



LUND UNIVERSITY

Control of Capacity-Constrained Networks

Agner, Felix

2025

Document Version:

Publisher's PDF, also known as Version of record

[Link to publication](#)

Citation for published version (APA):

Agner, F. (2025). *Control of Capacity-Constrained Networks*. [Doctoral Thesis (compilation), Department of Automatic Control]. Department of Automatic Control, Lund University.

Total number of authors:

1

General rights

Unless other specific re-use rights are stated the following general rights apply:

Copyright and moral rights for the publications made accessible in the public portal are retained by the authors and/or other copyright owners and it is a condition of accessing publications that users recognise and abide by the legal requirements associated with these rights.

- Users may download and print one copy of any publication from the public portal for the purpose of private study or research.
- You may not further distribute the material or use it for any profit-making activity or commercial gain
- You may freely distribute the URL identifying the publication in the public portal

Read more about Creative commons licenses: <https://creativecommons.org/licenses/>

Take down policy

If you believe that this document breaches copyright please contact us providing details, and we will remove access to the work immediately and investigate your claim.

LUND UNIVERSITY

PO Box 117
221 00 Lund
+46 46-222 00 00

Control of Capacity- Constrained Networks

in a District Heating Setting

FELIX AGNER

DEPARTMENT OF AUTOMATIC CONTROL | LUND UNIVERSITY



Control of Capacity-Constrained Networks

in a District Heating Setting

Felix Agner



LUND
UNIVERSITY

Department of Automatic Control

The cover page illustration is painted by Marie-Louise Agner.
Printed with permission.

Ph.D. Thesis TFRT-1146
ISBN 978-91-8104-288-7 (print)
ISBN 978-91-8104-289-4 (web)
ISSN 0280-5316

Department of Automatic Control
Lund University
Box 118
SE-221 00 LUND
Sweden

© 2025 by Felix Agner. All rights reserved.
Printed in Sweden by Media-Tryck.
Lund 2025

Till min Far.

Abstract

This thesis concerns control of capacity-constrained networks. These systems involve many agents interconnected by a resource distribution network. The capacity to generate and distribute this resource is constrained. This applies, for instance, to power grids, communication networks, smart surveillance camera networks, and district heating networks. District heating networks in particular are the main focus of this thesis. These systems distribute heat from producers to consumers through hot water pipelines. In this setting, the agents are the consumers in the network, who regulate the flow rate they receive from the network using control valves. Physical limitations limit these flow rates. Therefore, when the demand for heat is high, it may be impossible to satisfy the needs of all the agents. This can result in certain buildings becoming cold. This thesis presents several contributions in this setting.

Firstly: The nature of the flow rate constraints is investigated. In Paper I, it is shown that the set of feasible flow rates in a tree-structured district heating network is convex, allowing for convex optimization-based control structures. One such approach is proposed in the paper, in which the flow rates are distributed fairly between the agents. These control approaches require a model of the system hydraulics. In Paper V, a data-based method for establishing such a model is investigated in a laboratory environment.

Secondly: The limited network capacity should be utilized optimally. This is challenging in the multi-agent setting, where the agents regulate and actuate the flow of resources in a decentralized fashion. Papers II-IV concern controllers which not only asymptotically guide the network to an optimal resource distribution, but also function in the large-scale, multi-agent setting. These papers show that asymptotic optimality guarantees can be established using variations of standard proportional-integral control. Papers II and III concern a linear system setting with input saturation, which is extended to a nonlinear setting in Paper IV.

Acknowledgements

I am in the fortunate position to owe a long list of people my gratitude for their support during the last five years. Let me begin with supervisors. Anders Rantzer, thank you for the curiosity and encouragement you have given my research. You have always been supportive of me pursuing my own interests and allowed me to stay close to my intrinsic motivation. Richard Pates, thank you for always giving insightful feedback and perspective on my work. Finally, Pauline Kergus, you are probably the person who has spent the most time and effort on providing detailed feedback on my writing and my work. This has been invaluable and has taught me many of the writing skills that are otherwise hard to learn.

Thank you also, Pauline, for inviting me to Toulouse and subsequently sparking many core ideas in this thesis. Thank you professors Sophie Tarbouriech and Luca Zaccarian for the discussions and continued collaboration and thank you to all of the PhD students at the LAPLACE laboratory for the warm welcome.

Likewise, I want to thank the control department at Aalborg University for my visit to your lab. In particular Rafal Wisniewski for inviting me, my co-authors Christian Møller Jensen and Carsten Skovmose Kallesøe, and finally Saruch Rathore for guiding me around the department and local pubs.

I cannot thank all of the lovely people at the Department of Automatic control personally, as this would make too long of a list. You have all contributed to making this a great place to work for the last years, and I will sorely miss the discussion I have shared with you all over coffee. Thank you to our administrative staff for working so diligently to make our days run smoothly. Thank you to the technical staff for all of the times I came running to you when my computer or lab equipment malfunctioned. Thank you to the faculty members for your open doors and your insightful views. And finally thank you fellow PhD students for everything. Jonas, you deserve a particular thank you for being the PhD student to whom I have most often turned to discuss my research problems.

I would like to finish these acknowledgments by thanking my nearest and dearest. Thank you to my mother and father who have supported my educational interests from the day I started school to the day I defend this thesis. Thank you, Frida, for always listening, supporting and encouraging.

Acknowledgement of Funding

This work is funded by the European Research Council (ERC) under the European Union's Horizon 2020 research and innovation program under grant agreement No 834142 (ScalableControl).

The author of this thesis is a member of the ELLIIT Strategic Research Area at Lund University.

Contents

1. Introduction	11
1.1 Applications	12
1.2 Thesis Outline	13
2. Control Problem and Background	14
2.1 Motivating Examples	14
2.2 Problems Regarding Control of Large-Scale Systems	18
2.3 Related Control Literature and Background	20
3. District Heating and Cooling Networks	28
3.1 Background	29
3.2 Control of District Heating Networks	31
3.3 Hydraulic Models	33
3.4 Thermodynamic Models	38
4. Contributions	41
4.1 Included Papers	41
4.2 Additional Publications	45
5. Conclusion	46
Bibliography	49
Paper I. Combating District Heating Bottlenecks Using Load Control	57
1 Introduction	58
2 List of Notation	60
3 Problem and System Formulation	61
4 Control Strategies	67
5 Simulation and Results	71
6 Summary	75
7 Acknowledgements	76
References	78
Paper II. Anti-windup Coordination Strategy Around a Fair Equilibrium in Resource Sharing Networks	81
1 Introduction	82

2	System Description and Problem Formulation	83
3	Closed-Loop Equilibria	85
4	Optimality	90
5	Stability Properties	90
6	Numerical Example	92
7	Conclusion	93
	References	94
Paper III. Decentralized PI-control and Anti-windup in Resource Sharing Networks		97
1	Introduction	98
2	Problem Data and Proposed Controller	99
3	Main Results	101
4	Numerical Example	103
5	Proof of Equilibrium Existence and Uniqueness	105
6	Proof of Global Asymptotic Stability	107
7	Proof of Equilibrium Optimality	109
8	Conclusions	110
	References	112
Paper IV. On PI-control in Capacity-Limited Networks		115
1	Introduction	116
2	Problem Formulation	118
3	Considered Controllers and Main Results	119
4	Motivating Example - District Heating	122
5	Main Proofs	126
6	Conclusion	134
7	Acknowledgments	134
8	Remaining Proofs	135
	References	136
Paper V. Hydraulic Parameter Estimation for District Heating Based on Laboratory Experiments		141
1	Introduction	142
2	Hydraulic Parameter Estimation	144
3	Model Structures and Data Preprocessing	149
4	Experiment Description	151
5	Results	155
6	Conclusion	162
7	Future work	165
	References	169

1

Introduction

This thesis concerns control of systems where a large number of agents all share a common interconnection of limited capacity. This capacity represents e.g., the ability to deliver energy to consumers or send data to users. The concrete motivating example at the core of this thesis is district heating networks. A simplified yet typical way that these systems are operated is as follows: A large, central heating plant produces water of a desired temperature. This water is then circulated through a network of pipes to the consumers. These consumers are physically connected via the network, but they regulate the amount of hot water they receive through individual control valves. The amount of water that can be distributed to the consumers is limited due to the physical constraints of the components that control the flow rates: Valves cannot expand further once they are fully open, and pumps cannot exceed their maximum capacity. Therefore, when the outdoor temperature is critically low and the heat demand is high, it may be impossible to fully satisfy the needs of all consumers. In particular, high heat demands lead to high volume flow rates of water. Increased flow rates in the pipes induce frictional pressure losses. This is illustrated in Figure 1.1. These losses accumulate along the length of the network. Hence the pressure in the periphery of the network can drop to critically low levels. The peripheral consumers can counteract this effect in order to receive their desired flow rates by opening their valves, but this only works as long as the valves are not fully open. Once they are fully open, further pressure losses mean that the peripheral consumers cannot receive the flow rates that they require. This affects the indoor climate of these buildings: They become cold. In this sense, a district heating system is capacity-constrained; only a limited amount of water can be circulated. In extreme cases this will be insufficient to reject the disturbance (in this case the outdoor temperature) acting on the interconnected agents. It becomes impossible to stabilize the system at the desired set-point where all of the buildings are at comfort temperature.




Figure 1.1 A line-structure district heating network, where one central plant supplies the network with hot water. The pressure, indicated in the figure by the shaded color from red to blue, is higher near the distribution pump at the root of the network and lower in the periphery. This is caused by frictional pressure losses in the pipes. When the flow rates in the pipes increase, so do these losses. If the flow rates and hence the losses become too high, the pressure at the end of the network will be too low. Therefore, peripheral consumers will receive insufficient amounts of hot water.

1.1 Applications

Control problems of this or similar forms arise in many interesting engineering domains. Pioneering works considered the issue of congestion control in communication networks [Kelly et al., 1998; Low and Lapsley, 1999; Low et al., 2002; Kelly, 2003]. In this domain the issue lies in limited bandwidth, shared by a large number of connected users. Additionally, users must understand the current state of the network based solely on local measurements. Another area where problems of this form arise is energy transportation networks of different varieties. For instance, in power systems, optimal power flow concerns driving the system to an optimal equilibrium while leveraging constrained active and reactive power inputs [Molzahn et al., 2017; DallAnese and Simonetto, 2018; Biegel et al., 2012]. In heating, ventilation, and air conditioning (HVAC) systems, hydraulic cooling circuits can experience issues similar to those found in district heating systems. That is, the pressure in the system has to be managed such that all connected units receive sufficient flow rates for cooling purposes [Kallesøe et al., 2019; Kallesøe et al., 2020]. In modern distributed camera systems, limited communication bandwidth and, crucially, storage space for footage is shared between smart cameras. Hence this limited resource has to be distributed between the cameras [Martins et al., 2020; Martins and Årzén, 2021]. Finally, the main focus of this thesis is district heating networks. The above list of applications is necessarily non-extensive but highlights the applicability of considering control of capacity-constrained networks.

1.2 Thesis Outline

Chapter 2 introduces motivating control-theoretic examples along with associated control approaches and relevant mathematical background. Chapter 3 covers the necessary background on district heating systems and modeling of building thermodynamics. The contributions of this thesis are presented in Chapter 4. The introduction to this thesis concludes with a discussion and proposed future work in Chapter 5. After the introductory chapters, the five papers which constitute the contribution of this thesis are attached.

This work builds on the licentiate thesis [Agner, 2023] by the present author. Hence certain images and excerpts of text in this introduction are taken directly from the introduction of said licentiate thesis, potentially with minor variations.

Notation

The following basic notation is used in the introduction of this thesis. The set of real numbers is denoted \mathbb{R} , the set of real vectors of size n is denoted \mathbb{R}^n , and the set of real matrices of size $n \times m$ is denoted $\mathbb{R}^{n \times m}$. For a vector $v \in \mathbb{R}^n$ (or a matrix $M \in \mathbb{R}^{n \times m}$), we say that $v > 0$ (or $M > 0$) if v (or M) is element-wise positive. Likewise, $v \geq 0$ (and $M \geq 0$) denotes element-wise non-negativity. For a symmetric matrix $M \in \mathbb{R}^{n \times n}$, $M \succ 0$ denotes that M is positive definite. The vector $\mathbf{1}$ denotes a vector of all 1's, of suitable dimension as taken from context. The l_1 , l_2 , and l_∞ -norms of a vector $v \in \mathbb{R}^n$ are denoted $\|v\|_1 = \sum_{i=1}^n |v_i|$, $\|v\|_2 = \sqrt{\sum_{i=1}^n |v_i|^2}$, and $\|v\|_\infty = \max_{i \in 1, \dots, n} |v_i|$ respectively.

2

Control Problem and Background

This chapter seeks to introduce the control-theoretic perspective on the problems concerned in this thesis. To do so, we will begin with a small illustrative example. Then we will delve into the additional problems and design-constraints that emerge from the large-scale and multi-agent setting. This is followed by background on relevant and related control-approaches and mathematical concepts.

