transportation time have also the largest temperature loss in the system however it could be that some part of a network is older and the efficiency of the pipes insulation has been reduced. Therefore, having larger temperature loss even though they are closer than other areas. By satisfying the requirements of the consumers with the highest temperature loss (which are the critical points in the network), the other consumers' requirements are also therefore fulfilled. An exception is if other buildings have faulty service pipes into the houses or leakages, but these issues can be quickly discovered and fixed when investigating the smart meter data.

Chapter 3

Weather Forecasting for Cities

District heating is mostly applied in urban areas therefore in this section, we will introduce the climatic characteristics inside cities and highlight the effects of climate variables on heat consumption. Numerical Weather Predictions (NWPs) are also introduced as they are critical for district heating operations. They are needed to forecast the future, concerning demands, temperatures, and production planning, or in short; everything that district heating operates needs weather forecast as input in order to operate the systems efficiently. We will discuss the advantages of localizing NWPs to cities, and more specifically, we will look into enhancing short-term heat demand forecast in cities. The heat demand forecast accuracy is improved by correcting the short-term weather forecast using real-time measurements of the climate from a local station. Hence, increasing the accuracy of the short-term forecast is highly desirable for temperature control.

3.1 Urban Heat Island effects

Temperature optimization in the district heating network depends on obtaining reliable and relevant monitored outdoor air temperature data. The more accurate the air temperature around the district heating network and the more frequent the temperature measurements, the more accurate the temperature optimization model can be. Research has shown that the outdoor air temperature can vary across a large district heating network and it is therefore also important to obtain temperature data across the network - if possible Steeneveld et al. [14]. Outdoor temperature data are historically monitored in rural areas at sites where measuring the correct temperature has been easiest. Airports have been a good choice because they are in a rural setting, where the only impact is from the natural environment including lack of woody vegetation and directly exposed to natural rain, sun, and wind. Historically, temperature data from airports are often used as input for the temperature optimization in operating the district heating network, simply because they are available from the meteorological institutes. Recently data from other sources are becoming available, e.g. Danish Meteorological Institute recently started to give the public access to their climate station that are located everywhere in Denmark. However, it is important to notice that the air temperature measured in the airports may deviate from the temperature insides cities, where the air temperature is exposed to human activities and the built environment.

Research shows that the outdoor air temperature typically is higher in urban areas than in rural areas Steeneveld et al. [14]. The effect is termed urban heat island (UHI). An urban heat island (UHI) is an urban area that is warmer than its surrounding rural areas due to human activities or build human infrastructures. Research related to UHI has recently got more attention because of the concern that climate change with an average temperature increase of 2 to 3 K will cause more heat waves, becoming more severe in the future, causing significant stress to the urban population. That problem is however only relevant for hotter climates where there is no district heating or at least it happens outside of the heating season. It is though still relevant for the energy sector as a whole with significant cooling demand during a heat wave. Unfortunately, there is not the same interest in studying temperature differences between urban and rural areas during winter which

would be relevant for the district heating sector.

The variation of outdoor air temperature data is both spatial and temporal. A number of studies point towards a typical difference in urban and rural temperatures of 2-3 K. In a study in Barcelona the city centre was 2.9 K warmer than the airport during nighttime, but during the day the centre is slightly cooler than the periphery. Annually and overall, Barcelona centre is 1.4 K warmer than the airport. With regard to the average differences between the minimum at the two places, all are over 2.5 K, reaching 3 K in November and March [15]. Solecki et al. [16] examined the UHI mitigation potential of two highly urbanized places in the state of New Jersey, areas in and around the cities of Newark and Camden. Each city and surrounding suburbs included a set of neighborhoods with widely varying characters. The UHI effect in Newark is estimated to be on average about 3.0 K and for Camden between 1.0 and 1.5 K. Steeneveld et al. [14] has in a comprehensive study collected data from both private weather enthusiasts and weather stations to determine UHI in the Netherlands. They report a temperature difference of 2.5K. However, the paper focuses on UHI and its effects during warm seasons. There is no seasonal evaluation of the differences. It is therefore not certain that the same temperature difference occurs during the heating season. A review of research studies and data found that in the United States, the heat island effect results in daytime temperatures in urban areas about 1–7°F higher than temperatures in outlying areas and at nighttime temperatures are about 2–5°F higher. Humid regions (primarily in the eastern United States) and cities with larger and denser populations experience the greatest temperature differences [17].

For most cities, the difference in temperature between the urban and surrounding rural areas is largest at night. Throughout the daytime, particularly when the skies are cloudless, urban surfaces are warmed by the absorption of solar radiation. Surfaces in the urban areas tend to warm faster than those of the surrounding rural areas. By virtue of their high heat capacities, urban surfaces act as a giant reservoir of heat energy. As a result, the large daytime surface temperature within the UHI is easily seen via thermal remote sensing [18]. The typical temperature difference is several degrees between the center of the city and surrounding fields. The annual mean air temperature of a city with 1 million people or more can be 1.0–3.0 K warmer than its surroundings. In the evening, the difference can be as high as 12 K. [17]. This is also shown in the Barcelona study [15]. At night, the situation reverses. The absence of solar heating leads to the decrease of atmospheric convection and the stabilization of the urban boundary layer which traps urban air near the surface, and keeping surface air warm from the still-warm urban surfaces, resulting in warmer nighttime air temperatures within the UHI. Furthermore, the heat retention properties of urban areas, the nighttime maximum in urban canyons could also be due to the blocking of "sky view" during cooling: surfaces lose heat at night principally by radiation to the comparatively cool sky, and this is blocked by the buildings in an urban area. Radiative cooling is more dominant when wind speed is low and the sky is cloudless, and indeed the UHI is found to be largest at night in these conditions [19]. The outdoor air temperature changes from hour to hour, minute to minute, and even second to second. A change in wind and clouds can change the air temperature very rapidly.

During the last 100 years, cities have not been built with the UHI impact in mind. The main cause of the urban heat island effect is from the modification of land surfaces, which traps heat during the day. Waste heat is produced by energy usage as a secondary contributor. Dark surfaces such as roads and buildings absorb significantly more solar radiation, which causes increased heat absorption in cities more than suburban and rural areas during the day [16]. Materials commonly used in urban areas for pavement and roofs, such as concrete and asphalt, have significantly different thermal bulk properties (including heat capacity and thermal conductivity) and surface radiative properties (albedo and emissivity) than the surrounding rural areas. This causes a change in the energy budget of the urban area, often leading to higher temperatures than surrounding rural areas [20]. It is therefore also clear that mitigating strategies can be applied in city planning to reduce the UHI. Using lighter, more reflective materials in the built environment will reduce the UHI effect as well as planting trees will reduce the UHI effect [19].

The Barcelona study [15] also illustrated that the higher outdoor air temperature occurred at the city center. The further out the less dense the city is and the more trees are part of the build

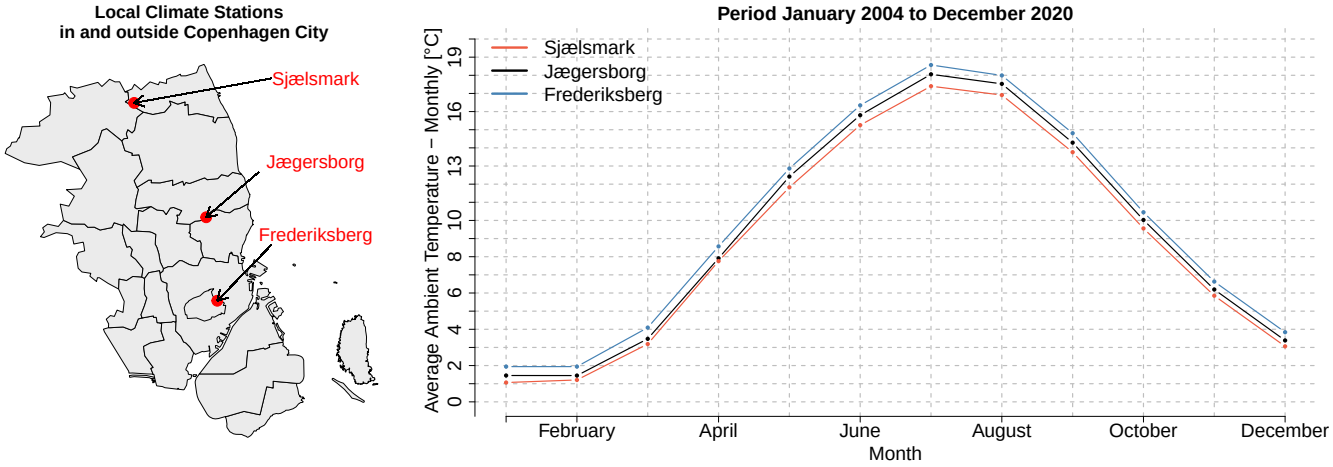

Figure 3.1: Urban Heat Island: Copenhagen demonstrated using three climate stations located with different proximity to the center of Copenhagen; The center (Frederiksberg), in the outskirts (Jægersborg), and in a rural area (Sjælsmark) as shown in the map and the plot showing the difference in the average monthly temperature for the three stations.

environment. In cities where the UHI effect is taken into account in the city planning, it is likely that the impact of the UHI is lower in the newer build districts.

Although as a rule of thumb the temperature is higher in urban areas than in rural areas where the outdoor temperature data is measured - it is also likely that the temperature difference between urban and rural is not so high in the outer districts of the city. It is therefore reasonable to have a number of temperature measurement stations implemented across the city.

3.2 Urban Heat Island: Copenhagen

As discussed, researchers have shown that there is a temperature difference within cities and between urban and rural areas. We would like to confirm this phenomenon in Copenhagen and investigate the magnitude of the difference. Our goal is also to discuss the impact on district heating operations due to different climates within the city. This will be done using three local climate stations. The data was extracted from the Open Meteorological Data provided by the Danish Meteorological Institute (DMI) [21]. Figure 3.1 shows a map of Copenhagen and the locations of the climate stations as red points, also the average monthly temperature from the stations is visualized in a plot. One climate station, Frederiksberg, is located very close to the city center, a densely populated area. The Jægersborg climate station is located in the outskirt of Copenhagen, while Sjælsmark is located in a rural place north of Copenhagen. The past hourly mean temperature was extracted from January 2004 to December 2020 from each climate station. The monthly average ambient temperature was then computed over the period. The temperature plot in Figure 3.1 demonstrates the UHI effect in Copenhagen. The temperature difference between the stations illustrates a significantly higher average temperature in stations that are located closer to the city, with more population and building mass. The difference is close to 1 K during the heating season and 1.5 K during summer periods. The difference between the season could be due to the solar radiation which heats up the buildings, streets while in the rural area, the terrain does not absorb as much solar gain, thus the temperature average increases over summer. The difference during heating seasons could be because of the heat from the buildings, transportation, and people as mentioned in Section 3.1.

Figure 3.2 shows hourly temperature average for four different months computed using same period. We can see that the climate has a time-varying process, both a diurnal variation and yearly. For example, comparing the result from July shows that the average temperatures in the mornings are very similar. However, as time progress, it differs, with higher temperature difference in the city and during the night it gets colder at the rural side, Sjælsmark. Thus, the city does not lose heat

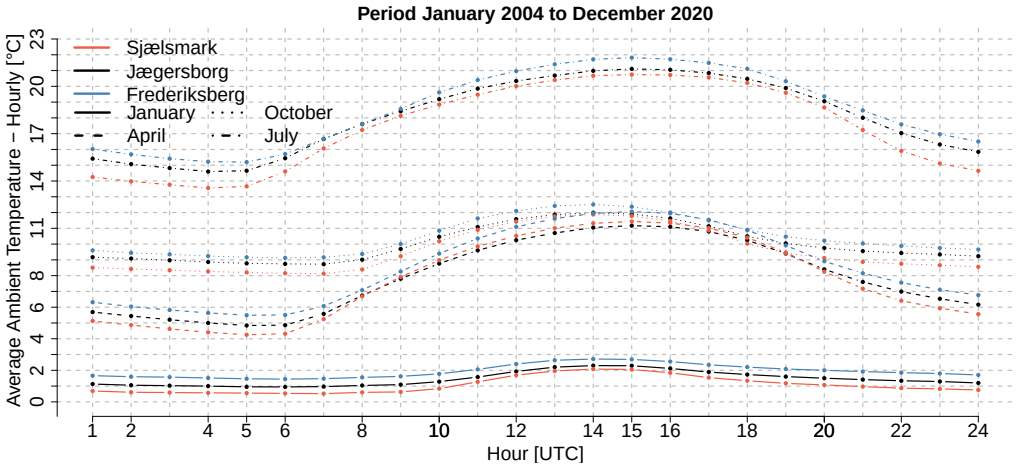

Figure 3.2: Urban Heat Island: Copenhagen, where four months are used to demonstrate time-varying climate for each station and between them. The plot shows the average temperature per hour for four months computed from the period, January 2004 to December 2020.

as fast as the rural part. January has quite a constant offset between the stations, except during the day when the temperature at Sjælmark is almost the same as at Jægersborg.

3.3 Numerical Weather Prediction in Cities

We have exemplified that the UHI phenomenon is relevant in Copenhagen where the temperature inside the city is different compared to the rural areas because of dense population and different environments (e.g. buildings, roads, etc) which entraps heat in cities. Three climate stations located in different areas in Copenhagen were used to demonstrate the phenomenon that exhibits higher temperatures closer to the city. Thus, the climate differs between locations. In Section 5, we will discuss that the temperature operation of district heating is heavily dependent on the most recent climate variables, where the ambient temperature is the most important. The climate where the district heating system is located needs to be analyzed to operate the network in an optimal setting. District heating is an efficient way to provide heat to buildings in densely populated areas. It therefore means that it is highly important to have localized climate data that can be used to analyze the heating demand dependency. For example, see Figure 4.2, to see heating demand vary over time. The variation can be explained both from the climate variation over time, e.g. high sea temperature in October contributes to higher ambient temperature during the night, see Figure 4.3. Hence, we have seen that ambient temperature differs depending on the location, building mass, sea temperature, i.e. there are many factors that contribute to the temperature. This holds also for other climate variables like wind and solar radiation. Climate variables are an important factor for analyzing heating consumption. In Nielsen and Madsen [22] and Madsen et al. [23] suggest that the climate variables: ambient temperature, solar radiation, and wind speed (including direction) have the most effect on the heating demand. They are also arranged in decrease importance. Nielsen and Madsen [22] give a detail description on how these climate variables interact with the heating consumption based on physical consideration, i.e. stationary relations. Here is a short summary of the findings;

- *Ambient Temperature:* The ambient temperature affects the indoor climate through heat conduction in the outer walls and windows, also through ventilation. It has been shown that the outdoor temperature affects the indoor temperature with a low-pass filter, a transfer function to model the variations in the outdoor temperature to variations in the indoor temperature.
- *Solar Radiation:* The solar radiation affects the indoor climate based on the angle of beams hitting the building, where the orientation of the beams through the windows and the window

area are most important. Basis functions are used to translate the non-linear dynamics of the solar radiation to its contribution to heating consumption.

- *Wind Speed:* The wind speed and the direction of the wind affect the indoor climate as natural ventilation, the effect is depending on the quality of the insulation. The wind speed also affects the convection heat coefficient on the outside of the buildings. It is therefore modeled as a low-pass filter as the contribution to the consumption.

Hence, to use these climate variables to describe the heating demand to estimate future supply temperatures then a forecast of them is needed. NWP is computed as a physical model of the atmosphere and ocean to predict climate variables. They are computed over a grid of the earth and are then interpolated together to a specific location where weather predictions are needed. However, NWPs have problems adjusting to the local climate in cities due to the local climate phenomenon. Thus, the models seem to have trouble adjusting to local heat contributions, e.g. solar heat in the street, heat from buildings, etc. District heating relies heavily on NWP to operate their system efficiently therefore it is important to correct them before using them as input. Especially, for temperature control of the system as it is done on a short-term horizon (between hourly and 24 hours) and is heavily dependent on the current local climate. Using local climate station to localize the NWP, corrects the short term NWP forecast by adapting them to the climate using real-time climate measurement [22]. Hence, this yields an optimal weather forecast for a certain area that can be used to operate the temperature control in the most optimal setting. This is discussed in more detailed in Section 4 where local climate station improved heat demand forecast and in Section 5 to improve the temperature control.

3.4 Localize Numerical Weather Prediction

We have discussed that it is important to localize numerical weather prediction to enhance the operation of district heating systems. Combining weather forecast to a local climate have been studied for many decades as it can be highly desired to have an accurate forecast to yield optimal operation. Incorporating certain local climate features into the forecasts is done to adjust the systematic errors from the NWP model. Glahn and Lowry [24] propose Model Output Statistics (MOS) to bind NWP to local climate stations observations, e.g. localize the forecasts. The MOS is a simple technique that uses linear regression where the observed climate variable is the response variable and predictors are the NWP variables which therefore bind the NWP to the local climate. It is a simple and frequently used method that will reduce systematic bias in the NWP if there is any. Crochet [25] propose using an adaptive method to reduce the systematic bias and lowering the RMSE of the NWP. A Kalman filter is used to localize NWP to the local area. It was demonstrated that the proposed methods decreased the systematic bias and reduced the error in areas where systematic bias is high. If there was no systematic bias then the Kalman filter does not significantly improve the forecast. The Kalman filter also gives useful information about the uncertainty of the local predictions when localizing the NWP to climate stations. The uncertainty from the weather forecast can be useful information for both temperature control of the network and production planning of the plant.

Chapter 4

Heat Demand Forecasts

Operating a district heating production facility and controlling the network efficiently is a difficult process. Both tasks need to consider multiple inputs to deliver a feasible production plan and accurate controller. One of these inputs is the heating demand of the consumers. Satisfying the consumers' heating consumption need is the most important requirement for the district heating utility. To meet these demands, the production planning needs to know the heating demand up to months ahead, e.g. scheduling biomass purchases to be able to deliver the required demand (long-term) [26]. While the temperature control of the network tries to meet these demands by regulating the supply temperature, $T_{s,t}$ and the mass flow of the water is indirectly varied to satisfy the demand, $\dot{m}_t$ as the heat energy is computed as the temperature difference times the mass flow and the specific heat constant,

$$Q_t = \dot{m}_t c_w (T_{s,t} - T_{r,t}), \quad c_w = 4.186 \text{J kg}^{-1}{}^{\circ}\text{C}^{-1} \tag{4.1}$$

to satisfy the heating demand, Q_t where the objective is to increase savings and reduce heat loss in the system. Thus, the controller needs to know the future demand between one hour ahead to the control horizon of the system which is usually the longest transportation time of the system. For large networks, the transportation time can be up to 24 hours. As the future demand is not known, a prediction of it needs to be available for the operators. The accuracy of the predictions is of high importance because the uncertainty of the forecast heating demand needs also be considered in the production and network operations. Along with the requirement of satisfying the heating demand of the consumers, the production of heat has become more complicated than firing up boilers using oil or natural gas. The shift from fossil fuels to renewable energy sources needs to be taken into account as these energy sources are not always available due to their weather dependency. Therefore, to utilize renewable energy sources, energy sector coupling is needed between the power and district heating market, i.e. smart energy system [27]. This makes accurate heat demand forecast highly valuable for the energy sector. Especially, as district heating plants usually also produce electricity with their CHP plants, and therefore need to plan their power production for the next market day. District heating utilities can also store energy as hot water in large thermal storage tanks when an excess of electricity is available.

Consequently, heat demand forecast is the first aspect that utilities need to have for operating the system efficiently. However, heating demand is an inherently non-linear and non-stationary process. The consumption has a non-linear relationship with the ambient air temperature because of the thermal mass of buildings [22, 5, 11]. The non-stationarity comes from the seasonal variation of the ambient temperature and social behavior of the consumers, i.e. time-varying demand. Other weather components than temperature also demonstrate some effect on the heating demand, it is however not as significant as the temperature and social behavior. Dotzauer [28] suggest a simple forecasting model that has a future insight into the heating demand for two different systems. Modeling the ambient air temperature dependency as a piecewise linear relationship to the heating demand is discussed. The social component was modeled by estimating a daily profile using the residual after having removed the dependency of the temperature from the heating demand.

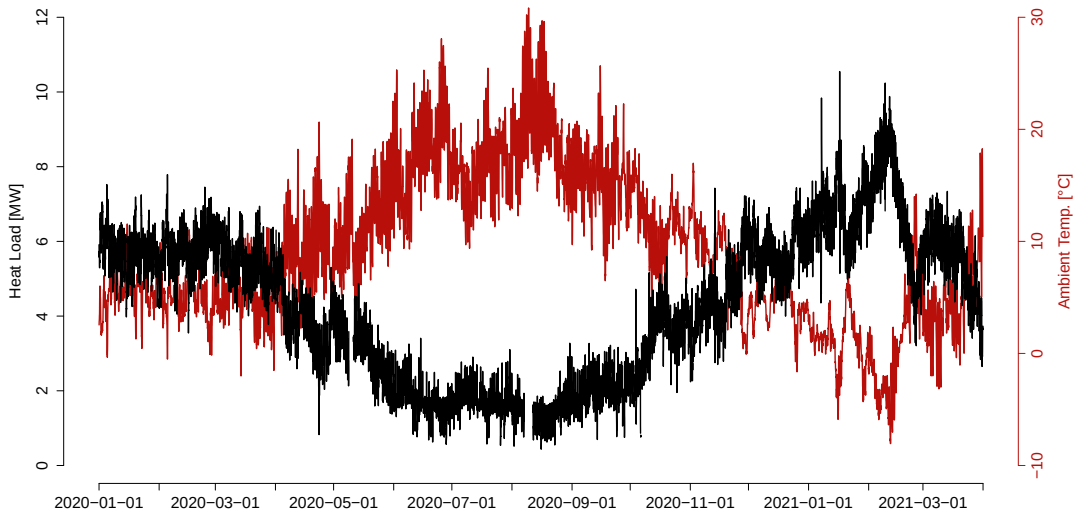

Figure 4.1: Heat load observation from Tingbjerg and ambient air temperature measured from a local climate station in Copenhagen, Station 06186 [21].

The model demonstrates adequate results however temperature measurements are used instead of the numerical weather forecast as input to the model. Also, the model does not handle the time-varying process as the parameters do not change over time as needed for heating demand systems. Nielsen and Madsen [22] suggest using adaptive methods to change parameters during the transition periods, i.e. from cold to warm season. They also propose to use on-line numerical weather predictions as inputs to the model and how to handle the nonlinear dependencies to be used in linear regression models. Their proposal of using grey-box modeling to describe the known physical relationship between heating demand and weather has high accuracy and has proven quite successful. Trying to use new sophisticated models to identify known physical relationships is time- and computer-demanding and is therefore undesirable.

In this project, ENFOR delivered the heat load forecast to be used for the temperature control in Tingbjerg. In the following, we shall describe the climate variables that influence the heat consumption, the consumption in Tingbjerg and compare the localized heat demand forecast from ENFOR to the scaled forecast that HOFOR used previously.

4.1 Data Exploration

Figure 4.1 shows the observed heat demand from Tingbjerg and the measured ambient air temperature at the local climate station in Copenhagen. The data from the climate station was extracted from DMI Open Data [21]. As we mentioned before, heat demand is a non-stationary process due to the time-varying demand following the climate and social behavior. The figures show that the heating demand follows the ambient temperature closely. During the heating season, the ambient temperature has a slow influence on the heating demand, and this influence is usually modeled using a low-pass filter [22]. As the temperature decreases or increases, the heating demand follows with a negative correlation as the plot demonstrates. In the summer season, when the temperature is above 17°C, there is no need for space heating. In these periods, district heating only needs to fulfill the need for domestic hot water usage, e.g. showering. One of the most difficult periods to predict the heat demand is during the transition periods from winter to summer and vice versa. The transition periods are when the temperature starts to increase or decrease, and at the same point in time, the solar gain starts to change.

In the spring, the solar radiation starts to warm up the buildings and therefore contributing to heating to maintain a comfortable indoor climate [29]. However, it has a complex relationship because the penetrating radiation onto the windows is related to the time-varying orientation of

solar radiation. The orientation changes during the day as the earth rotates around its axis and has yearly dependency as the earth rotates around the sun. The solar gain does not contribute as much to the heat consumption during the summer as the ambient temperature is the main driving contribution and is quite high during warm periods. However, during the fall similar effect appears as the angle changes and the ambient temperature has decreased. Thus, the climate variable, solar radiation influences heating consumption. The wind speed and the direction of the wind also contribute to the heating consumptions during the transition periods with natural ventilation. The weather is also known occasionally to change rather quickly during these periods.

The time-varying relationship between heat demand and ambient temperature is visualized in Figure 4.2 where the plots show the heat demand plotted against the temperature for each month. A reference curve is also plotted in each plot to highlight the difference between months, especially in the transition months. The piecewise linear reference curve was estimated by tuning the knots to fit the overall period in the top plot. When investigating the transition months from cold to warm (the months' March, April, and May) we can see that the demand tends to be more scattered below the reference curve between 0 and 10 degrees. In this period, solar radiation could have started to influence consumption even though the ambient temperature is still quite low. Comparing the spring transition period to fall, the demand in the fall is more constant, i.e. it has less spread around the curve. This could be due to the fact that Copenhagen is a coastal city. The sea temperature is higher in the fall period as it has been warmed up during the summer and it gives a quite constant heat to the city. Figure 4.3 shows the monthly average sea temperature in Copenhagen over 20 years. The plot shows that the fall periods have a higher temperature than spring. The data was extracted from the DMI Open Data platform [30]. The social behavior could also vary over time and that could drive the difference between the transition periods. People probably have a different perspective on ambient temperature when comparing the periods. Also, sea temperature could play a big part in keeping the climate milder in the fall.

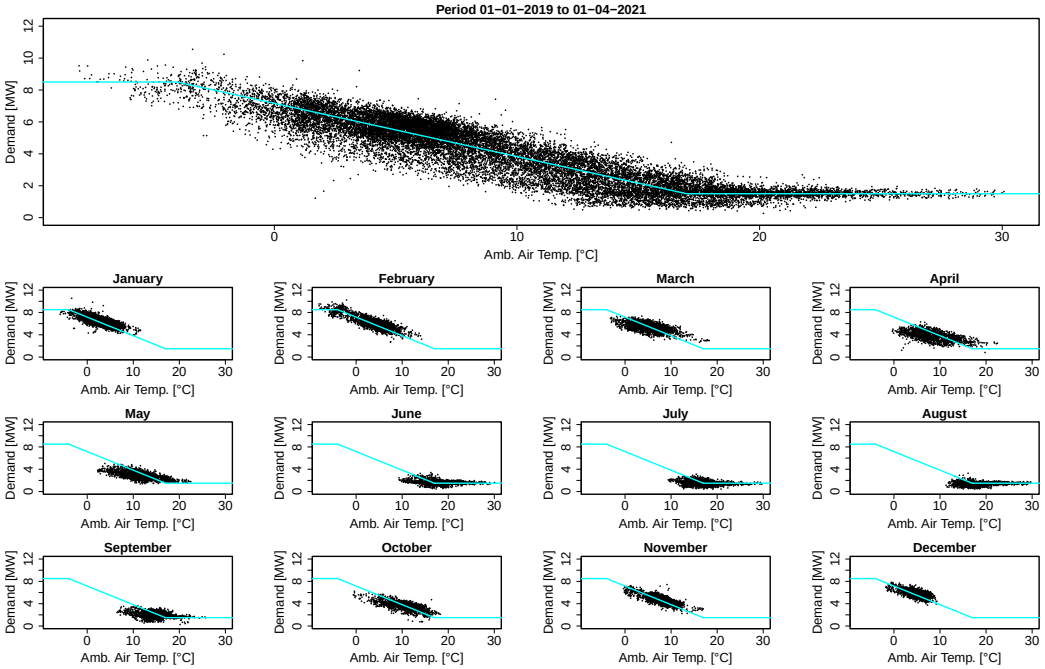

Figure 4.2: Time-varying relationship between heat demand and ambient air temperature demonstrated over different months.

As discussed, solar radiation is an interesting aspect regarding the heating demand during the transition periods. Figure 4.4 shows the mean solar radiation per month, the mean ambient air temperature per month, and the mean heating demand vs the mean solar radiation per month.

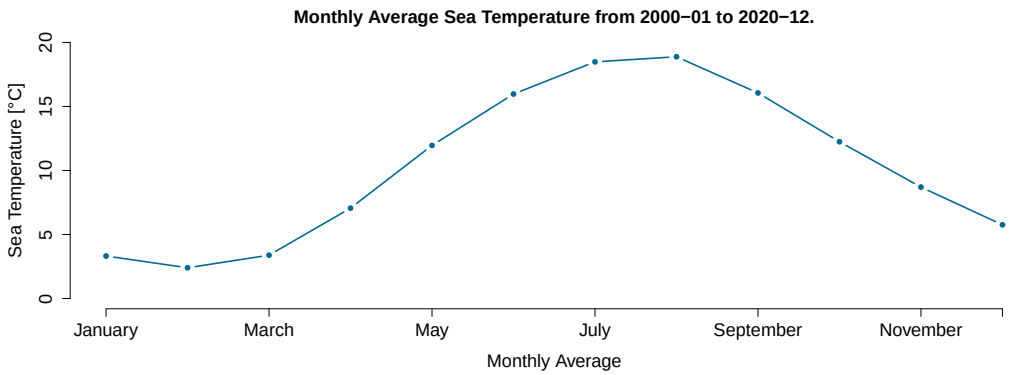

Figure 4.3: Monthly average sea temperature in Copenhagen from January 2000 to December 2020. Station 30336 from DMI Open Data [30]

The data was extracted from the DMI Open Data platform [21]. Investigating these plots together starting with the lowest plot, you would assume for the transition months that the solar radiation has less effect in spring. The spring months have similar solar radiation as the fall months but they require more heat. However, considering the mean ambient temperature it becomes evident that it is quite warm during the fall months than spring in Denmark. This could be explained by the fact that during spring and summer the ground and sea temperature have increased and have slow time variation, i.e. they react to changes slowly. Therefore, in the spring they have yet to be warmed up and during the fall they give away heat to the air temperature as we see in Figure 4.3.

We do not analyze the wind speed effect on the heating demand in-depth. However, it has an effect on the consumption, and largely when it is high and the ambient temperature is low [31, 32].

Considering these climate variables and how they influence heating consumption emphasizes the importance of having an accurate forecast of the heating demand to operate district heating systems. It shows that accurate numerical weather predictions of these variables and a time-varying model, that can react to these changes, is needed. Both the rapid changes in how the climate variables related to the consumption changes over the day and the slow variation in the heating demand can be explained by, for example, the change in social behavior, renovation of houses, new house connected to the network, so on and so forth.

4.2 Localized heat load forecast

In the previous subsection, we demonstrate that heat demand is a non-linear and non-stationary process due to dependencies on climate data and social components. We also learned in Section 3 that NWPs need correction to forecast the climate in specific areas to enhance the accuracy, especially for the short-term forecast. Therefore, we need to localize the NWP to a specific area using climate stations to handle the influence of the buildings, humans, cars, pollution that contribute to the climate. This was done for the heat demand forecast in Tingbjerg. A local climate station in Copenhagen that is located close to Tingbjerg was used to correct the NWPs. After the correction, the NWP was used as input to the heat demand forecast model. The model was therefore both localized to the heating demand in Tingbjerg by estimating the parameters of the model using historical demand from Tingbjerg and the climate as the NWPs are used as input.

In the previous operation in Tingbjerg, the heat demand forecast from a bigger area, Brønshøj, was scaled to match the demand in Tingbjerg. We will demonstrate the accuracy difference in heat demand forecasting by scaling the heating demand versus creating a localized heat demand forecast for Tingbjerg that uses numerical weather predictions that have been corrected to the climate in Tingbjerg. The scaled forecast is scaled by dividing the forecast for the large area by the percentage of Tingbjerg heating demand to the total demand. HOFORs estimation of the fraction

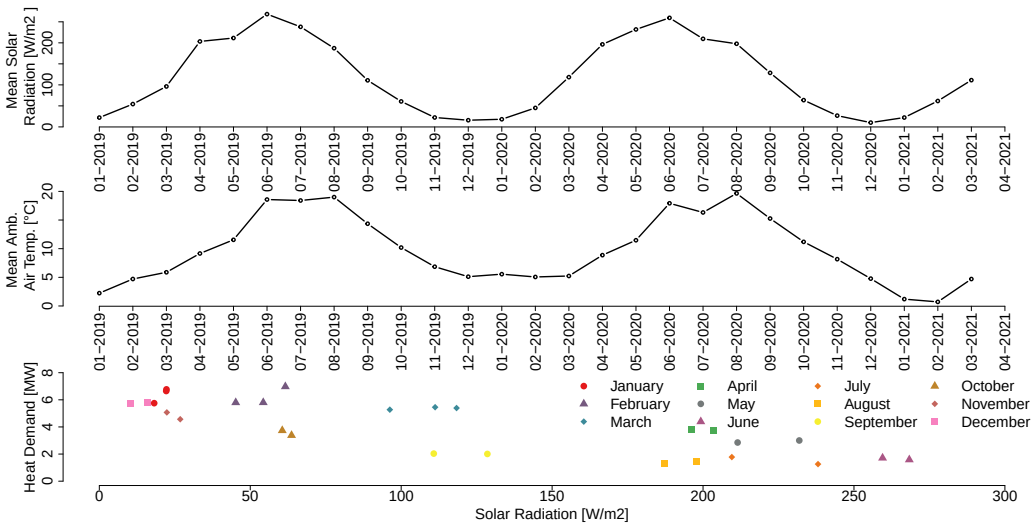

Figure 4.4: Climate data from the climate station in Copenhagen demonstrate the heating demand dependency on temperature and solar radiation. It shows the monthly average of these two climate variables.

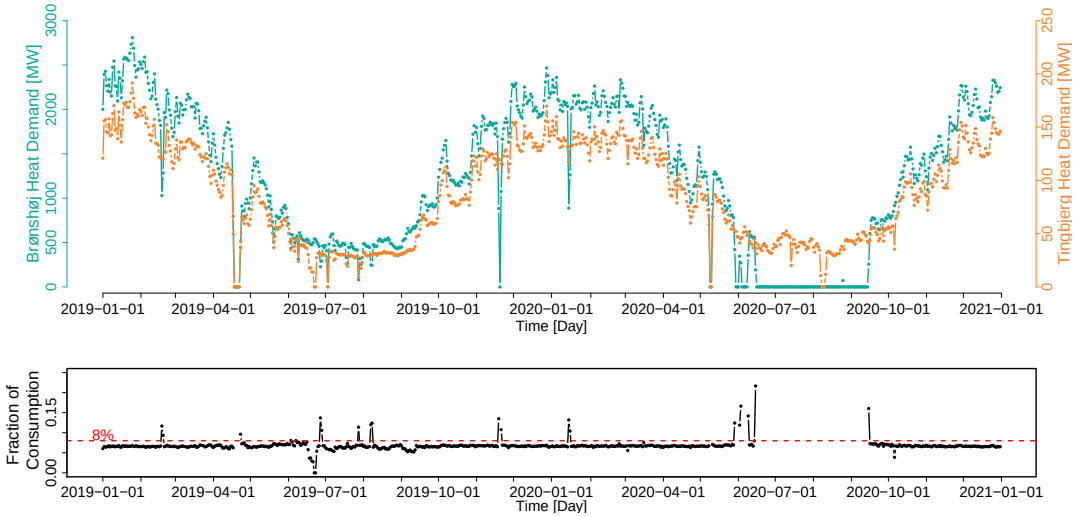

Figure 4.5: Daily heat demand from Brønshøj and Tingbjerg which is a small area inside of Brønshøj. Fraction of Tingbjerg daily consumption compared to Brønshøj consumption is illustrated in the bottom plot.