2.1 Motivating Examples

EXAMPLE 1

Consider the following dynamical system, which describes two interconnected agents where x_1 and x_2 are the states of each agent.

$$\begin{bmatrix} \dot{x}_1 \\ \dot{x}_2 \end{bmatrix} = - \begin{bmatrix} x_1 \\ x_2 \end{bmatrix} + \begin{bmatrix} 4 & -2 \\ -1.5 & 3 \end{bmatrix} \begin{bmatrix} \text{sat}(u_1) \\ \text{sat}(u_2) \end{bmatrix} + \begin{bmatrix} 1 \\ 1 \end{bmatrix} w \quad (2.1)$$

These states represent the offset from some reference value and should hence be regulated to zero. u_1 and u_2 represent the agents control actions, and w represents a disturbance. The saturation function $\text{sat}(u) = \max(\min(u, 1), -1)$ bounds the control actuation to the range $[-1, 1]$, representing physical limitations on the actuators. We can note that if agent 1 increases their control action, it has a positive impact on \dot{x}_1 but a negative effect on \dot{x}_2 , and vice versa. This reflects a form of competition for resources between the agents.

We now equip each agent with a proportional-integral controller with anti-windup (to be explained later in this chapter). Figure 2.1 shows simulations of this system, subject to disturbances $w = 1$, $w = 2$, and $w = 3$ respectively. With $w = 1$, the disturbance is fully rejected, and both agents reach the equilibrium $x_1 = x_2 = 0$. With $w = 2$, only agent one can reach $x_1 = 0$, whereas agent 2 converges to $x_2 = 0.5$, i.e., a constant control error. Finally, with $w = 3$, neither agent can reject the disturbance, and both agents align at a constant tracking error.





Figure 2.1 States x from Example 1, subjected to different disturbances w . With $w = 1$ (blue), both agents are able to reject the disturbance fully. With $w = 2$ (green), agent 1 can reject the disturbance, but agent 2 reaches a constant control error. With $w = 3$ (red), neither agent is able to reject the disturbance.

In this example, the agents are interconnected through their control actions. The control actions distribute a resource to the agents, where agent 1 receives the resource $4\text{sat}(u_1) - 2\text{sat}(u_2)$ and agent 2 receives the resource $3\text{sat}(u_2) - 1.5\text{sat}(u_1)$. Due to the saturation function, the capacity to provide the agents with their desired resource is constrained. Therefore, when the disturbance w becomes sufficiently large, the constrained capacity may be insufficient to stabilize the system such that $x_1 = x_2 = 0$. The new objective of the controllers in this setting should therefore be to stabilize the best possible equilibrium, given the current disturbance. The notion of the "best" equilibrium will be discussed momentarily. For the system (2.1), the set of all possible stable equilibria can be visualized. Figure 2.2a shows all such equilibria, i.e., states x that could solve (2.1) with $\dot{x} = 0$ given a corresponding control input u , subject to different disturbance levels w . Unsurprisingly, for the small disturbance levels $w = 0$ and $w = 1$, the origin $x_1 = x_2 = 0$ is found within the bounds of the gray and blue areas meaning that this is a possible equilibrium. For $w = 2$ and $w = 3$, however, this is not the case. The green and red areas, which represent these two disturbance levels, do not contain the origin. To determine the optimal equilibrium among these, we can compare them in terms of different objective functions $J(x)$, and placing preference on the equilibrium that minimizes said objective. Natural considerations for such objective functions are the l_1 , l_2 , and l_∞ -norms. In particular, $\|x\|_1$ can be interpreted as the *total* error experienced by all agents, and $\|x\|_\infty$ represents the worst error experienced by any agent. Figure 2.2b shows the equilibria that minimize these two objectives. In this example, we find that there is a trade-off between minimizing the total deviation (which minimally becomes $\|x\|_1 = 2.5$) and the worst-case deviation (which minimally becomes $\|x\|_\infty = 1.364$). We will now consider an example of a large-scale system consisting of many agents in order to further illustrate this equilibrium trade-off.



(a) The sets of all possible states x that can be stabilized for dynamics (2.1) in Example 1, given different disturbance levels w .

(b) A zoomed-in view of the red area of Figure 2.2a. The two highlighted equilibria $(1.0, 1.5)$ and $(1.364, 1.364)$ minimize $\|x\|_1$ and $\|x\|_\infty$ respectively.

Figure 2.2 A visualization of the possible equilibria for system (2.1).

EXAMPLE 2

Consider a system consisting of $n = 500$ agents. The dynamics of each agent $i = 1, \dots, n$ is given by

$$\dot{x}_i = -x_i + 1.1 \text{sat}(u_i) - \frac{i}{n^2} \sum_{j \neq i} \text{sat}(u_j) + w. \quad (2.2)$$

Let the disturbance $w = 1$. This disturbance acts equally on every agent. Here the agents share an interconnection through which they are granted a resource used to reject the disturbance w . For agent i , the term $-\frac{i}{n^2} \sum_{j \neq i} \text{sat}(u_j)$ encodes that when other agents j increase their control action, they utilize part of the capacity of the interconnection between the agents. This has a negative impact on the resource granted to agent i . This effect grows larger as i approaches n , meaning that agent n is heavily affected by the other agents, whereas agent 1 is not. Figure 2.3 displays the optimal equilibria of this system with respect to the norms $\|x\|_1$, $\|x\|_2$ and $\|x\|_\infty$, along with the corresponding control actuation $\text{sat}(u)$. For an agent i , the magnitude of $\text{sat}(u_i)$ dictates how much this agent should prioritize providing themselves with the shared resource ($\text{sat}(u_i)$ large) or avoid taking resource from others ($\text{sat}(u_i)$ small). In the equilibrium minimizing $\|x\|_\infty$, the resource is divided so that all agents receive the same error. This corresponds to control actions where agent 500 uses their full control capacity, and all agents proportionally reduce their control action to avoid taking resources from others. In the equilibrium minimizing $\|x\|_1$, all agents (almost) utilize their full control capacity, with the exception of


(a) Optimal states x .(b) Optimal control actuation $\text{sat}(u)$.

Figure 2.3 The optimal equilibria for system (2.2) with respect to $\|x\|_1$, $\|x\|_2$ and $\|x\|_\infty$ for $w = 1$. The figures are arranged by agent index on the horizontal axis.

agents 1-50. These agents already reach an equilibrium state of 0. The effect of all agents utilizing their maximum control capacity is that agents with higher index reach larger control errors. Finally, the equilibrium minimizing $\|x\|_2$ reaches a middle-ground between the first two. Here the largest errors are penalized the strongest, but no errors are fully disregarded. Hence all agents will reach at least a certain level of state error, with the purpose of leaving resources for the agents with the largest errors.

The objective function $J(x)$ represents a lens through which we view the optimality of a particular equilibrium. In Example 2, we can note that $\|x\|_\infty$ encodes a notion of fairness: The limited system capacity is allocated to ensure that the error of the agent which is most impacted by the current disturbance is minimized. On the other hand, the objective $\|x\|_1$ considers a form of system level optimum: The equilibrium in which total errors are minimized.

To summarize the point of these previous examples, this thesis concerns the design of controllers which stabilize equilibria in which a limited resource capacity is distributed optimally among a large number of agents. The resulting closed loop systems should be *asymptotically optimal*, loosely defined in the following sense.

DEFINITION 1—ASYMPTOTIC OPTIMALITY

Consider a dynamical system $\dot{x} = f(x, u, w)$ with states x and control input u , subject to a constant disturbance w . Then a (possibly dynamical) control law is asymptotically optimal with respect to f and an objective function $J(x)$ if, from any admissible initial condition, the closed-loop system converges to an equilibrium (x^0, u^0) such that for any other equilibrium pair (x^\dagger, u^\dagger) subject to $0 = f(x^\dagger, u^\dagger, w)$,

$$J(x^0) \leq J(x^\dagger). \quad (2.3)$$

This form of design-objective stands in contrast to large portions of traditional control literature, where the goal is to drive the system to its reference state as "cheaply" as possible, governed by some notion of "cheap" defined in terms of objectives placed on state deviations and control inputs. In our problem, it is impossible to reach the reference state at all, and therefore the question of which state to reach instead takes precedence. This does not mean that the system dynamics are unimportant and can be disregarded in favor of considering only equilibrium properties. The dynamics still inform whether a given equilibrium is stable, and whether the transient period leading to said equilibrium is well-behaved. These questions are clearly important for real applications.

Definition 1 concerns optimality under objectives $J(x)$ placed on the system state variables x , which will be the focus of this thesis. In many control applications it would be natural to penalize the control input instead, i.e., objectives on the form $J(u)$. This is reasonable when large control inputs are associated with a real cost, for instance high consumption of fuel, or significant wear-and-tear. An even more general problem formulation would be objectives on the form $J(x, u)$ concerning both the state variables and control inputs. In the setting of controlling heat allocation in traditional district heating networks, the control inputs u naturally represent valve positions, whereas x can represent building temperature deviations. Here, there is no real cost associated with increased valve openings, whereas a very real issue arises if building temperatures deviate heavily from desired set-points. While there is a real fuel cost associated with generating heat, this thesis specifically concerns the scenario when the maximum heat injection does not satisfy the current demand. A maximum heat injection should naturally be utilized, and the question is rather how this insufficient heat is to be distributed optimally to the agents.

2.2 Problems Regarding Control of Large-Scale Systems

In the large-scale networked setting concerned in this thesis, several challenges related to control naturally arise. This is because a large number of interconnected agents are affecting each other, without being able to access global information on the system. It is often infeasible to have each agent maintain a perfect model of the entire network. It is also unlikely that each agent can measure or know the state of every other agent in the system. Therefore, it is typically of interest to design

controllers bound by structural constraints inherent to the problem at hand. In this section we will discuss several controller structures which lend themselves nicely to systems of large-scale. The terminology for different types of scalable controller structures is not completely streamlined through the literature, and sometimes the terminology is used interchangeably. For the purpose of this thesis, the terminology defined here will be used.

Decentralized Control

A decentralized controller is implemented such that each agent in the network acts fully independently. They do so without relying on global information. In terms of the controller design, this means that each agent's controller should be designed based on models associated only with the agent's own system. In terms of actuation, the real-time signals used for actuation should only be measured locally. In the setting of district heating, these constraints are rather natural. The controller governing flow rates to each consumer may be designed with a detailed model of the building where it is installed, but not of every building connected to the network. The controller might also use local measurements from the building where it operates but should not require measurements from buildings in distant parts of the network. Notably, early works on congestion control in communication networks are implemented in a decentralized manner. Here the network-connected users update their communication rates based only on local information [Kelly et al., 1998; Low and Lapsley, 1999; Low et al., 2002; Kelly, 2003]. Another notable example is primary frequency control in power systems, where local frequency measurements are used to inform the power input of turbines [Machowski et al., 2020, p.358]. The proportional-integral control of Example 1 also falls under this category.

Low-rank Communication

In some applications, a certain degree of global communication can provide great benefits over a fully decentralized structure. Rather than designing a control system based on fully dense communication between all agents, it can often be sufficient to allow the agents to agree on a few global signals to use in their control computations. This type of communication becomes low rank in the sense that a large number of signals from the agents are condensed into a small number of globally coordinated signals. Consider for instance the implementation of the control signal

$$u = vk^{\top}x$$

where $u \in \mathbb{R}^n$ is the control action, $v \in \mathbb{R}^n$ and $k \in \mathbb{R}^n$ are controller gains, and $x \in \mathbb{R}^n$ is the system state. In actuating this signal, the agents need only agree on the scalar value of $k^{\top}x$. This can be achieved either by having one central node of communication, or by implementing a consensus-protocol, which can be done in a distributed fashion. A design method for this type of low-rank control structure was proposed in [Zecevic and Siljak, 2005], and rank-1 coordination was considered in the application of wind farm control in [Madjidian and Mirkin, 2014].

Distributed Control

A distributed controller allows communication between agents which are, in some sense, near each other. Analysis of this type of control structure often arises when the physical process lends itself to an interpretation of neighboring agents. Car platooning[Levine and Athans, 1966; Chu, 1974] and buffer network control[Iftar and Davison, 1990; Blanchini et al., 2016] are notable examples. In car platooning, the aim is to guide a fleet of cars to drive at the same speed and at a constant distance from each other. Here it may be reasonable to consider that the cars can measure their distance to their neighbors, and hence a decentralized control structure can naturally be implemented. In control of buffer networks, one considers buffer tanks connected by edges that guide a flow between the tanks. By implementing this flow as a function of the state in the connected buffer tanks, a decentralized strategy is achieved.

Summary of Scalable Control Architectures

In this thesis, we will mainly consider decentralized and low-rank control, as these structures lend themselves nicely to the application of district heating networks. While a distributed control structure may seem reasonable at first glance due to the obvious network structure, it is not trivial how such a control strategy should be implemented in practice. It is straight-forward to define two buildings as "neighbors" in a geographical sense, but this does not necessarily imply that the controllers that govern their heat load lend themselves to a notion of adjacency. For instance, it is not necessarily natural for the controllers in two neighboring buildings to communicate their current indoor temperatures with each other, only because they are neighbors.

2.3 Related Control Literature and Background

There are many fields of automatic control which tackle input constraints, scalability issues and optimality in different fashions. This section will cover related topics, along with necessary background on basic control concepts. The section concludes with an introduction to M-matrices, which is a central concept to papers II and III of this thesis.

PI-control with Anti-Windup

One of the most fundamental yet universally useful controllers in the toolbox of control systems is the proportional-integral-derivative (PID) controller. For a single-input, single-output (SISO) system, with a control error x (which should be kept at zero), such a controller is typically written on the form

$$u(t) = -k^P x(t) - k^I \int x(\tau) d\tau - k^D \frac{dx(t)}{dt}.$$




Figure 2.4 Examples of the saturation and dead-zone functions $\text{sat}(\cdot)$ and $\text{dz}(\cdot)$, with upper-and-lower limits 1 and -1 respectively.

Here k^P , k^I , and k^D are proportional, integral, and derivative controller gains respectively. This thesis concerns only PI controllers without derivative action. Such controllers can be rewritten on state-space form:

$$\begin{aligned}\dot{z} &= x \\ u &= -k^P x - k^I z\end{aligned}$$

Here z denotes the integral of the control error. In the context of this thesis, it is likely that a considered system is subject to input constraints and large disturbances. A typical model of input constraints in the SISO setting is the saturation function $\text{sat}(u)$, as considered in Examples 1 and 2. This model applies to many real-world actuators. For instance, valves which operate between fully opened and fully closed or car engines which can apply a limited range of torques. When saturation occurs such that $\text{sat}(u) \neq u$, a constant and nonzero control error x can occur, subsequently inducing an unwanted behavior termed *windup*. This implies that the constant error makes the integral state z "wind up", building an indefinitely increasing integrated error [Hägglund, 2023, p.48]. To counteract this behavior, there is a family of strategies denoted *anti-windup* compensation. This typically involves measuring the difference between the desired and the actuated control input $u - \text{sat}(u) = \text{dz}(u)$, denoted by the dead-zone function $\text{dz}(u)$, and using this to inform the update of the integral error. One simple anti-windup solution is to simply seize the update of the integral error z when in the saturated region $\text{dz}(u) \neq 0$ [Hägglund, 2023]. Another family of methods is direct, linear anti-windup [Galeani et al., 2009]. This corresponds to measuring $\text{dz}(u)$, and having this signal enter into the update of the controller state in a linear fashion. In the context of PI control, direct, linear anti-windup would be implemented as such:

$$\dot{z} = x + k^A \text{dz}(u) \tag{2.4a}$$

$$u = -k^P x - k^I z. \tag{2.4b}$$

Here k^A is a scalar anti-windup gain. The functions $\text{sat}(u)$ and $\text{dz}(u)$ are both visualized in Figure 2.4. The saturation function (and hence also the dead-zone function) can be generalized to work also with upper and lower limits distinct from 1 and -1 as shown in this example. A schematic overview of the complete closed-loop structure is shown in Figure 2.5. The intuition behind this approach is that when u grows large, the term $k^A \text{dz}(u)$ acts positively on the change in z , thus reducing u . We can see that when u is sufficiently large such that $k^A \text{dz}(u) = -x$, the integral error z ceases to grow. Modern approaches to anti-windup design have been extended well beyond the SISO case, into multiple-input-multiple-output scenarios [Galeani et al., 2009]. Historically, anti-windup compensation design has been considered to improve transient behavior and stability margins of the closed-loop system. In particular, the goal has been to ensure that when saturation occurs, the closed-loop system should behave close to how it would behave without saturation [Galeani et al., 2009]. Furthermore, anti-windup design can be used to maximize the size of the basin of attraction [Silva and Tarbouriech, 2005]. This becomes important when the natural dynamics of the system are unstable, and the control input is required for stabilization.

While anti-windup has historically focused on these transient properties, more recent research trends have drawn connections also between anti-windup and asymptotic optimality [Hauswirth et al., 2020a; Hauswirth et al., 2020b]. Informally, we can see the connection between anti-windup compensation and conditions of optimality by noting that $\text{dz}(u) \neq 0 \implies \text{sat}(u) = \pm 1$. In this sense, $\text{dz}(u)$ takes non-zero values only when the input constraints of the system are active. Hence, when the control loop is designed in an appropriate way, $\text{dz}(u)$ can play the role of a Lagrange multiplier when viewing the system from a primal-dual optimization algorithm perspective. This approach of viewing the system from an optimization perspective will soon be discussed further. In the later papers of this thesis, appropriately designed anti-windup will be shown to yield asymptotically optimal equilibria in certain settings.




Figure 2.5 Schematic illustrating a closed-loop system with anti-windup compensation. The output of the controller u is saturated. The dead-zone function $\text{dz}(u) = u - \text{sat}(u)$ is mapped back into the controller, and used to ensure that windup does not occur in the presence of persistent disturbances.

Extremum Seeking

Perhaps one of the oldest control branches concerned with asymptotic optimality is extremum seeking. This methodology traces back as far as the 1920's [Tan et al., 2010], and the research field grew during the 1950's and 60's in connection with the developments in adaptive control [Åström and Wittenmark, 2008, pp. 549-553]. In essence, this method relies on perturbing the system input with a slow, probing "dither-signal". This signal varies much slower than the natural system dynamics, gradually "learning" the local gradient of the system, thus driving the system towards an optimum. The separation in time scales between the system dynamics and the dither signal makes this method rather slow. While the method is powerful and essentially model free, it is most often applied to systems of low input-output dimensionality such as internal combustion engines, process control or optimal power-tracking [Tan et al., 2010]. The probing strategy at the heart of the method is difficult to apply in a multi-agent large-scale setting, especially when considering communication constraints. The large-scale context also further reduces the speed of the method, as each agent essentially corresponds to one new input direction that the dither signal needs to learn.

Real-Time Optimization

A more recent body of literature has investigated the potential of designing the closed-loop system to act practically as an optimization algorithm [Krishnamoorthy and Skogestad, 2022; Hauswirth et al., 2024]. Consider the problem of designing an asymptotically optimal controller for a dynamical system f with respect to some objective J as per Definition 1. Assume first of all that the natural dynamics of the system are stable, such that given a stationary input u and a stationary disturbance w , the system will converge to a unique equilibrium x defined by a steady-state map $x = h(u, w)$. If the control capacity is constrained such that $u(k) \in \mathcal{U}$ for some bounded set \mathcal{U} , the objective of the controller is to establish the input u which solves

$$\underset{u}{\text{minimize}} \quad J(h(u, w)) \quad (2.5a)$$

$$\text{subject to} \quad u(k) \in \mathcal{U} \quad (2.5b)$$

For now, disregard the capacity-constraint (2.5b), and consider w to be fully constant such that it can be disregarded, i.e., $h(u, w) = h(u)$. A sensible approach is then to implement the controller u by gradient flow:

$$\dot{u} = -\nabla h(u)^\top \nabla J(h(u))^\top. \quad (2.6)$$

With suitable assumptions on convexity, this controller should find the global optimum. This approach can then be extended to incorporate constraints such as $u \in \mathcal{U}$ with e.g., barrier function or projection methods. This type of method has shown promise for optimal operation in power systems [Ortmann et al., 2023]. While this

family of methods is fairly general and holds insightful interpretation, these methods come with several complications. Firstly, the gradient $\nabla h(u)$ of the steady state map must be at least approximately known. Furthermore, the controller structure inherits the sparsity pattern of this gradient, which may in general be dense. The methods also suffer from problems of time-scale separation the same way that extremum seeking methods do: The natural dynamics of the system have to converge much faster than the controller.

Model Predictive Control

Another broadly applied control method concerned with optimality is Model Predictive Control (MPC). This is a broad set of tools which are explicitly useful for handling objective functions and constraints, both in terms of control actions and states. MPC is bound to a discrete-time setting, where at each time instance t , the control action $u(t)$ is given by the solution to an optimization problem. For our problem setting, this optimization problem could naturally be formulated as

$$\underset{\hat{x}, u}{\text{minimize}} \quad \sum_{k=t}^{t+h} J(\hat{x}(k)) \quad (2.7a)$$

$$\text{subject to} \quad \hat{x}(k+1) = \hat{f}(\hat{x}(k), u(k)), \quad \forall k = t, \dots, t+h-1, \quad (2.7b)$$

$$u(k) \in \mathcal{U}, \quad \forall k = t, \dots, t+h-1, \quad (2.7c)$$

$$\hat{x}(t) = x(t) \quad (2.7d)$$

The constraint (2.7c) implies that the system input is bounded, meaning that a sufficiently large disturbance cannot be rejected. In (2.7b), the controller predicts the future states \hat{x} of the system with the use of a model \hat{f} of the dynamics, up to a prediction horizon h . This prediction is initialized such that $\hat{x}(t) = x(t)$ at the current time stamp t according to (2.7d). The objective of the algorithm is to minimize the expected total objective (2.7a) incurred over this time-frame MPC is a flexible tool, which can further be customized to the problem at hand. One could for instance consider time-varying objective functions $J(x, t)$ or constraints on the states x . However, MPC can be troublesome to use in the setting of this thesis for several reasons. Firstly, an explicit model \hat{f} is required. For a large-scale system, this requires knowledge of potentially thousands of individual components. The disturbance which might be unknown needs to be included in this model. Secondly even if \hat{f} is well modeled, the optimization problem itself may be hard to solve unless \hat{f} has exploitable properties such as linearity. Thirdly, the naive implementation of MPC is fully centralized, i.e., the control actions are governed by the central optimization algorithm which then needs to know the state of all of the agents in the system. There are however ways around this in specific contexts. MPC can be distributed such that each agent in the system solves its own local optimization problem, with communication between the agents ensuring that the resulting calculations align.

One can also have simple, local regulators for each agent, the set-points of which are dictated by a governing MPC controller in a hierarchical fashion. This structure is common in, for instance, the process industry.

Control of Buffer Networks

To find controllers which are well suited to a large-scale setting while ensuring asymptotic optimality is difficult. To establish such results, it is reasonable to consider specific instances of systems and objective functions. One such family of control application which displays many of the issues at the heart of this thesis is control of buffer networks [Bauso et al., 2013; Bürger et al., 2015; Blanchini et al., 2016; Trip et al., 2019; Blanchini et al., 2019]. In the basic form of this problem (of which variations and generalizations exist), the dynamics $f(x)$ is given by

$$\dot{x} = B \text{sat}(u) + w.$$

Here x represents the levels of some commodity stored in a set of n buffers. These buffers are connected through directed edges, through which the commodity can be transported. This transportation is governed by the constrained edge flow rates $\text{sat}(u)$, which is the controlled input to the system. Here w acts as a constant disturbance, which could for instance denote the commodity consumption at each buffer location. The matrix B encodes the network connection between buffers. If edge e leads from node i to node j then $B_{i,e} = -1$, $B_{j,e} = 1$, and $B_{k,e} = 0$ for all $k \neq i, j$. The control goal in this setting is to design controllers where the buffer levels x are guided to their desired set-points, utilizing flows which minimize some objective function $J(u)$. Several works have been published that solve instances of this problem or generalizations thereof [Bauso et al., 2013; Bürger et al., 2015; Blanchini et al., 2016; Trip et al., 2019; Blanchini et al., 2019]. What is remarkable in these works is that the controllers that achieve asymptotic optimality typically inherit the structure of B . The flow actuated by the edges is therefore only informed by the state of the connected buffers, and hence the control structure becomes distributed. As a particular example, the first of the cited works [Bauso et al., 2013] demonstrated that the control law $u = -\text{sat}(\gamma B^\top x)$ asymptotically minimizes the objective $J(u) = u^\top u$. Here the saturation function $\text{sat}(u)$ acts component-wise, and γ is a positive gain. This and other cited results highlight a family of control problems where structurally simple and distributed controllers can achieve interesting control goals in which they cannot be outperformed by complex, centralized controller alternatives. This form of theoretical result is what will be considered in papers II, III and IV of this thesis.

M-Matrices

This background section will now be concluded with an introduction to a certain class of matrices, denoted *M-matrices*. These matrices have properties which are exploited in papers III and II of this thesis. We here follow the definitions and properties as outlined in [Horn and Johnson, 1991, pp.113-115]. Firstly, let us define $\mathcal{Z}_n = \{M \in \mathbb{R}^{n \times n} \mid M_{i,j} \leq 0, \text{ if } i \neq j, i, j = 1, \dots, n\}$, as the set of all n by n matrices with non-positive off-diagonal elements. An M-matrix is then defined as:

DEFINITION 2

A matrix A is an M-matrix if $A \in \mathcal{Z}_n$ and all eigenvalues of A have a strictly positive real part.