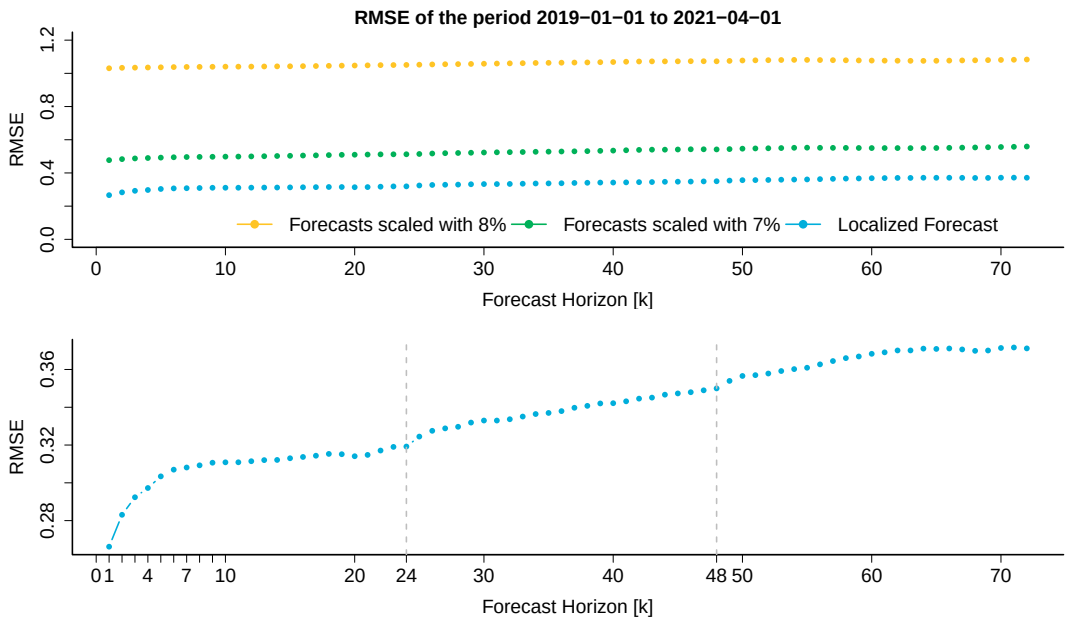

Figure 4.6: Performance of forecasts compared using RMSE for one to 72 steps ahead. The top plot visualize the results and bottom plot shows the localized forecast to demonstrate the affect of having local climate station to improve the short term forecasts, i.e. the forecast horizon between 1 and 8 hours.

of Tingjerg consumption was to be 8% of the total consumption in Brønshøj. Figure 4.5 shows the daily heat demand from Brønshøj and Tingbjerg in the top plot while the bottom plot shows the fraction of the Tingbjerg to the Brønshøj. Based on these data, the 8% that previously used, estimated from older historical data is off by 1% as for the years 2019 and 2020 the fraction of the consumption is around 7%. Notice, that the heat demand in Brønshøj during summer 2020 is zero. This could be the result of HOFOR supplying heat to Brønshøj from other heat sources than usually therefore the measurements are zero in this data-set.

The comparison between the forecasts is demonstrated in Figure 4.6. They are compared using the Root Mean Square Error (RMSE) metric,

$$\mathrm{RMSE}_k = \sqrt{\sum_{t=1}^{T} \frac{y_{t+k} - \hat{y}_{t+k|t}}{T}} \tag{4.2}$$

where the metric is computed for each prediction horizon, k. The heating demand observations are y and predictions are $\hat{y}$. This is computed over the whole period. We have also added the heat demand forecast scaled with the newer update on the fraction between Brønshøj and Tingbjerg demand than was estimated before. The upper plot demonstrates the RMSE over the first 72 steps horizon for all three forecasts, where the localized forecast significantly outperforms the scaled forecasts. This is not surprising as the model has both tuned the parameters to the area and the NWPs have been adjusted to the climate. Notice, the effect of localizing the weather forecast to the climate when comparing the localized forecast to the 7% scaled forecast in the first seven steps ahead. We see that there is a curvature for the localized which is the result of the corrections to the climate and the dynamic process of the forecast model. This is demonstrated in more detail in the lower plot in Figure 4.6 where the curvature is seen more easily. Usually, the RMSE over the horizon demonstrates a straight line with a slope over the horizon however by using the climate information in the area it is possible to enhance the short-term forecasts as demonstrated and using historical heat demand observation for the dynamical process of the model.

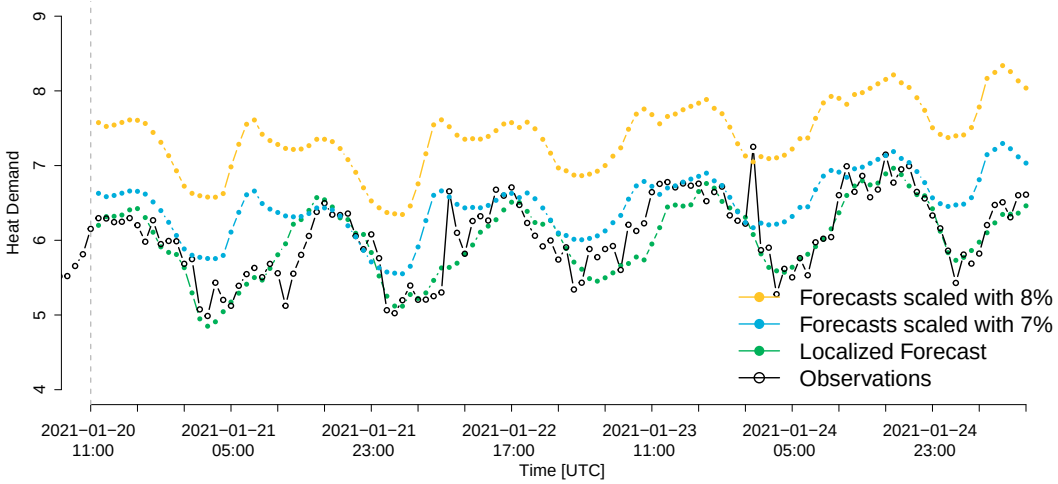

Figure 4.7: Figure demonstrate the performance of the three forecasting models. It shows the 72h steps ahead when generate at 2021-01-20 11:00.

Figure 4.7 shows the three forecasts created at 11:00 on the 2021-01-20 for the next 72 hours. We see that the localized forecast follows the observation significantly better than the other forecasts. Notice, that the scaled forecast does not resemble the diurnal variation in the Tingbjerg observation.

Chapter 5

Temperature Optimization and Control

We have demonstrated that district heating has a difficult task on forecasting the demand because of the local climate and social components of the consumers. An even more difficult task is to provide heat to the consumer without wasting heat and minimizing the cost. The process is a complex procedure as it needs to be delivered through the piping system to the consumer and more specifically it needs to arrive at the correct time to satisfy consumers demand. Hence, temperature control is an essential tool for the efficient operation of district heating. The production of heat has become more challenging as we move away from fossil fuels towards renewable energies. The goal of temperature control is to reduce heating production costs and heat losses in the network and at the same time, fulfill the requirements of the network and consumers. Nielsen [8] describes that the optimal operation of the district heating system is to be achieved by minimizing the production cost without violating these restrictions;

- A maximum allowable flow rate in the system
- A minimum guaranteed inlet temperature at the consumers
- A maximum allowable supply temperature
- Limited short term variation in the supply temperature
- Maximum allowable diurnal variations of the supply temperature

These restrictions are required to be satisfied by the temperature control and the controller also needs to reduce the operational costs without violating them. Therefore, it also needs to consider the heat loss in the pipe, the pumping costs, and maintenance costs of the system. As for all operational aspects of district heating, temperature control needs to know the future heating demand in order to minimize the operation costs for the given planning horizon. Thus, depending on the type of the plant, production will need to have 1) heat demand forecast, 2) supply temperature forecast, 3) future optimal scheduling of the productions, e.g. for a CHP plant needs future sales price for power, also heat and power production costs, 4) restriction in the system, e.g. hydraulics, minimum or maximum time-varying heat production 5) flexibility of the system. These factors need to be considered in advance to optimally achieve minimizing the operation costs. In this part of the report, we will only focus on finding the optimal future supply temperature to give the operators of the network, however in Nielsen [8] and Arvastson [4] further readings on how operating district heating in an optimal setting are given while considering the whole district heating system.

Benonysson et al. [3] formulate a mathematical model of a district heating system for estimating the optimal supply temperature. The model includes the production, the network, and the consumers where the objective is to minimize the operational cost while satisfying all requirements of the

system. They describe a number of items that need to be considered when modeling the dynamics of a district heating network. We have summarized them here below:

- *Time Delay:* The transportation time of the DH water from the production plant to the consumer. The transportation time varies for the individual consumer according to the distance from the plant and the flow velocity in the pipes. The heat capacity of the DH pipes also affects the time delays.
- *Heat loss:* The heat loss is approximately proportional to the difference between the temperature of the DH water and the surrounding temperature. Ground temperature varies over time, following the ambient air temperature with a slow change in undisturbed ground, i.e. not heated by the DH pipes.
- *Friction loss:* When the pumping energy transforms to heat energy due to the friction loss in the pipes. Most often negligible but when the flow velocity of the water is relatively high (extreme cases), the produced amount of heat can be of the same size as the heat loss to the surrounding ground.
- *Pressure:* The flow in the system changes spread in DH networks around 1000 times faster than temperature changes. This leads to to the fact that the dynamics of the flow in the network are of minor importance compared with the dynamics of the temperature changes, from an operational optimization point of view.

The dynamics of the network are therefore highly important to understand to be able to utilize these physical facts to enhance the operation of the network. They give the opportunity of keeping the supply temperature as low as possible when modeling them adequately along with accurate modeling of the consumers' dynamics, i.e. accurate heat demand predictions.

Initial controllers to operate the temperature optimally had a reference curve, i.e. a control schema that changes the set-point of the supply temperature for a given outside temperature. This is a good restriction on the supply temperature, it is however a naive control strategy. Firstly, it does not consider the time for the heat to reach the consumer, the transportation time. Secondly, the flow is usually kept at a low rate, thus the potential of keeping the flow close to the maximum hydraulic constraint is dismissed, i.e. higher flow results in lower supply temperature. Therefore, when the heat finally reaches the consumer, the outside temperature could have changed and the supply temperature is then either too high or too low than required. Also, as the reference curve only considers the ambient temperature as the meteorological factor that influences heat consumption, it must necessarily take into account the worst possible condition with respect to other meteorological factors [33]. This schema also frequently does not consider the social behavior of the consumer, the diurnal variation. Madsen et al. [5] propose a control schema that utilizes the relationship of supply temperature and flow in Eq. 4.1 and consider the dynamics of the network and social behavior of the consumer to increase savings potential. They predict the heating demand using historical heating demand, and the response of the network to change the supply temperature and keep the flow high to have the supply temperature as low as possible without violating any of the restrictions mentioned before. The response of the network is usually done by having measurements wells within the network, located where the lowest temperature is believed to be, a critical point. Thus, this critical point is used to give feedback on how the network is responding to changes at the production, makes it possible to estimate the time delay and heat loss in the network for given supply temperature and flow.

Madsen et al. [5] suggest controllers for the district heating to have an overall controller, supply temperature sub-controller, and a flow sub-controller to estimate the future supply temperature at the production. For each critical point, a sub-controller is developed to compute the lowest supply temperature from the plant satisfying the reference curve. Flow controller considers the variation of the heat demand, by varying either the mass flow or supply temperature. It utilizes the possibility of keeping the flow high as possible as the objective is to maintain as low a supply temperature as possible without violating the maximum flow constraint due to the hydraulic properties of the pipe

network. However, as the flow approaches the limit, the flow controller will start to increase its supply temperature predictions to meet the forecast heating demand. These controllers regulate the supply temperature and flow without violating any restriction with a certain probability, e.g. the supply temperature at the consumers needs to be above the reference curve 99% of the time. The overall controller at the productions then uses the set-points from the sub-controllers to select the highest supply temperature to generate the heat for the consumer.

In this report, a temperature controller from ENFOR was used, HeatTO™[1]. It is based on the methodology mentioned previously in this section, i.e. based on the articles and projects that have been discussed [8, 5]. Furthermore, in this project, HeatTO™ was extended to receive feedback from smart meters in large apartment buildings, instead of the normal measurements wells in the network. In Section 2.1 the smart-meter data and how they were selected to be used as feedback are described. Feedback from smart meters makes temperature control more desirable to use, especially for district heating networks that don't have measurement wells or are poorly calibrated. They can instead use feedback from smart meters to control the supply temperature at the production. Another advantage of using smart meters is the ability to change the location of critical points. District heating network dynamics changes over time due to many reasons, e.g. new areas are built that are further out in the system, old pipes are replaced with new, and old pipes getting older, i.e. the characteristics of the network vary over time. As the network dynamics vary, the critical point changes also, and therefore the critical feedback needs to be moved to give accurate feedback of the system to satisfy the requirements of the consumers. This is solved by selecting different smart meters to give feedback on the network. Bergsteinsson et al. [13] suggest how to use a group of single-family house smart-meters to establish an estimate of the temperature in the street temperature to be used as feedback. However, in this project, the area used has many large apartment buildings and it was demonstrated in Section 2.1 that some of the smart meters can be used to give feedback. Three meters were chosen to be used as feedback for the trial.

5.1 Trial at Tingbjerg

The demonstration of using additional data for the temperature control trial started on 1 November 2020 and lasted until 1 April 2021. The main task of the trial is to demonstrate that smart meters can be used as feedback of the network for temperature control. The trial was done in the Tingbjerg area using the HOFOR district heating system where the production unit is a heat exchanger that supplies the area with heat. Here the focus was on how district heating utilities can create value from digitalization and use it for control of the network. Three smart meters were selected to be used as feedback in Tingbjerg to control the distribution supply temperature at the heat exchanger.

Prior to the trial in Tingbjerg, an open-loop control was used to vary the supply temperature. It was operated by a hydraulic simulation of the system and using a reference curve at the consumers to estimate the optimal set-points for supply temperature. However, as there was no feedback of the network on how it reacted to changes or if the consumers were receiving what they are promised, i.e. the control was open-loop. It also used scaled heat demand forecast as was shown in Section 4 as input. The new controller in the trial is a feedback controller using data from the three smart meters, a reference curve at the location of each smart meter, localized heat demand forecast, and a flow controller to achieve the optimal future supply temperature to reduce the operational cost of the system by lowering the temperature.

Figure 5.1 illustrates the performance at the three smart meters, referenced as netpoints, for the old and new controller. The supply temperature for both operations is plotted versus the rolling average of the outside temperature for the past 24 hours. This is to stabilize the time series, smoothing out any small outliers. The grey solid lines show the reference curve that was used for the new controller and how it was believed to be for the old controller at the netpoints as it did not have feedback of the network before. All three plots demonstrate that the supply temperature from the new controller has less spread and rarely violates the desire reference curve at the netpoints thus the

[1] https://enfor.dk/services/heatto/

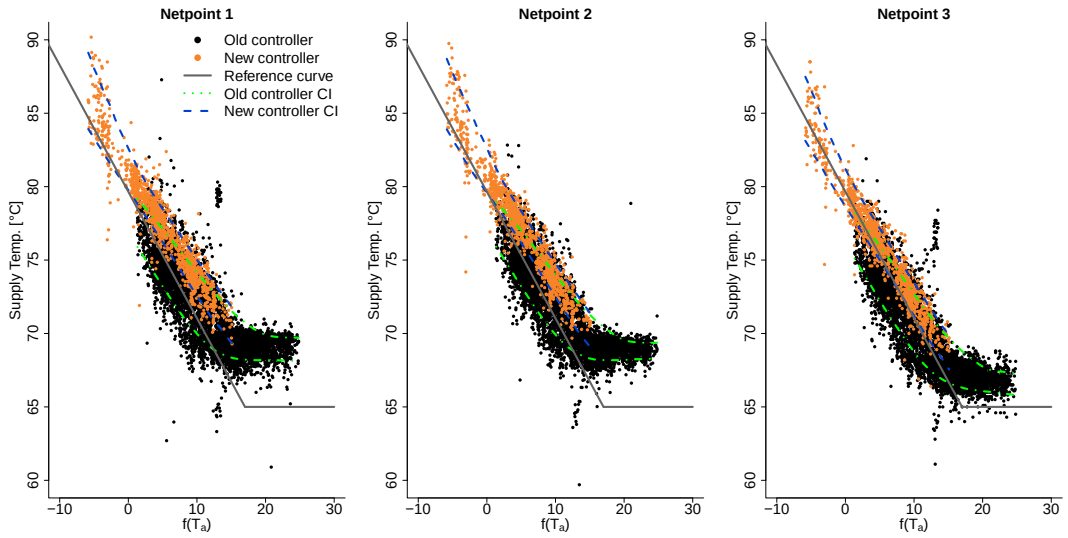

Figure 5.1: Figure compares the two controller performance at the netpoints. It also demonstrates the reference curve used and the estimated confidence interval of the supply temperature for both controllers

system was controlled with more precision. Hence, the new controller gives a better level of security than what was possible with the previous controller as the previous operation violates it rather frequently. This is visualized in more detail by comparing the confidence intervals (CI) between the two operations in the plots. The intervals were estimated using nonparametric quantile regression and using the 10th and 90th quantiles as the upper and lower bounds. These results indicate that the new controller can not compete with the old control of lowering the supply temperature as it can not violate the requirements as frequently as in the previous operation. This suggests that the reference curve for the new controller could have been lower, resulting therefore in a lower and more stable supply temperature. We notice that the reference curve can be adjusted 5°C lower when comparing supply temperature close to the reference curve for the new controller to the low group points at 10°C for the outside temperature. This suggests that the reference curve can be adjusted without breaking any requirements when comparing to the previous operation. The supply temperature will also be adjusted to investigate "what if" scenario when a more reasonable reference curve had been used. The suggested reference curve to be used with the new controller is demonstrated in Figure 5.2. The supply temperature for the trial has also been adjusted with the 5°C. By having more suitable reference curve for the new operation, it could have resulted in lower supply temperature without violating any constraints and thereby lowering the operational cost.

Figure 5.3 compares the old and the new controller operation at the heat exchanger. They were not in operation at the same time therefore the months the new controller was varying the temperature were also selected for the old controller, just one year earlier. The left plots are demonstrating the stability of the controllers, i.e. the variation in time of the supply temperature. This is important for the network, as large and frequent fluctuations in the supply temperature should be avoided as it increases the maintenance costs compared to more stable operation [34]. Comparing the old and new controller, it is noticeable for the period of the new controller, that the outside temperature is changing more dramatically. It is has a long period of very cold temperature and quite warm also. However, for the old controller, the outside temperature is more stable, there are never significant changes between warm or cold periods. Investigating the supply temperature in these two different periods it is evident that the new controller gives a more stable operation, i.e. fewer large fluctuating supply temperatures. This is exhibited in the bottom left plot where the difference series (subtracting past value with current value) of the supply temperatures are shown. The new controller difference series is more stable and the plot also show the variance of the differences where the old controller has around $\sigma^2_{old} = 0.633$ and new, $\sigma^2_{new} = 0.345$. Hence, with the new controller, there is less strain to the pipe system due to the slow variation in the supply temperature.

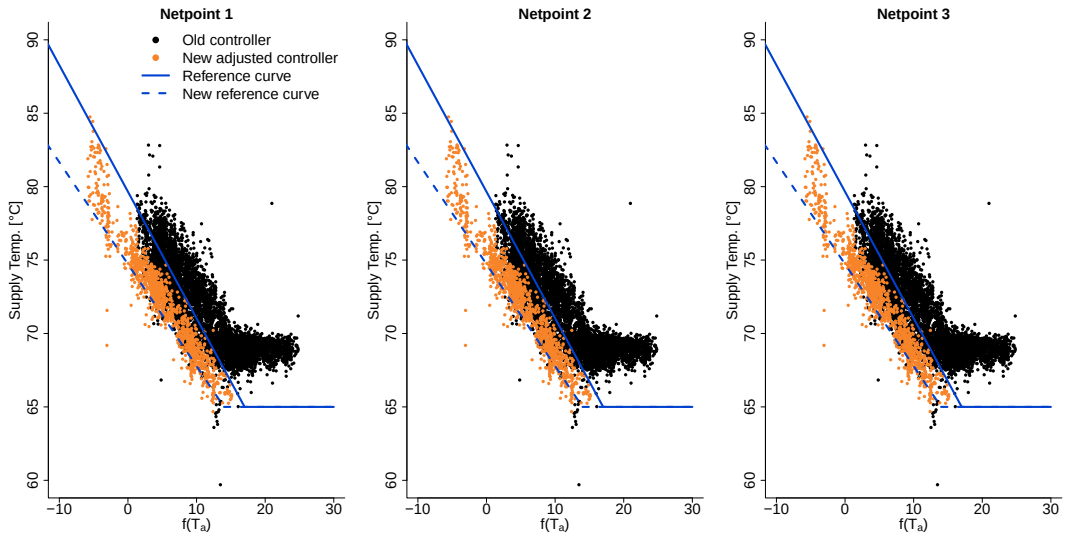

Figure 5.2: Figures shows the new adjusted reference curve and the supply temperature from the new controller periods using the 5°C adjustment.

The plots to the right in Figure 5.3 compare the performance of the controllers versus the degree days. Degree days are used to compare supply temperature between heating seasons when comparing different operations. The degree days, T^{dd} are computed by estimating the difference between the average ambient temperature, $\bar{T}_a$ an over one day, and using 17°C as the cut-off of heating demand from buildings,

$$T^{dd} = \max(0, 17 - \bar{T}_a) \tag{5.1}$$

The average supply and return temperature for each day is then computed and plotted against its corresponding degree day as shown in the top and bottom right plots. The top plots demonstrate the supply temperature performance of the controllers, and the adjusted supply temperature for the new controller as suggested before. We see that the new controller has quite a stable but higher supply temperature than the previous operation as expected because of the high reference curve as shown in Figure 5.1. Consequently, the new controller results in higher supply temperature at the production and thereby higher operation cost. However, it is not significantly higher than the previous operation even though it was penalized by higher restrictions. We see notably improved operation when the adjusted supply temperature is investigated, where the supply temperature is lower.

To compare these operations, a regression model using Ordinary Least Squares to estimate the parameters of a model with an intercept and slope have been fitted to each operation as shown in Figure 5.3 and rewritten here below,

$$\text{New controller: } T_{supply} = 68.48 + 0.71T^{dd} \tag{5.2}$$

$$\text{Old controller: } T_{supply} = 62.61 + 0.97T^{dd} \tag{5.3}$$

$$\text{New adjusted controller: } T_{supply} = 63.48 + 0.71T^{dd} \tag{5.4}$$

Hence, that the new controller has a lower slope, which indicates that it does not increase the supply temperature as fast when the outside temperature decreases compared to the previous operation. However, the intercept is quite higher, which translates to that the overall expectation of supply temperature is higher for the given degree day. The adjusted regression lines have a lower intercept as we have adjusted the data by 5°C however it still has a higher intercept than the previous operation.

Decreasing the supply temperature for operation leads to an increase in savings for the utility. Madsen et al. [33] suggest a rule of thumb for savings resulting from lowering the supply temperature

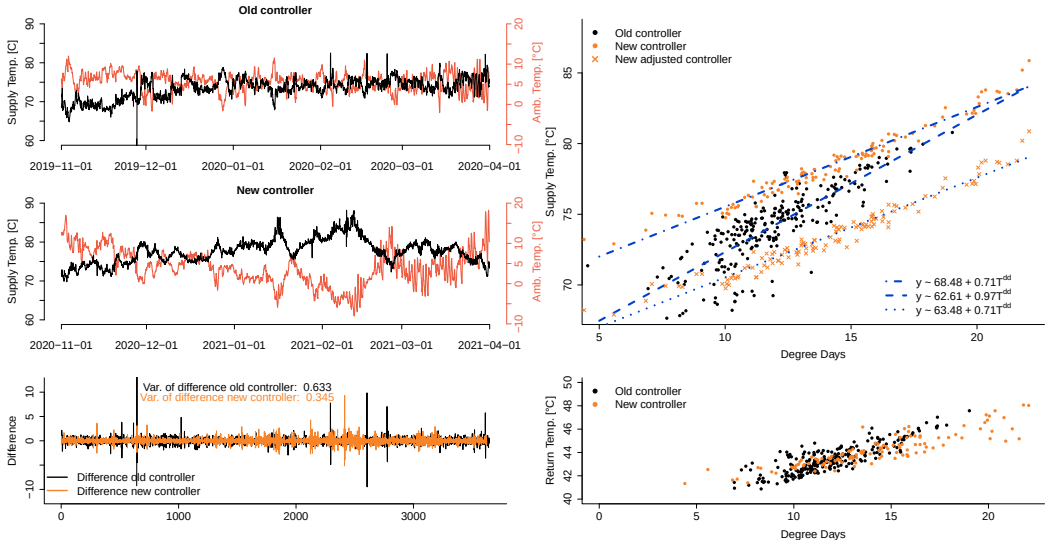

Figure 5.3: Comparing old and new controller at Tingbjerg. The left plots demonstrate the stability of the supply temperatures while the right plots compare the supply temperatures against degree days.

in CHP plant; For each degree lowered, the savings for the heat loss in the network is 0.5 % and at the production for more efficient production is 1 %, thus the savings can be compute as

$$\text{Savings } = (\text{Cost}_{\text{before}} * x\ [^{\circ}C] * 0.5\%) + \text{Shares}_{Production}(\text{Cost}_{\text{before}} * x\ [^{\circ}C] * 1\% \tag{5.5}$$

where x is the lowered supply temperature for the system. Thus, estimating from Figure 5.3 with the new adjusted controller, the supply temperature would be 3°C. The savings would be around 4.5% for the operation of the district heating network in Tingbjerg. The rule of thumb and how to compute the savings are heavily dependent on the system and how the heat is produced. Lowering the supply temperature for a CHP plant gives the highest savings for district heating. Decreasing the supply temperature results in an increase in the ratio of the power to heat output for CHP plant, and as electricity is more valuable than heat, a more profitable operation is achieved [5]. Thus, the equations are just a rule of thumb to demonstrate potential savings when sufficient production data is not available.

Concluding this section, we have demonstrated that having feedback of the network improves the stability of the system, i.e. few and smaller fluctuations in the supply temperature. We saw that the restrictions of the new operation were probably too high compared to the previous as it was allowed to violate the restrictions while the new controller tries to satisfy them within a certain probability. However, adjusting the result with the new controller, we could demonstrate a potential of 4.5% in savings.

Chapter 6

Conclusion

In this report, we have demonstrated how digitalization in district heating can improve the operation of the utility. Smart meters in the network can act as feedback of the system for temperature control to give response characteristics of the network. Accurate modeling of how the system will respond to changes at the production will increase the performance of temperature control and lowering the operation cost. The climate in cities was discussed and the climate in Copenhagen was used to visualize the temperature differences, i.e. different climates within the city. Therefore, a local climate station was used to bind weather forecasts to the area's climate, resulting in more accurate NWPs for the climate area. The NWPs improvements will result in a more accurate heat demand forecast, especially for the short-term forecast. An accurate short-term heat demand forecast is beneficial for the temperature control to deliver the desired consumer consumption and lowering the operational cost by intelligent control in the next hour.

A trial was conducted to demonstrate these benefits. The focus was to improve the heat demand forecast and temperature control of the area using additional data than typically is used. We illustrated the accuracy improvements of localizing heat demand forecast by using the area's historical demand and binding the NWP to the area using a local climate station. The local heat demand forecast was compared to the previous heat demand forecast in Tingbjerg where it was scaled from a forecast from a large area that contains the Tingbjerg area. This highlighted how crucial it is to localize heat demand forecast to an area where the temperature control is operating. The heat demand forecast is used for all operations of the utility, hence the desire of increasing the accuracy of the forecast that will improve the efficiency of the operations. Temperature control was in operation on-line during the trial to demonstrate that smart meters can be used as feedback for closed-loop control. The previous operation in the area was done with open-loop control using current ambient air temperature and hydraulic simulation of the system to operate the network. Unfortunately, the reference curve was placed quite conservatively in order to avoid complaints during the trial. Therefore, the trial operation demonstrated a higher supply temperature than the previous operation. However, we showed that the closed-loop control usually satisfied the reference control requirement at the consumer while the open-loop violated it frequently. The closed-loop operation also demonstrates that it results in higher precision of the supply temperature, e.g. it does not vary as much for the given ambient temperature compared to the open-loop operation. To summarize, the proposed data-driven methods lead to higher precision and that it has the potential of lowering the supply temperature by 5°C, i.e. savings potential of more optimal operation of the network.

In the report, all of these findings are demonstrated and highlighted, the importance of an accurate understanding of the area where heat is delivered. Three things are needed to be considered for efficient operation: 1) The local climate 2) The local social consumption behavior 3) The local response characteristics of the network. Accurate representation of the heating in the area and the local climate can be achieved by using data that has become available through the digitalization in district heating. We therefore conclude that digitalization in district heating will highly benefit the operation of district heating.

Bibliography

[1] D. F. Dominković, R. G. Junker, K. B. Lindberg, H. Madsen, Implementing flexibility into energy planning models: Soft-linking of a high-level energy planning model and a short-term operational model, Applied Energy 260 (2020) 114292. URL: https://www.sciencedirect.com/science/article/pii/S0306261919319798. doi:https://doi.org/10.1016/j.apenergy.2019.114292.

[2] A. Vandermeulen, B. van der Heijde, L. Helsen, Controlling district heating and cooling networks to unlock flexibility: A review, Energy 151 (2018) 103–115. URL: https://www.sciencedirect.com/science/article/pii/S0360544218304328. doi:https://doi.org/10.1016/j.energy.2018.03.034.

[3] A. Benonysson, B. Bøhm, H. F. Ravn, Operational optimization in a district heating system, Energy Conversion and Management 36 (1995) 297–314. URL: https://www.sciencedirect.com/science/article/pii/019689049598895T. doi:https://doi.org/10.1016/0196-8904(95)98895-T.

[4] L. Arvastson, Stochastic Modeling and Operational Optimization in District Heating Systems, Ph.D. thesis, Lund University, 2001.

[5] H. Madsen, K. Sejling, H. T. Søgaard, O. P. Palsson, On flow and supply temperature control in district heating systems, Heat Recovery Systems and CHP 14 (1994) 613 – 620. doi:https://doi.org/10.1016/0890-4332(94)90031-0.

[6] S. Grosswindhager, A. Voigt, M. Kozek, Predictive control of district heating network using fuzzy dmc, in: 2012 Proceedings of International Conference on Modelling, Identification and Control, 2012, pp. 241–246.

[7] H. Madsen, H. Nielsen, T. Nielsen, H. Søgaard, Control of Supply Temperature: EFP 1323/93-07, Institut for Matematisk Statistik og Operationsanalyse, 1996.

[8] T. S. Nielsen, Online prediction and control in nonlinear stochastic systems, Ph.D. thesis, Informatics and Mathematical Modelling, Technical University of Denmark, DTU, Richard Petersens Plads, Building 321, DK-2800 Kgs. Lyngby, 2002. URL: http://www2.compute.dtu.dk/pubdb/pubs/792-full.html.

[9] Directive 2012/27/eu of the european parliament and of the council of 25 october 2012 on energy efficiency, https://eur-lex.europa.eu/eli/dir/2012/27/oj, 2012. Accessed on 3 May 2021.

[10] M. H. Kristensen, S. Petersen, District heating energy efficiency of danish building typologies, Energy and Buildings 231 (2021) 110602. URL: https://www.sciencedirect.com/science/article/pii/S0378778820333880. doi:https://doi.org/10.1016/j.enbuild.2020.110602.

[11] C. Thilker, H. Bergsteinsson, P. Bacher, H. Madsen, D. Cali, R. Junker, Non-linear model predictive control for smart heating of buildings, in: Proceedings of Cold Climate HVAC and Energy 2021, 2021. URL: http://hvac2021.org, cold Climate HVAC amp; Energy 2021 ; Conference date: 20-04-2021 Through 21-04-2021.

[12] P. Bacher, P. de Saint-Aubain, L. Christiansen, H. Madsen, Non-parametric method for separating domestic hot water heating spikes and space heating, Energy and Buildings 130 (2016) 107–112. doi:10.1016/j.enbuild.2016.08.037.

[13] H. G. Bergsteinsson, T. S. Nielsen, J. K. Møller, S. B. Amer, D. F. Dominković, H. Madsen, Use of smart meters as feedback for district heating temperature control (2021). Submitted to the 17th International Symposium on District Heating and Cooling, DHC2021, 6–9 September 2021, Hamburg, Germany.

[14] G. Steeneveld, S. Koopmans, B. Heusinkveld, L. Van Hove, A. Holtslag, Quantifying urban heat island effects and human comfort for cities of variable size and urban morphology in the netherlands, Journal of Geophysical Research Atmospheres 116 (2011). URL: https://www.scopus.com/inward/record.uri?eid=2-s2.0-80455145345&doi=10.1029%2f2011JD015988&partnerID=40&md5=223f9ef2aaac5fc0473abb8e3db5e55e. doi:10.1029/2011JD015988, cited By 146.

[15] M. C. Moreno-garcia, Intensity and form of the urban heat island in barcelona, International Journal of Climatology 14 (1994) 705–710. URL: https://rmets.onlinelibrary.wiley.com/doi/abs/10.1002/joc.3370140609. doi:https://doi.org/10.1002/joc.3370140609. arXiv:https://rmets.onlinelibrary.wiley.com/doi/pdf/10.1002/joc.3370140609.

[16] W. D. Solecki, C. Rosenzweig, L. Parshall, G. Pope, M. Clark, J. Cox, M. Wiencke, Mitigation of the heat island effect in urban new jersey, Global Environmental Change Part B: Environmental Hazards 6 (2005) 39–49. URL: https://doi.org/10.1016/j.hazards.2004.12.002. doi:10.1016/j.hazards.2004.12.002. arXiv:https://doi.org/10.1016/j.hazards.2004.12.002.

[17] K. Hibbard, F. Hoffman, D. Huntzinger, T. West, Ch. 10: Changes in Land Cover and Terrestrial Biogeochemistry. Climate Science Special Report: Fourth National Climate Assessment, Volume I, Technical Report, 2017. URL: https://doi.org/10.7930/j0416v6x. doi:10.7930/j0416v6x, Referenced at: https://www.epa.gov/heatislands/learn-about-heat-islands#_ftn1, Accessed on 10.05.2021.

[18] H.-Y. Lee, An application of noaa avhrr thermal data to the study of urban heat islands, Atmospheric Environment. Part B. Urban Atmosphere 27 (1993) 1–13. URL: https://www.sciencedirect.com/science/article/pii/0957127293900414. doi:https://doi.org/10.1016/0957-1272(93)90041-4.

[19] C. J. Morris, Urban heat islands and climate change – melbourne, australia, 2006. URL: https://web.archive.org/web/20090310021108/http://www.earthsci.unimelb.edu.au/~jon/WWW/uhi-melb.html, accessed: 10.05.2021.

[20] T. R. Oke, The energetic basis of the urban heat island, Quarterly Journal of the Royal Meteorological Society 108 (1982) 1–24. URL: https://rmets.onlinelibrary.wiley.com/doi/abs/10.1002/qj.49710845502. doi:https://doi.org/10.1002/qj.49710845502. arXiv:https://rmets.onlinelibrary.wiley.com/doi/pdf/10.1002/qj.49710845502.

[21] Danish Meteorological Institute (DMI), Meteorological Observation (metObs), https://confluence.govcloud.dk/display/FDAPI/Danish+Meteorological+Institute+-+Open+Data, 2021. Accessed on 24.05.2021.

[22] H. Nielsen, H. Madsen, Predicting the Heat Consumption in District Heating Systems using Meteorological Forecasts, Informatics and Mathematical Modelling, Technical University of Denmark, DTU, 2000.

[23] H. Madsen, H. Søgaard, K. Sejling, O. Palsson, Models and Methods for Optimization of District Heating Systems.: Part I: Models and Identification Methods, Institut for Matematisk Statistik og Operationsanalyse, 1990.

[24] H. R. Glahn, D. A. Lowry, The use of model output statistics (mos) in objective weather forecasting, Journal of Applied Meteorology (1962-1982) 11 (1972) 1203–1211. URL: http://www.jstor.org/stable/26176961.

[25] P. Crochet, Adaptive kalman filtering of 2-metre temperature and 10-metre wind-speed forecasts in iceland, Meteorological Applications 11 (2004) 173–187. URL: https://rmets.onlinelibrary.wiley.com/doi/abs/10.

1017/S1350482704001252. doi:https://doi.org/10.1017/S1350482704001252. arXiv:https://rmets.onlinelibrary.wiley.com/doi/pdf/10.1017/S1350482704001252.