To avoid confusion, we can remark that this is equivalent to $-A$ being Metzler-Hurwitz. This is a matrix structure often encountered in the study of positive systems [Rantzer and Valcher, 2018]. M-matrices fulfill a list of equivalent and useful properties. The equivalences presented here are based on Theorem 2.5.3 in [Horn and Johnson, 1991, pp.114-115]. Said theorem contains further equivalences, which we omit here as they are not utilized in any of the papers of this thesis.

PROPOSITION 1

Let $A \in \mathcal{Z}_n$. Then the following statements are equivalent.

- (i) *A is an M-matrix.*
- (ii) *A is nonsingular and $A^{-1} \geq 0$.*
- (iii) *$Ax \geq 0 \implies x \geq 0$.*
- (iv) *There exists a diagonal positive matrix D such that DA and DAD^{-1} are strictly column-diagonally dominant.*
- (v) *There exists a vector $d > 0$ such that $d^\top A > 0$.*
- (vi) *DA and AD are M-matrices for any positive diagonal matrix D .*
- (vii) *There exists a positive diagonal matrix D such that $DA + A^\top D \succ 0$.*

Here, a matrix $M \in \mathbb{R}^{n \times n}$ is strictly column-diagonally dominant if $|M_{i,i}| > \sum_{j \neq i} |M_{j,i}|$ for all $i = 1, \dots, n$. In the context of this thesis, M-matrix properties will be used to analyze systems such as the following example.

EXAMPLE 3

Consider a network of n agents. Each agent $i \in 1, \dots, n$ is represented by a scalar state variable x_i , and subject to a constant disturbance w_i . The dynamics of agent i is governed by

$$\dot{x}_i = -x_i + B_i \text{sat}(u) + w_i. \quad (2.8)$$

where B_i is row i of the matrix $B \in \mathbb{R}^{n \times n}$ and B is an M-matrix.

Systems on this form are the topic of papers II and III of this thesis. Furthermore, at this point we can also note that both examples 1 and 2 concern systems on the form (2.8). The M-matrix property of B is well suited to model the notion of capacity-constraints in the context of this thesis. The non-positivity of the off-diagonal elements reflects competition between the agents, such that if agent j increases their control action, the resource granted to agent $i \neq j$ decreases. Properties (ii) – (iii) imply that increasing the resource granted to each agent ($B\text{sat}(u) \geq 0$) also requires a positive control action from each agent ($\text{sat}(u) \geq 0$). Properties (iv) – (v) together imply that if all agents increase their control actions ($\text{sat}(u) > 0$), the total (weighted) output of the system will be positive ($d^\top B\text{sat}(u) > 0$). Property (vi) implies that the system can be transformed slightly while maintaining the same core properties. Finally, (vii) is an exploitable property well suited for Lyapunov analysis.

3

District Heating and Cooling Networks

District heating and cooling networks are systems where hot or cold water produced at production facilities is distributed via pipelines to consumers. The consumers use the water for heating and cooling purposes. These systems play an important role in the energy systems of many countries. In 2016, district heating constituted 90 percent of energy use in Swedish multi-family dwellings [Energimyndigheten, 2017]. District cooling plays a counterpart to district heating, where the same form of infrastructure and connections are used. The results of this thesis are applicable also in this setting, but we will discuss district heating only so as to keep the discussion and examples concise.

This chapter will begin with background on district heating networks, their functionality, and developments in district heating technology. The chapter continues with background on control of district heating networks, beginning with control of current systems and continuing with newer research trends. We will finally consider mathematical models for hydraulics and thermodynamics.




Figure 3.1 An overview of a district heating network. At production sites, water from the less hot return network (blue) is heated and pumped out through the hot supply network (red) to reach consumers. Water subsequently returns through the return lines to continue this cycle.




Figure 3.2 A schematic of the substation of a building connected to a district heating network. Heat is delivered to the building through a heat exchanger (HEX). The volume flow rate on the primary side of the heat exchanger (left), i.e., the side connected to the network, is governed by a control valve. The radiator system of the building is connected to the secondary side (right) of the heat exchanger.

3.1 Background

District heating networks are systems where hot water produced at production facilities is distributed via pipelines to consumers. The consumers use the heat stored in the water for space heating or domestic hot water usage. A schematic of this system is shown in Figure 3.1. The transportation network consists of two layers: the supply network and the return network. "Cold" water (between around 40 to 60 degrees, depending on the system) from the return network is heated at production sites to a higher temperature (often to above 80 degrees, once again depending on the system). Large distribution pumps are then used to pump the hot water through the supply network out to consumers. The differential pressure generated by these pumps is what drives the flow of water. Consumers are connected to the network through *substations*. A schematic of a building with its substation is depicted in Figure 3.2. At such a substation, the water in the internal hydraulic system of each building is heated through a heat exchanger, utilizing the hot water from the district heating system. The substation has a *primary* and a *secondary* side. The primary side denotes the components on the side of the heat exchanger connected to the greater network. The secondary side denotes the internal heating system of the building. The heat supplied to the building is utilized for both space heating and domestic hot water usage. The connections for space heating and domestic hot water usage are typically separated into two individually controlled systems, which is not visible in Figure 3.2. Such a substation configuration is typical in Sweden, but there are many different variations [Frederiksen and Werner, 2013].

Generations of District Heating

The way in which district heating systems are constructed and operated has changed over the years, and development is ongoing. It is common to describe these systems as part of different generations to establish large progressions in technology [Fredriksen and Werner, 2013, cp. 8.1],[Lund et al., 2014]. The challenges and opportunities related to controlling these systems vary depending on which generation of district heating system is considered, making it important to distinguish between them.

The First Two Generations In the 1st generation of district heating systems, metal pipes installed in large ducts were used to distribute steam. This technology, which was developed in the United States, was largely driven by fossil fuel [Lund et al., 2021]. As the 2nd generation of district heating systems were developed in Europe there was a shift from using steam to using high temperature water as the energy carrier. The metal pipes remained however, along with the strong reliance on fossil fuel [Fredriksen and Werner, 2013].

The 3rd Generation The vast majority of existing district heating networks in the Nordics belong to the 3rd generation of district heating [Lund et al., 2014]. A reduction in water temperatures was implemented in the transition from 2nd generation technology. This allowed the use of a wider variety of fuel sources, as well as new piping technology [Fredriksen and Werner, 2013; Lund et al., 2021]. These 3rd-generation systems typically have the same structure as was described above for Figure 3.1: One or a few large heat suppliers produce all the heat in the network and utilize large, centrally operated pumps to distribute the water. Networks of this traditional, centralized structure form the main focus of this thesis.

The 4th Generation The current research, the discussion is largely centered around the 4th generation of district heating [Lund et al., 2014; Lund et al., 2018]. The implied technological changes from the 3rd generations are largely aimed at allowing district heating networks to act as a key part of a more sustainable and integrated energy system in the future. This includes a higher penetration of renewable heat sources such as solar, and new waste heat sources such as data centers. To be able to include such heat sources, as well as reducing network losses, 4th generation systems are designed with lower supply temperatures than previous generations. Furthermore, new challenges and opportunities related to control are found in this technological transition [Vandermeulen et al., 2018]. 4th generation district heating networks need to be smarter and more flexible than previous generations due to, among other issues, intermittent heat supply. However, smart metering and actuator technology allow for new control strategies and data driven modeling frameworks to be implemented.

The 5th Generation The 5th generation is perhaps a case of unfortunate naming procedures. Rather than being a strict progression from the 4th generation (which is

still not necessarily the current generation of technology), the 5th generation concerns a rather different setup entirely [Lund et al., 2021]. The term 5th generation typically refers to a form of ultra-low temperature network which supplies both heat and cold by utilizing bidirectional heat flow. In this setup, prosumers in the network can utilize water in the colder side of the network for cooling, and pump it into the warmer side, or use water from the warmer side and pump it into the cooler side. As the water is generally never of sufficiently high or low temperature to use directly for its intended purpose, each prosumer connected to this system is fitted with additional heat pumps [Lund et al., 2021].

3.2 Control of District Heating Networks

This section details the necessary background to understand the control systems which operate existing district heating networks. This is followed by new trends and developments in research on controlling district heating networks.

Control of Current Generation Networks

The majority of existing district heating networks are controlled by four independent control loops; *heat demand control*, *flow control*, *differential pressure control* and *supply temperature control* [Frederiksen and Werner, 2013; Vandermeulen et al., 2018].

The first two control loops are situated in customers' buildings. *Heat demand control* decides how much heat is used in the building, dictated by domestic hot water usage, as well as radiator valve settings. *Flow control* decides the flow rate of hot water going through the primary side of the building's substation. This is generally performed by a control valve, set to track a target temperature on the secondary side of the heat-exchanger. See Figure 3.2 for reference. If the water circulating on the secondary side of the heat-exchanger is too cold, the valve opens up to increase the primary-side volume flow rate. This process is typically governed by a PID controller. The set-point temperature for the PID controller is typically proportional to the outdoor temperature, which is the main disturbance acting on the building. The only measurements of consumer activity in these networks are typically supply-and-return temperatures at the substation level, as well as the associated flow rates. These measurements are generally not used for modeling or for control actuation and design. Rather, these measurements are used to calculate the energy consumption to be used for invoicing.

Differential pressure control and *supply temperature control* on the other hand are control loops governed by the network operator. *Supply temperature control* decides the temperature of the water supplied to the network and is typically an open-loop schedule based on the expected load for the coming day. The *differential pressure control* loop decides the pumping activity of the circulation pumps at the production facility. This control loop typically uses measurements of differential

pressure at the peripheral parts of the network. If this value drops too low, the flow control loop of the consumers will not work, as the consumer can only open their control valve up to a certain maximum point. Early research on control of district heating networks has often focused on these control loops which concern the heat production, e.g., [Benonysson et al., 1995].

The control strategies concerned in this thesis mirror and build on this existing infrastructure, in particular considering variations of the flow control loop. There are advantages of this approach: The present control structure is well proven in practice and hence stands the test of real-world deployment. Rather than fully rejecting the current structure in favor of a fully novel approach, introducing modular improvements to the system may increase the ease of implementation. Furthermore, the present control structure is already fully decentralized in terms of the flow control loop. Building on top of this architecture can allow partial preservation of this advantageous structure.

Research Trends in District Heating Control

The opportunities and challenges posed by transitions to 4th and 5th generations of district heating networks have sparked new research in district heating network control. The challenging multi-agent setting, and the seemingly "friendly" dynamical properties that allow scalable control architectures has led to a recent increase in interest in district heating networks within the control systems community. These include works focusing on modeling and stabilizing control of hydraulics [De Persis and Kallesøe, 2008; De Persis and Kallesøe, 2011; De Persis et al., 2014; Strehle et al., 2021; Strehle et al., 2022; Strehle et al., 2024; Machado et al., 2022a; Jeeninga et al., 2023], thermodynamics [Machado et al., 2023; Ahmed et al., 2023; Sibeijn et al., 2024; Simonsson et al., 2024a; Simonsson et al., 2024b], or both [Machado et al., 2022b]. These works have produced an impressive list of results in terms of passivity properties of the open-loop systems, and decentralized control solutions which stabilize desired system equilibria. However, a common assumption between the works which concern hydraulics is that the set of feasible volume flow rates is unbounded. This is implicitly assumed by considering pumps as ideal pressure sources without limitations. While this assumption is reasonable for analyzing the stability of the system under nominal operation, an extension is needed to understand how hydraulic capacity-constraints affect the system when the pumps are limited. Furthermore, all of the above works which focus on hydraulics except [Jeeninga et al., 2023] explicitly consider the case where each individual consumer substation is fitted with a pump. This corresponds to the type of configuration found in the 5th generation of district heating networks but is not typical for a majority of existing infrastructure. While this 5th generation scenario offers interesting control challenges due to the large number of actuators, systems within current generations of technology are still not fully understood.

Three of the five papers in this thesis concern actions which redistribute the

flow rate in the system in some fashion. Any such method which alters the nominal load pattern of the consumers is typically denoted *demand response*. This is a field of research which has recently grown in popularity within the district heating community, with several real-world implementations e.g., [Sweetnam et al., 2019; Ala-Kotila et al., 2020; Karasu et al., 2024]. These works do typically focus on technology in between the 3rd and 4th generation of district heating, such as is the focus of this thesis. However, unlike the literature originating from the control theory community, there is little focus on understanding exploitable properties for feedback solutions. Rather, these works tend to focus on detailed models, and price-based or optimization-based solutions, e.g., [Cai et al., 2020; Capone et al., 2021; Chaudhry et al., 2024]. For further details on demand response in district heating, see the review paper [Guelpa and Verda, 2021].

3.3 Hydraulic Models

In the following section we will discuss typical modeling procedures for district heating hydraulics. This entails models for the individual hydraulic components (e.g., individual pipes) and how to put these components together to analyze the network as a whole. Furthermore, we will discuss data-driven tuning of hydraulic models, as well as hydraulic constraints which limit the flow rate capacity in these systems.

Hydraulic Components

The most common way to model a district heating network is as a directed graph $\mathcal{G} = (\mathcal{V}, \mathcal{E})$, see for instance [De Persis et al., 2014; Machado et al., 2022b; Strehle et al., 2024]. Here the edges \mathcal{E} represent components such as valves, pipes and pumps, and the nodes \mathcal{V} represent intersections of edges. Each node $v \in \mathcal{V}$ is associated with a pressure p_i . Each edge $e \in \mathcal{E}$ is in turn associated with a volume flow rate q_e along the direction of the edge. This thesis concerns networks consisting of three types of components: *pipes*, *valves* and *pumps*. Most hydraulic models are built on the following common assumptions:

ASSUMPTION 1

The water circulating in the network is in a liquid state and is incompressible. The density of the water is approximately constant. There is negligible leakage.

Under these assumptions, we can model the individual components as follows.

Pipes Consider one edge $e \in \mathcal{E}$ representing a pipe, leading from node i to node j . A dynamic relation between the volume flow rate q_e through the pipe, and the nodal pressures p_i and p_j , can then be constructed through the control volume approach which yields a model on the form [De Persis et al., 2014]:

$$J_e \dot{q}_e = -s_e g(q_e) + p_i - p_j.$$

Here J_e corresponds to the inertial mass of the water in the pipe. The term $s_e g(q_e)$ models frictional losses in the pipe through the strictly monotonically increasing function g , and a positive resistance parameter s_e . The parameter s_e is unique to pipe e , whereas in this thesis we consider a shared function g to model all pipes. When considering a holistic view of the district heating network, the hydraulic transients are much faster than temperature dynamics. Therefore, this thesis concerns static models where $\dot{q}_e = 0$:

$$p_i - p_j = s_e g(q_e). \quad (3.1)$$

Furthermore, the resistive component $s_e g(q_e)$ is typically modeled using the Darcy-Weisbach equation [Frederiksen and Werner, 2013, p.442]. In this case

$$p_i - p_j = \frac{8f_{d,e}\rho_w L}{\pi^2 D_e^5} \cdot q_e |q_e|, \quad (3.2)$$

where D_e is the diameter of the pipe, ρ_w is the water density, L is the length of the pipe and $f_{d,e}$ is the Darcy friction factor. This factor varies with the Reynolds number and hence flow rate of the water but is often assumed constant in works concerning modeling and simulation of district heating networks (see e.g., [van der Heijde et al., 2017a; De Persis and Kallesøe, 2008; Machado et al., 2022b]). This assumption holds well when the water is turbulent [Frederiksen and Werner, 2013, p.442].

The Darcy-Weisbach model (3.2) can be compared to (3.1) by identifying $s_e = \frac{8f_{d,e}\rho_w L}{\pi^2 D_e^5}$ and $g(q_e) = q_e |q_e|$.

Note that consumer substations as depicted in Figure 3.2 typically consist of piping, a heat exchanger, and a valve, in series. The heat exchanger induces frictional losses in the same way that a pipe would and is therefore typically modeled with the same type of model as a pipe.

Valves The relation between volume flow rates and pressures for valves is often modeled as static, as the volume of water contained within the valve is relatively small. We model valves $e \in \mathcal{E}$ by

$$p_i - p_j = \theta_e k(q_e, u_e). \quad (3.3)$$

Here u_e denotes the controllable valve position, and θ_e is an edge-specific, positive resistance parameter similar to s_e for pipes. The continuous function $k(q_e, u_e)$ is strictly monotonically increasing in q_e , as higher flow rates increase pressure losses. It is strictly monotonically decreasing in the valve position u_e , representing that opening the valve decreases the resistance. The valve position is bounded due to physical limitations. We here consider the bounds $u_e \in]0, 1]$.

As u_e approaches fully closed, k satisfies that $\lim_{u_e \rightarrow 0} k(q_e, u_e) = \infty$ for any $q_e > 0$. This corresponds practically to removing this edge from the network. Consider as an example a valve with linear characteristics:

$$p_i - p_j = \theta_e \cdot \frac{1}{u_e^2} \cdot q_e |q_e|. \quad (3.4)$$

Linear valve characteristics are common in both district heating networks and process control [Frederiksen and Werner, 2013; Häggglund, 2023]. They are called "linear" as, given a constant pressure $p_i - p_j$ over the valve, the flow rate q_e becomes linear in the valve opening u_e . This can be shown by assuming a positive pressure $p_i - p_j > 0$ and rearranging (3.4):

$$q_e = u_e \sqrt{\frac{p_i - p_j}{\theta_e}}.$$

Pumps Lastly, pumps are the sources of pressure in the network. In many works, pumps are simply modeled as ideal pressure sources [De Persis and Kallesøe, 2008; De Persis and Kallesøe, 2011; De Persis et al., 2014; Strehle et al., 2021; Machado et al., 2022b; Machado et al., 2022a; Jeeninga et al., 2023], i.e., for a pump $e \in \mathcal{E}$ directed from node i to node j ,

$$p_i - p_j = -u_e,$$

where u_e is the pressure generated by the pump. Another approach is to model the pump through a pump head curve, where the generated pressure changes depending on the operating frequency r_e of the pump [Wang et al., 2017]. For a an e representing a pump, from node i to node j , an example of such a pump curve model is

$$p_i - p_j = c_1 q_e |q_e| - c_2 r_e - c_3 r_e^2$$

where c_1, c_2 and c_3 are positive parameters and $r_e \in [0, 1]$. However, the valve curve model can equivalently be viewed as an ideal pressure source connected in series with a pipe. This is found by simply choosing r_e appropriately by imposing $c_2 r_e + c_3 r_e^2 = u_e$.

Hydraulic Circuit Properties

We will now assemble the individual components in order to analyze the hydraulic properties of the full network. This thesis focuses on district heating networks that have certain structural properties:

ASSUMPTION 2

The supply-and-return networks in the district heating network are tree-structured and their topologies are identical. They contain only pipe-edges. There exists only one pump in the network, corresponding to a heat plant. This pump-edge connects the supply-and-return networks. Each other connection between the supply-and-return networks represents a valve.

Furthermore, we choose as convention and without loss of generality that the pump is directed from the return network to the supply network, valves are directed from the supply network to the return network, and pipes are directed away from the pump in the supply network and towards the pump in the return network. Figure 3.3

shows a schematic of a network which satisfies these qualities. Under these assumptions, there are several useful relations that can be formulated. Firstly, by removing the set of consumer valves which we can denote \mathcal{C} from \mathcal{E} , the remaining graph forms a spanning tree. By reintroducing any of these valves into the spanning tree, a fundamental loop is formed, which necessarily contains the pump. This property means that the flow rates $q_c, \forall c \in \mathcal{C}$, become independent values. Let the notation q without subscript denote the vector of these values. For instance, in the example of Figure 3.3, $q = [q_1, \dots, q_4]^T$. With this definition, q also corresponds to the set of flow rates which dictate the heat that is delivered to the consumers and hence lies at the heart of the topic of this thesis. Furthermore, under the assumption that there is no leakage in the network, we can introduce the following two fundamental concepts, analogous to Kirchhoff's circuit laws:

DEFINITION 3

- (i) *The sum of all flow rates leading in and out of each node $v \in \mathcal{V}$ must be 0.*
- (ii) *In traversing any loop in the network, the sum of pressure differences over each edge in this loop must be 0.*

Definition 3 (i) can be used to show that for all edges $e \in \mathcal{E}$, the flow rate q_e can be found linearly as

$$q_e = \sum_{i \in \mathcal{C}} l(e, i) q_i, \quad \forall e \in \mathcal{E} \quad (3.5)$$

where $l(e, i) = 1$ if e lies in fundamental loop i and $l(e, i) = 0$ otherwise. For instance, in the example of Figure 3.3 we find that $q_{12} = q_3 + q_4$. Furthermore, by utilizing (ii) on the fundamental loops of the network, we find that

$$\Delta p_{\text{pump}} = \sum_{e \in \mathcal{L}_i} s_e g(q_e) + \theta_i k(q_i, u_i), \quad \forall i \in \mathcal{C}. \quad (3.6)$$




Figure 3.3 A schematic of a district heating network, exemplifying the circuit-theoretical concepts employed to model district heating networks. The consumers are represented by valves, drawn as dotted edges 1-4. Edge 5, drawn in solid black, represents the pump. The supply-and return networks are represented by red edges 6-12 and blue edges 13-19 respectively. The green highlight displays the fundamental loop associated with edge 2. The figure is based on Figure 1 of Paper V.

where Δp_{pump} is the differential pressure over the pump, and \mathcal{L}_i denotes the set of pipes which are part of fundamental loop i . Hence in steady state, given a set of valve positions and a differential pressure actuated at the pump, the resulting flow rates q delivered to the consumers is the solution to the combination of (3.5) and (3.6).

Tuning of Hydraulic Models

There are several advanced control techniques that benefit from an explicit model of the hydraulic network. Examples include MPC controllers as introduced in Chapter 2, but also the model-based controller introduced in Paper I of this thesis. In this context, tuning a hydraulic model entails finding values of the parameters s_e and θ_e in (3.1) and (3.3) which match the observed behavior of the system. Several works considering such parameter estimation have been published in recent years [Wang et al., 2018; Liu et al., 2020; Luo et al., 2022; Zheng et al., 2024]. However, these studies have typically not been performed with control of flow rates in mind. These works typically assume that the differential pressure over each consumer substation can be measured and therefore the hydraulic properties of the consumer can be disregarded in the modeling effort. This means that no effort is placed on modeling the valves which regulate the flow rate. From a flow control perspective on the other hand, understanding the hydraulic behavior of the valves is key. This is the topic of Paper V of this thesis, where hydraulic parameter estimation is performed in a laboratory environment.

Hydraulic Capacity-Constraints

The flow rates q , and hence the feasible heat loads which can be delivered to the consumers, are bounded. This becomes evident by investigating the fundamental loop relations (3.6). Consider first that Δp_{pump} is upper bounded by some maximum pumping activity $\Delta p_{\text{pump}}^{\max}$. Note that both g and k are monotonically increasing in the volume flow rates q . Hence an increase in the flow rates q_i and q_e must be met by a corresponding increase in the valve position u_i , which reduces the pressure loss over the valve $k(q_i, u_i)$. However, the valve can only be opened to a limited extent, at which point the valve cannot be used to reduce this pressure loss further. Hence, given the bounds on the pump activity Δp_{pump} and the valve positions u , there is a bounded set \mathcal{Q} of flow rates $q \in \mathcal{Q}$ which can be delivered to the consumers. This bound implies that under peak load conditions in the network, it can be impossible to distribute the necessary heat to the consumers, i.e., the central topic of this thesis.

One of the key results of Paper I of this thesis is establishing that this set \mathcal{Q} is convex in the context of line-structured networks under common model assumptions for g and k . We will here informally show a minor extension of this result to the tree-structured case as described in the above section, based on the same form of relaxation as (16)-(17) as in Paper I. Consider relaxing the fundamental loop

relations (3.6) to find

$$\Delta p_{\text{pump}}^{\max} \geq \sum_{e \in \mathcal{L}_i} s_e g(q_e) + \theta_i k(q_i, 1), \quad \forall i \in \mathcal{C}. \quad (3.7)$$

Here, q_e is a positive sum of elements in q according to (3.5). Hence under assumption that $q > 0$, and assuming that g and k are convex in q_e for positive values of q_e and q_i , the set of q which satisfy (3.7) is convex. Consider any q which satisfies these convex constraints (3.7). As $k(q_i, u_i) \rightarrow \infty$ when $u_i \rightarrow 0$, there must exist a pump pressure $\Delta p_{\text{pump}} \leq \Delta p_{\text{pump}}^{\max}$ and corresponding feasible valve positions u which satisfy (3.6) exactly. Hence there exists a feasible configuration of the pump and valves which actuate this flow rate q .

These constraints on q , which are induced by the frictional pressure losses in the network, cause operational issues sometimes referred to as bottlenecks [Brange et al., 2019]: Under peak load conditions, it becomes difficult to provide sufficient flow rates to certain parts of the network, i.e., the main topic of this thesis. While this is a problem which already occurs in the current generation of district heating systems, it can potentially become even more important in 4th generation district heating networks when supply temperatures are lowered [Brange et al., 2019]. This is because, if the consumers cannot use the lower temperature water as efficiently, a higher flow rate is required in order to provide an equivalent heat load.

3.4 Thermodynamic Models

This section will cover thermodynamic modeling for district heating networks. We will first discuss modeling of the water temperatures in the network pipelines which is not a topic of focus in this thesis, followed by the modeling of temperatures in buildings connected to the network. It is common to partition the modeling of hydraulics and thermodynamics in the network into two layers, as we do here. This is generally justified by the time-scale separation between these two topics.

Network Models

The temperature of the water in the pipelines varies with many factors. The supply network temperatures change depending on the operation of the heat sources and losses to the ambient temperature. Furthermore, temperature fronts in the network will propagate according to the corresponding flow rates. There are also convection effects in the water, dissipating the heat. Additionally, the return temperature at each consumer substation depends on several factors connected to the substation itself. This involves the water temperature and flow rate on the secondary side of the heat exchanger. A full model involving all of these properties necessarily becomes complex. Therefore, it is a standing assumption in this thesis that the temperatures in the supply and return networks are constant in time. This assumption is clearly a simplification and is made in order to enable mathematical analysis and conclusions.