[26] D. Guericke, I. Blanco, J. M. Morales, H. Madsen, A two-phase stochastic programming approach to biomass supply planning for combined heat and power plants, Or Spectrum 42 (2020) 863–900.

[27] H. Lund, S. Werner, R. Wiltshire, S. Svendsen, J. E. Thorsen, F. Hvelplund, B. V. Mathiesen, 4th generation district heating (4gdh): Integrating smart thermal grids into future sustainable energy systems, Energy 68 (2014) 1–11. URL: https://www.sciencedirect.com/science/article/pii/S0360544214002369. doi:https://doi.org/10.1016/j.energy.2014.02.089.

[28] E. Dotzauer, Simple model for prediction of loads in district-heating systems, Applied Energy 73 (2002) 277–284. URL: https://www.sciencedirect.com/science/article/pii/S0306261902000788. doi:https://doi.org/10.1016/S0306-2619(02)00078-8.

[29] H. Madsen, J. Holst, Estimation of continuous-time models for the heat dynamics of a building, Energy and Buildings 22 (1995) 67–79. URL: https://www.sciencedirect.com/science/article/pii/037877889400904X. doi:https://doi.org/10.1016/0378-7788(94)00904-X.

[30] Danish Meteorological Institute (DMI), Oceanographic Observation (oceanObs), https://confluence.govcloud.dk/display/FDAPI/Danish+Meteorological+Institute+-+Open+Data, 2021. Accessed on 24.05.2021.

[31] K. Wojdyga, An influence of weather conditions on heat demand in district heating systems, Energy and Buildings 40 (2008) 2009–2014. URL: https://www.sciencedirect.com/science/article/pii/S0378778808001138. doi:https://doi.org/10.1016/j.enbuild.2008.05.008.

[32] P. Bacher, H. Madsen, H. Nielsen, B. Perers, Short-term heat load forecasting for single family houses, Energy and Buildings 65 (2013) 101–112. doi:10.1016/j.enbuild.2013.04.022.

[33] H. Madsen, H. Søgaard, K. Sejling, O. Palsson, Models and Methods for Optimization of District Heating Systems.: Part II: Models and Control Methods, Institut for Matematisk Statistik og Operationsanalyse, 1992.

[34] T. Nielsen, H. Madsen, J. Holst, H. Søgaard, Predictive control of supply temperature in district heating systems, 2002.

Appendix A

Temperature Optimization and Control at Svebølle Viskinge Fjernvarmeselskab

For the past couple of years, the district heating utility Svebølle Viskinge Fjernvarmeselskab has been improving its network operation by installing temperature sensors in the network (the critical points) to have feedback of the network. The previous operation of the supply temperature had been selected based on an open-loop system. Thus, without considering how the network response to changes and what supply temperature is received by the consumers. Along with getting feedback, they have also been using the temperature optimization software from ENFOR, the HeatTO™. The same software was installed and used in the Tingbjerg network.

The result of the changes made in the network operation can be seen in Figure A.1. In the figure, heating season is defined as the months; November, December, January, February, and March. The savings gain of using closed-loop temperature optimization can be seen when comparing the previous and new controller in the top plot where the average daily supply temperature against the degree days. The old controller was in operation during the heating season 2018/2019 and the new controller using the feedback and the HeatTO™ software during the heating seasons 2019/2020 and 2020/2021. The old controller was also kept running during the 2020/2021 heating season when the new controller was in operation as shown in the figure. It was only used in computing an alternative setpoints of the supply temperatures, while the new controller was operating the supply temperature in the network. The flow at the production from the previous and current operation is shown in the bottom plot. The plot shows that the new controller has a higher flow. Thus, it increases the flow until it reaches the physical flow maximum of the system before increasing the supply temperature.

Comparing the new operation to the heating season 2018/2019, it can be condluded that the supply temperature has been on average lowered by 8°C therefore the savings can be estimated using Eq. 5.5 to be 12%. However, comparing the new controller to the old controller during the same heat season (2020/2021), the temperature was decrease by 10°C, and the savings are 15%. Hence, investing in data-driven methods can increase the savings of the network.

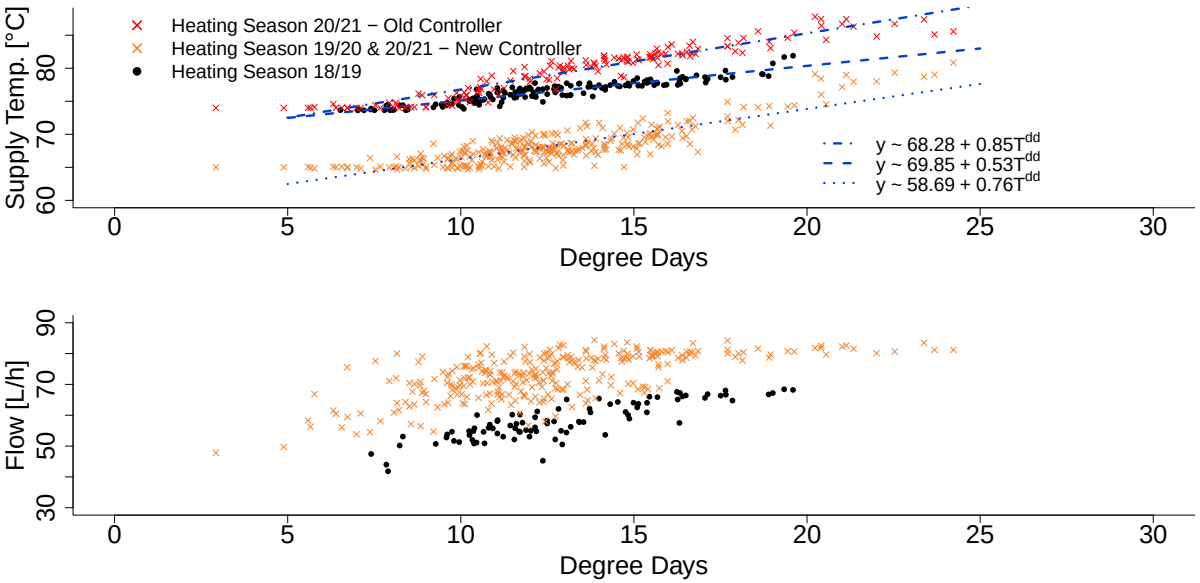

Figure A.1: Temperature optimization at Svebølle Viskinge for three heating seasons; 2018/2019, 2019/2020, and 2020/2021. The heating season months are November, December, January, February, and March. Average daily supply temperature and flow are plotted against the degree days in the top and bottom plots.

Danmarks
Tekniske
Universitet

Richard Petersens Plads, Building 324
2800 Kgs. Lyngby
Tlf. 4525 3031

https://www.compute.dtu.dk/

Paper H

QUANTIFICATION OF HEAT DEMAND FORECAST ACCURACY IMPROVEMENTS BY LOCALIZED WEATHER FORECAST

Authors:
Hjörleifur G. Bergsteinsson, Jan Kloppenborg Møller and Henrik Madsen.

Submitted To:
International Journal of Sustainable Energy.

Quantification of Heat Demand Forecast Accuracy Improvements by Localized Weather Forecast

Hjörleifur G. Bergsteinsson[a], Jan Kloppenborg Møller[a] and Henrik Madsen[a]

[a]Technical University of Denmark , Anker Engelunds Vej 1, Building 101A, 2800 Kongens Lyngby, Denmark

ARTICLE HISTORY
Compiled October 21, 2022

ABSTRACT
Weather forecasts are essential for district heating (DH) utility operations as they prepare the utility for future consumption so that optimal operation is achieved by supplying sufficient heat while keeping costs low. Weather forecasts are usually converted into heat demand forecasts which are used for production planning and control of the temperatures in the DH network. Hence, increasing the accuracy of weather forecasts will lead to improvements in the system's operational performance. However, numerical weather predictions (NWPs) are obtained by complex dynamical equations that are computed over the earth as grid values, and NWPs are designed for rural areas, not urban areas. Thus, for DH applications, weather forecasts do not give an optimal representation of the local climate that drives the local heat demand in a city. In this paper, we propose a method to localize weather forecasts and we demonstrate the benefits by quantifying heat demand forecast accuracy improvements.

KEYWORDS
Localized forecast; recursive estimation; heat demand; weather forecast; NWP;

1. Introduction

The role of district heating is changing dramatically, with more sector coupling and increasing amounts of energy being generated by renewable energy sources (RES). Traditionally, the role of district heating was to provide heat to consumers when needed, and usually by burning fossil fuels. Sometimes, district heating production plants also produce electricity. These are called cogeneration plants. However, with awareness of the global climate crisis due to increasing CO2 concentrations in the air, more climate-friendly operation of energy production is needed. Hence, district heating needs to shift away from fossil fuels towards RES, and due to its flexibility (energy storage), district heating has become a crucial element of efficient operation of the overall integrated energy system. To maximize the potential of district heating, more advanced methods are needed for its operation. Advanced methods for production optimization of plants have been proposed for optimal scheduling of heating units and for bidding into the electricity markets to increase the share of RES and lower overall system costs [1]. Advanced methods for delivering the heat to consumers have also been proposed. For instance, Madsen et al. [2] propose advanced optimization

CONTACT Hjörleifur G. Bergsteinsson Email: hgbe@dtu.dk

methods to control supply temperature and flow in district heating networks in order to reduce heat production costs and heat losses in the network.

Heat demand forecasts are needed for efficient operation of production planning and for temperature control of the network. For some applications, like for optimal use of storage systems, the forecasts are most conveniently provided as scenarios. Hence, accurate future scenarios of heat consumption are desirable to lower the overall cost of operating a district heating system. Heat consumption is directly related to the climate, and therefore heat demand forecast models usually use forecasts of weather variables as inputs. Clearly, the accuracy of heat demand can be increased by improving the weather forecasts. District heating is typically applied in urban areas where the climate is generally different from rural areas due to human activities, buildings, and infrastructures. This effect is termed Urban Heat Island (UHI). Therefore, the operation of district heating needs to account for this local climate to achieve efficient operation. Heating consumption is highly correlated with climate variables, especially the ambient air temperature. Predictions of the weather are therefore desired by district heating companies to ensure that heat production is sufficient to meet future heating consumption and to ensure efficient production. Numerical Weather Prediction (NWP) is used to forecast the weather by simulating physics-based partial differential equations of the atmospheric processes [3]. Most often the NWP does not represent the UHI effect and hence the weather inside cities, and consequently a systematic bias between the NWP and local weather stations is often observed [4]. It would therefore be beneficial to localize the weather forecast to remove the bias and hence reduce the prediction uncertainty of the weather forecasts for district heating applications.

Several studies have investigated the possibility of increasing the forecasting capabilities of NWP by including the UHI effects in the model. See e.g. Ronda et al. [5] for their suggestion of incorporating UHI in NWP for the city of Amsterdam.

2. Method

Heat demand forecast models have been studied extensively for the past decades and both grey-box and black-box models have been proposed. Black-box models are methods where the modeller feeds data through the model, trains or calibrates the model parameters using these data, and finally produces forecasts without considering any physical knowledge of the system. As an example, Idowu et al. [6] propose four different methods to forecast heat demand. Grey-box models are based on a combination of physical knowledge and information embedded in data. The methods use statistical methods to estimate the parameters and reduce the model dimensions compared to a purely physical model of the system. Aalborg and Madsen [7] have discussed a physical-based model of the heat consumption in a district heating network and they suggested a simplified version of the full (complex) model. They demonstrate that a rather simple model, based on the physical nature but calibrated using data on heat consumption, is adequate for forecasting heat demand.

The forecasting model that will be used to produce heat demand forecasts is generated using recursive and adaptive techniques, as implemented in the R package, *onlineforecast* [8]. The proposed forecasting method in this paper uses physical knowledge on how heating consumption and weather are linked as suggested by Aalborg and Madsen [7]. An optimal model was identified using the forward selection approach as shown in Bacher et al. [9], where inputs are added one-by-one to the model, and for each input, it is investigated if the forecasting performance increases significantly. The

Root-Mean-Square-Error (RMSE) is used to compare the forecasting performance for each k-step horizon. Different transformations of the inputs are also investigated. Models for heat demand are usually based on weather forecast inputs (e.g., air temperature and wind) and social components (e.g., Fourier Harmonics). NWPs are created from a system of nonlinear differential equations created from the laws of physics, and are used to predict the weather by describing the physical processes in the atmosphere. These differential equations cannot be solved for each point on the earth, and therefore a spatial grid of the earth is used as points for the equations to be solved at. However, these systems model urban areas with the same formulations as rural areas, hence discarding the Urban Heat Island effect [10].

The raw forecast of the ambient air temperature does therefore not consider the local climate in the area where the heat demand forecast is needed. However, if there is a local climate station that measures the temperature inside the area, then it will be possible to adjust the temperature forecast using such local measurements. Glahn and Lowry [11] proposed the Model Output Statistics method to localize NWPs using a regression model where the dependent variable is the observed weather variable at the local weather station, and the explanatory variable is the forecast weather variable,

$$y = \beta_0 + \beta_1 T_a^{\text{nwp}} + \epsilon. \tag{1}$$

The forecasts of the regression model are then used as inputs for the heat demand forecast model to quantify the accuracy improvement by localizing the NWP. This method provides the possibility to locally adjust the NWP and therefore update it with information on the local climate and hence provide better information for predicting the heat demand. The method also makes it possible to estimate the uncertainty of the NWP for this location.

3. Results

Figure 1. Figure shows the location of the climate station and Tingbjerg, where the district heating is located. The distance between the two places is approximately 6 km. [2]

[2]Figure created using https://maps.google.com/.

The demo case used in this project is a small area, Tingbjerg, in Copenhagen and it mainly consists of apartment buildings. The heat demand data is provided by the district heating utility company, HOFOR, which operates the network. The heat demand is used to establish the forecasting model and validate the results. The NWPs are forecasts obtained using ECFMW and they are updated twice a day; at noon and midnight[3]. The local climate station is located in Copenhagen, and the data is provided by the Open Data platform by the Danish Meteorological Institute (DMI)[4]. Unfortunately, the local climate station is not located at the same place as the district heating network, but it is quite close and also located in the city (Figure 1). It is therefore expected that the use of data from this city climate station will lead to improved forecasts of the local climate. However, it would have been optimal to have the climate station located inside the area where the desired heat demand forecasts are needed.

The optimal model found for forecasting the heat demand was,

$$\begin{aligned}\hat{y}_{\mathrm{t+k|t}} =&\beta_{0,\mathrm{k}} + \beta_{1,\mathrm{k}} y_{\mathrm{t}} + \mu_{\mathrm{k}}(t, n_{\mathrm{har}}, \alpha_{\mathrm{diu}}) + \beta_{3,\mathrm{k}} H(q) T^{\mathrm{obs,nwp}}_{a,t+k|t} + \\ &\beta_{4,\mathrm{k}} H(q) W^{\mathrm{nwp}}_{s,t+k|t} + \beta_{5,\mathrm{k}} H(q) G^{\mathrm{nwp}}_{t+k|t},\end{aligned} \tag{2}$$

where β are the coefficients of the model, $T^{\mathrm{obs,nwp}}_{a,t+k|t}$ is a combined sequence of measured and forecast ambient air temperature (°C) including current measurements and NWP of the ambient air temperature, i.e. $T^{\mathrm{obs,nwp}}_{a,t+k|t} = \{\ldots, T^{\mathrm{obs}}_{a,t-1}, T^{\mathrm{obs}}_{a,t}, T^{\mathrm{nwp}}_{a,t+1|t}, T^{\mathrm{nwp}}_{a,t+2|t}, \ldots, T^{\mathrm{nwp}}_{a,t+K|t}\}$, $W^{\mathrm{nwp}}_{s,t+k|t}$ is the NWP of wind speed (m/s), $G^{\mathrm{nwp}}_{t+k|t}$ is the NWP of global radiation (W/m2), $\mu_{\mathrm{k}}(t, n_{\mathrm{har}}, \alpha_{\mathrm{diu}})$ describes the diurnal curve using Fourier harmonic series where t is the time of the day (hour), n_{har} is the number of harmonics, and finally α_{diu} is a vector consisting of the coefficients for the included harmonics. H(q) is a transfer function that acts as a low-pass filter with a stationary gain equal to one,

$$H(q) = \frac{1-a}{1-q^{-1}}, \tag{3}$$

where a is the time constant describing how the buildings are affected by changes in the corresponding climate variable, e.g. ambient air temperature.

The NWPs are localized by estimating parameters in Equation 1 during 2019, and then used out-of-sample to forecast the weather variables during 2020. The parameters of the model are estimated and the forecasts of Equation 1 are obtained using the R-package `onlineforecast`. This R-package for forecasting uses either least squares (constant parameters over time) or recursive least squares with exponential forgetting factor (adaptive parameters, i.e. varying over time). The measured temperature comes from a climate station located at the Landbohøjskolen in Frederiksberg, roughly 6 km from Tingbjerg. The performance of the raw NWPs and the NWPs that have been locally adjusted to the climate of the ambient air temperature can be seen in Figure 2. Three different localized NWPs were created: 1) Parameters estimated using least squares (LS) estimation, i.e., constant parameters (lm), 2) Parameters estimated using recursive and adaptive estimation, with forgetting factor ($\lambda = 0.994$) (rls), and 3)

[3]European Centre for Medium-Range Weather Forecasts (ECMWF), HRES - High-Resolution Forecast.

[4]Danish Meteorological Institute (DMI), Meteorological Observation, Station ID: 06186. Available: https://confluence.govcloud.dk/display/FDAPI/Meteorological+Observation [Accessed: 05-02-2022]

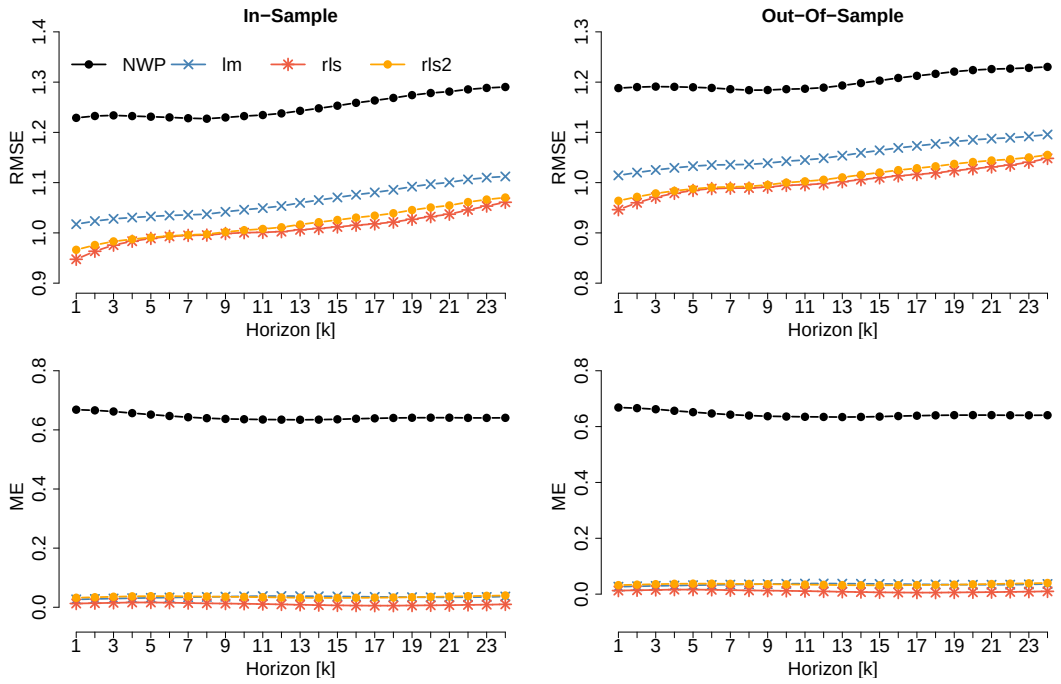

Figure 2. Figure shows the accuracy of the NWPS in Tingbjerg. The left plots show the in-sample and the right plots show the out-sample. The RMSE is shown in the top plots and ME in the bottom plots.

Parameters estimated using recursive and adaptive estimation, with forgetting factor ($\lambda = 0.998$), i.e., making the parameters change slower over time than case 2 (rls2). The forgetting factor for the second method, rls, was found by minimising the RMSE of the prediction on horizons, $k = \{1, \ldots, 24\}$ and the forgetting factor in rls2 was selected by hand. The results from the in-sample period are shown in the left plots, while the right plots show the out-of-sample results. Two error scores are used; the RMSE in the top plots and mean error (ME) in the bottom plots. The plots demonstrate the suggested method for locally adjusted forecasts to perform much better than the forecasts based directly on the NWP. From the ME, it is seen that the NWP is biased; it underestimates the temperature in the city, as expected due to the UHI effect. The rls method with optimal selection of the forgetting factor shows the best performance.

These four NWPs of the ambient air temperature are then used in Equation 2 to forecast the heat demand in Tingbjerg. The heat demand forecast accuracy improvement was validated by computing the Relative Root Mean Square Error (RRMSE) by computing the localized forecasting RMSE against the heat demand forecast RMSE-BASE using the raw NWP as inputs,

$$\left(\frac{\text{RMSE}}{\text{RMSE}_{\text{BASE}}} - 1\right) * 100 \tag{4}$$

The results shown in Figure 3 show that all three localized forecasts improved the accuracy of the heat demand forecasts when compared to the raw NWP model. The rls method which adapts faster has the best performance overall, especially for horizons up to 12 hours ahead, which are the most important horizons for temperature optimization, where the objective is to lower the supply temperature. For the last horizons, the rls performs similarly to the other localized forecasts, although it still demonstrates around 1.5% higher accuracy than using the raw NWPs as input.

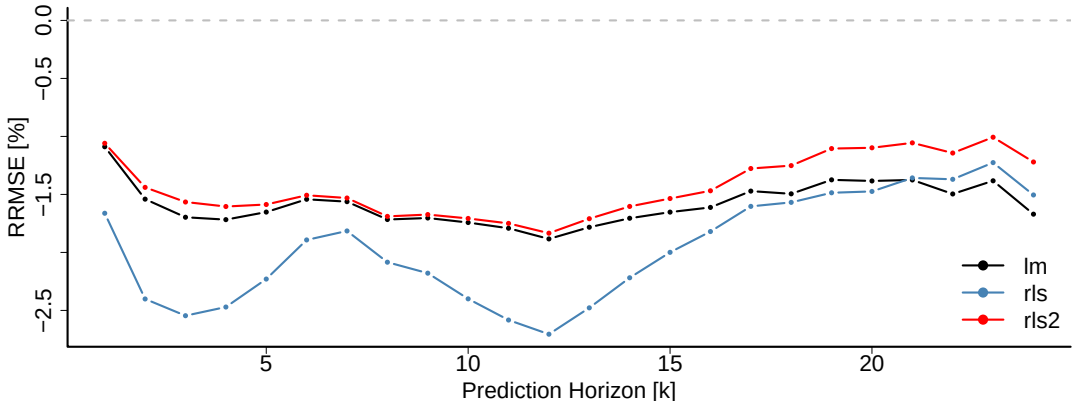

Figure 3. Accuracy improvements in heat demand forecast by localizing the meteorological forecast of the ambient air temperature to a local climate station. The RRMSE for the three localized methods compared against the forecasting model with the raw NWP as inputs.

4. Conclusion

In this paper, we demonstrated that NWP delivers a biased forecast of the temperature for the location of the climate station, and on average it predicts a lower temperature than is measured. Therefore, locally combining the NWP to measurements of the temperature from a climate station in Copenhagen would reduce the bias and give a more accurate representation of the local climate. Three different methods were proposed to locally adjust the NWP, one where the parameters are constant and two methods where the parameters vary over time with different forgetting factors.

Secondly, it is shown that locally combining weather forecasts leads to higher accuracy of the heat demand forecasts. The raw and three locally combined NWPs of the ambient air temperature were used to forecast heat demand in Tingbjerg (Copenhagen). The results show that the use of local climate data combined with NWP outperforms the forecasts based solely on the raw NWP for all forecasting horizons. This suggests that having a better understanding and forecasting of the local climate will improve the heat demand forecast, which in turn will contribute to cost savings achieved due to better operation of the district heating system.

We conclude that, in order to optimize the operation of the district heating systems, it would be beneficial to install and use a local climate station to localize the weather forecast to enhance the heat demand forecasts.

5. Acknowledgement

This work is funded by Innovation Fund Denmark through the project HEAT 4.0 (8090-00046B), TOP-UP (9045-00017B), FED (8090-00069B), and SEM4Cities (0143-00032B). The IDASC (18012745) project funded by Region Hovedstaden and the IEA TS4 annex on Digitalisation of District Heating: Optimised Operation (and Maintenance) of District heating and Cooling Schemes via Digital Processes Management funded by the EUDP Denmark, have also supported this work. The authors would like to thank HOFOR for making their data available and for their support.

References

[1] I. Blanco, D. Guericke, A. N. Andersen, H. Madsen, Operational planning and bidding for district heating systems with uncertain renewable energy production, Energies 11 (2018) 3310.

[2] H. Madsen, K. Sejling, H. T. Søgaard, O. P. Palsson, On flow and supply temperature control in district heating systems, Heat Recovery Systems and CHP 14 (1994) 613–620.

[3] P. Bauer, A. Thorpe, G. Brunet, The quiet revolution of numerical weather prediction, Nature 525 (2015) 47–55.

[4] P. Crochet, Adaptive kalman filtering of 2-metre temperature and 10-metre wind-speed forecasts in iceland, Meteorological Applications 11 (2004) 173–187.

[5] R. Ronda, G. Steeneveld, B. Heusinkveld, J. Attema, A. Holtslag, Urban finescale forecasting reveals weather conditions with unprecedented detail, Bulletin of the American Meteorological Society 98 (2017) 2675–2688.

[6] S. Idowu, S. Saguna, C. Åhlund, O. Schelén, Applied machine learning: Forecasting heat load in district heating system, Energy and Buildings 133 (2016) 478–488.

[7] H. Aalborg, H. Madsen, Predicting the heat consumption in district heating systems using meteorological forecasts (2000).

[8] P. Bacher, H. G. Bergsteinsson, L. Frölke, M. L. Sørensen, J. Lemos-Vinasco, J. Liisberg, J. K. Møller, H. A. Nielsen, H. Madsen, onlineforecast: An r package for adaptive and recursive forecasting, arXiv preprint arXiv:2109.12915 (2021).

[9] P. Bacher, H. Madsen, H. A. Nielsen, B. Perers, Short-term heat load forecasting for single family houses, Energy and buildings 65 (2013) 101–112.

[10] A. Baklanov, G. Sue, M. Alexander, M. Athanassiadou, Meteorological and air quality models for urban areas, volume 184, Springer, 2009.

[11] H. R. Glahn, D. A. Lowry, The use of model output statistics (mos) in objective weather forecasting, Journal of Applied Meteorology and Climatology 11 (1972) 1203–1211.

Paper I

ONLINEFORECAST: AN R PACKAGE FOR ADAPTIVE AND RECURSIVE FORECASTING

Authors:
Peder Bacher, Hjörleifur G. Bergsteinsson, Linde Frölke, Mikkel L. Sørensen, Julian Lemos-Vinasco, Jon Liisberg, Jan Kloppenborg Møller, Henrik Aalborg Nielsen and Henrik Madsen.

Submitted To:
The R Journal.

onlineforecast: An R Package for Adaptive and Recursive Forecasting

by Peder Bacher, Hjörleifur G. Bergsteinsson, Linde Frölke, Mikkel L. Sørensen, Julian Lemos-Vinasco, Jon Liisberg, Jan Kloppenborg Møller, Henrik Aalborg Nielsen and Henrik Madsen

Abstract Systems that rely on forecasts to make decisions, e.g. control or energy trading systems, require frequent updates of the forecasts. Usually, the forecasts are updated whenever new observations become available, hence in an online setting. We present the R package onlineforecast that provides a generalized setup of data and models for online forecasting. It has functionality for time-adaptive fitting of dynamical and non-linear models. The setup is tailored to enable the effective use of forecasts as model inputs, e.g. numerical weather forecast. Users can create new models for their particular applications and run models in an operational setting. The package also allows users to easily replace parts of the setup, e.g. using new methods for estimation. The package comes with comprehensive vignettes and examples of online forecasting applications in energy systems, but can easily be applied for online forecasting in all fields.

Introduction

Time series analysis and forecasting are of indispensable importance to numerous practical fields such as business, finance, science and engineering (Cryer and Chan, 2008). Time series analysis is the process of statistical modelling of time series, i.e. data which is sampled at different points in time over a period – often with a constant distance in time, i.e. equidistant. Classical time series models for a single equidistant time series use past values of the response variable (model output) as predictors (inputs). In this way, appropriate models describing the inherent auto-correlation structure of the time series can be realized. Such models are exponential smoothing (e.g. Holt-Winters), AutoRegressive (AR) and Moving Average (MA), and usually the combination of the latter two as ARMA models. When multiple correlated time series are at hand, they can be used as model inputs to improve forecasts. They are then called eXogenous variables and the classical model becomes an ARMAX – hence the *X* indicates that input variables are included. ARMAX models are optimal for forecasting the output of linear time invariant (LTI) systems, however for most forecasting application models for non-linear systems are needed. A wide range of techniques for modelling non-linear systems exists, either based on some approach of input transformation or some local fitting method. The onlineforecast package implements an advanced model setup for modelling and forecasting the output of non-linear time varying systems. The setup was developed for applications such as forecasting wind power (Nielsen et al., 2002) and thermal loads in district heating (Nielsen and Madsen, 2006). The significance of the package is in the "online" term, indicating that at each sampling point the model parameter estimates are updated in an effective way for generating multi-step forecasts.

The use of ARMAX models and variations thereof for forecasting has been and is still widespread (De Gooijer and Hyndman, 2006), especially for energy systems due to the high dependency between e.g. weather, load, renewable generation and periodic phenomena. Load forecasting is an obvious example. A nice overview for electric load forecasting is given by Alfares and Nazeeruddin (2002) and Hong and Fan (2016), and for heat load by Dotzauer (2002) who demonstrates the dependency between the response variable, heat load, and the predictor, ambient temperature, using a piecewise linear function. It is also proposed to model the daily and weekly diurnal using hours of the week as inputs. For solar power forecasting (Kleissl, 2013) the improvement from an autoregressive (AR) to an AR with eXogenous input (ARX), where the ARX model uses numerical weather predictions (NWPs) as inputs, is demonstrated by Bacher et al. (2009). The ARX model uses past observations and NWPs of global irradiance to forecast the power production from PV systems and the ARX model obtains higher accuracy than the AR model. Bacher et al. (2013) identifies exogenous variables that are suitable for forecasting the heat load of a building, with similar models.

Energy systems are time-varying systems as they usually change over time due to wear and contamination, like dirt on solar panels or changes in usage. For example, with new tenants in a house the dependency between the heat load and other variables, like calendar and temperature, changes. Therefore, a forecast model needs to adapt – the model coefficients are not optimal if they are constant, they need to be updated and allowed to change over time. The Recursive Least Square (RLS) method provides a recursive estimation scheme for the coefficients in regression models, where they are updated at each step when new data becomes available. Introduction of a forgetting factor in RLS allows

control on how fast the coefficients can change over time – this is referred to as adaptive recursive estimation, with exponential forgetting, in linear regression and autoregressive models. The method is described by Ljung and Söderström (1983), for advances that has been made since then see e.g. (Engel et al., 2004).

The objective of the onlineforecast package is to make it easy to set up and optimize non-linear models for generating online multi-step forecasts. The package contains functionalities not directly available elsewhere:

- Use of forecasts, e.g. NWPs, as input to multi-step forecast models.
- Optimal tuning of models for multi-step horizons.
- Recursive estimation for tracking time-varying systems.

The package provides a framework for handling data and setting up models, which makes it easy to apply it in a wide range of forecasting applications.

Time series modelling and forecasting in R

A wide range of existing software useful for time series forecasting is currently available – all have their suitable applications (Chatfield and Xing, 2019; Siebert et al., 2021). In the following an overview is given of the most relevant R packages for forecasting at the time of writing – generally, the same functionalities are available in Python packages.

Classical ARMAX models can be fitted with the `arima()` function from the stats package and the `Arima()` function from the forecast package (Hyndman and Khandakar, 2008) provides automatic model selection with `arima()`. R Packages like marima (Spliid, 1983), KFAS, sysid and dlm (Petris, 2010) can also be used for fitting ARMAX models. Spliid (1983) proposed a very fast and simple method for parameter estimation in large multivariate ARMAX models with a pseudo-regression method that repeats the regression estimation until it converges. The other packages represent time series and regression models as state-space models and use a Kalman or Bayesian filter to include exogenous variables in the model, and optimally reconstruct and predict the states. The advances of the onlineforecast models compared to ARMAX models are first of all a recursive fitting scheme, which allows for much faster and time adaptive fitting. Furthermore, model coefficients are tuned as a function of the forecast horizon. This optimize the use of multi-step forecasts as models inputs, such functionality is not available for ARMAX models.

State-space modelling is frequently used to describe time series data from a dynamical system, e.g. a falling body, see (Madsen, 2007). The dynamical system can in such cases be written as differential equations or difference equations. State-space models use filter techniques to optimally reconstruct and predict the states, e.g. the Kalman filter, the extended Kalman filter or other Bayesian filters. This gives the possibility of tracking the coefficients over time, i.e. time-varying parameter estimation. The KFAS package (Helske, 2017) provides state-space modelling, where the observations come from the exponential family, e.g. Gaussian or Poisson. The ctsm-r package provides a framework for identifying and estimating partially observed continuous-discrete time state space models, referred to as grey-box models. This modelling approach bridges the gap between physical and statistical modelling using Stochastic Differential Equations (SDEs) to model the system equations in continuous time and the measurement equations in discrete time. Packages for discrete time state-space modelling are: dlm for Bayesian analysis of dynamic linear models, MARSS and SSsimple for fitting multivariate state-space models. The onlineforecast models are basically fitted using a Kalman filter, as explained in Section .2, thus existing packages could be applied. However, the use of forecasts as model inputs would be very cumbersome and is made very easy with the onlineforecast setup.

For non-parametric time series models, the number of available packages is growing rapidly. NTS provides simulation, estimation, prediction and identification for non-linear time series data. It also includes threshold autoregressive models (e.g. self-exciting threshold autoregressive models) and neural network estimation. tsDyn provides methods for estimating non-parametric time series models, including neural network estimation. Neural network, deep learning and machine learning methods are available in R. Recurrent neural networks are in the rnn, the keras and tensorflow packages. Additive time series models, where non-linear trends are fitted with seasonality patterns, are in prophet. Time adaptive neural networks, i.e. with recursive updating, can be implemented in various ways (Yang et al., 2019), however currently no effective implementation is available.

Some packages can be useful for forecast evaluation, e.g. ForecastTB presented in (Bokde et al., 2020). Packages like forecastML and modeltime (Alexandrov et al., 2020) provide functionality that simplifies

the process of multi-step-ahead forecasting with machine learning algorithms. This purpose of handling multi-step-ahead forecasts is also a key feature of the onlineforecast package. The classical time series models, such as ARMAX and Exponential Smoothing models, are mostly optimal for modelling Linear Time Invariant (LTI) systems however most systems are not LTI. Furthermore, since a model is always a simplification of reality, optimal multi-step forecasting is often not possible with the classical models, especially when using exogenous inputs. For optimal multi-step ahead forecasting the models must be tuned for each horizon – which is exactly what the onlineforecast package does.

Functionality of onlineforecast

A model is an approximation to the real world, thus it will always be a simplification and can never predict perfectly. One of the main challenges of identifying a good forecast model is to find the most informative input variables and the best structure of the model. The onlineforecast package provides functionality for defining, validating and selecting models in a systematic way.

To introduce the onlineforecast models consider the simplest model with one input. It is the linear model for the k'th horizon

$$Y_{t+k|t} = \beta_{0,k} + \beta_{1,k} u_{t+k|t} + \varepsilon_{t+k|t} \tag{1}$$

where $Y_{t+k|t}$ is the response variable and $u_{t+k|t}$ is the input variable. The coefficients are $\beta_{0,k}$ and $\beta_{1,k}$, note that they are subscripted with k to indicate that they are estimated for individually for every horizon. The error $\varepsilon_{t+k|t}$ represents the difference between the model prediction and the observed value for the k-step horizon. The interpretation of the subscript notation $t+k|t$ on a variable is, that it is the k-step prediction calculated using only available information at time t, usually referred to either "conditional on time t" or "given time t".