Future studies where thermodynamic models are included can build on existing literature on thermodynamic simulation [Simonsson et al., 2024a; Simonsson et al., 2024b] and analysis of dynamic properties [Machado et al., 2022b; Machado et al., 2023; Sibeijn et al., 2024].

Building Models

This thesis concerns understanding how buildings are affected by insufficient heat supply under large disturbances. For this purpose, thermodynamic building models are required. A common approach to modeling thermodynamics in buildings is through the use of *RC-models*. These models play a thermodynamics counterpart to electrical circuits composed of resistors and capacitors. In this form, the building is represented by a set of temperatures T_i , $i = 1, \dots, k$, representing k different components. Note that in this section we are only describing one building. Hence the indexing in this section no longer refers to the different buildings in the full network, but rather to components in one individual building. The temperature T_i could be the temperature of a wall, a room, a radiator, or another component of the building. Each such component has a thermal mass C_i , corresponding to a capacitance in the electrical counterpart. If two components in the building are connected, say two rooms with temperatures T_i and T_j connected by a wall, then the heat transfer between these components is represented by a resistance $R_{i,j}$. The resistances form a lumped representation of heat transfer in the form of both conductance, convection and radiation (where the nonlinear parts of the latter are assumed negligible) [Frahm et al., 2022]. In literature concerning modeling and control of buildings, one can find numerous examples of this type of model being employed [Bacher and Madsen, 2011; Saletti et al., 2022; Frahm et al., 2022; van der Heijde et al., 2017b]. Depending on the application, one can establish models of varying complexity, ranging from modeling the temperature in individual walls, radiators and rooms [Bacher and Madsen, 2011] to lumping the thermal energy of the entire building into one state [Saletti et al., 2022]. In essence, a model of this form would look as follows. The dynamics of each temperature state is given by

$$C_i \dot{T}_i = \sum_{j \in \mathcal{N}(i)} \frac{1}{R_{i,j}} (T_j - T_i) - \frac{1}{R_{i,\text{ext}}} (T_i - T_{\text{ext}}) + P_i.$$

Here \mathcal{N}_i is the set of all other model components directly connected to component i . In the case where state i is connected to the exterior, $\frac{T_i - T_{\text{ext}}}{R_{i,\text{ext}}}$ denotes the heat transfer between state i and the exterior. Finally, P_i denotes direct energy injection into state i , for instance through the substation heat exchanger. To exemplify further, Paper I of this thesis utilizes a model with two temperature states T_{in} and T_{hs} representing

the indoor temperature and the temperature of the water in the radiator system:

$$\begin{aligned} C_{\text{in}} \dot{T}_{\text{in}} &= \frac{1}{R_{\text{hs}}} (T_{\text{hs}} - T_{\text{in}}) + \frac{1}{R_{\text{ext}}} (T_{\text{ext}} - T_{\text{in}}) \\ C_{\text{hs}} \dot{T}_{\text{hs}} &= \frac{1}{R_{\text{hs}}} (T_{\text{in}} - T_{\text{hs}}) + P. \end{aligned}$$

Here the interior temperature is affected by heat injection from the heating system and heat losses to the exterior. The heating system temperature is affected by the heat transferred to the interior, and the heat injection P delivered to the building through the substation heat exchanger:

$$P = c_p \rho_w T_{\Delta} q$$

where c_p is the specific heat capacity of water and ρ_w is the water density, assumed to be constant. T_{Δ} is the difference in water temperature between the supply and return layers on the primary side of the heat exchanger and q is the volume flow rate on the primary side of the heat exchanger. The relation between P and q is in general not linear, and the level of nonlinearity varies with the sizing of the heat exchanger. In practice, this means that the difference between supply temperature T_{sup} and return temperature T_{ret} varies with the volume flow rate. However, substation heat exchangers are often generously sized, making the relationship close to linear [Frederiksen and Werner, 2013, p.401], which is the assumption maintained in this thesis.

4

Contributions

In this chapter, the five publications and their contributions which make out this thesis are presented. Section 4.2 lists two additional papers by the present author which are excluded in order to improve the overall coherence of the thesis.

A Note on Notation The notation is not consistent between all five papers. This is largely but not exclusively due to the different target audiences to which the different papers were directed. Papers I and V were published in journals for energy research, whereas Papers II, III and IV were aimed towards a control theory audience. Therefore, the reader is referred to the *notation*-section of each paper for clarifications.

4.1 Included Papers

Paper I

F. Agner et al. (2022). “Combating district heating bottlenecks using load control”. *Smart Energy* **6**. ISSN: 2666-9552. DOI: 10 . 1016 / j . segy . 2022 . 100067

Scientific Summary In this paper, we characterize the constraint on district heating flow rates in the context of the current generation of district heating networks. As discussed in Section 3.3, the flow rates which can be delivered to the consumers are constrained. Denoting these flow rates q , we can say that $q \in \mathcal{Q}$. Under a typical Darcy-Weisbach model formulation for pipes, we find that the set \mathcal{Q} for a line-structured district heating network is convex and parameterized by the hydraulic parameters of the network components. This corresponds to the values of s_e and θ_e in (3.1) and (3.3).

We exploit the convexity of \mathcal{Q} for scalable control design, where the heat load in the system is coordinated via a central optimization algorithm. This allows the heat to be distributed in a more fair fashion than in the scenario where each agent locally decides their heat load freely. This is demonstrated through a simulation. This chronologically first paper sets the stage for the later manuscripts in this thesis.

Author Contribution F. Agner derived the result that \mathcal{Q} can be expressed as a convex set, proposed the final coordinating control strategy, and developed the simulation. P. Kergus contributed to writing the first manuscript draft. The whole author group contributed to scientific discussions during the project and editing of the manuscript.

Paper II

F. Agner et al. (2023b). “Anti-Windup Coordination Strategy Around a Fair Equilibrium in Resource Sharing Networks”. *IEEE Control Systems Letters* **7**, pp. 2521–2526. ISSN: 2475-1456. DOI: 10.1109/LCSYS.2023.3287252, ©2023 IEEE. Reprinted with permission.

Scientific Summary Paper I leaves open the question if coordination between the agents can be imposed without requiring an explicit model of the network hydraulics. In this paper, we approach this problem by considering a multi-agent model on the form

$$\dot{x} = -x + B \text{sat}(u) + w \quad (4.1)$$

akin to Example 3. Here each entry of x corresponds to the state of one agent, to be regulated to zero. w is a large, constant disturbance and B is an M-matrix. This matrix B is not explicitly known to the control designer. We propose a control-strategy where each agent uses a local PI-controller, with a coordinating rank-one anti-windup term. This becomes a variation of the PI-controller (2.4) where, for each agent i , the controller takes the form

$$\dot{z}_i = x_i + \beta \mathbf{1}^\top \text{dz}(u) \quad (4.2a)$$

$$u_i = -k_i^P x - k_i^I z. \quad (4.2b)$$

The coordination between the agents is here imposed through the term $\beta \mathbf{1}^\top \text{dz}(u)$, where β is a positive gain. Recall that $\mathbf{1}$ denotes a vector of all 1’s.

We establish a closed-form expression for disturbances w for which the proposed closed-loop system admits an equilibrium. This occurs when w is approximately parallel to $\mathbf{1}$, i.e., the disturbance acts similarly on all agents. Most importantly, such equilibria are fair, in the sense that they minimize the objective $\|x\|_\infty$. We furthermore demonstrate passivity of the closed-loop system in the linear domain where $\text{sat}(u) = u$.

Author Contribution The model structure and controllers of this paper were developed by F. Agner, under discussions with P. Kergus, L. Zaccarian and S. Tarbouriech. The final anti-windup compensator was suggested by S. Tarbouriech as a simplified version of an initially more complicated, dynamic version suggested by F. Agner. L. Zaccarian contributed significantly to the results regarding equilibrium existence, and proposed the theoretical connection to buffer networks, e.g., the paper [Bauso et al., 2013] and following works. F. Agner derived the theoretical

statements in the paper, along with their associated proofs and developed the simulation of the paper. A. Rantzer provided the proof of linear domain passivity of the closed-loop system. The initial manuscript was written by F. Agner with assistance from P. Kergus. Subsequent editing was mainly performed by F. Agner and assisted by all co-authors.

Paper III

F. Agner et al. (2024a). “Decentralized PI-control and anti-windup in resource sharing networks”. *European Journal of Control*, p. 101049. ISSN: 0947-3580. DOI: 10.1016/j.ejcon.2024.101049

Scientific Summary This paper complements Paper II by considering the same form of system model (4.1) and basic controller structure. The system model is slightly extended to consider a scaling factor on the stabilizing term, i.e.,

$$\dot{x}_i = -a_i x_i + B_i \text{sat}(u) + w_i, \quad \forall i = 1, \dots, n. \quad (4.3)$$

This paper also extends the model from considering only $\text{sat}(\cdot)$ to a richer class of nonlinear functions. In this paper, unlike Paper II, the considered control law requires no coordination between the agents. The PI-controllers of the agents are designed with fully decentralized anti-windup compensation and take exactly the form of the (2.4). This control strategy requires no communication between the agents neither in design nor actuation. This strategy resembles the form of controller typically used in current generation district heating systems. We show that this strategy globally, asymptotically stabilizes a unique equilibrium, and that this equilibrium minimizes a weighted objective $\|\Gamma x\|_1$. Here Γ is a diagonal matrix which depends on the properties of the matrix B , and the individual stabilizing parameters a_i . This result highlights a class of problems where the fully decentralized strategy cannot be outperformed by more complex controller options.

Author Contribution The main concepts of this paper build on the results of Paper II, and hence the conceptualization should be accredited to the same discussions between F. Agner, P. Kergus, L. Zaccarian and S. Tarbouriech. F. Agner proposed analyzing the fully decentralized anti-windup compensator. F. Agner developed the theoretical results of the paper, with assistance in conceptualizing and validating the proofs from J. Hansson. A. Rantzer suggested the change of variables which allowed the proof of global, asymptotic stability. F. Agner wrote the initial manuscript version and the included simulation. All authors contributed to manuscript editing.

Paper IV

Submitted to Automatica November 20th, 2024. Preprint available: F. Agner and A. Rantzer (2024). *On PI-control in capacity-limited networks*. arXiv: 2411.14077 [eess.SY].

Scientific Summary This paper extends the results of papers II and III. The same two controllers are considered, but now in a setting where the agents are interconnected through a fully nonlinear function b . This corresponds to a variation of Example 3, where the M-matrix is now replaced by the nonlinear function b :

$$\dot{x}_i = -a_i x_i + b_i(\text{sat}(u)) + w_i, \quad \forall i = 1, \dots, n. \quad (4.4)$$

The function b is assumed to have similar input-output properties as an M-matrix, corresponding approximately to negative off-diagonal properties of b , and (iv) of Proposition 1.

The extension to this nonlinear setting is crucial for representing models of district heating hydraulics, such as considered in Paper I. We are able to extend the main results of stability, equilibrium existence and equilibrium optimality from papers II and III to this nonlinear setting. Finally, the two considered control strategies are demonstrated in a simulation of a district heating network.

Author Contribution All theoretical results of the paper were proposed and proven by F. Agner, with guidance and suggestions from A. Rantzer regarding simplifications and exposition of the results. F. Agner developed the simulation included in the paper. The manuscript was authored fully by F. Agner, with input and guidance from A. Rantzer regarding disposition and notation.

Paper V

F. Agner et al. (2024b). “Hydraulic parameter estimation for district heating based on laboratory experiments”. *Energy* **312**, p. 133462. ISSN: 0360-5442. DOI: <https://doi.org/10.1016/j.energy.2024.133462>

Scientific Summary This paper concerns experiments for identification of hydraulic model parameters, i.e., pipe resistances s and valve curve parameters θ as per Section 3.3. Such models are important for advanced control, simulation and analysis of district heating systems. For instance, such a model is central to the proposed method of Paper I. The experiments were performed in the Smart Water Infrastructure Laboratory in Aalborg, Denmark. Here we extended the fully theoretical framework considered in [Agner et al., 2023a], and suggested model improvements to capture real-world complications in the form of unknown valve curve characteristics and valve hysteresis. The experimental data was also published in an openly available repository for future studies.

Author Contribution F. Agner planned and performed the experiments in this paper with technical assistance from C.M. Jensen and C.S. Kallesøe. The subsequent modeling and parameter estimation was performed by F. Agner, in discussion with the whole group of authors regarding model improvements and design. The initial manuscript was authored by F. Agner, with assistance in writing the description of the laboratory equipment from C.M. Jensen. Editing of the manuscript was performed by the full group of authors.

4.2 Additional Publications

The following publications and works in progress were excluded from the thesis to improve the disposition and coherence.