The package offers to estimate the coefficients using either the Least Squares (LS) or Recursive Least Squares (RLS) method. In the LS method, the coefficients are constant, while the in RLS method the coefficients can change over time

$$Y_{t+k|t} = \beta_{0,k,t} + \beta_{1,k,t} u_{t+k|t} + \varepsilon_{t+k|t} \tag{2}$$

as indicated with the t on the coefficients. This allows for tracking changes occurring over time.

The package allows for easy definition of transformations and thus the possibility to fit non-linear models e.g.

$$Y_{t+k|t} = \beta_{0,k,t} + \beta_{1,k,t} f(u_{t+k|t}; \alpha) + \varepsilon_{t+k|t} \tag{3}$$

where the function $f(u_{t+k|t}; \alpha)$ is some non-linear function of the input $u_{t+k|t}$ with parameter α, e.g. a low pass filter on the outdoor temperature to model building heat dynamics. The package sets up tuning of the non-linear function parameters, e.g. if the parameter α determines the degree of low-pass filtering it can be tuned with an optimizer to match the system dynamics inherent in the data at hand.

An example of generated forecasts can be appreciated in Figure 1. Hourly forecasts up to 36 steps ahead of heat load in a single building are shown for three consecutive steps. This is the typical structure of forecasts generated with the package. It can be seen how the forecasts change slightly as they are updated in each step, e.g. around 12:00 the second day, hence horizon $k = 23$ in the upper plot, which corresponds to $k = 21$ in the lower plot.

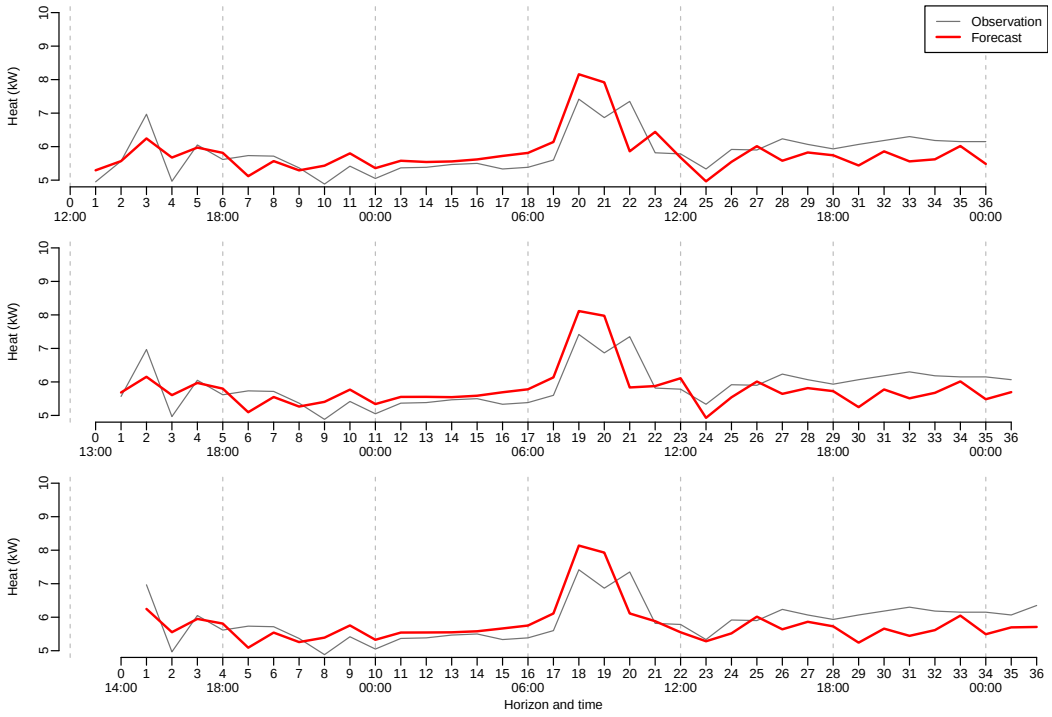

Figure 1: Example of hourly load forecasts at three consecutive time steps. The upper is calculated at 12:00, the middle is calculated at 13:00 and the lower at 14:00. It can be seen how the forecasts change slightly as they are updated in each step, most clearly seen around 12:00 on the second day.

Vignettes

A great way to get actual hands-on experience is through vignettes. They are available when installing the package and on the website onlineforecasting.org, where also examples of different forecast applications can be found. The package vignettes are:

- setup-data covers how data must be set up. The vignette goes into detail on how observations and model inputs (forecasts) are set up. The vignette also focuses on the importance of aligning forecasts correctly in time.
- setup-and-use-model focus on how to set up a model and use it to generate forecasts.
- model-selection demonstrates how model selection can be carried out.
- forecast-evaluation covers the evaluation of forecasts, and how to use this information to improve a model.
- online-updating demonstrates how to update a model in actual operation when new observations become available. This functionality is not described in the R examples in the present paper.

Furthermore, one vignette is available only on the website:

- nice-tricks provides some useful tips on how to make the workflow easier with the package.

Paper structure

The structure of the paper is the following: In Section 2.2 the notation used in the paper and how to set up data is introduced. The core methodology is presented in Section 2.3 and important aspects of forecast modelling are outlined in Section 2.4. In Section 2.5 examples with R code are presented to provide a short hands-on tutorial. The paper ends with a summary and conclusions in Section 2.6.

In addition, three appendixes are included in the paper. In Appendix .1 some guidelines on mathematical notation of forecast models are provided. In Appendix .2 the regression schemes are covered in full detail.

Notation and forecast matrices

The notation in this article follows Madsen (2007) as close as possible. All time series considered are equidistantly sampled and the sampling period is normalized to 1. Hence, the time t is simply an integer indexing the value of a variable at time t. The same goes for k which indexes the forecast horizon k steps ahead. In the onlineforecast setup, forecasts are calculated at time t for each horizon up to n_k steps ahead. To achieve the desired notation that can deal with overlapping time series, a two dimensional index is required. The notation used is

$$u_{t+k|t} \tag{4}$$

which translates to: the value of variable u at time $t+k$ *conditional* on the information available at time t. The conditional term is indicated by the bar $|$. Thus, for $k > 0$ this is a forecast available at t and k is the horizon. When writing a forecast model the following convention is used, here a simple example

$$Y_{t+k|t} = \beta_{0,k} + \beta_{1,k} u_{t+k|t} + \varepsilon_{t+k|t} \tag{5}$$

where $Y_{t+k|t}$ is the model output, $\beta_{0,k}$ and $\beta_{1,k}$ are the coefficients and $\varepsilon_{t+k|t}$ with $\mathrm{Var}(\varepsilon_{t+k|t}) = \sigma_k^2$ is the error. The error process and variance σ_k^2 is thus separate for each horizon. Note, that the model is fitted separately for each horizon, so the coefficients take different values for each horizon, and the predictions and errors are separated for each horizon. This was a simplified example, see Appendix .1 on how to write the full forecast models.

Forecast matrix

A forecast matrix is the format of forecast data in the onlineforecast setup. See examples in the setup-data vignette. Data must have this format in order to be used as model input, and the forecasts generated are in this format. The forecast matrix holds for any past time *the latest available forecast along the row* for the corresponding time

$$\boldsymbol{u}_n = \begin{array}{c} \begin{array}{ccccc} \mathbf{k0} & \mathbf{k1} & \mathbf{k2} & \dots & \mathbf{k}n_k \end{array} \\ \begin{pmatrix} u_{1|1} & u_{2|1} & u_{3|1} & \dots & u_{1+n_k|1} \\ u_{2|2} & u_{3|2} & u_{4|2} & \dots & u_{2+n_k|2} \\ \vdots & \vdots & \vdots & & \vdots \\ u_{t-1|t-1} & u_{t|t-1} & u_{t+1|t-1} & \dots & u_{t-1+n_k|t-1} \\ u_{t|t} & u_{t+1|t} & u_{t+2|t} & \dots & u_{t+n_k|t} \\ \vdots & \vdots & \vdots & & \vdots \\ u_{n|n} & u_{n+1|n} & u_{n+2|n} & \dots & u_{n+n_k|n} \end{pmatrix} \end{array} \begin{array}{c} \rightarrow \textbf{horizon/time} \downarrow \\ 1 \\ 2 \\ \vdots \\ t-1 \\ t \\ \vdots \\ n \end{array} \tag{6}$$

where

- t is the counter of time for equidistant time points with sampling period 1 (note that t is not included in the matrix, it is simply the row number).
- n is the number of time points in the matrix. Hence, the data is available and can be used as model input at time $t = n$.
- n_k is the longest forecasting horizon.
- The column names (in R) are indicated above the matrix, they are simply a '**k**' concatenated with the value of k, e.g. n_k in the last column.

Note, that the **k0** column holds values with forecast horizon $k = 0$, which could be real time observations. Usually, only the horizons to be forecasted should be included, hence often **k0** is not needed.

For example with a prediction horizon $n_k = 24$ at $t = 100$, we will have the forecast matrix

$$\boldsymbol{u}_{100} = \begin{array}{c} \begin{array}{ccccc} \textbf{k0} & \textbf{k1} & \textbf{k2} & \dots & \textbf{k24} \end{array} \\ \begin{pmatrix} u_{1|1} & u_{2|1} & u_{3|1} & \dots & u_{25|1} \\ u_{2|2} & u_{3|2} & u_{4|2} & \dots & u_{26|2} \\ \vdots & \vdots & \vdots & & \vdots \\ u_{99|99} & u_{100|99} & u_{101|99} & \dots & u_{123|99} \\ u_{100|100} & u_{101|100} & u_{102|100} & \dots & u_{124|100} \end{pmatrix} \end{array} \begin{array}{c} \rightarrow \textbf{horizon/time} \downarrow \\ 1 \\ 2 \\ \vdots \\ 99 \\ 100 \end{array} \tag{7}$$

In Section 2.5.1 examples of how data and forecast matrices are set up in R are given.

Two-stage modelling procedure

A widespread approach to modelling non-linear functional relations between inputs and output is a two-stage modelling procedure. See, e.g. Breiman and Friedman (1985) and Weisberg (2005) for direct transformation of predictor variables, and Hastie et al. (2009) for non-parametric transformation techniques (basis functions). Using transformations allows for fitting complex models with robust and fast estimation techniques. In the first stage, the *transformation stage*, the inputs are mapped by some function – potentially into a higher dimensional space. In the second stage, the *regression stage*, a linear regression model[1] is applied between the transformed inputs and the output. An exemplification of this is presented in the following.

As an example a model with two inputs is presented. In this model the transformation stage consists of generating an intercept and mapping the two inputs (they are set up as forecast matrices $\boldsymbol{u}_{1,t}$ and $\boldsymbol{u}_{2,t}$)

$$\text{Intercept:} \quad x_{0,t+k|t} = 1 \tag{8}$$

$$\text{Input 1:} \quad x_{1,t+k|t} = f_1(u_{1,t+k|t}, \boldsymbol{\alpha}_1) \tag{9}$$

$$\text{Input 2:} \quad \boldsymbol{x}_{2,t+k|t} = f_2(u_{2,t+k|t}, \boldsymbol{\alpha}_2) \tag{10}$$

where the f's are transformation functions that map the inputs to regressors. Note, that the intercept is simply a constant passed on to the regression. The transformations result in multiple inputs for the regression – the latter actually as multiple variables indicated by the bold font notation. In the regression stage the linear model

$$Y_{t+k|t} = \beta_{0,k} x_{0,t+k|t} + \beta_{1,k} x_{1,t+k|t} + \boldsymbol{\beta}_{2,k}^T \boldsymbol{x}_{2,t+k|t} + \varepsilon_{t+k|t} \tag{11}$$

is fitted. The regression is carried out separately for each horizon k. Thus, the combined model has:

- An intercept
- Two inputs: $u_{1,t+k|t}$ and $u_{2,t+k|t}$
- Output: $Y_{t+k|t}$
- Transformation functions: f_1 and f_2
- Transformation parameters: $\boldsymbol{\alpha}_1$ and $\boldsymbol{\alpha}_2$
- Regression coefficients: $\beta_{0,k}$, $\beta_{1,k}$ and $\boldsymbol{\beta}_{2,k}$

Some transformation parameters should be optimized for the data at hand, e.g. a low-pass filter coefficient depends on the system dynamics. The same goes for some parameters related to the regression scheme, e.g. the forgetting factor (introduced below). We will refer to them together as ***offline parameters***. The onlineforecast package provides a setup, where the offline parameters can be optimized using a heuristic optimization (e.g. a BFGS quasi-Newton method). The default score, which

[1] In the remaining of the text, when the term "regression" is used it is implicit that it is "linear regression".

is minimized, is the Root Mean Square Error (RMSE) of the predictions – hence offline parameters in the model above, given data from the period, $t = 1, 2, \ldots, n$, are optimized by solving

$$\min_{\boldsymbol{\alpha}_1, \boldsymbol{\alpha}_2} \quad \frac{1}{n-k} \sum_{t=1}^{n-k} (y_{t+k} - \hat{y}_{t+k|t}(\boldsymbol{\alpha}_1, \boldsymbol{\alpha}_2))^2 \tag{12}$$

Naturally, other scores can be minimized (e.g. MAE or the Huber psi-function, however the regression schemes should be modified accordingly, which is not trivial).

The regression coefficients are calculated with a closed-form scheme: either with the Least-Squares (LS) or the Recursive Least-Squares (RLS) scheme – in the latter the coefficients are allowed to vary over time. In both schemes the coefficients are gathered in the vector $\boldsymbol{\beta}_k$ and calculated separately for each horizon k. In Appendix .2 both schemes are presented in full detail.

In the LS scheme the coefficients are constant during the entire period. The output vector is $\boldsymbol{y}_{k,n}$ and for a given value of the transformation parameters (i.e. here $\boldsymbol{\alpha}_1$ and $\boldsymbol{\alpha}_2$) the transformed data is calculated and set up in the design matrix $\boldsymbol{X}_{k,n}$. The LS coefficients are then calculated by

$$\hat{\boldsymbol{\beta}}_k = (\boldsymbol{X}_{k,n}\boldsymbol{X}_{k,n})^{-1}\boldsymbol{X}_{k,n}\boldsymbol{y}_{k,n} \tag{13}$$

and the predictions calculated by

$$\hat{\boldsymbol{y}}_{k,n} = \boldsymbol{X}_{k,n}\hat{\boldsymbol{\beta}}_k \tag{14}$$

where $\hat{\boldsymbol{y}}_{k,n} = \begin{bmatrix} \hat{y}_{1+k|1} & \hat{y}_{2+k|2} & \cdots & \hat{y}_{n|n-k} \end{bmatrix}^T$ are the predictions. Note, that for the LS scheme the predictions are "in-sample", since data from the entire period is used for the coefficient calculation.

In the RLS scheme the coefficients are calculated recursively, meaning that they are updated in every time step – the RLS is actually a Kalman filter with the model coefficients in the state vector. At each time t the coefficients are updated by

$$\boldsymbol{R}_{k,t} = \lambda \boldsymbol{R}_{k,t-1} + \boldsymbol{x}_{k,t}\boldsymbol{x}_{k,t}^T \tag{15}$$

$$\hat{\boldsymbol{\beta}}_{k,t} = \hat{\boldsymbol{\beta}}_{k,t-1} + \boldsymbol{R}_{k,t}^{-1}\boldsymbol{x}_{k,t}(y_t - \boldsymbol{x}_{k,t}^T\hat{\boldsymbol{\beta}}_{k,t-1}) \tag{16}$$

and the predictions by

$$\hat{y}_{t+k|t} = \boldsymbol{x}_{t+k|t}\hat{\boldsymbol{\beta}}_{k,t} \tag{17}$$

where $\boldsymbol{x}_{k,t}$ is the data available for horizon k at time t (a row in the design matrix $\boldsymbol{X}_{k,n}$), and $\boldsymbol{x}_{t+k|t}$ is the k'th horizon transformed input forecast available at time t, see the appendix for all details. The coefficients adapts to data over time and the level of adaptivity is controlled by setting the forgetting factor λ to a value between 0 and 1. For $\lambda = 1$ all past data at t is equally weighted. For $\lambda < 1$ higher weight is put on recent data – the smaller the value the faster the model adapts to recent data. By optimizing the forgetting factor as an offline parameter the model adaptivity can be tuned.

An important point to notice is, that the offline parameters are always constant for the given period, hence all predictions are essentially "in-sample". However, depending on the regression scheme there is a difference: with the LS scheme the regression coefficients are calculated once using all data, thus the predictions are (fully) "in-sample", where as with the RLS scheme they adapt through the period and the predictions are "out-of-sample" (except for the offline parameters). This makes a difference, since model over-fitting is less of a problem when using the RLS scheme.

The typical onlineforecast setup is to optimize the (usually few) offline parameters in an "offline" setting, but calculate the regression coefficients adaptively with the RLS. This has the advantage that the model adapts and tracks the systematic changes in input-output relations, while keeping the setup computationally very effective – updating the coefficients and calculating a forecast at each time t takes few operations. The optimization of offline parameters can be carried when computational resources are available (e.g. every week for hourly forecasts).

Transformations

In the transformation stage the inputs are mapped using some function as demonstrated above, for more examples, see the setup-and-use-model vignette. The onlineforecast package has functions available for most common use, however it is easy to write and use new functions as they are simply R functions. The main functionality they have to fulfil is to return a forecast matrix (or a list of them),

which is also the reason why some of the regular R functions and operators has been extended for the multi-horizon setup, e.g. for splines as explained below. For R examples, see Section 2.5.2 and the vignettes. The currently available transformation functions are:

- Low-pass filtering, `lp()`: A low-pass filtering for modelling linear dynamics as a simple RC-model. See e.g. Nielsen and Madsen (2006) for further information.
- Basis splines, `bspline()`: Use the `bs` function for calculating regression splines basis functions.
- Periodic basis splines, `pbspline()`: Use the `pbs` function for calculating periodic regression splines basis functions.
- Fourier series, `fs()`: Fourier series as periodic regression basis functions.
- Auto-regressive, `AR()`: For including Auto-Regressive (AR) terms.
- Intercept, `one()`: Generates a forecast matrix of ones, i.e. intercept.

In the following section, the low-pass filtering is shortly described below. For more examples of transformations, see the package vignettes.

The implementation in onlineforecast allows all parameters, which are used in some way (except the regression coefficients), to be included in an optimization – using any available optimizer i R. This includes e.g. the RLS forgetting factor, knot points or order of splines – hence both continuous and integer variables. This functionality is achieved using a simple syntax as explained in Section 2.5.

Low-pass filtering

When modelling time series from linear dynamical systems, the classical ARMAX model is often the optimal choice (Madsen, 2007). However, for multi-step forecasting this is often not the case, especially for longer horizons. In the onlineforecast setup, where the regression model is fitted for each horizon, a "trick" can be used for modelling linear dynamics: simply apply a filter on the input and then use the filtered input in the regression stage. For example, dynamics between ambient air temperature and heat demand are slow due to the thermal mass of the building. Thus they can be modelled using a low-pass filter, see Nielsen and Madsen (2006) for modelling heat load in district heating and Bacher et al. (2013) for forecasting single buildings heat load.

In the package the simple low-pass filter

$$x_{t+k|t} = \frac{(1-a)u_{t+k|t}}{1 - ax_{t-1+k|t-1}} \tag{18}$$

is implemented. The filter coefficient a must take a value between 0 and 1 and should be tuned to match the time constant optimal for the particular data. When the current implemented low-pass filter is applied in the transformation stage, on some forecast matrix $\boldsymbol{u}_{t+k|t}$, the filter is applied on each column. Hence, independently for each horizon k. More advanced filters can be implemented.

Model selection and validation

Model selection

In statistics, different model selection procedures are used (Madsen and Thyregod, 2010). Essentially, a backward or a forward selection procedure can be applied, or some combined approach. In the onlineforecast package both procedures are implemented, as well as a combined approach, see the model-selection vignette for examples.

In each step of the selection process two properties of the model can be modified:

- Model inputs: In each step, inputs can either be removed or added.
- Integer offline parameters: In each step integer parameters, such as the number of knot points in a basis spline or the number of harmonics in a Fourier series can be counted one up or down.

In each step of the process, the offline parameters are first optimized to minimize the score for each modified model (in most cases the appropriate score is the RMSE in Equation (12) summed for selected

horizons). Then the scores of the modified models are compared with the score of the currently selected model and the model with the lowest score is selected for the next step. This continues until no further improvement of the score is achieved and the model with the lowest score is selected. It is important to note, that the implemented procedure should only be used with the RLS scheme, with the LS scheme the score is calculated fully in-sample leading to over-fitting.

Model validation

The most important aspects of validation of forecast models are discussed in this section, see the forecast-evaluation vignette for examples.

Training and test set

One fundamental caveat in data-driven modelling is over-fitting. This can easily happen when the model is fitted (trained) and evaluated on the same data. There are essentially two ways of dealing with this: Penalize increased model complexity (regularization) or divide the data into a training set and test set (cross-validation) (Tashman, 2000). In most forecasting applications the easiest and most transparent approach is some cross-validation – many methods for dividing into sets are possible. In the onlineforecast setup, when a model is fitted recursively using the RLS only past data is used when calculating the regression coefficients, so there is little need for dividing into a training and a test set.

The offline parameters (like the forgetting factor and low-pass filter coefficients) are optimized on a particular period, hence over-fitting is possible, however it is most often very few parameters compared to the number of observations – so it is very unlikely to over-fit a recursive fitted model in this setup.

Scoring

Scoring forecasts can be done in many ways, however in the onlineforecast, where the conditional mean is estimated and when using the RLS scheme, it is straightforward to choose the Root Mean Square Error (RMSE) in Equation (12) as the best score to use. When using the LS scheme it can be favourable to include regulation to avoid over-fitting, hence AIC or BIC is preferable. One important point when comparing forecasts is to only include the complete cases, i.e. forecasts at time points with no missing values across all horizons and across all evaluated models. A function for easy selection of only complete cases given multiple forecasts is implemented, see the examples in Section 2.5.4.

Residual analysis

Analysing the residuals is an important way to validate that a model cannot be further improved or learn how it can improved. The main difference from classical time series model validation, where only the one-step ahead error is examined, is that multiple horizons should be included in the analysis. The two most important analysis:

- Plot residual time series to find where large forecast errors occur.
- Plot scatter plots of the residuals vs. other variables to see if any apparent dependencies are not described by the model.

In order to dig a bit more into the result of a recursive estimation, the regression coefficients can be plotted over time. In this way, it is possible to learn how the relations between the variables in the model evolve over time. If drastic changes are found in some periods it might be worthwhile to zoom into those periods to learn what causes these changes and potentially how to improve the model. In case auto-correlation is left in the residuals, an error model can be used to improve the forecasts by applying an auto-regressive model on the residuals. This is somewhat equivalent to include an MA part in the original model.

As summarizing measures for validation of how well dynamics are modelled:

- Plot the auto-correlation function (ACF) of the one-step residuals.

- Plot cross-correlation functions from one-step residuals to other variables, see (Bacher et al., 2013).

Systematic patterns found in these functions lead to direct knowledge on how to improve the model, see for example the table on page 155 in Madsen (2007).

Example with R code

A short introduction to the functionalities and steps in setting up a model is given in the following – for more details, see the vignettes listed in Section 2.1.3 and the website onlineforecasting.org.

First, a few remarks on the implementation. onlineforecast models are set up using an object-oriented R6 class. The main reason for this is that R6 objects are pointed to per reference, which allows to make minimum changes in computer memory when updating a model fit with new data – this would not be possible with the regular S3 class objects, as they are always copied in memory when changed inside a function.

Furthermore, it is noted, that model inputs and transformations simply are defined using R code. The regular `formula` class is not used, since it cannot operate as needed on the multi-horizon forecast matrices. The provided code for the inputs defines the transformations etc. and is executed for each input to generate the data used for regression.

Setup of data

Data must be set as variables in a list, here we have loaded `D` with the data for the examples:

```
class(D)


## [1] "data.list" "list"
```

As seen its class is `data.list`, which is inherited from the `list` class. Hence, it is simply a `list` extended with some modified and new functions (can be listed with `methods(class="data.list")`).

All inputs to be used must be formatted as forecast matrices and set in the list as `data.frames`. For example the ambient temperature forecasts:

```
class(D$Ta)


## [1] "data.frame"


head(D$Ta[ ,1:8], 4)


##         k1       k2      k3      k4       k5       k6       k7       k8
## 1 -2.82340 -3.20275 -3.1185 -3.0896 -3.13200 -3.16130 -3.16645 -3.08885
## 2 -2.90405 -3.11850 -3.0896 -3.1320 -3.16130 -3.16645 -3.08885 -2.77165
## 3 -2.93590 -3.08960 -3.1320 -3.1613 -3.16645 -3.08885 -2.77165 -2.32185
## 4 -2.89315 -3.11285 -3.0484 -3.1090 -3.11600 -2.80990 -2.36895 -2.00945
```

The time must set in a `POSIXct` vector named `t`:

```
D$t[1:4]


## [1] "2010-12-15 01:00:00 UTC" "2010-12-15 02:00:00 UTC"
## [3] "2010-12-15 03:00:00 UTC" "2010-12-15 04:00:00 UTC"
```

Observations are simply set in a numeric vector:

```
D$heatload[1:4]


## [1] 5.916667 5.850000 5.850000 5.883333
```

For more details on the `data.list` class, see the setup-data vignette – which demonstrates useful functions for manipulating, validating and exploring forecast data.

Defining a model

Models are set up using the R6 class `forecastmodel`. An object of the class is instantiated by:

```
model <- forecastmodel$new()
```

It holds variables and functions for representing and manipulating a model.

If we want to forecast the observed `heatload` variable in the data list `D`, we set that as the model output by:

```
model$output <- "heatload"
```

The model inputs must then be defined. We can add an input as a linear function by:

```
model$add_inputs(Ta = "Ta")
```

The code given as text simply evaluates into the `Ta` forecast matrix, which will lead to the k-step forecast of ambient temperature (i.e. a column in `Ta`) will be set directly into the design matrix for the k horizon regression (more explanation of this is given in the end of the current section).

Adding an intercept to a model can be done by:

```
model$add_inputs(mu = "one()")
```

where the function `one()` evaluates into a forecast matrix of 1's, which will be inserted in the design matrix, see details in Appendix .2.

Functions for a range of useful transformations were already listed in Section 2.3.1. Dynamics can be modelled using filters, for example low-pass filtering of a variable with:

```
model$add_inputs(Ta = "lp(Ta, a1=0.9)")
```

will apply a low-pass filter along each column of `Ta` and return a forecast matrix with the modified data. The filter coefficient is set to $a = 0.9$. To illustrate the effect of this, see the vignette setup-and-use-model.

Non-linear effects can be modelled using basis functions. For mapping an input to basis splines the function `bspline()` is provided. It is a wrapper of the `bs()` function from the splines package and has the same arguments. To e.g. include a non-linear function of the ambient temperature:

```
model$add_inputs(Ta = "bspline(Ta, df=5)")
```

where `df` is the degrees of freedom of the spline function.

Functions can be nested, e.g. first a low-pass filter before mapping to basis splines:

```
model$add_inputs(Ta = "bspline(lp(Ta, a1=0.9), df=5)")
```

Varying-coefficient models can be realized with multiplication of inputs (Hastie and Tibshirani, 1993). For more details, see the solar-power-forecasting example on the website. For more examples of input transformations, e.g. fourier-series and auto-regressive inputs, see the setup-and-use-model vignette.

Execution of the input transformations

As seen above, R code is given for each input. The code is given as text and will simply be executed to calculate the data for regression. This is carried out with the function `model$transform_data(data)`, inside which:

```
eval(parse(text=frml), data)
```

is executed for each input. The `frml` is the R code for the input (as a character e.g. as `"Ta"` or `"bspline(Ta,df=5)"` in the examples) and `data` is the list containing the variables used in the `frml` code (as `D` in the examples).

This way of defining the input formulas simply as code is very flexible. It also allows for easy debugging, for example a function used in the code can be set for debug and it is possible to step through its execution during the transformation – or even by setting `"browser();"` directly into in the input code to stop and step through the execution.

The only constraint to an inputs code is that it must just return a forecast matrix as a `data.frame` (or list of them). The regression can then be carried out *separately for every horizon* by, for the k horizon taking the k column from each of the returned matrices and bind these together into a design matrix, on which LS or RLS can be applied to calculate the k horizon forecasts.

Model fitting and offline parameter tuning

After setting up a model it can be fitted to data by carrying out the transformation and regression. To simply the coding functions of fitting a model are provided. Different functions implement different regression schemes. The two currently available fitting functions are `lm_fit()` and `rls_fit()` – they take offline parameters as a vector, fits a model and returns the RMSE score (summed for all fitted horizons).

To demonstrate this we replace the inputs on the model defined above with two inputs by:

```
model$inputs <- NULL
model$add_inputs(mu = "one()",
                 Ta = "lp(Ta, a1=0.9)")
```

We also have to set the "score period", which is simply a logical vector that specifies the observations to be included in the score calculation. It is useful for defining a burn-in period and dividing in test and training sets. For the current linear regression we simply include all points by:

```
D$scoreperiod <- rep(TRUE, length(D$t))
```

Now the summed RMSE for the horizons 1 to 6 steps ahead can be obtained by:

```
model$kseq <- 1:6
lm_fit(c(Ta__a1=0.8), model, D, scorefun=rmse, returnanalysis=FALSE)


## [1] 5.039611
```

The function can be passed to an optimizer, which can then find the parameter value(s) which minimizes the score.

Any optimizer function in R can be used, but, again to simplify the code, wrappers for `optim()` are included – similar wrappers can easily be made for other optimizers. The parameter(s) to optimize within the wrapper function is defined by:

```
model$add_prmbounds(Ta__a1 = c(min=0.8, init=0.95, max=0.9999))
```

Note, the double underscore syntax. The double underscore separates the input name and the name of the parameter. So in the case above the value of `a1` in the R code for input `Ta` will be optimized, starting at an initial value of 0.95 and staying within the specified bounds.

We can then run the optimization calculating scores (to save computational time run on only a horizons 3 and 18 steps ahead):

```
lm_optim(model, D, kseq=c(3,18))
```

The `lm_optim()` function is a wrapper for the R optimizer function `optim()`. It returns the result from `optim()` and sets optimized parameters in:

```
model$prm


##    Ta__a1
## 0.8926346
```

i.e. a lower low-pass coefficient than the initial value of 0.95.

For more details, see the setup-and-use-model vignette.

Calculating forecasts

While developing models it is most convenient to use the fit functions for calculating predictions, e.g.:

```
model$kseq <- 1:24
fit <- lm_fit(model$prm, model, D)
```

will return a list holding the forecasts (in the forecast matrix `fit$Yhat`) and other useful information. Forecasts can also be calculated directly with a predict function:

```
lm_predict(model, model$transform_data(D))
```

will return a forecast matrix using the input data in `D`.

Evaluation

Finally, it is time to evaluate the forecasts and potentially get inspired to improve the model. For a more comprehensive introduction, see the forecast-evaluation vignette.

First, a plot of the forecasts is always a good idea to learn how the model work:

```
D$Yhat <- fit$Yhat
plot_ts(subset(D,D$scoreperiod), "heatload$|Yhat", kseq=c(1,5,24), p=p)
```

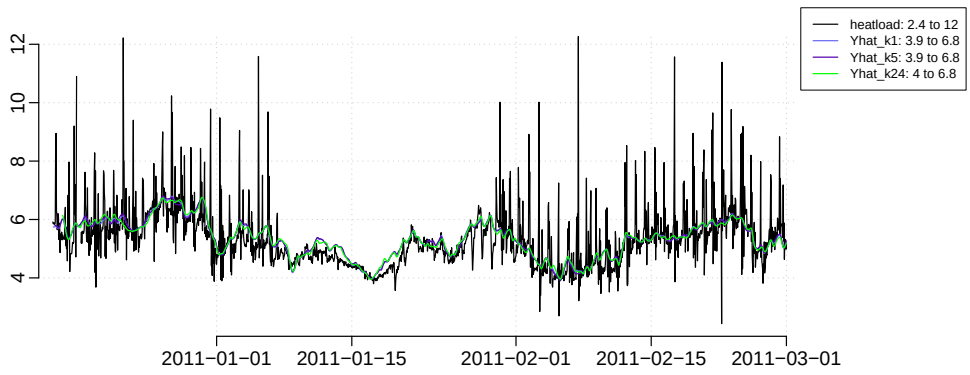

We can plot the ACF of the one-step residuals by:

```
acf(residuals(fit)$h1, na.action=na.pass, lag.max=96, main="")
```

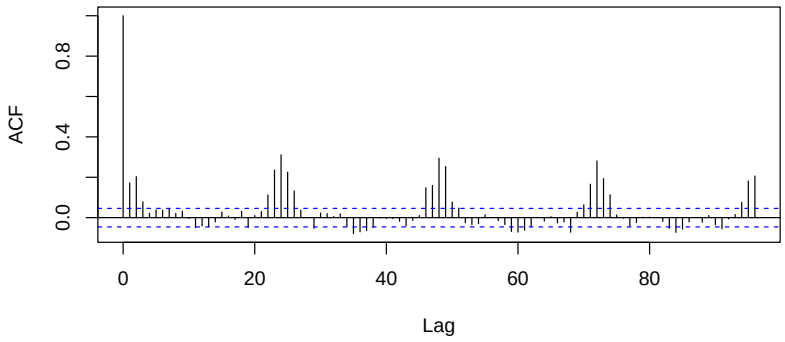

The ACF plot suggests that there remains a diurnal pattern to be modelled. It can be achieve by adding a diurnal curve to the model, e.g. with Fourier series basis functions. This is demonstrated in the vignette setup-and-use-model.

We also want to calculate the score as a function of the horizon:

```
inscore <- D$scoreperiod & complete_cases(fit$Yhat)
RMSE <- score(residuals(fit), scoreperiod = inscore)
plot(RMSE, ylim=c(0.75,0.88), xlab="Horizon")
```

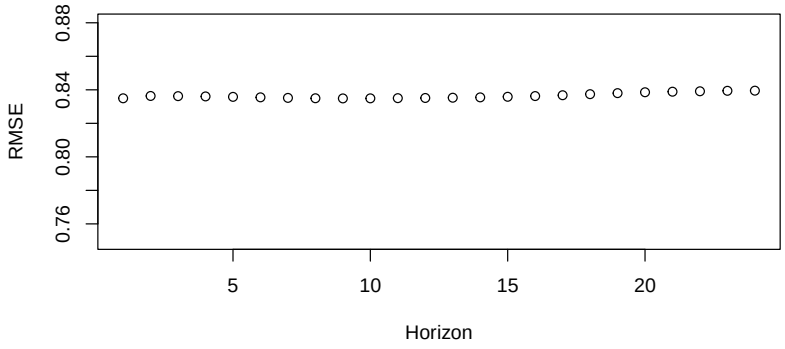

Which is relatively constant, since the model is very simple. The offline parameters were optimized for $k = 3$ and $k = 18$, which can explain why it is not monotonic increasing with the horizon.

Discussion and conclusion

Extending functionality

The current package is designed to make it easy to implement new transformation functions and regression schemes, as well as using other optimizers for tuning parameters.

Implementing a new transformation function is straight forward. It must receive either a forecast matrix or a list of forecast matrices and return either after processing. Furthermore, when used in an operational online setup, where the transformation is executed when new data arrives, it is possible to save state information inside a transformation function, such that next time the function is called, the state can be read and used. See the `lp()` function for inspiration when writing a new transformation function.

A new regression scheme, e.g. a kernel or quantile regression, can be implemented. A fitting function should be implemented in similar way as `lm_fit()` and `rls_fit()`, such that the first argument is the parameter vector and it returns a score value, which can be passed to an optimizer.

It is very easy to use other optimizers. The current fitting functions can simply be passed to any optimizer in R, which follows the `optim()` way of receiving a function for optimization, see the code

in `lm_optim()`.

In future versions new regression techniques, e.g. kernel regression (local fitting) and quantile regression, might be added. The latter opens up the possibilities to calculate probabilistic forecasts, see (Nielsen et al., 2006) and (Bjerregaard et al., 2021), as well as carry out normalization and Copula transformations, which can be very useful for spatio-temporal forecast models, see (Tastu et al., 2011) or (Lemos-Vinasco et al., 2021).

Summary and conclusion

This paper provides an entry point and reference for working with the onlineforecast package. The paper covers version 1.0 of the package, which has been available on CRAN in almost one year at the time of writing.

The main contribution of the package is to make it easy to generate online multi-step forecasts in a flexible way. The package contains functionalities not directly available elsewhere, such as:

- Enabling the use of input variables given as forecasts, e.g. NWPs, in an easy and flexible way.
- Optimal tuning of non-linear models for multi-step horizons.
- Recursive estimation for tracking time-varying systems computationally efficient for multiple horizons.

The onlineforecast package has a significant value for anyone who needs to carry out operational online forecasting. For example, in energy scheduling, where recursive updated forecasts are needed as input to optimal decision making and real-time control of systems. It can also be very useful for companies that need online forecasts for other monitoring and real-time applications – especially the functionality for model updating with very little computational costs when new data becomes available, is a unique feature of the package.

Computational details

We have tried to make the onlineforecast package depend on as few other packages as possible. Only a few additional packages are used in the core functionalities: **R6** for the "usual" OOP functionalities and **Rcpp** (Eddelbuettel and Balamuta, 2018) with **RcppArmadillo** (Eddelbuettel and Sanderson, 2014) for easy integration of fast compiled code. For extending the modelling possibilities the **splines** and **pbs** packages are essential, and for nice caching the **digest** package. We acknowledge the **devtools** and **knitr** (Xie, 2015), **rmarkdown** (Xie et al., 2018), **R.rsp**, **testthat** (Wickham, 2011) packages, which are indispensable for developing a package. We acknowledge the R community and the amazing work behind R done by many people over the years!

The results in this paper were obtained using R 4.1.3. R itself and all packages used are available from the Comprehensive R Archive Network (CRAN) at `https://CRAN.R-project.org/`.

Acknowledgments

The software has been developed with funding from multiple projects: PTXHeatUtilisation (Energy Cluster Denmark), Flexible Energy Denmark, Heat 4.0 and Decision support tools for smart home energy management systems (Innovation Fund Denmark, No. 9045-00017B, 8090-00046B and 8053-00156B), TOP-UP (Innovation Fund Denmark and ERA-NET, No. 9045-00017B), Digital-twin, IEA Annex 71 and 83 Danish participation (EUDP, No. 64019-0570, 64017-05139 and 64020-1007), and finally SCA+ (EU Interreg, No. 20293290).

Bibliography

A. Alexandrov, K. Benidis, M. Bohlke-Schneider, V. Flunkert, J. Gasthaus, T. Januschowski, D. C. Maddix, S. S. Rangapuram, D. Salinas, J. Schulz, et al. Gluonts: Probabilistic and neural time series modeling in python. *J. Mach. Learn. Res.*, 21(116):1–6, 2020. [p2]

H. K. Alfares and M. Nazeeruddin. Electric load forecasting: Literature survey and classification of methods. *International Journal of Systems Science*, 33(1):23–34, 2002. doi: 10.1080/00207720110067421. URL https://doi.org/10.1080/00207720110067421. [p1]

P. Bacher, H. Madsen, and H. A. Nielsen. Online short-term solar power forecasting. *Solar Energy*, 83 (10):1772 – 1783, 2009. ISSN 0038-092X. doi: https://doi.org/10.1016/j.solener.2009.05.016. [p1]

P. Bacher, H. Madsen, H. A. Nielsen, and B. Perers. Short-term heat load forecasting for single family houses. *Energy and buildings*, 65:101–112, 2013. [p1, 8, 10]

M. B. Bjerregaard, J. K. Møller, and H. Madsen. An introduction to multivariate probabilistic forecast evaluation. *Energy and AI*, page 100058, 2021. [p15]

N. D. Bokde, Z. M. Yaseen, and G. B. Andersen. Forecasttb—an r package as a test-bench for time series forecasting—application of wind speed and solar radiation modeling. *Energies*, 13(10):2578, 2020. [p2]

L. Breiman and J. H. Friedman. Estimating optimal transformations for multiple regression and correlation. *Journal of the American statistical Association*, 80(391):580–598, 1985. [p6]

C. Chatfield and H. Xing. *The analysis of time series: an introduction with R*. Chapman and hall/CRC, 2019. [p2]

J. D. Cryer and K.-S. Chan. *Time series analysis: with applications in R*, volume 2. Springer, 2008. [p1]

J. G. De Gooijer and R. J. Hyndman. 25 years of time series forecasting. *International journal of forecasting*, 22(3):443–473, 2006. [p1]

E. Dotzauer. Simple model for prediction of loads in district - heating systems. *Applied Energy*, 73(3-4): 277–284, 2002. doi: 10.1016/S0306-2619(02)00078-8. [p1]

D. Eddelbuettel and J. J. Balamuta. Extending extitR with extitC++: A Brief Introduction to extitRcpp. *The American Statistician*, 72(1):28–36, 2018. doi: 10.1080/00031305.2017.1375990. [p15]

D. Eddelbuettel and C. Sanderson. Rcpparmadillo: Accelerating r with high-performance c++ linear algebra. *Computational Statistics and Data Analysis*, 71:1054–1063, March 2014. URL http://dx.doi.org/10.1016/j.csda.2013.02.005. [p15]

Y. Engel, S. Mannor, and R. Meir. The kernel recursive least-squares algorithm. *IEEE Transactions on signal processing*, 52(8):2275–2285, 2004. [p2]

T. Hastie and R. Tibshirani. Varying-coefficient models. *Journal of the Royal Statistical Society: Series B (Methodological)*, 55(4):757–779, 1993. [p11]

T. Hastie, R. Tibshirani, and J. Friedman. *The elements of statistical learning: data mining, inference, and prediction*. Springer Science & Business Media, 2009. [p6]

J. Helske. KFAS: Exponential family state space models in R. *Journal of Statistical Software*, 78(10):1–39, 2017. doi: 10.18637/jss.v078.i10. [p2]

T. Hong and S. Fan. Probabilistic electric load forecasting: A tutorial review. *International Journal of Forecasting*, 32(3):914–938, 2016. [p1]

R. J. Hyndman and Y. Khandakar. Automatic time series forecasting: the forecast package for R. *Journal of Statistical Software*, 26(3):1–22, 2008. URL https://www.jstatsoft.org/article/view/v027i03. [p2]

J. Kleissl. *Solar energy forecasting and resource assessment*. Academic Press, 2013. [p1]

J. Lemos-Vinasco, P. Bacher, and J. K. Møller. Probabilistic load forecasting considering temporal correlation: Online models for the prediction of households' electrical load. *Applied Energy*, 303: 117594, 2021. [p15]

L. Ljung and T. Söderström. *Theory and practice of recursive identification*. MIT press, 1983. [p2]

H. Madsen. *Time series analysis*. CRC Press, 2007. [p2, 5, 8, 10]

H. Madsen and P. Thyregod. *Introduction to general and generalized linear models*. CRC Press, 2010. [p8]

H. A. Nielsen and H. Madsen. Modelling the heat consumption in district heating systems using a grey-box approach. *Energy and buildings*, 38(1):63–71, 2006. [p1, 8]

H. A. Nielsen, H. Madsen, and T. S. Nielsen. Using quantile regression to extend an existing wind power forecasting system with probabilistic forecasts. *Wind Energy: An International Journal for Progress and Applications in Wind Power Conversion Technology*, 9(1-2):95–108, 2006. [p15]

T. S. Nielsen, H. A. Nielsen, and H. Madsen. Prediction of wind power using time-varying coefficient functions. In *Proceedings of the XV IFAC World Congress*, 2002. [p1]

G. Petris. An r package for dynamic linear models. *Journal of Statistical Software*, 36(1):1–16, 2010. [p2]

A. H. Sayed and T. Kailath. A state-space approach to adaptive rls filtering. *IEEE signal processing magazine*, 11(3):18–60, 1994. [p21]

S. L. Shah and W. R. Cluett. Recursive least squares based estimation schemes for self-tuning control. *The Canadian Journal of Chemical Engineering*, 69(1):89–96, 1991. [p21]

J. Siebert, J. Groß, and C. Schroth. A systematic review of python packages for time series analysis. *arXiv preprint arXiv:2104.07406*, 2021. [p2]

H. Spliid. A fast estimation method for the vector autoregressive moving average model with exogenous variables. *Journal of the American Statistical Association*, 78(384):843–849, 1983. [p2]

L. J. Tashman. Out-of-sample tests of forecasting accuracy: an analysis and review. *International journal of forecasting*, 16(4):437–450, 2000. [p9]

J. Tastu, P. Pinson, E. Kotwa, H. Madsen, and H. A. Nielsen. Spatio-temporal analysis and modelling of short-term wind power forecast errors. *Wind Energy*, 14(1):43–60, 2011. [p15]

S. Weisberg. *Applied linear regression*, volume 528. John Wiley & Sons, 2005. [p6]

H. Wickham. testthat: Get started with testing. *The R Journal*, 3:5–10, 2011. URL https://journal.r-project.org/archive/2011-1/RJournal_2011-1_Wickham.pdf. [p15]

Y. Xie. *Dynamic Documents with R and knitr*. Chapman and Hall/CRC, Boca Raton, Florida, 2nd edition, 2015. URL https://yihui.org/knitr/. ISBN 978-1498716963. [p15]

Y. Xie, J. Allaire, and G. Grolemund. *R Markdown: The Definitive Guide*. Chapman and Hall/CRC, Boca Raton, Florida, 2018. URL https://bookdown.org/yihui/rmarkdown. ISBN 9781138359338. [p15]

C. Yang, J. Qiao, Z. Ahmad, K. Nie, and L. Wang. Online sequential echo state network with sparse rls algorithm for time series prediction. *Neural Networks*, 118:32–42, 2019. [p2]

Forecast model notation

In this section it is shown how to write onlineforecast models in mathematical notation. Both in a full description and how to write a shorter summarized description. Note, that when variables are noted in bold font it indicates that they are multi-variate.

A model can be described in full detail as presented in the following.

The transformation stage

$$\begin{aligned}
\text{Intercept:} \quad & \mu_{t+k|t} && = 1 && (19)\\
\text{Periodic:} \quad & \boldsymbol{x}_{\text{per},t+k|t} && = f_{\text{fs}}(t; n_{\text{har}}) && (20)\\
\text{Part 1:} \quad & x_{1,t+k|t} && = H(B;a)u_{1,t+k|t} && (21)\\
\text{Part 2:} \quad & \boldsymbol{x}_{23,t+k|t} && = f_{\text{bs}}(u_{2,t+k|t}; n_{\text{deg}})u_{3,t+k|t} && (22)\\
\text{Part 3:} \quad & x_{4,t+k|t} && = u_{4,t} && (23)
\end{aligned}$$

and the regression stage

$$Y_{t+k|t} = \beta_{0,k}\mu_{t+k|t} + \boldsymbol{\beta}_{1,k}^T\boldsymbol{x}_{\text{per},t+k|t} + \beta_{2,k}x_{1,t+k|t} + \boldsymbol{\beta}_{3,k}^T\boldsymbol{x}_{23,t+k|t} + \beta_{4,k}x_{4,t+k|t} + \varepsilon_{t+k|t} \quad (24)$$

Thus the model inputs are:

- t is simply the time value.
- $u_{1,t+k|t}$ some forecast input (e.g. NWP variable).
- $u_{2,t+k|t}$ some forecast input (e.g. could be a deterministic value, e.g. time of day which is always know (the $|t$ could be omitted)).
- $u_{3,t+k|t}$ some forecast input (e.g. NWP variable).
- $u_{4,t}$ some value at time t (e.g. an observation variable).

The functions which maps the inputs (u's) to the regression inputs (x's) are:

- $f_{\text{fs}}(t; n_{\text{har}})$ is a function generating Fourier series of some implicit period length.
- $H(B;a)$ is a low-pass filter.
- $f_{\text{bs}}(u_{2,t+k|t}; n_{\text{deg}})$ is a function generating basis splines.

Their parameters are the transformation parameters:

- n_{har} is the number of harmonics.
- a is the low-pass filter coefficient.
- n_{deg} is the degrees of freedom of the spline function.

which must be set or optimized.

The regression coefficients are

$$\begin{aligned}
\boldsymbol{\beta}_k &= \begin{bmatrix}\beta_{0,k} & \boldsymbol{\beta}_{1,k}^T & \beta_{2,k} & \boldsymbol{\beta}_{3,k}^T & \beta_{4,k}\end{bmatrix}^T && (25)\\
&= \begin{bmatrix}\beta_{0,k} & \beta_{1,1,k} & \beta_{1,2,k} & \cdots & \beta_{1,2n_{\text{har}},k} & \beta_{2,k} & \beta_{3,1,k} & \beta_{3,2,k} & \cdots & \beta_{3,n_{\text{deg}},k} & \beta_{4,k}\end{bmatrix}^T && (26)
\end{aligned}$$

If the model is fitted with a recursive scheme, thus the coefficients change over time, it should be indicated by adding a t to the subscript, e.g. $\beta_{0,k,t}$. Furthermore, other parameters can exist, which can enter an optimization at the transformation stage, e.g. the RLS forgetting factor λ. The parameters which are optimized in the transformation stage should be presented together.

Specifying the model in all details can be cumbersome to include in some texts, so it makes sense to simplify the notation. When using a simpler notation, as suggested below, it should be stated, what is implicit (e.g. the regression stage). Referencing the present text should be sufficient when using a simpler notation. Naturally, all inputs, functions, etc., should be described in some way.

A model can be specified in a simpler way, e.g. the model above in one equation

$$Y_{t+k|t} = \beta_{0,k} + \boldsymbol{\beta}_{1,k}^T f_{\text{fs},k}(t; n_{\text{har}}) + \beta_{2,k} H_k(B; a) u_{1,t+k|t} + \boldsymbol{\beta}_{3,k}^T f_{\text{bs},k}(u_{2,t+k}; n_{\text{deg}}) u_{3,t+k|t} + \beta_{4,k} u_{4,t} + \varepsilon_{t+k|t} \tag{27}$$

or writing the regression stage implicitly by removing the regression coefficients where it is meaningful

$$Y_{t+k|t} = \mu_k + f_{\text{fs},k}(t; n_{\text{har}}) + H_k(B; a) u_{1,t+k|t} + f_{\text{bs},k}(u_{2,t+k|t}; n_{\text{deg}}) u_{3,t+k|t} + \beta_k u_{4,t} + \varepsilon_{t+k|t} \tag{28}$$

It is then implicit that the functions are different from the previous stated functions, since they include the regression coefficients. Again, if fitted with a recursive scheme, then it can be indicated by adding a t subscript, e.g. $f_{\text{fs},k,t}(t; n_{\text{har}})$.

To simplify further the k on the functions can be implicit

$$Y_{t+k|t} = \mu + f_{\text{fs}}(t; n_{\text{har}}) + H(B; a) u_{1,t+k|t} + f_{\text{bs}}(u_{2,t+k|t}; n_{\text{deg}}) u_{3,t+k|t} + \beta u_{4,t} + \varepsilon_{t+k|t} \tag{29}$$

and similarly the transformation parameters can be implicit

$$Y_{t+k|t} = \mu + f_{\text{fs}}(t) + H(B) u_{1,t+k|t} + f_{\text{bs}}(u_{2,t+k|t}) u_{3,t+k|t} + \beta u_{4,t} + \varepsilon_{t+k|t} \tag{30}$$

Then the functions and their parameters, and the fitting scheme (i.e. with either LS or RLS for each horizon) should be described in some other way.

Finally, the most simplified notation would be to even remove the time indexing

$$Y = \mu + f_{\text{fs}}(t) + H(B) u_1 + f_{\text{bs}}(u_2) u_3 + \beta u_4 + \varepsilon \tag{31}$$

after making clear how all the variables are defined.

Regression

In this section the two regression schemes implemented in onlineforecast are described. When fitting a model, thus estimating the regression coefficients, data from a period $t \in (1, 2, \ldots, n)$ is used and passed on to either: the `lm_fit()` function which implements the Least Squares (LS) scheme, or the `rls_fit()` function, which implements the Recursive Least Squares (RLS) scheme.

One important difference between the two implementations is that in the LS the coefficients are estimated using data from the entire period, thus they are constant during the period and the calculated predictions are "in-sample". This is opposed to the RLS, where the coefficients are updated through the period using only past data at each time t. In that case the coefficients vary over time and the calculated predictions are "out-of-sample".

This difference is explained in the following and indicated by subscripting the coefficient vector with t only for the RLS.

Least squares

The regression coefficients for the k'th horizon is set in the vector

$$\boldsymbol{\beta}_k = \begin{bmatrix} \beta_{0,k} & \beta_{1,k} & \ldots & \beta_{p,k} \end{bmatrix}^T \tag{32}$$

Note, that t is not included in the subscript.

The input data for the k horizon is the design matrix

$$\boldsymbol{X}_{k,t} = \begin{bmatrix} x_{0,1+k|1} & x_{1,1+k|1} & \cdots & x_{p,1+k|1} \\ x_{0,2+k|2} & x_{1,2+k|2} & \cdots & x_{p,2+k|2} \\ \vdots & \vdots & & \vdots \\ x_{0,n-1|n-1-k} & x_{1,n-1|n-1-k} & \cdots & x_{p,n-1|n-1-k} \\ x_{0,n|n-k} & x_{1,n|n-k} & \cdots & x_{p,n|n-k} \end{bmatrix} \tag{33}$$

The output observations are in the vector

$$\boldsymbol{y}_{k,n} = \begin{bmatrix} y_{1+k} & y_{2+k} & \cdots & y_{n-1} & y_n \end{bmatrix}^T \tag{34}$$

The LS estimates of the coefficients are

$$\hat{\boldsymbol{\beta}}_k = (\boldsymbol{X}_{k,n}\boldsymbol{X}_{k,n})^{-1}\boldsymbol{X}_{k,n}\boldsymbol{y}_{k,n} \tag{35}$$

The predictions are "in-sample" and calculated by

$$\hat{\boldsymbol{y}}_{k,n} = \boldsymbol{X}_{k,n}\hat{\boldsymbol{\beta}}_k \tag{36}$$

and returned when fitting a model with `lm_fit()`.

The estimated coefficients may now be used for "out-of-sample" prediction (for $t_{\text{new}} \geq n$), with the input

$$\boldsymbol{x}_{t_{\text{new}}+k|t_{\text{new}}} = \begin{bmatrix} x_{0,t_{\text{new}}+k|t_{\text{new}}} & x_{1,t_{\text{new}}+k|t_{\text{new}}} & \cdots & x_{p,t_{\text{new}}+k|t_{\text{new}}} \end{bmatrix}^T \tag{37}$$

by

$$\hat{y}_{t_{\text{new}}+k|t_{\text{new}}} = \boldsymbol{x}_{t_{\text{new}}+k|t_{\text{new}}}\hat{\boldsymbol{\beta}}_k \tag{38}$$

This can be done by providing new data to the `lm_predict()` function.

Recursive least squares

In the RLS scheme the coefficients are recursively updated through the period. Time t steps from 1 to n and in each step the "newly" obtained data at t is used for calculating updated coefficients. The coefficient vector has the same structure as for LS

$$\boldsymbol{\beta}_{k,t} = \begin{bmatrix} \beta_{0,k,t} & \beta_{1,k,t} & \cdots & \beta_{p,k,t} \end{bmatrix}^T \tag{39}$$

The only difference is that we now subscript with t because it varies over time.

Only the most recent input data at t (the row at t from the LS design matrix in Equation (33)) is used in each update

$$\boldsymbol{x}_{k,t} = \begin{bmatrix} x_{0,t|t-k} & x_{1,t|t-k} & \cdots & x_{p,t|t-k} \end{bmatrix}^T \tag{40}$$

and similarly the most recent output observation y_t. At each time t the coefficients are updated by

$$\boldsymbol{R}_{k,t} = \lambda\boldsymbol{R}_{k,t-1} + \boldsymbol{x}_{k,t}\boldsymbol{x}_{k,t}^T \tag{41}$$

$$\hat{\boldsymbol{\beta}}_{k,t} = \hat{\boldsymbol{\beta}}_{k,t-1} + \boldsymbol{R}_{k,t}^{-1}\boldsymbol{x}_{k,t}(y_t - \boldsymbol{x}_{k,t}^T\hat{\boldsymbol{\beta}}_{k,t-1}) \tag{42}$$

Hence, when applying RLS for data from the period $t \in (1, 2, \ldots, n)$ the RLS provides a new value of the coefficients for each time t (opposed to LS).

The predictions are calculated recursively as well by using the updated coefficients at each time t.

Given the inputs

$$\boldsymbol{x}_{t+k|t} = \begin{bmatrix} x_{0,t+k|t} & x_{1,t+k|t} & \dots & x_{p,t+k|t} \end{bmatrix}^T \tag{43}$$

the prediction is

$$\hat{y}_{t+k|t} = \boldsymbol{x}_{t+k|t}\hat{\boldsymbol{\beta}}_{k,t} \tag{44}$$

Only past data has been used when calculating the predictions through the period, hence they are "out-of-sample" predictions (these predictions are returned by `rls_fit()`).

The initial value of R is set simply set to a zero matrix with diagonal $1/10000$ and β set to a zero vector.

An alternative updating scheme, which is actually the implemented scheme (gives the same results as the scheme above), is the Kalman gain scheme (Sayed and Kailath, 1994), where matrix inversion is avoided

$$\boldsymbol{K}_{k,t} = \frac{\boldsymbol{P}_{k,t-1}\boldsymbol{x}_{k,t}}{\lambda + \boldsymbol{x}_{k,t}^T\boldsymbol{P}_{k,t-1}\boldsymbol{x}_{k,t}} \tag{45}$$

$$\hat{\boldsymbol{\beta}}_{k,t} = \hat{\boldsymbol{\beta}}_{k,t-1} + \boldsymbol{K}_{k,t}(y_t - \boldsymbol{x}_{k,t}^T\hat{\boldsymbol{\beta}}_{k,t-1}) \tag{46}$$

$$\boldsymbol{P}_{k,t} = \frac{1}{\lambda}\left(\boldsymbol{P}_{k,t-1} - \boldsymbol{K}_{k,t}\boldsymbol{x}_{k,t}^T\boldsymbol{P}_{k,t-1}\right) \tag{47}$$

This actually opens up the possibilities for self-tuned variable forgetting (Shah and Cluett, 1991).

Peder Bacher
Dynamical Systems
Department of Applied Mathematics and Computer Science
Technical University of Denmark
Asmussens Allé, Building 303B
2800 Kgs. Lyngby, Denmark
E-mail: pbac@dtu.dk
URL: https://www.compute.dtu.dk/english/research/research-sections/dynsys/

Hjörleifur G. Bergsteinsson
Dynamical Systems
Department of Applied Mathematics and Computer Science
Technical University of Denmark
Asmussens Allé, Building 303B
2800 Kgs. Lyngby, Denmark
E-mail: hgbe@dtu.dk
URL: https://www.compute.dtu.dk/english/research/research-sections/dynsys/

Linde Frölke
Dynamical Systems
Department of Applied Mathematics and Computer Science
Technical University of Denmark
Asmussens Allé, Building 303B
2800 Kgs. Lyngby, Denmark
E-mail: hgbe@dtu.dk
URL: https://www.compute.dtu.dk/english/research/research-sections/dynsys/

Mikkel L. Sørensen
Dynamical Systems
Department of Applied Mathematics and Computer Science
Technical University of Denmark
Asmussens Allé, Building 303B
2800 Kgs. Lyngby, Denmark
E-mail: mliso@dtu.dk
URL: https://www.compute.dtu.dk/english/research/research-sections/dynsys/

Julian Lemos-Vinasco

Dynamical Systems
Department of Applied Mathematics and Computer Science
Technical University of Denmark
Asmussens Allé, Building 303B
2800 Kgs. Lyngby, Denmark
E-mail: jlvi@dtu.dk
URL: https://www.compute.dtu.dk/english/research/research-sections/dynsys/

Jon Liisberg
Dynamical Systems
Department of Applied Mathematics and Computer Science
Technical University of Denmark
Asmussens Allé, Building 303B
2800 Kgs. Lyngby, Denmark
E-mail: jlvi@dtu.dk
URL: https://www.compute.dtu.dk/english/research/research-sections/dynsys/

Jan Kloppenborg Møller
Dynamical Systems
Department of Applied Mathematics and Computer Science
Technical University of Denmark
Asmussens Allé, Building 303B
2800 Kgs. Lyngby, Denmark
E-mail: jkmo@dtu.dk
URL: https://www.compute.dtu.dk/english/research/research-sections/dynsys/

Henrik Aalborg Nielsen
ENFOR A/S
Røjelskær 11, 3.
2840 Holte, Denmark
E-mail: han@enfor.dk
URL: https://www.enfor.dk

Henrik Madsen
Dynamical Systems
Department of Applied Mathematics and Computer Science
Technical University of Denmark
Asmussens Allé, Building 303B
2800 Kgs. Lyngby, Denmark
E-mail: hmad@dtu.dk
URL: https://www.compute.dtu.dk/english/research/research-sections/dynsys/

Paper J

HEAT LOAD FORECAST ACCURACY IMPROVEMENT WITH SPATIAL HIERARCHIES

Authors:
Hjörleifur G. Bergsteinsson, Mikkel Lindstrøm Sørensen, Jan Kloppenborg Møller and Henrik Madsen.

To Be Submitted

Spatial Hierarchical Forecasting for Heat Load

Hjörleifur G. Bergsteinsson[a,*], Mikkel Lindstrøm Sørensen[a], Jan Kloppenborg Møller[a], Henrik Madsen[a]

[a]*Technical University of Denmark, Denmark*

Abstract

District heating is an efficient method of distributing heat in densely populated areas at a low cost. The heat is usually produced at central production plants and then distributed to consumers through large networks of pipes. However, district heating is gradually becoming more decentralised with additional heat sources, e.g. heat pumps, solar thermal farms, and industrial waste heat connected to the network. Therefore, the system is changing from a system with centralised heat sources to a more decentralised system with several different heat sources within the network, including also still a large production area. Operationally this is more complex than the previous setup, especially in terms of temperature optimisation. Typically, the temperature must be adjusted for each area in order to work efficiently with the decentralised heat sources, so a forecast of the local heat demand is required. It is relatively easy to make a forecast for each area, but they are usually made independently and are therefore not necessarily coherent. This paper proposes a methodology to spatially reconcile hierarchies of individual localised heat demand forecasts with a coherency constraint. This results in coherent reconciled forecasts. The suggested methodology has been demonstrated in several case studies and has been shown to be able to enhance forecast accuracy significantly. This paper will use two case studies to illustrate the proposed method. One case study has a few areas, while the other case study will have more areas, and here it is proposed to add a new level of aggregation to the hierarchy to increase accuracy. The results in this paper show that the reconciled forecast, where information is shared between areas through the spatial hierarchy, improves forecast accuracy by 1% to 20%, depending on the prediction horizon.

Keywords: Forecast Reconciliation, Adaptive Estimator, Spatial Hierarchy, District Heating

1. Introduction

New methods for the operation of district heating are developing quite rapidly. The supply temperature in the network needs to be reduced to increase efficiency and maximise the flexibility potential of district heating [1]. Improving the efficiency of district heating and maximising flexibility are important as they can be seen as a solution to the occasional electricity surplus that is expected to grow over time with the increase in the share of intermittent renewable energy sources (RES) [2]. District heating can increase the efficiency of the overall energy system due to its unique ability to store heat over long periods of time and its high efficiency in converting electricity into heat, e.g. through heat pumps [3]. In addition, due to the increasing number of intermittent heat sources, more thermal energy storage (TES) systems are being installed [4, 5]. These new changes with additional heat sources and TES systems lead to more complex systems than traditional district heating systems. Future district heating systems will become more decentralised, with more RES sources and a closer coupling with the power sector.

However, to make this transition more feasible and efficient, the supply temperature of the network needs to be lowered to make the integration of the new heat sources, e.g. heat pumps, into the district heating system more effective [3, 6]. In addition, more detailed knowledge of the system needs to be acquired as district heating becomes more decentralised with heat units and TES in the network. For instance, if the district heating system consists of a large production area with transmission lines and two heat exchangers supplying two distribution networks for heat. The distribution networks could be equipped with heat pumps and TES. Therefore, the supply temperature for each distribution network area and the transmission line needs to be optimised precisely in order for the system to work efficiently. Temperature optimisation requires information about future demand to obtain optimal supply temperature set points [7, 8]. It takes time to send one unit of hot water from a production plant to the users (the time delay). Therefore, information about the future heat load and network characteristics (e.g., time delay, temperature loss) are needed to ensure that sufficient temperature reaches the consumers [9]. Therefore, an accurate heat load forecast must be available. In practice, however, district heating companies usually only produce a forecast for a large area, or they distribute a forecast for a large area to smaller areas by scaling it according to a heat consumption ratio determined in the past [10]. This practice introduces a significant error and bias in the consumption forecast, as the dynamics of the heat load can change rapidly due to climate or changing behaviour.

Accurate forecasts of the local heat load are therefore important for district heating utilities for controlling the supply

*Corresponding Author
Email address: hgbe@dtu.dk
Postal Address: Anker Engelunds vej 1, Building 101A, 2800 Kongens Lyngby, Denmark

temperature. However, the district heating utility needs a total heat load forecast for its production optimisation and planning. Therefore, forecasts for the total heat load and the local areas are needed. Ideally, these forecasts must be coherent, i.e. the aggregation of the local forecast must be consistent with the total forecast. However, this is not guaranteed if the forecasts are produced independently of each other. Coherence can be achieved by reconciling forecast hierarchies [11]. Hierarchy constraints ensure that forecasts are coherent but have also been shown to lead to more accurate forecasts [12] due to a transfer of information between hierarchies. Constraints between hierarchies are used in the reconciliation process, where it uses incoherent forecasts and produces coherent forecasts, which are called reconciled forecasts. This not only solves the problem of incoherent forecasts but also improves forecast accuracy [13].

1.1. Heat Load Characteristics and Forecasting

Energy forecasting is evolving rapidly, especially for renewable energy sources [14]. Energy forecasts are valuable because they are needed to optimise future strategies for the sectors. For instance, in a district heating system, it could be decided if the TES system is charged when electricity prices are expected to be low during low heat demand or if the Combined Heat and Power plant (CHP) is operated at full capacity when electricity prices are expected to be high. Even if the heat demand is low during the operation of the CHP, the TES system could be charged. Therefore, a heat load forecast is needed to support decision-making and production optimisation when selecting heat units [15], charging/discharging the TES system [4, 16], or temperature optimisation [8, 17, 18].

Heat load forecast is therefore often studied, and its popularity has increased in recent decades due to the growing importance of unlocking the flexibility of district heating for energy systems with a high share of fluctuating renewables. It is important to understand the characteristics of heat in order to build an appropriate and robust forecasting model. Heat load can be divided into two main categories: Space heating and domestic hot water usage with the addition of heat losses in the systems. Space heating is about heating the consumer's home to maintain thermal comfort, while domestic hot water usage is about meeting the consumer's other needs, such as showering. The heat demand can also be divided into a physical and a social heat load, as described in Gadd and Werner [19]. The physical part is determined by the weather and the thermal insulation of the building, where the ambient air temperature is a significant variable for driving the heat consumption [20]. The insulation of the building envelope acts as a resistance in the heat transfer between the ambient air and indoor temperature, and the heat consumption is used to maintain a certain indoor temperature. Other weather variables also influence heat consumption, e.g. solar radiation and wind speed. Solar radiation reduces heat consumption as the solar beams enter through the window and heat the floor. The effect depends, for example, on the strength of the radiation, the angle of the rays and the size of the window. The wind influences the natural ventilation in buildings, and a higher wind speed increases heat consumption. The social effect affects heat consumption when hot water is used, e.g. when taking a shower in the morning before work. The heat load is not stationary, as heat consumption also changes due to weather changes. For example, as the ambient air temperature rises at the beginning of summer, less space heating is then needed until a certain threshold temperature is reached, at which space heating is no longer required. Social behaviour also changes over time, e.g. during the summer holidays when fewer people are at home. More detailed information about physical and social heat load dependencies can be found in Nielsen and Madsen [20] and Gadd and Werner [19].

Knowledge as outlined above can then be used to identify significant input variables for the model and how they should be treated (e.g., instantaneous or transformed). Several prediction model methods have been proposed that can be used to predict heat load. Dotzauer [21] propose a linear regression model where the relationship between the heat and the input variables is investigated, e.g. the dependence on the ambient air temperature is treated by creating a piecewise linear function. Dahl et al. [22] also uses a linear regression model with seasonal lags. The inputs are the weather forecasts, where the aim was to investigate the gain from uncertainty in heat load forecast by using an ensemble of weather forecasts as inputs. Grosswindhagera et al. [23] proposes using Seasonal Autoregressive Integrated Moving Average (SARIMA) formulated in a state-space form to produce online forecasts using the Kalman filter. These methods are simple and neglect the non-linearity and non-stationarity of the heat load. In Dahl et al. [24] and Idowu et al. [25], linear regression, neural network and support vector regression (SVR) are compared, which in Idowu et al. [25] also is compared with regression tree. The SVR method provides the best predictive performance with the lowest prediction error. The SVR and the neural network can deal with non-linearity because they are nonlinear models, whereas the SVR can use nonlinear kernel values to model the relationship, and the neural network has nonlinear activation functions. However, these models have problems with the non-stationary heat load and would therefore often need to be recalibrated frequently. Nielsen and Madsen [26] propose to use the grey-box method to predict the heat load where physical insights and statistical methods are used to find an optimal model. A linear regression model is used for predicting heat load, where the inputs are transformed to deal with non-linearity. Also, the coefficients are estimated using recursive least squares (RLS) with exponential forgetting to deal with non-stationarity, allowing the coefficients to adapt to changes as new observations become available and exponentially down-weighting the older ones. Recurrent neural networks (RNN) and convolutional neural networks, and long short-term memory (CNN-LSTM) have been proposed to extend the neural network to deal with non-stationarity, as shown in Kato et al. [27] and Song et al. [28].

1.2. Hierarchical Forecasting and Reconciliation

Reconciliation is the process of making forecasts coherent according to their hierarchical structure. The individual forecasts are usually incoherent, and in many applications, these forecasts need to be coherent as specified by their hierarchy. Energy production planning, for example, needs coherent fore-

casts to make optimal decisions based on forecasts for different horizons [29]. The reconciliation process not only makes the individual base forecast coherent but has also been shown to increase the accuracy of the forecasts [12, 13, 29]. Van Erven and Cugliari [30] prove that reconciliation forecasts should perform at least as well as the base forecast on average. However, this depends on the quality of the weights in the reconciliation process, as shown in Nystrup et al. [31], where it is shown that no improvements can be achieved if the errors of the base forecasts, which are used to estimate the weights, are too highly correlated. Therefore, the independent base forecast at the different levels in the hierarchy cannot come from the same model using the same information (e.g. input variables) at all levels, as they have nothing to exchange between levels. Hollyman et al. [32] show that the reconciliation process is a special case of a combination forecasting method by reformulating it into a combination of direct forecasts using linear coherent constraints.

Several recent studies have demonstrated the improvement in accuracy using the reconciliation process, and the benefits of coherent forecasts in the energy sector [33]. Nystrup et al. [29] demonstrates the improvement for short-term electricity load forecasts. Jeon et al. [34] demonstrate the reconciliation process for probabilistic forecasting of wind power and electric load to ensure coherence, resulting in higher accuracy of the forecasts. Ber [12] demonstrate that using temporal hierarchies with the hierarchy of all-natural levels from one hour ahead to daily time resolutions leads to 15% higher accuracy compared with state-of-the-art hourly heat demand forecasting.

Thus, reconciling an independent base forecast of heat demand between different areas inside the district heating network can result in accuracy improvements for all forecasts. These improvements will then lead to higher cost savings for the operation of the system since a more precise forecast will improve the production optimization, temperature optimization and the operation of the decentralized heat sources, e.g. heat pumps. It will also increase the possibility of reducing supply temperature in the network, reducing cost, and reducing heat losses in the system [8]. This will also increase the efficiency of the heat sources in the system, e.g. cogeneration plant [35]. More importantly, it will increase the feasibility of power-to-heat units, and the flexibility of district heating is then leading to a flexibility of the power system [3].