Excluded Publication I

F. Agner et al. (2023a). “Hydraulic Parameter Estimation in District Heating Networks”. *IFAC-PapersOnLine*. 22nd IFAC World Congress **56**:2, pp. 5438–5443. ISSN: 2405-8963. DOI: 10.1016/j.ifacol.2023.10.194

Excluded Publication II

F. Agner et al. (2024c). “A data-based comparison of methods for reducing the peak flow rate in a district heating system”. *Smart Energy*. ISSN: 2666-9552. DOI: <https://doi.org/10.1016/j.segy.2024.100168>. In press.

5

Conclusion

This thesis addresses control of capacity-constrained networks. The results are grounded in the application of district heating networks. In this context, the capacity to deliver heat to consumers is limited by hydraulic constraints.

We demonstrate how the feasible volume flow rates in a line-structured network are convex. This result can rather trivially be extended also to tree-structured networks. The convexity of this constraint was utilized to formulate a load coordination strategy to provide a fair allocation of system flow rates. For this purpose, an explicit hydraulic model of the system is required. This thesis demonstrates in a laboratory environment how parameters for such a model can be estimated from operational data. Future works could employ such an explicit hydraulic model for more advanced control strategies, such as model predictive control. Outlooks for future work also include investigating the conditions under which the set of feasible flow rates remains convex, even under meshed network structures, and structures which involve decentralized pumping. This decentralized pumping structure is likely to become more common with the transition to future generations of district heating.

While data-driven hydraulic modeling is considered in this thesis, maintaining a model of a full district heating network consisting of consumer agents counted in the hundreds or thousands remains a challenging task. The challenge further increases if the aim is to use this large and complex model for designing the controller of each individual agent. Hence this thesis also considered control structures that do not explicitly use such a model but rather exploit expected input-output properties of the agent interconnection. In particular, the thesis considers two forms of proportional-integral control, using either fully decentralized anti-windup compensation, or a form of coordinating anti-windup action which balances the control actions of the agents. This thesis demonstrates that such an approach can be globally asymptotically stable in the fully decentralized case, and that the closed-loop equilibria to these closed-loop systems will be optimal. This is in the sense that the decentralized and coordinating control strategies minimize objectives $J(x)$ of equilibrium control errors, in the form of $J(x) = \|\Gamma x\|_1$ and $J(x) = \|x\|_\infty$ respectively. Here Γ is a diagonal matrix embodying certain system parameters. The considered design goal of asymptotically minimizing these objectives could be extended in fu-

ture works. In the district heating setting for instance, a popular consideration is introducing load flexibility over time, to increase or decrease demand peaks in the network.

This thesis has completely omitted the analysis of thermodynamics in the network, which clearly are affected by changes in flow rates. A first step towards bridging this gap could be to extend the results of Paper IV to time-varying interconnections between the agents. In this sense, the thermodynamics would not be explicitly modeled but rather assumed to vary within a bounded rate. Another reasonable step would be to evaluate the proposed control methods of this thesis in more high-fidelity simulations.

More detailed building models than what is considered in papers II- IV are necessary in order to move towards practical applicability. In a real building, more than one temperature state is clearly necessary. Furthermore, the actual indoor temperature is rarely used directly for feedback in the control loop governing the flow rate. Rather, the flow control loop typically utilizes measurements of the temperature in the building heating system, which itself is governed by e.g., thermostatic valves.

The coordinating controller subject to papers II and IV can be developed further. With the current implementation, equilibrium existence is not guaranteed for certain disturbances. What can occur is that the controllers of two (or more) different agents wind up in different directions (positive and negative), which will be invisible in the coordination term $\mathbf{1}^\top \text{dz}(u)$. Simple, heuristic extensions of the coordination rule can solve this issue (e.g., if agent i is currently leveraging maximally negative control action ($\text{dz}(u_i) < 0$), they ignore positive values of the coordination term $\mathbf{1}^\top \text{dz}(u) > 0$). However, such approaches can introduce new complications, such as cluttered notation and complicating the analysis of equilibrium optimality.

In the transition from papers II and III to Paper IV, an extension from linear agent interconnections modeled by M-matrices to a fully nonlinear interconnection was introduced. Only two assumptions beyond continuity were placed on this interconnection, corresponding roughly to properties found in M-matrices. In the case of M-matrices, a whole list of equivalent properties can be listed, i.e., Proposition 1. It would be interesting to investigate such a list of equivalent properties in the nonlinear domain as well. A list of such equivalent properties can be utilized as a tool to investigate if a particular system falls under the assumptions of Paper IV. Furthermore, the explicit connection between these assumptions and district heating hydraulics are not shown in this thesis. For tree structured networks, the assumptions can be proven to hold based on (3.6). This result is omitted from this thesis in order to maintain novelty for a potential future publication.

This thesis concerns issues which would be interesting to view from a game-theoretic perspective. Game theory is a common way to view systems where several agents all interact and subsequently receive some utility from their shared interaction. In this setting, one could consider market design mechanisms to incentivize the fair distribution of the limited resource capacity. This stands in contrast to the approach of this thesis, where the coordination between agents is not incentivized,

but rather mandated.

On the topic of agent coordination, further work could also tackle the issue of certain agents in the system redesigning their controllers to cheat the coordination scheme. Such agents can perhaps trick the system into allocating more of the limited resource capacity to them.

A significant outstanding issue is general stability guarantees for the coordinating closed-loop system analyzed in Paper II and IV. As the main author working with these papers, I have performed a vast number of simulations that lead me to believe that under reasonable assumptions on controller parameters, and assuming the existence of an equilibrium, these systems are stable. However, a proof of this statement is nonetheless not established at the time of writing this thesis.

Finally, a key take-away from this thesis: By leveraging an understanding of system properties, complex control challenges can be tackled with simple feedback-based solutions.

Bibliography

- Agner, F. (2023). *On Hydraulic Constraints in Control of District Heating Systems*. Supervisors: Anders Rantzer, Richard Pates, Pauline Kergus. Licentiate Thesis. Department of Automatic Control, Lund University, Lund, Sweden. URL: <https://portal.research.lu.se/en/publications/on-hydraulic-constraints-in-control-of-district-heating-systems>.
- Agner, F., J. Hansson, P. Kergus, A. Rantzer, S. Tarbouriech, and L. Zaccarian (2024a). “Decentralized PI-control and anti-windup in resource sharing networks”. *European Journal of Control*, p. 101049. ISSN: 0947-3580. DOI: 10.1016/j.ejcon.2024.101049.
- Agner, F., C. M. Jensen, A. Rantzer, C. S. Kallesøe, and R. Wisniewski (2024b). “Hydraulic parameter estimation for district heating based on laboratory experiments”. *Energy* **312**, p. 133462. ISSN: 0360-5442. DOI: <https://doi.org/10.1016/j.energy.2024.133462>.
- Agner, F., P. Kergus, R. Pates, and A. Rantzer (2022). “Combating district heating bottlenecks using load control”. *Smart Energy* **6**. ISSN: 2666-9552. DOI: 10.1016/j.segy.2022.100067.
- Agner, F., P. Kergus, R. Pates, and A. Rantzer (2023a). “Hydraulic Parameter Estimation in District Heating Networks”. *IFAC-PapersOnLine*. 22nd IFAC World Congress **56**:2, pp. 5438–5443. ISSN: 2405-8963. DOI: 10.1016/j.ifacol.2023.10.194.
- Agner, F., P. Kergus, A. Rantzer, S. Tarbouriech, and L. Zaccarian (2023b). “Anti-Windup Coordination Strategy Around a Fair Equilibrium in Resource Sharing Networks”. *IEEE Control Systems Letters* **7**, pp. 2521–2526. ISSN: 2475-1456. DOI: 10.1109/LCSYS.2023.3287252.
- Agner, F. and A. Rantzer (2024). *On PI-control in capacity-limited networks*. arXiv: 2411.14077 [eess.SY].

- Agner, F., U. Trabert, A. Rantzer, and J. Orozalieva (2024c). “A data-based comparison of methods for reducing the peak flow rate in a district heating system”. *Smart Energy*. ISSN: 2666-9552. DOI: <https://doi.org/10.1016/j.segy.2024.100168>.
- Ahmed, S., J. E. Machado, M. Cucuzzella, and J. M. A. Scherpen (2023). “Control-oriented modeling and passivity analysis of thermal dynamics in a multi-producer district heating system”. *IFAC-PapersOnLine*. 12th IFAC Symposium on Nonlinear Control Systems NOLCOS 2022 **56**:1, pp. 175–180. ISSN: 2405-8963. DOI: [10.1016/j.ifacol.2023.02.030](https://doi.org/10.1016/j.ifacol.2023.02.030).
- Ala-Kotila, P., T. Vainio, and J. Heinonen (2020). “Demand response in district heating market-results of the field tests in student apartment buildings”. *Smart Cities* **3**:2, pp. 157–171. ISSN: 2624-6511. DOI: <https://doi.org/10.3390/smartcities3020009>.
- Åström, K. and B. Wittenmark (2008). *Adaptive Control*. 2nd ed. Dover Publications Inc., Mineola, New York. DOI: [doi:10.1515/9783111104959](https://doi.org/10.1515/9783111104959).
- Bacher, P. and H. Madsen (2011). “Identifying suitable models for the heat dynamics of buildings”. en. *Energy and Buildings* **43**:7, pp. 1511–1522. ISSN: 03787788. DOI: [10.1016/j.enbuild.2011.02.005](https://doi.org/10.1016/j.enbuild.2011.02.005).
- Bauso, D., F. Blanchini, L. Giarre, and R. Pesenti (2013). “The linear saturated decentralized strategy for constrained flow control is asymptotically optimal”. *Automatica* **49**:7, pp. 2206–2212. ISSN: 0005-1098. DOI: [10.1016/j.automatica.2013.03.029](https://doi.org/10.1016/j.automatica.2013.03.029).
- Benonysson, A., B. Bøhm, and H. F. Ravn (1995). “Operational optimization in a district heating system”. en. *Energy Conversion and Management* **36**:5, pp. 297–314. ISSN: 01968904. DOI: [10.1016/0196-8904\(95\)98895-T](https://doi.org/10.1016/0196-8904(95)98895-T).
- Biegel, B., P. Andersen, J. Stoustrup, and J. Bendtsen (2012). “Congestion Management in a Smart Grid via Shadow Prices”. *IFAC Proceedings Volumes*. 8th Power Plant and Power System Control Symposium **45**:21, pp. 518–523. ISSN: 1474-6670. DOI: [10.3182/20120902-4-FR-2032.00091](https://doi.org/10.3182/20120902-4-FR-2032.00091).
- Blanchini, F., D. Casagrande, F. Fabiani, G. Giordano, and R. Pesenti (2019). “Network-decentralised optimisation and control: An explicit saturated solution”. *Automatica* **103**, pp. 379–389. ISSN: 0005-1098. DOI: [10.1016/j.automatica.2019.02.009](https://doi.org/10.1016/j.automatica.2019.02.009).
- Blanchini, F., E. Franco, G. Giordano, V. Mardanlou, and P. L. Montessoro (2016). “Compartmental flow control: Decentralization, robustness and optimality”. *Automatica* **64**, pp. 18–28. ISSN: 0005-1098. DOI: [10.1016/j.automatica.2015.10.046](https://doi.org/10.1016/j.automatica.2015.10.046).
- Brange, L., K. Sernhed, and M. Thern (2019). “Decision-making process for addressing bottleneck problems in district heating networks”. en. *International Journal of Sustainable Energy Planning and Management* **20**, pp. 37–50. DOI: <https://doi.org/10.5278/ijsepm.2020.19.4>.