1.3. Contribution

The purpose of this study is to develop a spatial hierarchical framework that handles the dynamics of heat load to improve the accuracy of individual forecasts for each area in the hierarchy for district heating demand. The proposed method makes individual base forecasts coherent. It improves accuracy by using the reconciliation process with an empirical covariance matrix estimator estimated from the base forecast errors. Due to the nature of the heat demand observations, it is proposed to use a recursive and adaptive covariance estimator, i.e. the estimator can be easily updated and weights down previous observations to give more importance to new observations. We propose using exponential smoothing with a forgetting factor to estimate the covariance matrix as suggested in Ber [12].

We propose to estimate an individual covariance matrix for each prediction horizon and use an optimal forgetting factor for each horizon. We document the accuracy improvements by using a covariance matrix estimated for each horizon compared to only using forecast errors from one-step-ahead predictions. Also, the accuracy difference between using the same forgetting factor for all horizons and the optimal factor for each horizon is also investigated. These results are demonstrated using two different case studies, one with few areas and another with many areas. For the case study with many areas, it was possible to add a new level, and it is shown that adding a new level to the hierarchy increases the accuracy. An operational state-of-the-art heat load forecasting system is used to compare and illustrate the accuracy improvements possible through the reconciliation process. The paper also includes a simulation study on heat load forecasting to investigate the accuracy improvements when the spatial hierarchy is applied in the reconciliation process.

The paper is organised as follows. The data from the two case studies are presented in Section 2. Section 3 discusses the base forecasting model used by the authors to forecast heat demand. The spatial hierarchies and the reconciliation process are also presented. The results are presented in Section 4 and discussed in Section 6. The paper is then concluded in Section 7.

2. Data

The data used in this study is heat load from two district heating utilities in Denmark; *Fjernvarme Fyn* and *Brønderslev Forsyning*. Both utilities produce and deliver heat to their consumers in Fyn and Brønderslev, respectively. *Fjernvarme Fyn*'s heat load data consists of total heat consumption and the twelve areas into which the total heat consumption is divided, while *Brønderslev Forsyning* consists of total consumption and three areas. The data for both utilities have an hourly resolution from 1 January 2019 to 1 December 2021 for Fyn, while Brønderslev has data from 1 February 2020 to 1 June 2022. The heat load is denoted by

$$\left\{Y_t^{\mathrm{i}};\ t = 1, \ldots, T,\ i = \mathrm{Total}, \mathrm{Area}_1, \ldots, \mathrm{Area}_{\mathrm{g}}\right\}, \qquad (1)$$

where T is the total number of observations, and g is the corresponding total number of areas. The measurements are visualised in Figure A.19 for Fyn over the whole period and in Figure A.18 for Brønderslev.

The heat consumption from both utilities shows seasonal variations in heat load as described in Gadd and Werner [19], correlating with ambient air temperature, i.e. high in cold periods and low in warm periods. The daily heat demand profile also correlates with ambient air temperature, but the social behaviour of each area also shapes the profile. In addition, there are some missing values in all series. The magnitude of demand is not the same in all areas, and the volatility is also different. The dynamics in the areas at Fyn, as seen in Figure A.19, differ more than the dynamics between the areas in Brønderslev, as seen in Figure A.18. The areas in Brønderslev have similar

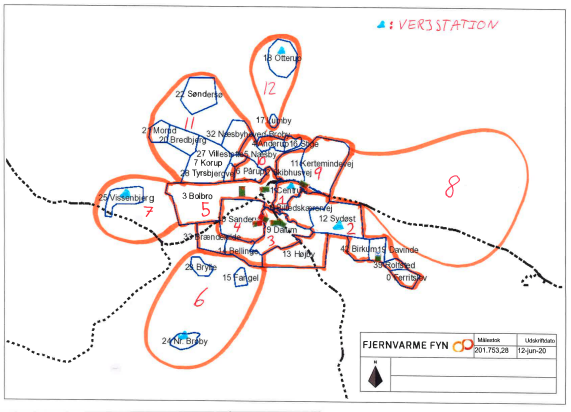

Figure 1: Layout of the areas at *Fjernvarme Fyn*.

dynamics with slight differences in magnitude but almost identical. The areas on Fyn have both different dynamics and magnitudes. However, some areas are similar and could therefore possibly be grouped together. Figure 1 shows the individual areas that *Fjernvarme Fyn* supplies heat to.

The total heat load in this study is not the total heat load for the utilities, as the heat loss is discarded here. This approach is used since otherwise, heat loss would also have to be included as an area in the hierarchy, and a forecasting model would have to be created to predict heat loss. It can be difficult to accurately predict heat loss as it varies depending on the pipe's supply temperature and flow and the pipes' surrounding temperature. Hence, it was decided not to consider this in this study for simplicity. Therefore, the total heat load in this study is only the aggregated load from all areas.

2.1. Operational heat load forecasts

Both *Brønderslev Forsyning* and *Fjernvarme Fyn* use forecast with an hourly resolution for several forecast horizons, which are used to optimise their operations. A commercial forecast provider delivers the hourly heat load forecasts to both utilities. The HeatFor™[1] solution provides the heat load forecasts. The heat load forecast, $\hat{Y}^{\mathrm{i}}_{t+k|t}$, is updated every hour for k steps in advance for each group i, as shown below,

$$\left\{\hat{Y}^{\mathrm{i}}_{t+k|t}, t = 1,\dots,T,\ k = 1,\dots,K, i = \text{total}, \text{area}_1,\dots,\text{area}_{\mathrm{g}}\right\}. \quad (2)$$

State-of-the-art operational forecasts are referred to throughout the text as *operational* base forecasts. This work aims to improve the accuracy of the one-step to 24-step forecast. Thus, the forecast horizon of interest for improving hourly forecast accuracy is $k = 1, 2, \dots, 24$.

2.2. Numerical Weather Prediction

The numerical weather predictions (NWPs) used as input to the forecast models were provided by the MetFor™[2]; i.e. by the same commercial forecast provider as for the heat load. The NWPs consist of climate variables with an hourly resolution, updated every hour and available for forecasting heat load $\hat{Y}^{\mathrm{i}}_{t+k|t}$.

An example of an NWP for the kth prediction horizon is the predicted ambient temperature [in °C] denoted by

$$\left\{T^{\mathrm{a,NWP}}_{t+k|t}, t = 1,\dots,T,\ k = 1,\dots,K\right\}. \quad (3)$$

The NWPs are used to produce the heat load forecasts required in addition to the *operational* base forecast.

3. Methods

This section introduces the methods used for generating the forecasts for the study. In section 3.1, the forecasting methodology used in creating the base forecasts is presented, which are needed additionally with the *operational* base forecasts. A grey-box method is used to generate the base forecasts, which will be referred to as *simple* base forecast throughout the text. Section 3.2 presents the methods of hierarchical forecasting and the linear constraints they impose on spatial aggregation. Lastly, Section 3.3 defines the reconciliation process, which is done when the base forecast does not satisfy the coherency constraints by the hierarchy.

3.1. Heat Load Base Forecast

Since the *operational* base forecast is not available for each level of aggregation, additional forecasts for heat load are needed. The forecasting method used here is a linear regression model where the coefficients are estimated using adaptive RLS with a forgetting factor, where past observations are exponentially weighted downwards as suggested by Ljung and Söderström [36]. This forecasting model has shown promising results for several energy forecast studies. For instance, Bacher et al. [37] and Bacher et al. [38] use the proposed method to predict the electricity generation from PV systems and heat load for single-family homes, Ber [12] to predict the heat load for a district heating system to be used in temporal hierarchical forecasting, RAS [39] to predict the electrical load for supermarket refrigeration, and Jónsson et al. [40] to predict electricity spot prices.

The forecasts are made using the `R`- package, `onlineforecast` [41]. The package provides an ideal forecasting framework for heat load forecasting as it provides the tools to deal with the non-stationarity and non-linearity of the heat load time series. Thus, account for the time-varying dynamics and the non-linear relationship between the load and input variables, such as NWP and social behaviour. The forecasting methodology of the package is only briefly introduced here. For a more detailed introduction, see Bacher et al. [41] and Bergsteinsson et al. [42]. The package is based on a regression model that models the output variable as a linear combination of the input variables. However, it also contains a possibility for mapping the input variables to handle non-stationarities and non-linearities of the output and input variables. The method consists of a two-stage modelling procedure as proposed in [43] and [39]. This procedure consists of a *Transformation Stage* and a *Regression Stage*. In the *Transformation Stage*, the input

[1] `https://enfor.dk/services/heatfor/`
[2] `https://enfor.dk/services/metfor/`

variables are transformed either directly by a function or non-parametrically, e.g. using splines. A linear model is then created using the transformed data to predict the heat load. Subsequently, in the *Regression Stage*, the model's coefficients are estimated recursively using the RLS method. The two-stage modelling procedure thus deals with non-linearity and non-stationarity using the RLS method by recursively estimating the model coefficients and transforming the input variables appropriately.

3.2. Hierarchical Time Series Forecasting

Hierarchies for time series define the structure from the most granulated series to the most aggregated. Hierarchies are divided into three types; temporal hierarchies (see e.g. [44]), spatial hierarchies, in literature sometimes called cross-sectional hierarchies (see e.g. [45]), and spatio-temporal hierarchies which combines the two [46].

In *spatial* hierarchies, the structure is related to either spatial aggregation or some grouping aggregation, e.g. Yang et al. [47] uses a geographical hierarchy to make PV forecasts coherent across the whole energy system from the inverters on the PV, subsystems, PV plants and all the way up to the transmission zones. All these different groups require individual base forecasts to operate their system and are therefore not necessarily coherent. They argue that if the forecasts are constrained to be coherent, then the decision-makers in the electricity grid can improve their planning of the grid due to the intermittent individual power injection to the grid. Hence, a more detailed and coherent overview of when and where the power will enter the grid is obtained.

Temporal hierarchies consider aggregation in the time domain. An example is described in Ber [12] reconciles temporal hierarchies forecasts to make different heat load resolution forecasts coherent. This allows for better alignment of the decision-making for production scheduling which has a higher temporal resolution (e.g. day ahead), while temperature optimization needs heat load forecasts on lower resolution (e.g. hourly). Yagli et al. [48] demonstrates a spatio-temporal hierarchy of PV systems, where the focus is on simultaneously making the forecasts spatially and temporally coherent. Petropoulos et al. [49] discuss this more in detail and refer to multiple studies that use spatial, temporal or both in their research.

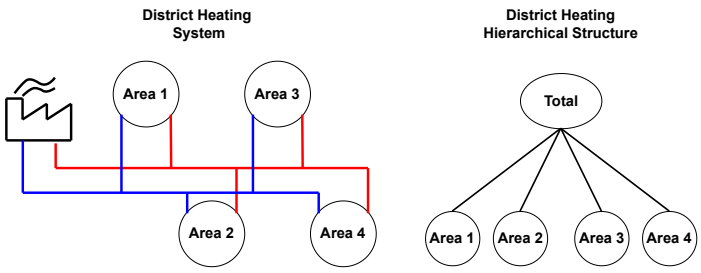

Figure 2: Example of a simple spatial hierarchy structure for heat load forecasting.

This study considers spatial hierarchies of the areas within district heating networks. In these systems, the total heat consumption should be coherent with the area forecasts, i.e. the aggregation of area forecasts should equal the forecast total consumption. For example, a district heating system that has one production plant, one transmission line and four distribution systems. Assuming no heat loss, the heat produced at the production plant is the total heat consumption, and the four distribution systems are the bottom areas. Figure 2 illustrates this type of district heating system on the left and shows the corresponding hierarchy on the right. Thus, the total of the four areas should aggregate to the total demand.

A summation matrix $\boldsymbol{S}$ is used to describe the structure of a hierarchy. The general definition of the summation matrix for any balanced hierarchy structure (as shown in Figure 3) is given in Nystrup et al. [29] as

$$\boldsymbol{S} = \begin{bmatrix} I_{\mathrm{m}/\ell_1} \otimes 1_{\ell_1}^{\mathrm{T}} \\ \vdots \\ I_{\mathrm{m}/\ell_{\mathrm{L}}} \otimes 1_{\ell_{\mathrm{L}}}^{\mathrm{T}} \end{bmatrix} \tag{4}$$

where $\otimes$ denotes the Kronecker product, $I_{m/\ell}$ is an identity matrix of order m/ℓ, and 1_ℓ is an $\ell-$vector of ones. The aggregation levels are a factor of m, which is the sampling frequency of the lowest level. In the example above, $\ell_1 = m$, $\ell_L = 1$, and m/ℓ is the number of observations at the aggregation level ℓ. Using the hierarchy structure in Figure 2 to illustrate this, the hierarchy has aggregation levels $\ell_1 = 4$ and $\ell_2 = 1$ with $m = 4$ and the dimension of the base forecasts in the structure is $n = 5$. This results in the following summation matrix corresponding to Figure 2

$$\boldsymbol{S} = \begin{bmatrix} 1 & 1 & 1 & 1 \\ 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 \end{bmatrix} \tag{5}$$

It is also possible to add additional layers to the hierarchy. This adds to the complexity of the problem, but assuming that the forecast errors are not too highly correlated, then the added information should improve the accuracy of the reconciled forecasts [31]. Figure 3 shows the hierarchy from Figure 2 with an additional layer.

Here a layer has been added between the area forecasts and the total forecast, which aggregates the area forecasts in pairs; however non-symmetrical hierarchies are equally valid. Adding a layer in this manner results in the following summation matrix corresponding to Figure 3.

$$\boldsymbol{S} = \begin{bmatrix} 1 & 1 & 1 & 1 \\ 1 & 1 & 0 & 0 \\ 0 & 0 & 1 & 1 \\ 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 \end{bmatrix} \tag{6}$$

This requires $m = 4$ bottom-level forecasts and $n = 7$ base forecasts.

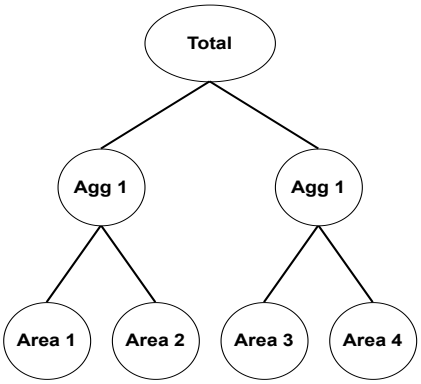

Figure 3: Example of hierarchy with added aggregation layer between area and total forecasts

3.3. Reconciling Forecasts

The reconciliation process describes the transformation of individual time series to be coherent according to the defined hierarchical structure. In relation to heat demand forecasts, it results in making them coherent such that the aggregation of the lower levels match the higher levels and vice versa. Hence, the process is defined by linear constraints. Thus, for the reconciliation process, a projection matrix is needed that projects the individual base forecast to a coherent subspace as defined by the linear constraints [50]. For example, making independent base forecast reconciled using the bottom-up method,

$$\tilde{Y} = \boldsymbol{SG}\hat{Y}, \tag{7}$$

$$\boldsymbol{G} = \left[0_{m\times(n-m)}|I_m\right], \tag{8}$$

where $\tilde{Y}$ is the reconciled forecast, $\boldsymbol{G}$ is a extraction matrix of order $m \times n$, which extracts the m bottom-level forecasts, $\hat{Y}$ is the base forecasts. Here, $\boldsymbol{G}$ extracts the bottom forecasts and aggregates them up according to the summation matrix and the $\boldsymbol{SG}$ is the projection matrix in the reconciliation process. The $\boldsymbol{G}$ matrix is the mapping of the base forecasts to the bottom-level reconciled forecasts therefore the accuracy of the reconciled forecasts depends on $\boldsymbol{G}$. Simply choosing $\boldsymbol{G}$ according to the bottom-up method is inefficient as it disregards all information in the higher levels.

Therefore, a more optimal approach is needed where all information in the hierarchy is shared between the forecasts using the $\boldsymbol{G}$ and hence the projection matrix to reduce the error of the reconciled forecast. Hyndman et al. [11] propose a regression approach to estimate the mapping matrix $\boldsymbol{G}$ using generalized least square estimation, minimising the coherency errors, i.e. the error between the base forecast and the reconciled forecast, subject to the coherency constraint. In this formulation, the base forecasts are then written in regression form,

$$\hat{Y}_{t+k|t} = \boldsymbol{S}\beta(k) + \epsilon(k), \tag{9}$$

where $\beta(k) = \mathbf{E}[Y_{\ell,t+k}|Y = y_1, \ldots, y_t]$ is the unknown conditional mean of the future values of the most granular observed series, i.e. the reconciled forecasts. The $\epsilon(k)$ represents the error between the base forecasts and their expected value, the coherency error $\hat{Y}-\tilde{Y}$. The error $\epsilon(k)$ is assumed to have zero mean and covariance matrix, Σ. Hence, the generalized least squares estimation of $\beta(k)$ in Eq. (9). If Σ is assumed to be known and the base forecasts are unbiased, the reconciled forecasts can be estimated by

$$\tilde{y} = S(S^{\mathrm{T}}\Sigma^{-1}S)^{-1}S^{\mathrm{T}}\Sigma^{-1}\hat{y}, \tag{10}$$

where the matrix $\boldsymbol{G} = (S^{\mathrm{T}}\Sigma^{-1}S)^{-1}S^{\mathrm{T}}\Sigma^{-1}$. The issue with this approach is that the covariance of the coherency errors Σ is not identifiable, as shown by Wickramasuriya et al. [13]. Consequently, multiple authors have suggested possible alternatives. Hyndman et al. [11] argued that Σ would be difficult to estimate and replaced it with an identity matrix, thus placing equal weights on all base forecasts. Hyndman et al. [45] proposed estimating Σ by using weighted least squares where the variance of the one-step ahead base forecast is used in place of Σ. Athanasopoulos et al. [44] proposed three different structures of the estimator based on the in-sample base forecast errors. However, all of the estimators disregard variance between groups and levels in the hierarchy. Wickramasuriya et al. [13] introduce what they call the minimum trace (MinT) reconciliation, which uses the full variance-covariance matrix of the base forecast errors. In Nystrup et al. [29], it is also proposed to use the full variance-covariance matrix to maximise the accuracy potential of the reconciliation process. While other methods have since been proposed, e.g. the Combined Conditional Coherent forecasts (CCC) by Hollyman et al. [32], the MinT has seemingly become somewhat of a go-to method for forecast reconciliation. Hence, the full covariance of the base forecast errors will be used for the reconciliation process in this study.

Operational forecasts often need to be re-estimated multiple times per day according to the operation it is needed for. Therefore, the reconciled forecast also needs to be re-estimated with the same frequency as the base forecast. This requires fast computations such that the forecasts are available when they are needed. For example, temperature optimization in district heating networks is usually done on an hourly resolution based on new heat demand observations for the past hour. The data needs to be sent between servers and needs to be quality-checked before it is used. Input variables also need to be available as soon as possible, but they also take time. Then the base forecast needs to be updated and used e.g., for recalculation of the optimal set point of the temperature optimization. Hence, the computation of the reconciliation forecasts needs to be fast.

To ensure rapid computations Ber [12] propose using the exponential smoothing method to estimate the covariance,

$$\hat{\Sigma}_{\mathrm{t}} = \lambda\hat{\Sigma}_{\mathrm{t-1}} + (1-\lambda)e_{\mathrm{t}}e_t^{\mathrm{T}}. \tag{11}$$

where e_{t} is the newest base forecast error at time t, λ is the forgetting factor and $\hat{\Sigma}_{\mathrm{t-1}}$ is the previously estimated covariance matrix. Hence, updating the estimator with new information is quick, and only the previous estimator and current error need to be stored. This method also makes the estimator adaptive due to the forgetting factor where past information is exponentially down-weighted, i.e. newer observations have more influence on the estimation. This is useful, for instance, when the system is non-stationary, and forecasts need to adapt quickly to changes

that occur. We therefore propose to estimate the covariance estimator using the exponential smoothing method, as the forecasts in this paper are used operationally and are updated on an hourly basis to handle the non-stationary heat load.

The accuracy improvements achieved using the proposed reconciliation process will be demonstrated using the Relative Root Mean Squared Error (RRMSE),

$$\text{RRMSE} = \frac{\text{RMSE}}{\text{RMSE}^{\text{base}}} - 1, \tag{12}$$

where $\text{RMSE}^{\text{base}}$ is the Root Mean Squared Error (RMSE) of the base forecasts, and RMSE is of the reconciled forecast. The result demonstrates either improvement or decline in performance compared to the base forecast, where negative values correspond to improvements in accuracy over the base forecast. The RRMSE is frequently used to compare forecasts between different methods due to its interpretability of the relative measure [51].

4. Results

This section presents the results of applying the reconciliation process to the heat demand forecasts for the spatial hierarchy. Section 4.1 presents an example of the *simple* base forecast model and demonstrates its performance. Section 4.2 shows the optimisation of the forgetting factor for the covariance estimator. The forgetting factor must be optimised over an in-sample period to achieve optimal improvement in accuracy over the out-of-sample period. In Section 4.3 the improvements for the case study *Brønderslev Forsyning* are presented. Furthermore, in Section 4.3.1, the difference in accuracy for the reconciliation process when using the one-step prediction errors versus using the corresponding prediction horizon for each horizon for the estimation of the covariance matrix is investigated. Finally, Section 4.4 presents the improvements in *Fjernvarme Fyn* and demonstrates the difference in accuracy when using different hierarchy structures.

4.1. Base Forecast

The model identification and validation will follow the steps described in Bacher et al. [38], find the optimal model by extending the model by investigating if there are any missing dynamics left by residual analysis of the one-step error as well as investigating if the errors and other inputs are correlated using the cross-correlation function (CCF) for each modelling step. The CCF is used to identify any remaining dynamics that an input variable can explain or if different transformations for the current inputs can be used to improve the model. The models are thereby constructed by using the forward selection principle, i.e. adding new inputs sequentially and examining which input had the lowest error score during the scoring-period period. The same approach will be applied here to find the optimal models, and using the Root Mean Squared Error (RMSE) will be used to compare the performance of models. However, this process will not be shown in this work as explaining the process is tedious. Similar models have been proposed. For instance, Ber [12] uses the same method to establish a forecasting model for each temporal level.

An example of a model created for this study is the total forecast in *Brønderslev Forsyning*. The training period is from 2020-02-01 to 2020-05-01, with the first month used as a "burn-in" period, i.e., discarded when calculating the error score. The scoring period is then 2020-03-01 to 2020-05-01. The final model is,

$$\begin{aligned}\hat{Y}_{t+k|t} = \theta_{0,k} + \theta_{1,k} Y_t + \theta_{2,k} H_{a_\text{T}}(q) T^{\text{a,NWP}}_{t+k|t} + \\ \theta_{3,k} H_{a_\text{W}}(q) W^{\text{a,NWP}}_{t+k|t} + \theta_{4,k} H_{a_\text{G}}(q) G^{\text{a,NWP}}_{t+k|t} + \\ \theta_{5,k} T^{\text{a,NWP}}_{t+k|t} + \theta_{6,k} W^{\text{a,NWP}}_{t+k|t} + \theta_{5,k} G^{\text{a,NWP}}_{t+k|t},\end{aligned} \tag{13}$$

where the filters are of the from

$$H_a(q) = \frac{1-a}{1-aq^{-1}}, \tag{14}$$

with the backward shift operator ($[q^{-1}]$). The offline coefficients (i.e. the constants for the model which are estimated in the in-sample period) of the model are: the time constants for ambient air temperature ($[a_\text{T}]$), wind speed ($[a_\text{W}]$), global radiation ($[a_\text{G}]$) and the forgetting factor ($[\lambda]$). The off-line coefficients are estimated during the in-sample period from 01/02/2020 to 01/05/ 2020, removing the first month in the scoring period, i.e. the first month is treated as a burn-in period, while the other errors are used to calculate the RMSE, i.e. the scoring period. The forecast horizons used for the estimation of the offline coefficients are $k = \{3, 6, 12, 18, 24\}$. The offline coefficient estimates are presented in Table 1. The θ coefficients are estimated recursively as new observations become available, see Bacher et al. [37] for further details.

Table 1: Estimates of the offline coefficients for the total heat demand forecast at *Brønderslev Forsyning*

a_T	a_W	a_G	λ
0.937	0.816	0.980	0.992

The RMSE for each horizon is shown in Figure 4. The top left and right plots show the performance of the model for the in-sample and out-of-sample periods. The estimation of the offline coefficients shows that they are valid estimates, as the RMSE for the out-of-sample period is similar to that for the in-sample period and shows no significant increase in error. The bottom graph in Figure 4 shows a realisation of the model prediction to demonstrate the performance of the model prediction. The grey line in the plot indicates when the prediction was made, at time t = 2020-11-09 23:00 for the next 24 hours.

4.2. Optimization of hyperparameters

The forgetting factor λ is used to update the empirical covariance estimator and needs to be optimised, as discussed in Section 3.3, to achieve the optimal accuracy improvements. The forgetting factor is determined by minimising the RMSE of the reconciliation forecasts for the total and the areas in the hierarchy during the training period. It is estimated either as one optimal forgetting factor for all horizons by minimising the RMSE

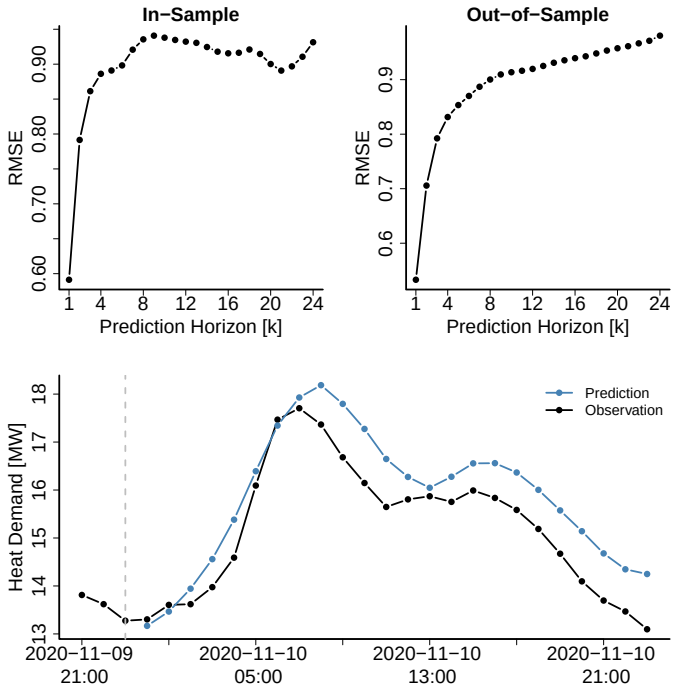

Figure 4: The top plots shows the RMSE for the in-sample period on the left and the out-of-sample period on the right for each horizon. The bottom plot shows one forecasting realisation for the total heat load forecast from one to 24 step-ahead created at 2020-11-09 23:00, as highlighted with a vertical dashed grey line.

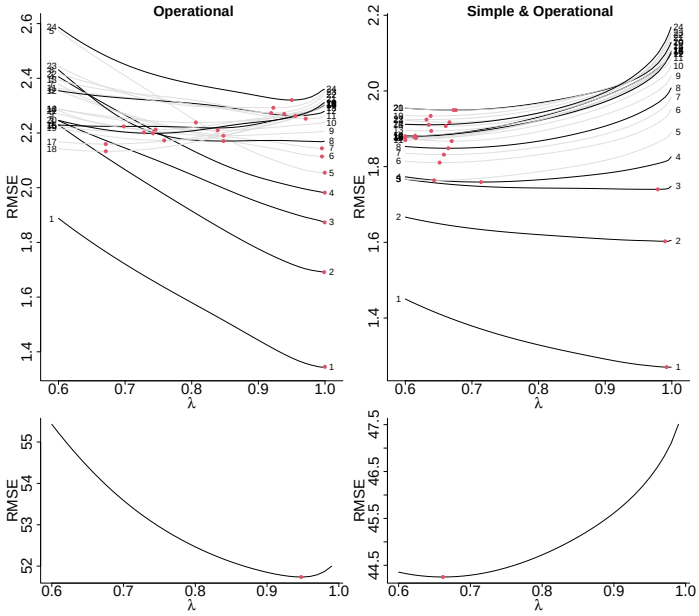

Figure 5: Figure shows RMSE of the prediction errors versus the forgetting factor for each horizon in the top plots while the bottom plots show it for all horizons for Brønderslev. The shorter the prediction horizon, the longer memory is needed; as the horizon increases, the memory decreases. The red dots show the optimal λ, i.e. lowest RMSE for each horizon.

for all horizons or as a single optimal forgetting factor for each horizon by minimising the RMSE for each horizon. The value of the forgetting factor may vary for the different horizons, as shown in Nielsen and Madsen [26]. It is therefore important to investigate whether the optimal forgetting factor is different for each horizon or whether a global forgetting factor is sufficient. In this study, horizons from one to 24 hours ahead are investigated. Therefore, the covariance estimator will have either one forgetting factor or 24 forgetting factors.

The result of the investigation of the forgetting factor for the *Brønderslev Forsyning* utility is shown in Figure 5. First, two cases are examined: 1) The *operational* base forecast is used for every level in the hierarchy (case one). The results for case one are shown in the left plots, 2) The *operational* base forecast is used for three areas only, and the total forecast uses the *simple* base forecast from the forecast model in Eq. 13 (case two). The results for case two are shown in the right plots. The top plots illustrate the performance of the reconciliation forecasts using a covariance estimator that is estimated for each horizon with a forgetting factor that was optimised for each horizon, while the bottom plots illustrate the performance using a single global forgetting factor for all horizons for the covariance estimation. To investigate the effect of the forgetting factor on the reconciliation accuracy performance, the profile of the RMSE is calculated for different forgetting factors and visualised in Figure 5. The forgetting factor on the interval $[0.6, 0.99]$ with a step of 0.01 was calculated. Also, two additional forgetting factors were added for the profile calculation, 0.995 and 0.999. The optimal value was then optimised using `nlminb()` in `R` [52], shown as a red dot for each horizon.

The results show that for the short horizons, the forgetting factors are found to be optimal with a high forgetting factor where the RMSE is minimised, as can be seen in the upper plots in Figure 5. For instance, the first seven horizons for the first case and the first three for the second case, where the *simple* base forecasts are used for the total. At higher horizons, the forgetting factor is optimal for both cases with a rather small memory; however, there is a rather flat curve, i.e. the optimum is not very well defined, especially for the second case. For example, a forgetting factor of 0.8 translates to effective memory of $N_{\text{eff}} = 5$. This is quite a small number of effective observations that can be used to estimate the empirical covariance matrix, which could be prone to very large prediction errors.

When comparing the optimal factors between the two cases, it can be seen that the RMSE is significantly lower when using the author's base forecast for the total demand. This can be seen in more detail in Table 2, where the optimal forgetting factor and corresponding RMSE are shown for both optimisations, a single forgetting factor for each horizon and one for all horizons which are shown inside the parenthesis. So, these results suggest that finding the optimal forgetting factor for each horizon should lead to greater accuracy improvements than using a single forgetting factor. At the very least, use different forgetting factors for the lower and higher horizons, as they tend to cluster together, and the RMSE curves are quite flat around their optimal points.

Table 2: The optimal forgetting factors for each horizon and the optimal forgetting factors for all horizons for both *operationa* and *simple* base forecast are shown in the table with the corresponding RMSE value. The values in the parenthesis are the result of using a global forgetting factor. Notice that the forgetting factor and RMSE for the *simple* base forecast are significantly lower than for the *operational*

Horizon	*Operational*		*Operational & Simple*	
[k]	Forgetting Factor [λ]	RMSE	Forgetting Factor [λ]	RMSE
1	0.9999 (0.9475)	1.345 (1.38)	0.9922 (0.6617)	1.27 (1.403)
2	0.9987 (0.9475)	1.691 (1.723)	0.9898 (0.6617)	1.603 (1.646)
3	0.9995 (0.9475)	1.873 (1.908)	0.979 (0.6617)	1.74 (1.754)
4	0.9997 (0.9475)	1.982 (2.008)	0.7136 (0.6617)	1.759 (1.762)
5	0.9996 (0.9475)	2.055 (2.08)	0.6431 (0.6617)	1.764 (1.764)
6	0.9955 (0.9475)	2.114 (2.131)	0.6515 (0.6617)	1.811 (1.81)
7	0.9953 (0.9475)	2.144 (2.158)	0.6579 (0.6617)	1.832 (1.831)
8	0.9999 (0.9475)	2.17 (2.172)	0.6645 (0.6617)	1.849 (1.848)
9	0.8475 (0.9475)	2.19 (2.199)	0.6696 (0.6617)	1.867 (1.867)
10	0.839 (0.9475)	2.21 (2.226)	0.616 (0.6617)	1.876 (1.877)
11	0.971 (0.9475)	2.252 (2.254)	0.6 (0.6617)	1.87 (1.876)
12	0.9999 (0.9475)	2.28 (2.263)	0.6148 (0.6617)	1.882 (1.885)
13	0.9999 (0.9475)	2.293 (2.27)	0.6389 (0.6617)	1.894 (1.895)
14	0.8062 (0.9475)	2.238 (2.262)	0.6 (0.6617)	1.88 (1.886)
15	0.7282 (0.9475)	2.202 (2.264)	0.6 (0.6617)	1.878 (1.883)
16	0.8363 (0.9475)	2.218 (2.258)	0.6 (0.6617)	1.878 (1.884)
17	0.671 (0.9475)	2.16 (2.257)	0.6136 (0.6617)	1.878 (1.881)
18	0.671 (0.9475)	2.133 (2.25)	0.6607 (0.6617)	1.907 (1.907)
19	0.7587 (0.9475)	2.173 (2.249)	0.638 (0.6617)	1.934 (1.934)
20	0.7418 (0.9475)	2.199 (2.256)	0.6723 (0.6617)	1.949 (1.949)
21	0.7458 (0.9475)	2.213 (2.27)	0.6759 (0.6617)	1.949 (1.949)
22	0.9192 (0.9475)	2.273 (2.279)	0.6318 (0.6617)	1.924 (1.925)
23	0.8679 (0.9475)	2.293 (2.295)	0.6663 (0.6617)	1.917 (1.917)
24	0.9501 (0.9475)	2.32 (2.32)	0.6352 (0.6617)	1.911 (1.911)

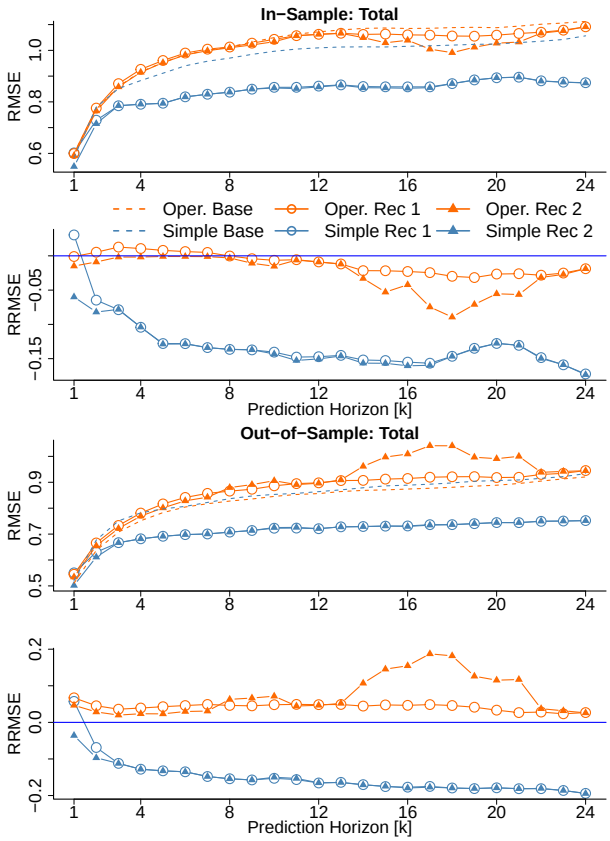

Figure 6: The upper plots in the in-sample or out-of-sample period illustrate the RMSE of the total forecast error for the two cases 1) operational for all levels, 2) *simple* at the top and *operational* at the bottom. Case 1 is shown in orange, and case 2 is in black. The base forecast errors for both cases are marked with a dashed line. Also, the use of either one forgetting factor for each horizon with a line and dot and a unique forgetting factor for each horizon with a line and a triangle

4.3. Accuracy Improvements in Brønderslev

Using the forgetting factor found in Section 4.2 to adaptively update the covariance estimator for the reconciliation process to produce reconciled forecasts as new information becomes available. The RMSE of the base and the reconciliation forecasts for the two cases are calculated for the total area level and the RRMSE to demonstrate the accuracy performance for each horizon in percentage. The results for the *Brønderslev Forsyning* demo case are shown in Figure 6. The accuracy improvements by using forgetting factors optimised for each horizon or a single forgetting factor for all horizons are also shown in the plots. The two upper plots show the result for the in-sample period, while the two lower plots show the result for the out-of-sample period. The lower plots for in-sample and out-of-sample plots show the improvement in accuracy between the base forecast and the reconciled forecast using the RRMSE score. Figure A.17 shows the same plots for the bottom three areas.

Case two, where the *simple* base forecast was used for the total area in the hierarchy, shows significant improvements compared to case one, where the *operational* base forecast was used at all levels. Case one shows a lower accuracy for the reconciliation forecast compared to the base forecast for almost all areas and horizons, the longer horizon however demonstrate a slight improvement in accuracy), as can be seen in Figure 6 and Figure A.17. We can also see that using the optimal forgetting factor for each horizon results in greater improvements for the first four steps ahead than the single forgetting factor. Using a single forgetting factor for all horizons usually results in a worse accuracy than the base forecast for the one-step horizon. After the first horizon, the improvements in accuracy are similar to the other horizons. This is consistent with the result of optimising the forgetting factor in Figure 5. For example, in case two, the forgetting factor for the first three horizons was high, and the single forgetting factor was low and similar to the forgetting factor for the other horizons.

These results become even more apparent when calculating the cumulative sum of squared errors for the base forecast and the reconciled forecast with the optimal forgetting factor for each horizon in the total heat load. This is shown in Figure 7 for one step and 24 steps ahead for both the in-sample and out-of-sample forecasts. This shows a significant improvement in the accuracy using the proposed method. It can also be observed that the gains in accuracy are most significant in the colder periods. In contrast, in the warm period (summer), the error slope is similar (flat) when comparing the base and reconciled errors.

The correlation matrices shown in Figure 8 are calculated from the covariance estimator at the time shown in the plots from the 24-ahead forecast errors for both cases, using only the *operational* forecast (plots on the left) and the *operational* and *simple* forecasts (the plots on the right) as base forecast. The upper and lower plots show instances from the summer

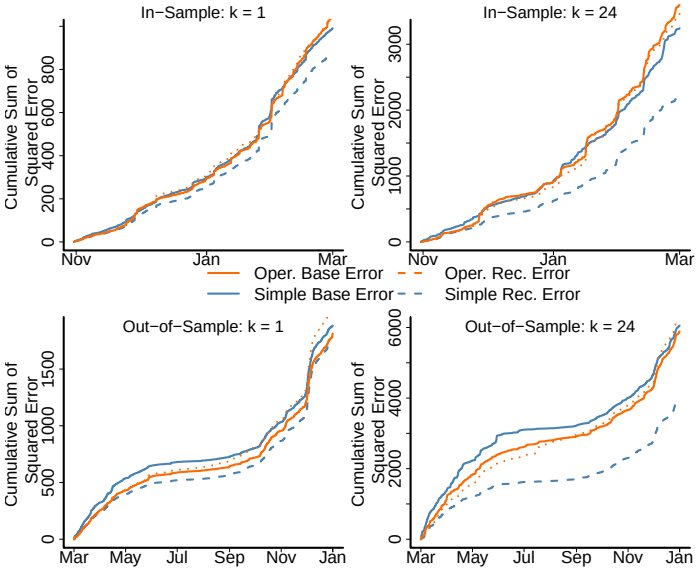

Figure 7: Figure shows the cumulative sum of squared error for the base and reconciled forecast of the total heat load where the reconciled forecast estimated using optimal forgetting for each horizon using the authors base forecast of the total aggregation.

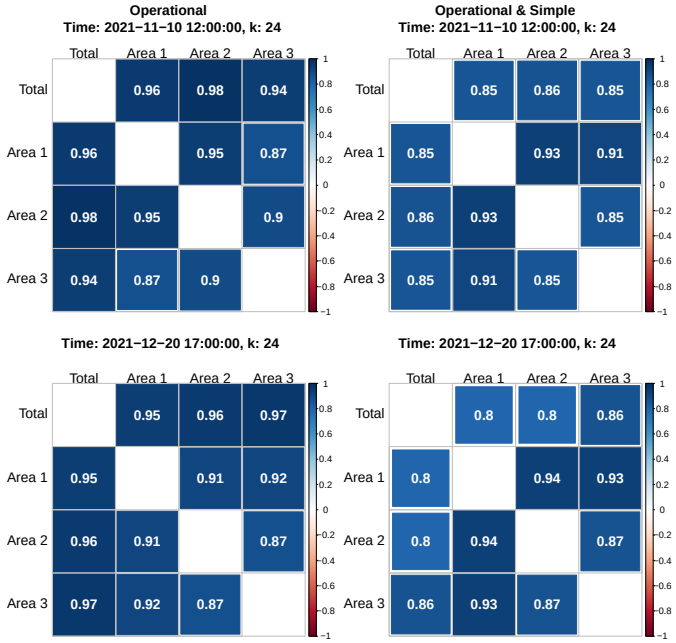

Figure 8: Figure illustrates the correlation between the areas when using only *operational* base forecast in the left plots and using *operational* and *simple* forecast as the base in the right plots. The plots show correlation instances at periods in summer in the top plots and lower plots in winter.

and winter periods. Only using the *operational* forecast results in a higher correlation, while using the author forecast results in a lower correlation. This could indicate that using forecasts produced with the same model leads to too similar errors and therefore cannot share any useful information in the reconciliation process. Similar results are found and discussed in Nystrup et al. [31].

Figure 9 shows an example of the realisation of the prediction for both the base and reconciled forecast of the one to 24 steps ahead. Each plot shows a different area, e.g. the top left plot shows total demand. The top plot shows both the base forecast of *operational* and *simple* base forecast and the corresponding reconciled forecast with a unique forgetting factor for each horizon to estimate the covariance matrix. From these plots, it is difficult to tell which forecast performs best; however, the base forecasts for the total heat load are quite different.

4.3.1. One-Step Ahead Empirical Covariance Matrix

In the literature, one-step errors are usually used to tune models, i.e. to estimate the model's coefficients. These estimates are then used for multi-step predictions. However, this can lead to sub-optimal results because the correlation between the output and input variables can change depending on the horizon. The same is true for the reconciliation forecast, as the empirical covariance matrix is estimated using forecast errors. The accuracy improvements from the reconciled process result from the base forecast errors from the recursively estimated empirical covariance matrix. Therefore, having a covariance estimator for each horizon might be beneficial. We used an empirical covariance matrix for each horizon calculated from the corresponding k-step prediction error in the previous section. However, we would like to investigate whether there is a

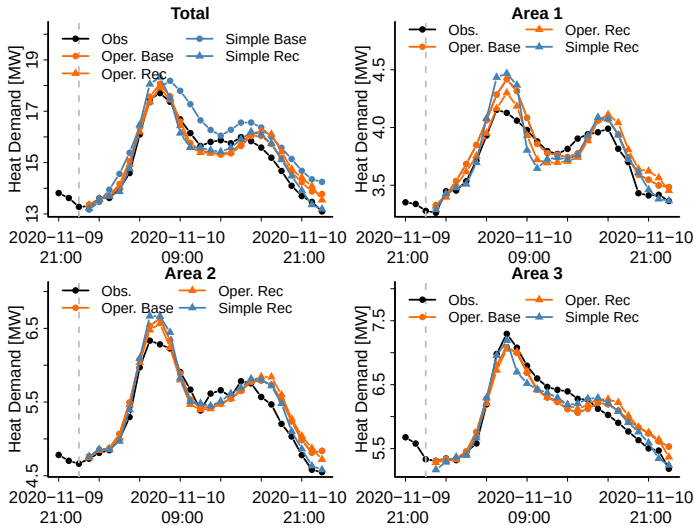

Figure 9: Realisations of base and reconciled forecast for each area in the *Brønderslev Forsyning* hierarchy.

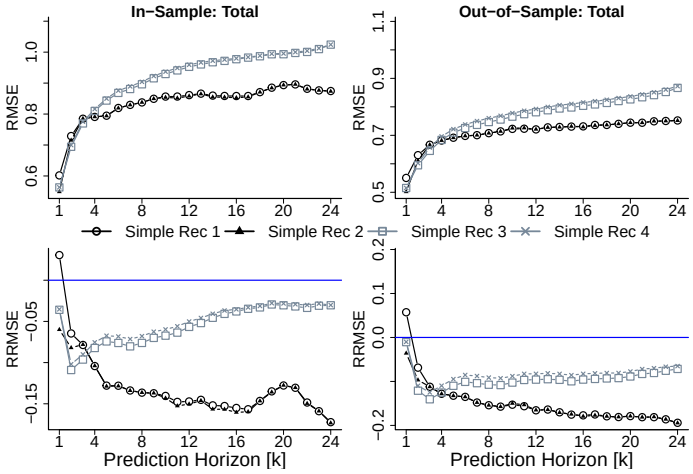

Figure 10: The figure shows the difference in accuracy between using errors from all prediction horizons and using errors only one step ahead when estimating the empirical covariance matrix. *Simple Rec 1* and *Simple Rec 2* show the results of using errors from each prediction horizon when a single forgetting factor is optimised for all horizons or each horizon has a unique forgetting factor. *Simple Rec 3* and *Simple Rec 4* show the result of optimising the single forgetting factor for all horizons or each horizon has a unique forgetting factor using only the error from the one step ahead to estimate the covariance matrix for all horizons. The base forecasts from case two were used here.

significant difference in accuracy when we use either the one-step error for all horizons or the k-step error for the covariance estimator. Here we will only use case 2 with the *simple* total forecast and the *operation* forecast for the bottom areas.

The results are shown in Figure 10, which shows the accuracy improvements when using the one-step or multi-step prediction errors to estimate the covariance matrix. The top plots show the RMSE score, and the bottom plot shows the RRMSE, while the left and right plots show the in-sample and out-of-sample results. For simplicity, only the results for the total area are shown. *Simple Rec 1* is the result of using errors from all horizons and only one forgetting factor, while *Simple Rec 2* uses errors from all horizons and has an optimal forgetting factor for each horizon. *Simple Rec 3* and *Simple Rec 4* show the same, but only using the base error from the one-step prediction with one forgetting factor or optimal for each horizon. The improvements using only the errors from the one-step forecast are significantly lower than those using multi-step prediction errors. You can see the optimised forgetting factors for each horizon and the forgetting factor for all horizons in Table A.3.

4.4. Accuracy Improvements in FYN and Hierarchy Structure Investigation

The current hierarchy at Fjernevarme Fyn is the total and the 12 different areas however due to the high number of areas and how they are linked together as seen in Figure 1. They are quite grouped together, with some areas located far from the central point. Hence, it gives the opportunity to create more aggregation levels to extend the hierarchy structure and hopefully further enhance the accuracy improvements. An hourly average of the heat demand for each weekday for the four seasons is shown in Figure A.16. The four heating demand seasons in Denmark are listed below,

1. ***Winter*** *(December, January, February, March)* during cold periods when the heating demand is high due to the high consumption of space heating to keep the indoor climate comfortable.

2. ***Spring*** *(April, May)* during the transition period from cold to warm with the influence of solar irradiance of warming houses, thus lowering the amount of space heating needed.

3. ***Summer*** *(June, July, August, September)* when space heating is usually not needed in Denmark, only domestic hot water is needed (e.g. hot tap water and showering).

4. ***Fall*** *(October, November)* during the transition period when space heating is required again due to lowering ambient air temperatures.

Comparing the areas' heating dynamics together makes it evident that each area has unique heat dynamics; however, there are some possibilities to group some of them depending on their behaviour from Figure A.16 when comparing their shape. For instance, Areas 6, 8 and 12 exhibit similar shapes. This analysis can give an idea of which areas should be aggregated together. Based on these results, new aggregate levels have been added. Three new aggregations have been added, which aggregated four areas together. Aggregation 1 is the aggregation of { Area 1, Area 2, Area 3, Area 4 }, Aggregation 2 is {Area 5, Area 9, Area 10, Area 11 }, and Aggregation 3 is {Area 6, Area 7, Area 8, Area 12} based on from the result in Figure A.16 and the layout in Figure 1.

The base forecast for the new aggregation level is then created (*simple*) while still using the *operational* base forecast for the total and the 12 areas. The forecasting model used for the new aggregation level is similar to the model for the total for *Brønderslev Forsyning* as shown in Eq. 13. The prediction horizon will also be the same, one- to 24-steps ahead. We will use both hierarchies, a total of 12 areas and a total with three aggregations of the bottom 12 areas, to compare the results using a simple hierarchy and a hierarchy with an additional aggregation level. Thus, investigate when more information is added to the hierarchy and how it will affect the accuracy improvements of the reconciled forecasts.

The accuracy improvements for the total aggregation are shown in Figure 11 where it shows that the reconciled forecast for both cases improves the accuracy in almost all cases, especially in the longer horizons as was the case for *Brønderslev Forsyning* study. The reconciled forecasts with *Oper. Rec 1* and *Oper. Agg Rec 1* shows the reconciled forecast where the same forgetting factor is optimised for all horizons while *Oper. Rec 2* and *Oper. Agg Rec 2* shows the result for optimising the forgetting factor for each horizon. *Oper. Rec* is where the simple hierarchy is used and *Oper. Agg Rec* is when the additional aggregation level has been added to the hierarchy. In the top plots, the RMSE of the forecasts are shown, and in the bottom, the accuracy improvements as the RRMSE. The left and right plots show the result in the in-sample and out-of-sample periods.

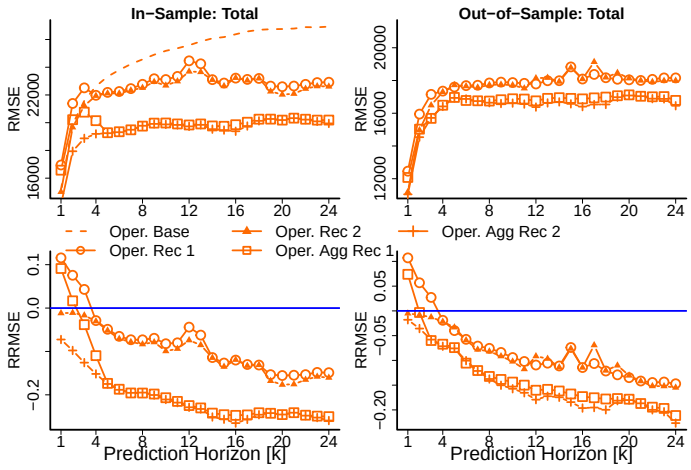

Figure 11: The plots show the accuracy improvements of the base forecast in *Fjernvarme Fyn*, where two different hierarchies are investigated. The top plots show the RMSE, and the bottom plots show the improvements using the RRMSE, while the left and right plots show the result in the in-sample and out-of-sample periods.

The results in the plots in Figure 11 demonstrate that the reconciled forecast for both hierarchy structure cases improves the accuracy of the base forecasts. The first three horizons needed to have their own forgetting factor, the same as we saw for the *Brønderslev Forsyning* case. The longer horizon otherwise improves the accuracy for both setups. Adding an additional aggregation level to the hierarchy structure demonstrates significant accuracy improvement, almost double the improvement.

5. Simulation Study

To show the effect of the forgetting factors in a controlled environment, a simulation study is performed. The setup for the simulation study will consist of three areas and a total of one year of hourly measurements each. The simulated load from the three areas are produced by simple Auto-Regressive (AR) models of lag one with different offsets.

$$Y_{\text{Area1,t}} = 0.7Y_{\text{Area1,t}-1} + 140 + \epsilon_{1,t}, \tag{15}$$
$$Y_{\text{Area2,t}} = 0.8Y_{\text{Area2,t}-1} + 150 + \epsilon_{2,t}, \tag{16}$$
$$Y_{\text{Area3,t}} = 0.9Y_{\text{Area3,t}-1} + 158 + \epsilon_{3,t}, \tag{17}$$

Where total is the aggregation of the areas and $\epsilon_{i,t}$ are white noise

$$Y_{\text{Total,t}} = Y_{\text{Area1,t}} + Y_{\text{Area2,t}} + Y_{\text{Area3,t}}. \tag{18}$$

Level shifts are added to the simulated heat load at random intervals to introduce some additional complexity and simulate spikes in heat demand. These are added at randomly sampled time points. At each sampled time point, the load is increased for 48 hours by a random amount sampled from the distribution $N(30, 2)$ and rounded to the nearest integer. The resulting simulated heat load is visualized in Figure 12 where the raw

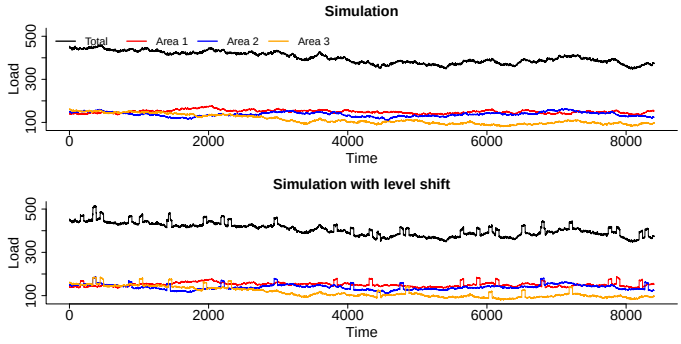

Figure 12: Simulated heat load measurements. Top: raw simulated measurements from AR(1) models. Bottom: simulated measurements with levels shifts.

simulation is shown in the top plot while the lower plot shows the simulation with level shifts.

To reconcile the spatial hierarchy, base forecasts are, of course, needed. These are again created using the `R` package `onlineforecast`, where the model coefficients are estimated using the Ordinary Least Squares method. The models that are constructed are AR(1) processes for both the areas and the total, where the coefficients are re-estimated every time a new observation becomes available. A unique model is created for each prediction horizon. Hence,

$$Y_{\text{Area1,t}} = \phi_{1,\text{t,k}} Y_{\text{Area1,t}-1} + \epsilon_{1,\text{t}}, \tag{19}$$
$$Y_{\text{Area2,t}} = \phi_{2,\text{t,k}} Y_{\text{Area2,t}-1} + \epsilon_{2,\text{t}}, \tag{20}$$
$$Y_{\text{Area3,t}} = \phi_{3,\text{t,k}} Y_{\text{Area3,t}-1} + \epsilon_{3,\text{t}}, \tag{21}$$
$$Y_{\text{Total,t}} = \phi_{4,\text{t,k}} Y_{\text{Total,t}-1} + \epsilon_{4,\text{t}}. \tag{22}$$

Where the subscript t is the time and k is the prediction horizon. The $\phi_{\text{i,t,k}}$ is the AR(1) coefficient of the model and $\epsilon_{\text{i,t}}$ is the error term.

When reconciling this simulated hierarchy, an investigation is performed into the memory used for updating the empirical covariance matrix when new information is available. This is done by running the reconciliation process twice, once with a short memory and once with a longer memory. The short memory will have the forgetting factor set as $\lambda = 0.8$; thus, the effective memory is the past five time steps, i.e. $N_{\text{eff}} = 5$. The long memory will have the forgetting factor as $\lambda = 0.99$, i.e. $N_{\text{eff}} = 100$. The RMSE for each horizon for the base forecast and the two scenarios of memory for the reconciled forecast are shown in Figure 13 for each area. The long memory demonstrates slight improvements compared to the base forecast, while the short memory improves it significantly, except for the first three horizons, where it has lower accuracy.

Figure 13 shows the accuracy on the one-hour prediction horizon. This result indicates that a short memory might be preferable; however this should be investigated further. The top rows of figures 14 and 15 show a segment of the data along with the base forecast as well as the two reconciled forecasts for prediction horizons $k = 1$ and $k = 12$ respectively. The bottom rows of these figures show the corresponding cumulative sum of squared error for each forecast. It is evident from Figure 14 that

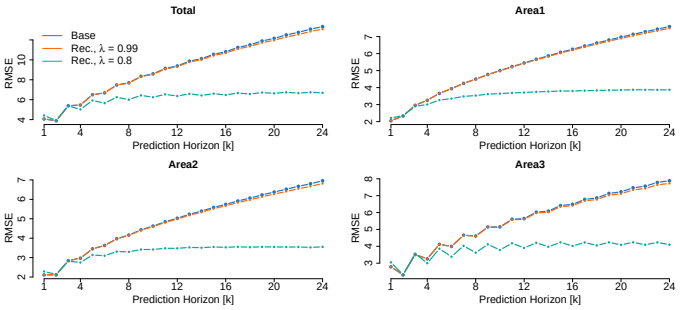

Figure 13: Accuracy of the base and reconciled forecasts with the two forgetting factors being investigated for the three areas and the total.

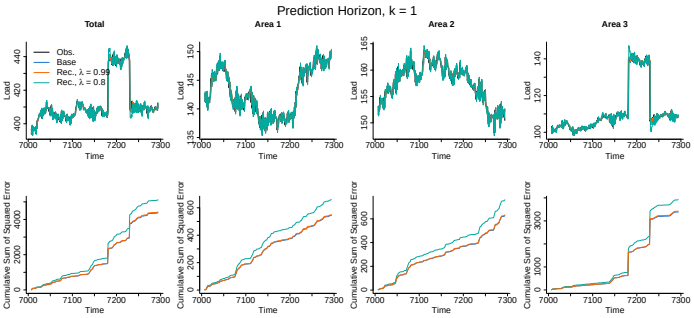

Figure 14: Top row: Segment of the simulated observations along with base forecasts and reconciled forecasts with memories $\lambda = 0.8$ and $\lambda = 0.99$ for the one-hour prediction horizon. Bottom row: cumulative sum of squared errors for the three forecasts for each area.

the short memory is very volatile, resulting in lower accuracy than the base forecast. The longer memory seems to perform similarly to the base forecast. Short memory significantly outperforms the base forecast for the longer prediction horizon, as was also seen in Figure 13. Figure 15 demonstrates that the significant accuracy improvements come from sudden changes in the data. Thus, during the level shift around 7200 hours in Area 3, the short memory is able to track the changes faster, thus resulting in more accurate forecasts. Even though there are no significant level shifts in Area 1 and Area 2, the short memory outperforms the other forecasts. Notice the level jump in the cumulative sum of squared errors in Area 1 and Area 2 for the base forecast and long memory reconciled forecast, while this is not noticeable in the short memory forecast. Hence highlighting the power of being able to react quickly to changes in the data. The short memory reconciled forecasts do however show a significant overshoot when changes in the base forecast happen quickly, as seen in Area 3 when the level shift occurs. This is likely due to the weights used in the reconciliation process being estimated based on base forecasts from before the change. It is likely that this could be addressed by improving the base forecasts or tweaking the memory a bit. Improving base forecast to react to changes in data in a more timely manner would result in weights not being so far out of scale. Alternatively, the memory should be tweaked to be a bit longer, thus proving a middle ground between the two examples shown.

It is clear from this simulation study that choosing a fitting memory for the proposed reconciliation process is vital to

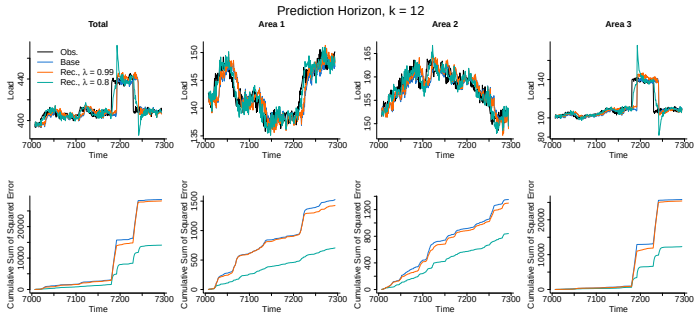

Figure 15: Top row: Segment of the simulated observations along with base forecasts and reconciled forecasts with memories $\lambda = 0.8$ and $\lambda = 0.99$ for the 12-hour prediction horizon. Bottom row: cumulative sum of squared errors for the three forecasts for each area.

ensuring good performance. Too high memory can result in reconciled forecasts not being reactive enough, while too short memory can result in overreacting to changes in the base forecast. As such, the memory should be set based on the data behaviour to ensure optimal performance.

6. Discussion

This paper proposes to improve state-of-the-art operational forecast accuracy by using spatial hierarchies with an adaptive covariance matrix in the reconciliation process to produce coherent forecasts for district heating systems. The coherency is introduced by linear constraints defined by the hierarchy, where the lower levels are forced to be aggregated equal to the upper level. The independent and not necessarily coherent base forecasts are first created and then projected onto the coherent subspace using a projection matrix. The projection matrix is created from the summation matrix of the hierarchy, and the covariance estimator is estimated from the errors of the in-sample base forecasts errors. Due to the non-stationary nature of the heat load, the covariance matrix must be able to adapt as the heat load changes over time. Therefore, adaptive and recursive covariance estimation is performed using exponential smoothing as proposed by Ber [12]. Additional forecasts were needed to investigate whether the *operational* base forecast was too similar. Also, investigating if adding a new level of aggregation to the hierarchy would enhance the accuracy improvements even further. A linear regression model was proposed to forecast the additional heat load forecasts. The two-stage forecasting framework was used, where the input variables are first transformed, and then the coefficients are estimated using RLS with exponential forgetting. This framework is ideal for the inherent non-linearity and non-stationary heat load and for producing robust online forecasts with high accuracy.

We initially defined the spatial hierarchy of district heating with the total as the top level and the individual areas as the lowest levels. This study disregarded heat losses from production to areas for simplicity, as heat losses can be challenging to forecast with high accuracy. The first attempt to make reconciliation forecasts for the *Brønderslev Forsyning* utility showed almost no improvements for the in-sample period. It also re-

sulted in worse accuracy than the base forecast in the out-of-sample period. It was found that the similarity of the *operational* base forecasts between levels was too high by investigating the correlation matrix. Therefore, no gain in accuracy could be achieved if the correlation between the base forecasts is too high as they did not exchange any useful information, as shown in Nystrup et al. [31]. Replacing the forecast with the *simple* base forecast showed high accuracy improvements in both the in-sample and out-of-sample periods, about 15% depending on the horizon.

We concluded that the forgetting factor needs to be optimised for each horizon, especially for the short horizons, which tend to be very high. It was also discovered that the optimum of the forgetting factor for higher horizons was low and not very well defined. The low forgetting factor allows the covariance matrix to react quickly to sudden changes when the base forecast performs poorly. The ability to react quickly to changes, e.g. level shifts, can be useful when sudden and unexpected changes in the heat load occur. This was demonstrated in a simulation study where it was investigated where the accuracy improvements could come from. The base forecast model used in this work tends to have a high forgetting factor, which could be the reason why the low forgetting factor leads to such a significant improvement in accuracy. However, choosing a higher forgetting factor could lead to a more robust estimate and reduce the risk of high prediction errors. It is also argued that a higher forgetting factor enables more robust prediction by reducing overshooting/undershooting of the prediction by having too low a forgetting factor.

The errors used in the covariance estimation were investigated. Frequently, the one-step errors are used to estimate coefficients of forecasting models, which are then used for forecasting multiple-step horizons ahead. It could be more beneficial to use a specific forecast model for each k-step ahead forecast and use the appropriate errors to estimate the coefficients. We demonstrated that it is important to use the k-step errors when estimating the covariance estimator to predict the k-step ahead heat load by comparing the accuracy to when only using the one-step errors for all horizons when estimating the covariance estimator. By using the corresponding k-step errors, almost twice the accuracy is obtained.

Another case study was carried out using the heat load forecast from the *Fjernevarme Fyn* utility to validate the results of the case study of *Brønderslev Forsyning*. It showed similar accuracy improvements and highlighted the importance of finding the optimal forgetting factor for each horizon. We also showed that adding an additional level of aggregation with the *simple* base forecasts to the hierarchy significantly improved the accuracy. Thus, adding more information to the voting process shows that more knowledge is shared between levels about the covariance matrix. This is in line with current literature that suggests using full hierarchies to achieve optimal improvements.

We therefore conclude that spatial hierarchies improve the accuracy of state-of-the-art operational heat load forecasting, as demonstrated in two different case studies with different hierarchy structures. This will be important for future energy systems as an improved heat load forecast increases the flexibility potential of district heating by providing more information for decision-making. Furthermore, a coherent forecast for district heating operations is important for temperature, and production optimisation as district heating systems are becoming more decentralized with multiple areas with local heat units or TES systems. We also believe that selecting a robust forgetting factor that reduces the probability of large forecast errors will be important. Too similar forecasts could be problematic in not improving the accuracy, leading to worse accuracy than the base forecast. Hence, different forecasting methods for other areas could be more beneficial and should be investigated in more detail.

7. Conclusion

We propose a novel method to increase the accuracy of operational state-of-the-art heat load forecasts by exploiting information between different areas through the spatial hierarchy. The improved heat load forecasts were computed using a reconciliation process where the individual base forecasts are forced to be coherent using the predefined hierarchy structure, and the information is exchanged between them using the proposed covariance estimator. It is proposed to estimate the covariance estimator recursively and make it adaptive using the exponential smoothing formula. Based on two case studies, it was shown that the proposed method significantly increases the accuracy compared to operational state-of-the-art heat load forecasts for all horizons used in this work. This is highly desired by district heating utilities as future systems will be more decentralised and therefore more local heat load forecasts will be needed. Improving the accuracy of the local heat load forecast will improve the operation of the district heating by optimising the network's temperature levels and production.

It was concluded that the covariance estimator needs to be estimated for each forecast horizon to achieve the highest accuracy improvements for each horizon. Therefore, using information from the same horizon will be more optimal than using only the one-step-ahead prediction errors, as is the tendency to do in forecasting.

It was found that optimising the forgetting factor for each horizon is crucial for covariance estimation to achieve optimal improvement in accuracy. Optimal forgetting factors were found to be quite small compared to what is commonly used in research for forecasts, usually around 0.65 compared to the usual 0.999 for forecasting models. However, a higher forgetting factor of around 0.99 was found for shorter horizons. A low forgetting factor allows the reconciliation process to shift the weights when the errors of the base forecast start to increase, e.g. due to level shifts in the heat load. We argue that it might reasonable to select a higher forgetting factor than found through optimisation. Higher forgetting will lead to a more robust estimate, as a low forgetting factor might be prone to large prediction errors.

Even more improvements in accuracy were demonstrated by adding an additional level of aggregation in the hierarchy.

This provides more information to share in the hierarchy, leading to higher improvements than a hierarchy with fewer levels of aggregation. The definition of the aggregation levels for the spatial hierarchy is trivial compared to a more straightforward hierarchy structure for the temporal hierarchy. An investigation of the optimal hierarchy structure for the spatial hierarchy could therefore be explored in future research. There are several possible directions for future research, but a more detailed investigation of where these significant accuracy improvements come from at such a low forgetting factor would be essential before this can be used operationally.

8. Acknowledgement

This work is funded by Innovation Fund Denmark through the projects Heat 4.0 (8090-00046B), CITIES (1305-00027B), TOP-UP (9045-00017A), Flexible Energy Denmark (8090 -00069B), ARV (EU H2020 101036723), and the Norwegian FME-ZEN project funded by the ZEN partners and the Research Council of Norway (257660).

We thank the energy forecasting provider ENFOR, especially Torben Skov Nielsen, for his support and supplying the forecast and information. We would also like to thank *Fjernvarme Fyn* and *Brønderslev Forsyning* for making their heat demand data available.

References

[1] A. Vandermeulen, B. van der Heijde, L. Helsen, Controlling district heating and cooling networks to unlock flexibility: A review, Energy 151 (2018) 103–115.

[2] International Energy Agency (IEA), Renewables 2021, Analysis and forecasts to 2026, EA, Paris, France, 2021. https://www.iea.org/reports/renewables-2021.

[3] A. David, B. V. Mathiesen, H. Averfalk, S. Werner, H. Lund, Heat roadmap europe: Large-scale electric heat pumps in district heating systems, Energies 10 (2017).

[4] E. Guelpa, V. Verda, Thermal energy storage in district heating and cooling systems: A review, Applied Energy 252 (2019) 113474.

[5] IRENA, Innovation outlook: thermal energy storage, International Renewable Energy Agency, Abu Dhabi, 2020. https://www.irena.org/-/media/Files/IRENA/Agency/Publication/2020/Nov/IRENA_Innovation_Outlook_TES_2020.pdf.

[6] A. Bloess, W.-P. Schill, A. Zerrahn, Power-to-heat for renewable energy integration: A review of technologies, modeling approaches, and flexibility potentials, Applied Energy 212 (2018) 1611–1626.

[7] T. S. Nielsen, Online prediction and control in nonlinear stochastic systems, Ph.D. thesis, Informatics and Mathematical Modelling, Technical University of Denmark, DTU, 2002. URL: http://www2.compute.dtu.dk/pubdb/pubs/792-full.html.

[8] On flow and supply temperature control in district heating systems, Heat Recovery Systems and CHP 14 (1994) 613 – 620.

[9] P. Pinson, T. S. Nielsen, H. A. Nielsen, N. K. Poulsen, H. Madsen, Temperature prediction at critical points in district heating systems, European Journal of Operational Research 194 (2009) 163–176.

[10] H. Bergsteinsson, S. Ben Amer, P. Nielsen, H. Madsen, Digitalization of District Heating, Technical University of Denmark, 2021.

[11] R. Hyndman, R. Ahmed, G. Athanasopoulos, H. Shang, Optimal combination forecasts for hierarchical time series, Computational Statistics and Data Analysis 55 (2011).

[12] Heat load forecasting using adaptive temporal hierarchies, Applied Energy 292 (2021) 116872.

[13] S. L. Wickramasuriya, G. Athanasopoulos, R. J. Hyndman, Optimal forecast reconciliation for hierarchical and grouped time series through trace minimization, Journal of the American Statistical Association 114 (2019) 804–819.

[14] T. Hong, P. Pinson, Y. Wang, R. Weron, D. Yang, H. Zareipour, Energy forecasting: A review and outlook, IEEE Open Access Journal of Power and Energy 7 (2020) 376–388.

[15] I. Blanco, D. Guericke, A. Andersen, H. Madsen, Operational planning and bidding for district heating systems with uncertain renewable energy production, Energies 11 (2018).

[16] F. Verrilli, S. Srinivasan, G. Gambino, M. Canelli, M. Himanka, C. Del Vecchio, M. Sasso, L. Glielmo, Model predictive control-based optimal operations of district heating system with thermal energy storage and flexible loads, IEEE Transactions on Automation Science and Engineering 14 (2016) 547–557.

[17] T. S. Nielsen, H. Madsen, Control of supply temperature in district heating systems, in: Proceedings of the 8th International Symposium on District heating and Cooling, 2002.

[18] D. Quaggiotto, J. Vivian, A. Zarrella, Management of a district heating network using model predictive control with and without thermal storage, Optimization and Engineering 22 (2021) 1897–1919.

[19] H. Gadd, S. Werner, Daily heat load variations in swedish district heating systems, Applied Energy 106 (2013) 47–55.

[20] H. Nielsen, H. Madsen, Modelling the heat consumption in district heating systems using a grey-box approach, Energy and Buildings 38 (2006) 63–71.

[21] E. Dotzauer, Simple model for prediction of loads in district - heating systems, Applied Energy 73 (2002).

[22] M. Dahl, A. Brun, G. Andresen, Using ensemble weather predictions in district heating operation and load forecasting, Applied Energy 193 (2017).

[23] S. Grosswindhagera, A. Voigtb, M. Kozeka, Online short-term forecast of system heat load in district heating networks, Tsp 1 (2011).

[24] M. Dahl, A. Brun, O. S. Kirsebom, G. B. Andresen, Improving short-term heat load forecasts with calendar and holiday data, Energies 11 (2018) 1678.

[25] S. Idowu, S. Saguna, C. Åhlund, O. Schelén, Applied machine learning: Forecasting heat load in district heating system, Energy and Buildings 133 (2016) 478–488.

[26] H. Nielsen, H. Madsen, Predicting the Heat Consumption in District Heating Systems using Meteorological Forecasts, Informatics and Mathematical Modelling, Technical University of Denmark, 2000.

[27] K. Kato, M. Sakawa, K. Ishimaru, S. Ushiro, T. Shibano, Heat load prediction through recurrent neural network in district heating and cooling systems, in: 2008 IEEE international conference on systems, man and cybernetics, IEEE, 2008, pp. 1401–1406.

[28] J. Song, L. Zhang, G. Xue, Y. Ma, S. Gao, Q. Jiang, Predicting hourly heating load in a district heating system based on a hybrid cnn-lstm model, Energy and Buildings 243 (2021) 110998.

[29] P. Nystrup, E. Lindström, P. Pinson, H. Madsen, Temporal hierarchies with autocorrelation for load forecasting, European Journal of Operational Research 280 (2020) 876–888.

[30] T. Van Erven, J. Cugliari, Game-theoretically optimal reconciliation of contemporaneous hierarchical time series forecasts, in: A. Antoniadis, J.-M. Poggi, X. Brossat (Eds.), Modeling and Stochastic Learning for Forecasting in High Dimensions, Springer International Publishing, Cham, 2015, pp. 297–317.

[31] P. Nystrup, E. Lindström, J. K. Møller, H. Madsen, Dimensionality reduction in forecasting with temporal hierarchies, International Journal of Forecasting 37 (2021) 1127–1146.

[32] R. Hollyman, F. Petropoulos, M. E. Tipping, Understanding forecast reconciliation, European Journal of Operational Research 294 (2021) 149–160.

[33] M. L. Sørensen, P. Nystrup, M. B. Bjerregård, J. K. Møller, P. Bacher, H. Madsen, Recent developments in multivariate wind and solar power forecasting, Wiley Interdisciplinary Reviews: Energy and Environment (2022) e465.

[34] J. Jeon, A. Panagiotelis, F. Petropoulos, Probabilistic forecast reconciliation with applications to wind power and electric load, European Journal of Operational Research 279 (2019) 364–379.

[35] L. Arvastson, Stochastic Modeling and Operational Optimization in Dis-

trict Heating Systems, Ph.D. thesis, Centre for Mathematical Sciences, Lund University, 2001.
[36] L. Ljung, T. Söderström, Theory and practice of recursive identification, volume 4 of *The MIT press series in signal processing, optimization, and control*, MIT Press, 1983.
[37] P. Bacher, H. Madsen, H. A. Nielsen, Online short-term solar power forecasting, Solar energy 83 (2009) 1772–1783.
[38] P. Bacher, H. Madsen, H. Nielsen, B. Perers, Short-term heat load forecasting for single family houses, Energy and Buildings 65 (2013) 101–112.
[39] Load forecasting of supermarket refrigeration, Applied Energy 163 (2016) 32–40.
[40] T. Jónsson, P. Pinson, H. A. Nielsen, H. Madsen, T. S. Nielsen, Forecasting electricity spot prices accounting for wind power predictions, IEEE Transactions on Sustainable Energy 4 (2012) 210–218.
[41] P. Bacher, H. G. Bergsteinsson, L. Frölke, M. L. Sørensen, J. Lemos-Vinasco, J. Liisberg, J. K. Møller, H. A. Nielsen, H. Madsen, onlineforecast: An r package for adaptive and recursive forecasting, arXiv preprint arXiv:2109.12915 (2021).
[42] H. G. Bergsteinsson, J. K. Møller, C. A. Thilker, D. Guericke, A. Heller, T. S. Nielsen, H. Madsen, Data-driven methods for efficient operation of district heating systems, in: Handbook of Low Temperature District Heating, Springer, 2022, pp. 129–163.
[43] T. Nielsen, H. Madsen, H. Nielsen, P. Pinson, G. Kariniotakis, N. Siebert, I. Marti, M. Lange, U. Focken, L. Von Bremen, P. Louka, G. Kallos, G. Galanis, Short-term wind power forecasting using advanced statistical methods, volume 3, 2006, pp. 2482–2490.
[44] G. Athanasopoulos, R. Hyndman, N. Kourentzes, F. Petropoulos, Forecasting with temporal hierarchies, European Journal of Operational Research 262 (2017) 60–74.
[45] R. J. Hyndman, A. J. Lee, E. Wang, Fast computation of reconciled forecasts for hierarchical and grouped time series, Computational Statistics and Data Analysis 97 (2016) 16 – 32.
[46] N. Kourentzes, G. Athanasopoulos, Cross-temporal coherent forecasts for australian tourism, Annals of Tourism Research 75 (2019) 393–409.
[47] D. Yang, H. Quan, V. R. Disfani, L. Liu, Reconciling solar forecasts: Geographical hierarchy, Solar Energy 146 (2017) 276–286.
[48] G. M. Yagli, D. Yang, D. Srinivasan, Reconciling solar forecasts: Sequential reconciliation, Solar Energy 179 (2019) 391–397.
[49] F. Petropoulos, D. Apiletti, V. Assimakopoulos, M. Z. Babai, D. K. Barrow, S. B. Taieb, C. Bergmeir, R. J. Bessa, J. Bijak, J. E. Boylan, et al., Forecasting: theory and practice, International Journal of Forecasting (2022).
[50] A. Panagiotelis, G. Athanasopoulos, P. Gamakumara, R. J. Hyndman, Forecast reconciliation: A geometric view with new insights on bias correction, International Journal of Forecasting 37 (2021) 343–359.
[51] R. Hyndman, A. Koehler, Another look at measures of forecast accuracy, International Journal of Forecasting 22 (2006) 679–688.
[52] R Core Team, R: A Language and Environment for Statistical Computing, R Foundation for Statistical Computing, Vienna, Austria, 2019.

Appendix A.

Table A.3: The optimal forgetting factors for each horizon and the optimal forgetting factors for all horizons for using only one-step ahead errors to estimate the empirical covariance matrix. Notice how the forgetting factor increases again for higher prediction horizons

	Operationa & *Simple*	
Horizon [k]	Forgetting Factor [λ]	RMSE
1	0.999	1.2702
2	0.7149	1.5505
3	0.7803	1.7078
4	0.8059	1.7925
5	0.8259	1.8616
6	0.7917	1.9041
7	0.7824	1.9231
8	0.7838	1.9527
9	0.7785	1.9884
10	0.7765	2.0125
11	0.7829	2.0358
12	0.7955	2.0608
13	0.8184	2.0781
14	0.8353	2.0932
15	0.8644	2.1037
16	0.9958	2.1101
17	0.9961	2.1153
18	0.9954	2.1254
19	0.9955	2.1353
20	0.9948	2.1353
21	0.9931	2.1445
22	0.9908	2.1538
23	0.9913	2.1595
24	0.9945	2.1728

Table A.4: The optimal forgetting factors for each horizon and the optimal forgetting factors for all horizons for both ENFOR and authors base forecast are shown in the table with the corresponding RMSE value. Notice that the forgetting factor and RMSE for the *simple* base forecast are significantly lower than for the *operational*. It is not possible to compare the RMSE between the two case studies in this table as ENFOR Agg has another aggregation level, thus a higher RMSE.

	Operational		*Operational* Agg	
Horizon [k]	Forgetting Factor [λ]	RMSE	Forgetting Factor [λ]	RMSE
1	0.9861	38718.26	0.9878	88435.2
2	0.9767	49146.8	0.9829	97964.3
3	0.9649	52476.19	0.9753	99853.03
4	0.8701	54040.31	0.9542	99052.69
5	0.693	54598.05	0.9095	96424.64
6	0.6851	54702.81	0.8797	94701.16
7	0.6844	54994.56	0.8618	93483.41
8	0.6596	55289.8	0.8615	93578.51
9	0.659	55938.88	0.8426	93094.56
10	0.6204	55599.19	0.853	93088.3
11	0.6	55853.17	0.8456	93184.57
12	0.8662	57567.87	0.8553	93069.96
13	0.8475	57528.59	0.8753	94143.76
14	0.7152	56412.53	0.8699	93725.18
15	0.6539	55621.23	0.8729	93569.79
16	0.7652	56567.24	0.8846	94219.94
17	0.6842	56327.91	0.8688	94575.05
18	0.7358	56720.07	0.8723	95835.68
19	0.6376	55265	0.8962	98122.15
20	0.603	55002.29	0.9227	99076.89
21	0.6	55236.32	0.9224	99671.8
22	0.6633	55841.8	0.7555	106038.7
23	0.6431	55896.72	0.9059	97152.2
24	0.6392	55781.87	0.8975	96414.18

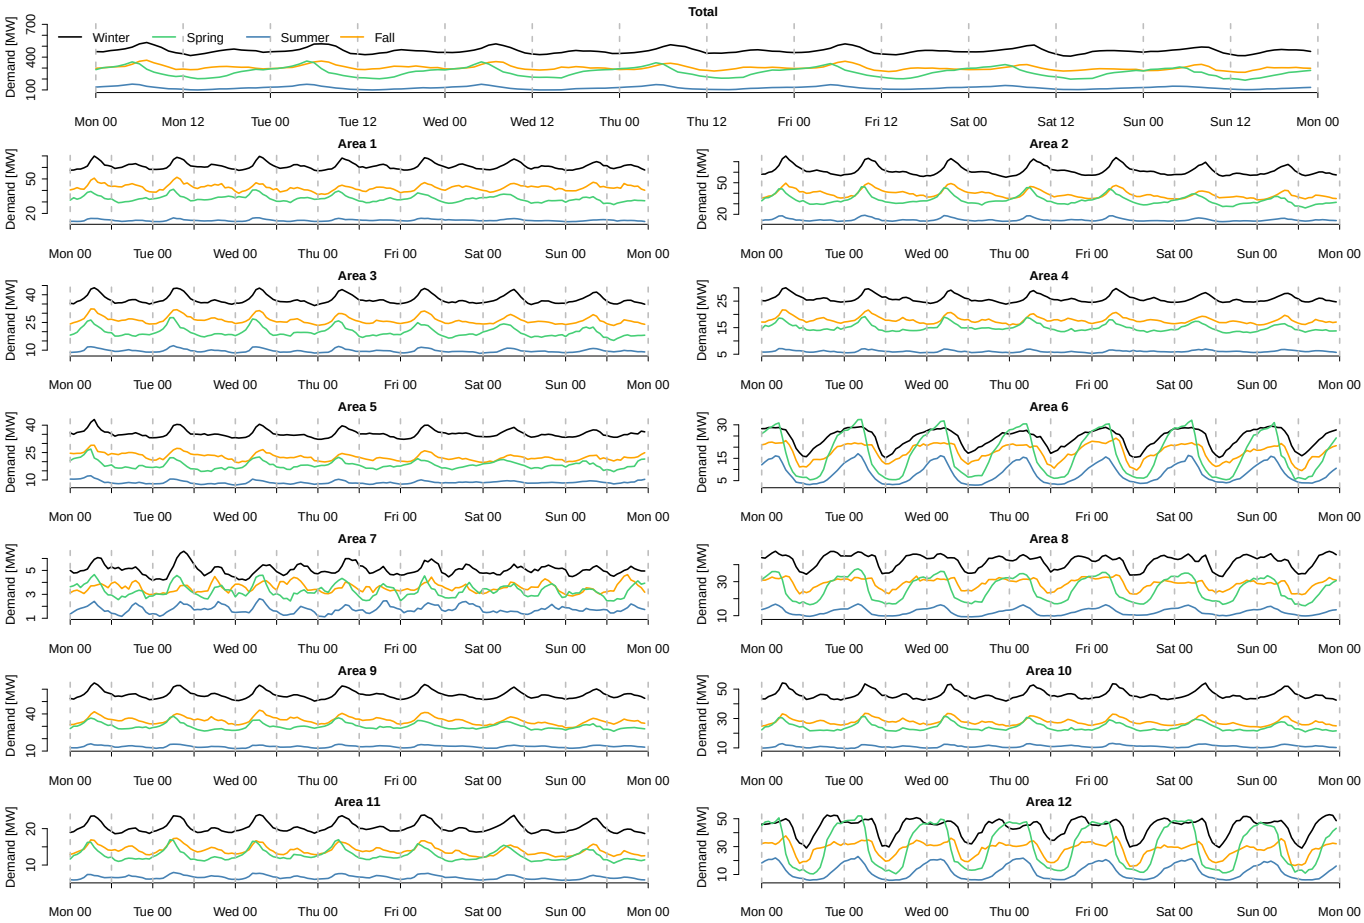

Figure A.16: The weekday seasonality for the areas at *Fjernvarme Fyn* is visualized here by splitting it up into four seasons; Winter (December, January, February, Mars), Spring (April, May), Summer (June, July, August, September), and Fall (October, November).

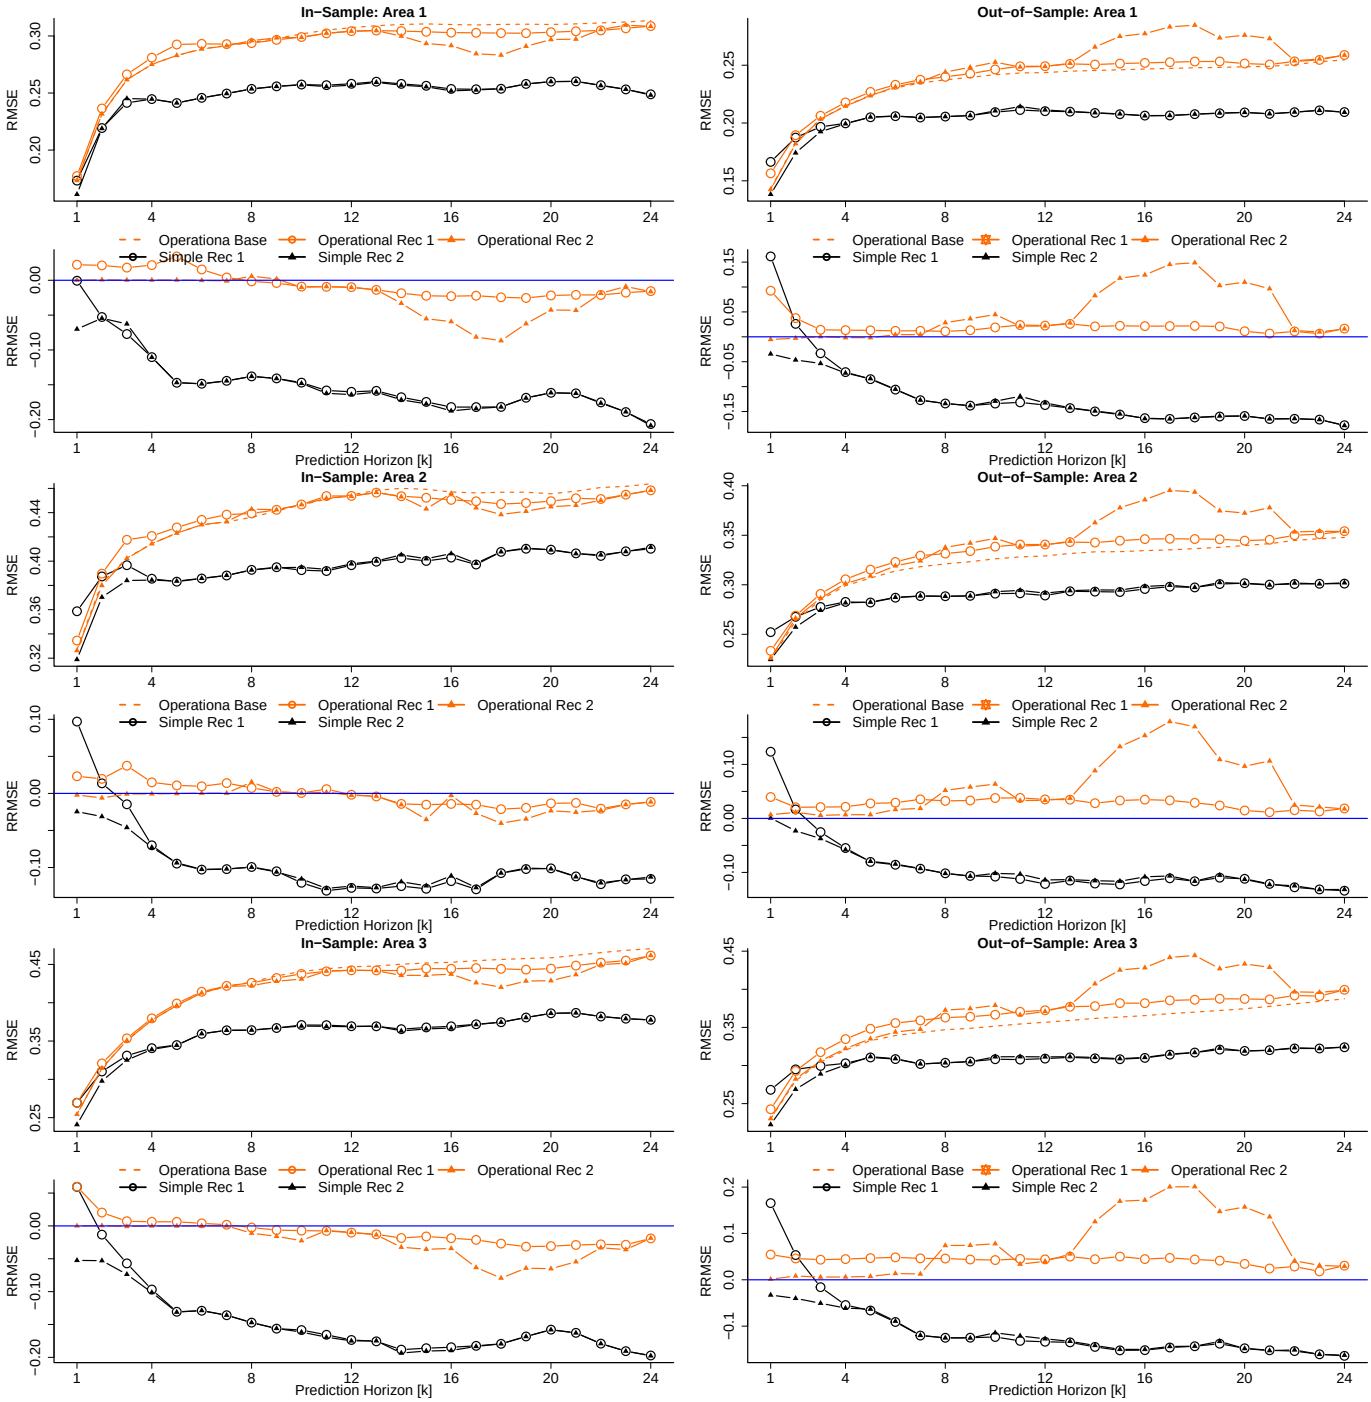

Figure A.17: Figure illustrates accuracy improvements of the reconciled forecast using aggregation hierarchy for in-sample and out-of-sample for all areas and total.

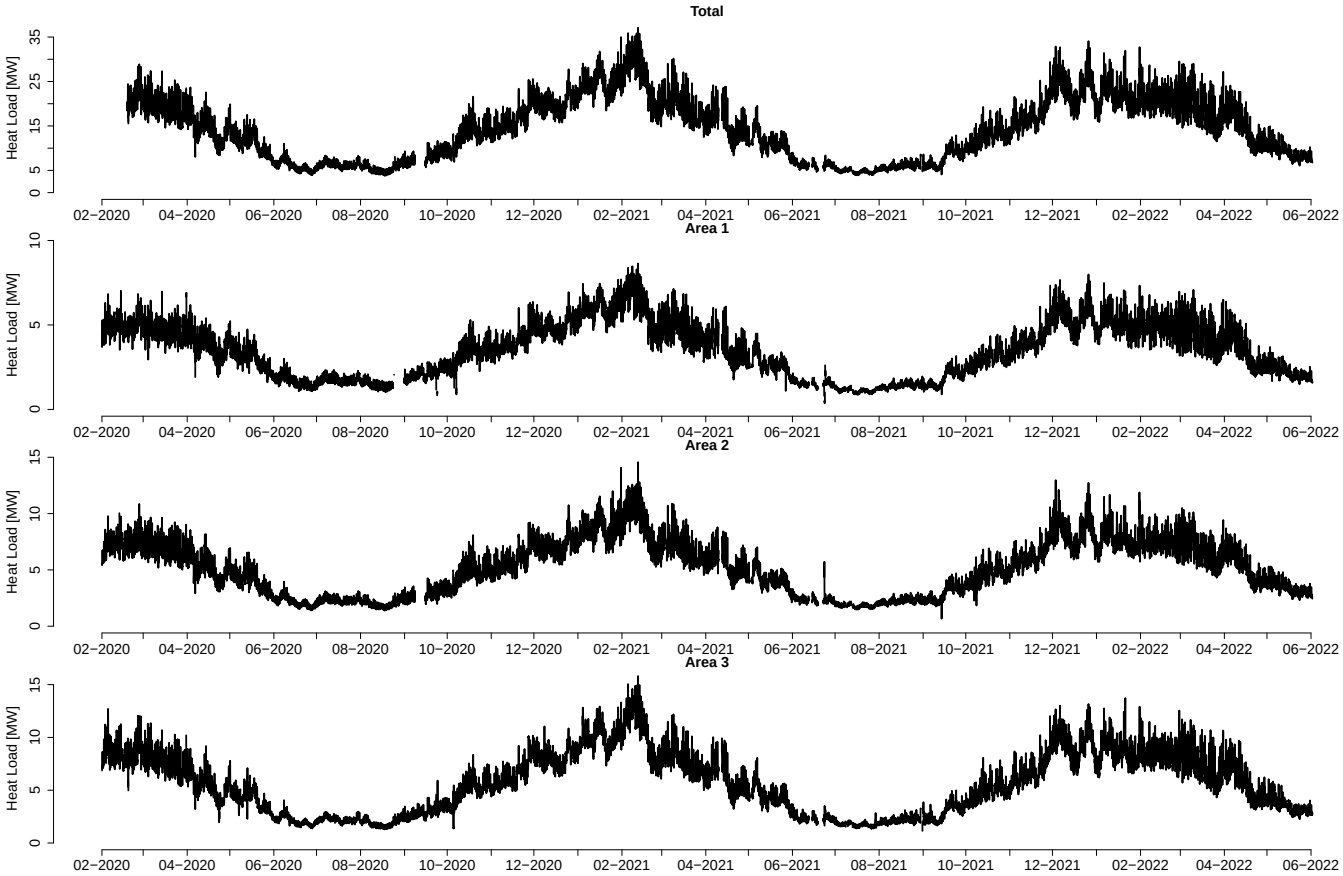

Figure A.18: Time series plot of the heat load over the two and half year period for each group inside the *Brønderslev Forsyning* system.

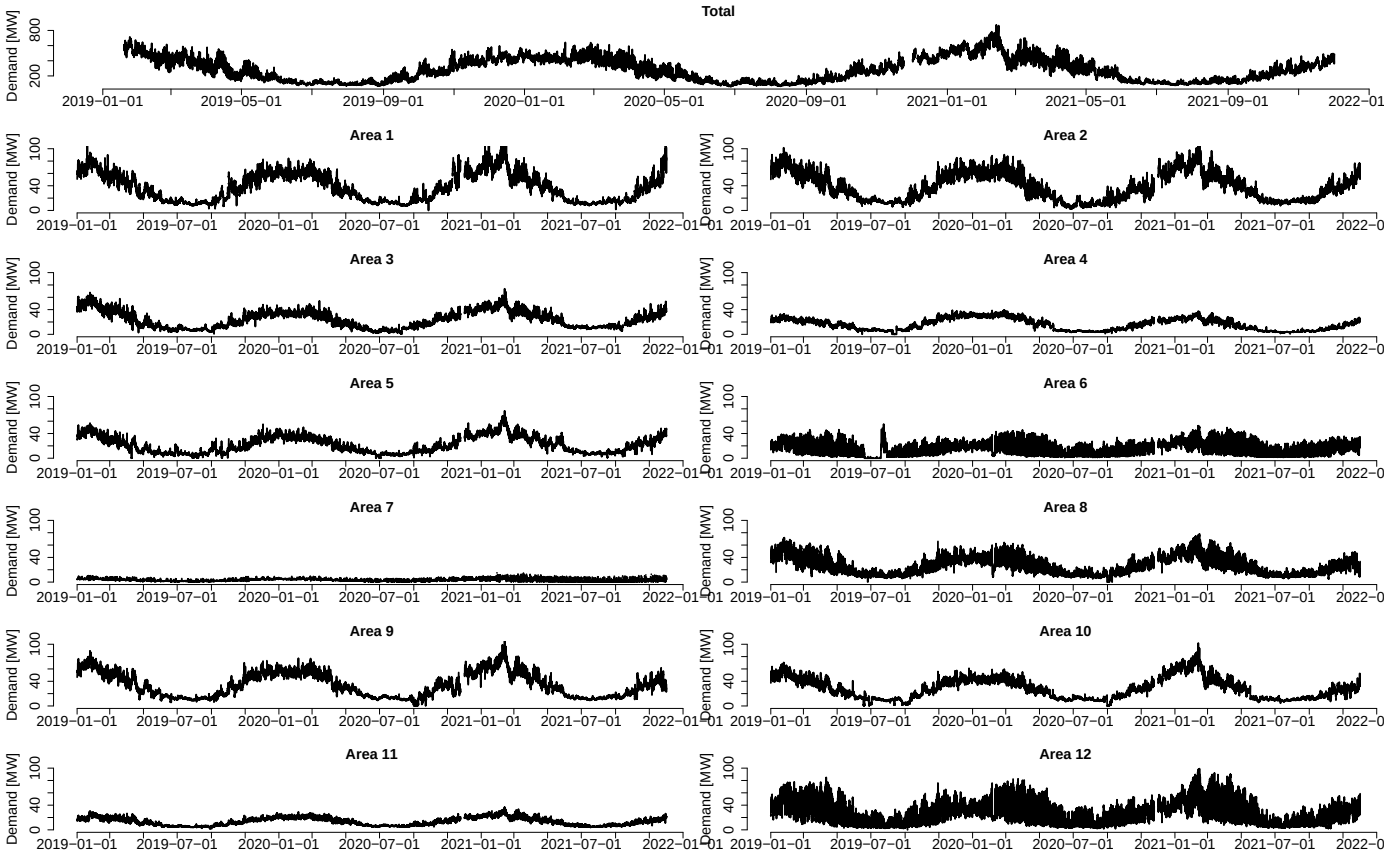

Figure A.19: Time series plot of the heat load over the two and half year period for each group inside the *Fjernvarme Fyn* system.

Appendix A Code Example

The code snippet below shows the `onlineforecast` heat load forecast from Section 3.1.1.

```
#### Data ####
names(D)

#### Periods ####
## ForecastPeriod is to slice the data frame
ForecastPeriod <- c(as.POSIXct("2020-02-01 00:00", tz = "UTC"), as.POSIXct("
    2021-12-31 23:00:00", tz = "UTC"))
Dfit <- subset(D, ForecastPeriod)

## TrainDates is the training period
TrainDates <- c(as.POSIXct("2020-02-01 00:00", tz = "UTC"),as.POSIXct("
    2020-12-31 23:00:00", tz = "UTC"))
Dtrain <- subset(Dfit, TrainDates)

## The training period
Scoreperiod_start <- as.POSIXct("2020-03-01 00:00", tz = "UTC")
Scoreperiod_end <- as.POSIXct("2020-12-31 23:00", tz = "UTC")
Dtrain$scoreperiod <- in_range(Scoreperiod_start, Dtrain$t,Scoreperiod_end)

## Out-Of-Sample period
Scoreperiod_start <- as.POSIXct("2021-01-01 00:00:00", tz = "UTC")
Scoreperiod_end <- as.POSIXct("2021-12-31 23:00:00", tz = "UTC")
Dfit$scoreperiod <-  in_range(Scoreperiod_start, Dfit$t,Scoreperiod_end)

#### Model 1
model <- forecastmodel$new()
model$output <- "Y" # The response variable, Y (heat load)

## Inputs
model$add_inputs(Ta = "Ta", # Air temperature, NWP
                 mu_tday = "fs(tday/24, nharmonics = 3)", # harmonics
                 mu = "one()")  # Intercept

## Bounds for the lambda in the optimiser
model$add_regprm("rls_prm(lambda=0.9)")
model$add_prmbounds(lambda = c(0.8,0.99,0.9999))

## Which horizons to use for the optimiser, i.e. only uses error from these
    horizons
model$kseq <- c(6,12,24,48)
model$prm <- rls_optim(model, Dtrain, printout = FALSE)$par

## Use the optimise forgetting factor to produce forecast for all horizon
model$kseq <- 1:72
val1 <- rls_fit(model$prm, model, Dfit, returnanalysis = TRUE)
## Save the predictions
Dfit$Yhat <- val1$Yhat

#### Model 2
model2 <- forecastmodel$new()
model2$output <- "Y"  # The response variable, Y (heat load)
## Inputs
model2$add_inputs(Ta = "lp(Ta, a1 = 0.9)", # now low pass the temperature
                 mu_tday = "fs(tday/24, nharmonics = 3)", # harmonics
                 mu = "one()") # Intercept
```

```
## Bounds for the lambda in the optimiser
model2$add_regprm("rls_prm(lambda=0.9)")
model2$add_prmbounds(Ta__a1 = c(0.5,0.99,0.9999),
                     lambda = c(0.5,0.99,0.9999))
## Which horizons to use for the optimiser, i.e. only uses error from these horizons
model2$kseq <- c(6,12,24,48)
model2$prm <- rls_optim(model2, Dtrain, printout = FALSE)$par
(model2$prm)
## Use the optimise forgetting factor to produce forecast for all horizon
model2$kseq <- 1:72
val2 <- rls_fit(model2$prm, model2, Dfit, returnanalysis = TRUE)
## Save the predictions
Dfit$Yhat2 <- val2$Yhat


### Produce the plots
kRanal <- 1
Residuals1 <- cbind(t = Dfit$t, residuals(Dfit$Yhat, Dfit$Y))
Residuals2 <- cbind(t = Dfit$t, residuals(Dfit$Yhat2, Dfit$Y))


par(mfrow = c(1,3))
par(cex = 1.5, cex.lab = 1.5, cex.axis = 1.5, lwd = 3)
par(mar=c(4.5,4.5,1.5,1)) #bottom, left, top and right
acf(Residuals1$h1[Dfit$trainScore], na.action = na.pass, xlab = "Lag (hours)")
boxplot(Residuals1$h1[Dfit$trainScore] ~ Dfit$tday$k0[Dfit$trainScore], xlab = "Hours", ylab = "One-Step Error", pch = 16, cex = 0.5)

x <- Dfit$TaObs[Dfit$trainScore] ## Input to the model
y <- Residuals1$h1[Dfit$trainScore] ## Output to the model
ccf(x, y, na.action = na.pass, lag.max = 48, main = "", xlab = "Lag (hours)", ylab = "CCF")
mtext(expression("T"["a,obs"]), side = 3, line = -4.5, outer = TRUE, col = 1, cex = 2, adj = 0.87)
mtext("Model 1 One-Step Error Analysis", side = 3, line = -1.2, outer = TRUE, col = 1, cex = 2)

par(mfrow = c(2,1))
par(cex = 1.5, cex.lab = 1.2, cex.axis = 1, lwd = 3)
par(mar=c(4,4.5,0.5,1)) #bottom, left, top and right

plot(Dfit$t[Dfit$scoreperiod], val1$Lfitval$k4[Dfit$scoreperiod,ncol(val1$Lfitval$k4)], type = "l",
     xlab = "", ylab = expression(theta[0]), axes = F, ylim = c(0,12), col = "#F17720")
lines(Dfit$t[Dfit$scoreperiod], val2$Lfitval$k4[Dfit$scoreperiod,ncol(val2$Lfitval$k4)], col = "#0474BA")

legend("bottomleft", legend=c("Model 1", "Model 2"), ncol = 2,
       col=c("#F17720", "#0474BA"), lty=1, cex=0.8, bty = "n")

axis(2, seq(0,12,2))


plot(Dfit$t[Dfit$scoreperiod],val1$Lfitval$k4[Dfit$scoreperiod,1], type = "l",
     xlab = "", ylab = expression(theta[1]), axes = F, ylim = c(-0.7,0), col = "#F17720")
lines(Dfit$t[Dfit$scoreperiod], val2$Lfitval$k4[Dfit$scoreperiod,1], col = "#0474BA")
axis(2, seq(-0.7,0,0.1))

```

```
105
106 axis.POSIXct(side = 1, x = Dfit$t[Dfit$scoreperiod],  xaxt="s",
107               at = seq(from = Dfit$t[Dfit$scoreperiod][1] ,
108                        to = Dfit$t[Dfit$scoreperiod][length(Dfit$t[Dfit$
                             scoreperiod])] + 60*60*24*2,
109                        by = "month"), format = "%Y-%m-%d",  mgp=c(3,-0.1,-1))
110
111 Residuals1 <- cbind(t = Dfit$t, residuals(Dfit$Yhat2, Dfit$Y))
112
113 par(mfrow = c(1,3))
114 par(cex = 1.5, cex.lab = 1.5, cex.axis = 1.5, lwd = 3)
115 par(mar=c(4.5,4.5,1.5,1)) #bottom, left, top and right
116 acf(Residuals1$h1[Dfit$trainScore], na.action = na.pass, xlab = "Lag (hours)")
117 boxplot(Residuals1$h1[Dfit$trainScore] ~ Dfit$tday$k0[Dfit$trainScore], xlab =
      "Hours", ylab = "One-Step Error", pch = 16, cex = 0.5)
118
119 x <- Dfit$TaObs[Dfit$trainScore] ## Input to the model
120 y <- Residuals1$h1[Dfit$trainScore] ## Output to the model
121 ccf(x, y, na.action = na.pass, lag.max = 48, main = "", xlab = "Lag (hours)",
    ylab = "CCF")
122 mtext(expression("T"["a,obs"]), side = 3, line = -4.5, outer = TRUE, col = 1,
    cex = 2, adj = 0.93)
123 mtext("Model 2 One-Step Error Analysis", side = 3, line = -1.2, outer = TRUE,
    col = 1, cex = 2)
124
125
126 par(mfrow = c(2,1))
127 par(cex = 1.5, cex.lab = 1.2, cex.axis = 1, lwd = 3)
128 par(mar=c(4,4.5,0.5,1)) #bottom, left, top and right
129
130 plot(Dfit$t[Dfit$scoreperiod], val1$Lfitval$k4[Dfit$scoreperiod,ncol(val1$
    Lfitval$k4)], type = "l",
131      xlab = "", ylab = expression(theta[0]), axes = F, ylim = c(0,12), col = "
         #F17720")
132 lines(Dfit$t[Dfit$scoreperiod], val2$Lfitval$k4[Dfit$scoreperiod,ncol(val2$
    Lfitval$k4)], col = "#0474BA")
133
134 legend("bottomleft", legend=c("Model 1", "Model 2"), ncol = 2,
135        col=c("#F17720", "#0474BA"), lty=1, cex=0.8, bty = "n")
136
137 axis(2, seq(0,12,2))
138
139
140 plot(Dfit$t[Dfit$scoreperiod],val1$Lfitval$k4[Dfit$scoreperiod,1], type = "l",
141      xlab = "", ylab = expression(theta[1]), axes = F, ylim = c(-0.7,0), col =
          "#F17720")
142 lines(Dfit$t[Dfit$scoreperiod], val2$Lfitval$k4[Dfit$scoreperiod,1], col = "
    #0474BA")
143 axis(2, seq(-0.7,0,0.1))
144
145
146 axis.POSIXct(side = 1, x = Dfit$t[Dfit$scoreperiod],  xaxt="s",
147               at = seq(from = Dfit$t[Dfit$scoreperiod][1] ,
148                        to = Dfit$t[Dfit$scoreperiod][length(Dfit$t[Dfit$
                             scoreperiod])] + 60*60*24*2,
149                        by = "month"), format = "%Y-%m-%d",  mgp=c(3,-0.1,-1))
```

Technical University of Denmark
Department of Applied Mathematics and Computer Science

Richard Petersens Plads, Building 324
2800 Kgs. Lyngby
Tel.: +45 4525 3031

www.compute.dtu.dk