- Bürger, M., C. De Persis, and F. Allgöwer (2015). “Dynamic Pricing Control for Constrained Distribution Networks With Storage”. *IEEE Transactions on Control of Network Systems* **2**:1, pp. 88–97. ISSN: 2325-5870. DOI: 10.1109/TCNS.2014.2367572.
- Cai, H., S. You, and J. Wu (2020). “Agent-based distributed demand response in district heating systems”. *Applied Energy* **262**, p. 114403. ISSN: 0306-2619. DOI: 10.1016/j.apenergy.2019.114403.
- Capone, M., E. Guelpa, and V. Verda (2021). “Multi-objective optimization of district energy systems with demand response”. *Energy* **227**, p. 120472. ISSN: 0360-5442. DOI: <https://doi.org/10.1016/j.energy.2021.120472>.
- Chaudhry, A. M., M. Delvaux, P. Z. Csurcsia, S. Chicherin, J. Hachez, and S. Bram (2024). “Enhancing operational optimization of district heating substations through refined estimations of network campus buildings heat demands to achieve a low return from the network”. en. *Energy and Buildings* **313**, p. 114241. ISSN: 03787788. DOI: 10.1016/j.enbuild.2024.114241.
- Chu, K.-c. (1974). “Decentralized control of high-speed vehicular strings”. *Transportation Science* **8**:4, pp. 361–384. ISSN: 00411655, 15265447.
- Dall’Anese, E. and A. Simonetto (2018). “Optimal power flow pursuit”. *IEEE Transactions on Smart Grid* **9**:2, pp. 942–952. DOI: 10.1109/TSG.2016.2571982.
- De Persis, C., T. N. Jensen, R. Ortega, and R. Wisniewski (2014). “Output Regulation of Large-Scale Hydraulic Networks”. *IEEE Transactions on Control Systems Technology* **22**:1, pp. 238–245. ISSN: 1558-0865. DOI: 10.1109/TCST.2012.2233477.
- De Persis, C. and C. S. Kallesøe (2008). “Proportional and Proportional-Integral Controllers for a Nonlinear Hydraulic Network”. *IFAC Proceedings Volumes*. 17th IFAC World Congress **41**:2, pp. 319–324. ISSN: 1474-6670. DOI: 10.3182/20080706-5-KR-1001.00054.
- De Persis, C. and C. S. Kallesøe (2011). “Pressure Regulation in Nonlinear Hydraulic Networks by Positive and Quantized Controls”. *IEEE Transactions on Control Systems Technology* **19**:6, pp. 1371–1383. ISSN: 1558-0865. DOI: 10.1109/TCST.2010.2094619.
- Energimyndigheten (2017). *Summary of energy statistics for dwellings and nonresidential premises for 2016*. Tech. rep. Statens energimyndighet. URL: <https://www.energimyndigheten.se/49abd2/globalassets/statistik/officiell-statistik/statistikprodukter/energistatistik-i-smh-fbhlok/rapporter/energistatistik-for-smahus-flerbostadshus-och-lokaler-2016.pdf>.
- Frahm, M., F. Langner, P. Zwickel, J. Matthes, R. Mikut, and V. Hagenmeyer (2022). “How to derive and implement a minimalistic RC model from thermodynamics for the control of thermal parameters for assuring thermal comfort in buildings”, pp. 1–6. DOI: 10.1109/OSMSES54027.2022.9769134.

- Frederiksen, S. and S. Werner (2013). *District heating and cooling*. Studentlitteratur.
- Galeani, S., S. Tarbouriech, M. Turner, and L. Zaccarian (2009). “A Tutorial on Modern Anti-windup Design”. *European Journal of Control* **15**:3, pp. 418–440. ISSN: 0947-3580. DOI: 10.3166/ejc.15.418-440.
- Guelpa, E. and V. Verda (2021). “Demand response and other demand side management techniques for district heating: a review”. *Energy* **219**, p. 119440. ISSN: 0360-5442. DOI: <https://doi.org/10.1016/j.energy.2020.119440>.
- Hägglund, T. (2023). *Process Control in Practice*. De Gruyter, Berlin, Boston. ISBN: 9783111104959. DOI: doi:10.1515/9783111104959.
- Hauswirth, A., F. Dorfler, and A. Teel (2020a). “On the Robust Implementation of Projected Dynamical Systems with Anti-Windup Controllers”. In: *2020 American Control Conference (ACC)*. IEEE, Denver, CO, USA, pp. 1286–1291. ISBN: 9781538682661. DOI: 10.23919/ACC45564.2020.9147378.
- Hauswirth, A., F. Dörfler, and A. Teel (2020b). “On the Differentiability of Projected Trajectories and the Robust Convergence of Non-Convex Anti-Windup Gradient Flows”. *IEEE Control Systems Letters* **4**:3, pp. 620–625. ISSN: 2475-1456. DOI: 10.1109/LCSYS.2020.2988515.
- Hauswirth, A., Z. He, S. Bolognani, G. Hug, and F. Dörfler (2024). “Optimization algorithms as robust feedback controllers”. en. *Annual Reviews in Control* **57**, p. 100941. ISSN: 13675788. DOI: 10.1016/j.arcontrol.2024.100941.
- Horn, R. A. and C. R. Johnson (1991). *Topics in Matrix Analysis*. Cambridge University Press. DOI: 10.1017/CB09780511840371.
- Iftar, A. and E. J. Davison (1990). “Decentralized robust control for dynamic routing of large scale networks”. In: *1990 American Control Conference*, pp. 441–446. DOI: 10.23919/ACC.1990.4790775.
- Jeeninga, M., J. E. Machado, M. Cucuzzella, G. Como, and J. Scherpen (2023). “On the Existence and Uniqueness of Steady State Solutions of a Class of Dynamic Hydraulic Networks via Actuator Placement”. In: *2023 62nd IEEE Conference on Decision and Control (CDC)*. ISSN: 2576-2370, pp. 3652–3657. DOI: 10.1109/CDC49753.2023.10384050.
- Kallesøe, C. S., B. K. Nielsen, A. Overgaard, and E. B. Sørensen (2019). “A Distributed Algorithm for Auto-Balancing of Hydronic Cooling Systems”. In: *2019 IEEE Conference on Control Technology and Applications (CCTA)*. IEEE, Hong Kong, China, pp. 655–660. ISBN: 9781728127675. DOI: 10.1109/CCTA.2019.8920702.
- Kallesøe, C. S., B. K. Nielsen, and A. Tsouvalas (2020). “Heat Balancing in Cooling Systems using Distributed Pumping”. en. *IFAC-PapersOnLine* **53**:2, pp. 3292–3297. ISSN: 24058963. DOI: 10.1016/j.ifacol.2020.12.1139.

- Karasu, M. B., . Yankolu, A. Aykut, D. Ay, and . Özpoyraz (2024). “Thermal request optimization of a smart district heating system”. en. *Energy and Buildings* **324**, p. 114841. ISSN: 0378-7788. DOI: 10.1016/j.enbuild.2024.114841.
- Kelly, F. P., A. K. Maulloo, and D. K. H. Tan (1998). “Rate control for communication networks: shadow prices, proportional fairness and stability”. en. *Journal of the Operational Research Society* **49**:3, pp. 237–252. ISSN: 0160-5682, 1476-9360. DOI: 10.1057/palgrave.jors.2600523.
- Kelly, F. (2003). “Fairness and stability of end-to-end congestion control*”. *European Journal of Control* **9**:2, pp. 159–176. ISSN: 0947-3580. DOI: <https://doi.org/10.3166/ejc.9.159-176>.
- Krishnamoorthy, D. and S. Skogestad (2022). “Real-Time optimization as a feedback control problem A review”. en. *Computers & Chemical Engineering* **161**, p. 107723. ISSN: 00981354. DOI: 10.1016/j.compchemeng.2022.107723.
- Levine, W. and M. Athans (1966). “On the optimal error regulation of a string of moving vehicles”. *IEEE Transactions on Automatic Control* **11**:3, pp. 355–361. DOI: 10.1109/TAC.1966.1098376.
- Liu, Y., P. Wang, and P. Luo (2020). “Pipe hydraulic resistances identification of district heating networks based on matrix analysis”. *Energies* **13**:11. ISSN: 1996-1073. DOI: 10.3390/en13113007.
- Low, S. and D. Lapsley (1999). “Optimization flow control. I. Basic algorithm and convergence”. *IEEE/ACM Transactions on Networking* **7**:6, pp. 861–874. ISSN: 10636692. DOI: 10.1109/90.811451.
- Low, S., F. Paganini, and J. Doyle (2002). “Internet congestion control”. *IEEE Control Systems Magazine* **22**:1, pp. 28–43. DOI: 10.1109/37.980245.
- Lund, H., S. Werner, R. Wiltshire, S. Svendsen, J. Thorsen, F. Hvelplund, and B. Mathiesen (2014). “4th generation district heating (4GDH)”. *Energy* **68**, pp. 1–11. ISSN: 03605442. DOI: 10.1016/j.energy.2014.02.089.
- Lund, H., P. A. Østergaard, M. Chang, S. Werner, S. Svendsen, P. Sorknæs, J. E. Thorsen, F. Hvelplund, B. O. G. Mortensen, B. V. Mathiesen, C. Bojesen, N. Duic, X. Zhang, and B. Möller (2018). “The status of 4th generation district heating: research and results”. *Energy* **164**, pp. 147–159. ISSN: 0360-5442. DOI: <https://doi.org/10.1016/j.energy.2018.08.206>.
- Lund, H., P. A. Østergaard, T. B. Nielsen, S. Werner, J. E. Thorsen, O. Gudmundsson, A. Arabkoohsar, and B. V. Mathiesen (2021). “Perspectives on fourth and fifth generation district heating”. *Energy* **227**, p. 120520. ISSN: 0360-5442. DOI: <https://doi.org/10.1016/j.energy.2021.120520>.
- Luo, P., H. Wang, Y. Liu, Q. Du, and Z. Zhang (2022). “Resistance characteristic parameters estimation of hydraulic model in heating networks based on real-time operation data”. *Buildings* **12**:6. ISSN: 2075-5309. DOI: 10.3390/buildings12060743.

- Machado, J. E., M. Cucuzzella, N. Pronk, and J. M. A. Scherpen (2022a). “Adaptive Control for Flow and Volume Regulation in Multi-Producer District Heating Systems”. *IEEE Control Systems Letters* **6**, pp. 794–799. ISSN: 2475-1456. DOI: 10.1109/LCSYS.2021.3085702.
- Machado, J. E., M. Cucuzzella, and J. M. A. Scherpen (2022b). “Modeling and passivity properties of multi-producer district heating systems”. *Automatica* **142**, p. 110397. ISSN: 0005-1098. DOI: 10.1016/j.automatica.2022.110397.
- Machado, J. E., J. Ferguson, M. Cucuzzella, and J. M. A. Scherpen (2023). “Decentralized Temperature and Storage Volume Control in Multiproducer District Heating”. *IEEE Control Systems Letters* **7**, pp. 413–418. ISSN: 2475-1456. DOI: 10.1109/LCSYS.2022.3189321.
- Machowski, J., Z. Lubosny, J. W. Bialek, and J. R. Bumby (2020). *Power System Dynamics: Stability and Control*. 3rd. John Wiley & Sons. ISBN: 9781119526360.
- Madjidian, D. and L. Mirkin (2014). “Distributed control with low-rank coordination”. *IEEE Transactions on Control of Network Systems* **1**:1, pp. 53–63. DOI: 10.1109/TCNS.2014.2309712.
- Martins, A. and K.-E. Årzén (2021). “Dynamic Management of Multiple Resources in Camera Surveillance Systems”. In: *2021 American Control Conference (ACC)*. IEEE, New Orleans, LA, USA, pp. 2061–2068. ISBN: 9781665441971. DOI: 10.23919/ACC50511.2021.9482666.
- Martins, A., M. Lindberg, M. Maggio, and K.-E. Årzén (2020). “Control-Based Resource Management for Storage of Video Streams”. en. *IFAC-PapersOnLine* **53**:2, pp. 5542–5549. ISSN: 24058963. DOI: 10.1016/j.ifacol.2020.12.1564.
- Molzahn, D. K., F. Dörfler, H. Sandberg, S. H. Low, S. Chakrabarti, R. Baldick, and J. Lavaei (2017). “A Survey of Distributed Optimization and Control Algorithms for Electric Power Systems”. *IEEE Transactions on Smart Grid* **8**:6, pp. 2941–2962. ISSN: 1949-3061. DOI: 10.1109/TSG.2017.2720471.
- Ortmann, L., C. Rubin, A. Scozzafava, J. Lehmann, S. Bolognani, and F. Dörfler (2023). “Deployment of an Online Feedback Optimization Controller for Reactive Power Flow Optimization in a Distribution Grid”. In: *2023 IEEE PES Innovative Smart Grid Technologies Europe (ISGT EUROPE)*. IEEE, Grenoble, France, pp. 1–6. DOI: 10.1109/ISGTEUROPE56780.2023.10408057.
- Rantzer, A. and M. E. Valcher (2018). “A Tutorial on Positive Systems and Large Scale Control”. en. In: *2018 IEEE Conference on Decision and Control (CDC)*. IEEE, Miami Beach, FL, pp. 3686–3697. ISBN: 978-1-5386-1395-5. DOI: 10.1109/CDC.2018.8618689.

- Saletti, C., N. Zimmerman, M. Morini, K. Kyprianidis, and A. Gambarotta (2022). “A control-oriented scalable model for demand side management in district heating aggregated communities”. *Applied Thermal Engineering* **201**, p. 117681. ISSN: 1359-4311. DOI: <https://doi.org/10.1016/j.applthermaleng.2021.117681>.
- Sibeiijn, M., S. Ahmed, M. Khosravi, and T. Keviczky (2024). “Dissipativity Analysis for Economic Nonlinear MPC of District Heating Networks”. In: *2024 European Control Conference (ECC)*, pp. 1111–1118. DOI: 10.23919/ECC64448.2024.10590838.
- Silva, J. da and S. Tarbouriech (2005). “Antiwindup design with guaranteed regions of stability: an LMI-based approach”. *IEEE Transactions on Automatic Control* **50**:1, pp. 106–111. ISSN: 1558-2523. DOI: 10.1109/TAC.2004.841128.
- Simonsson, J., K. T. Atta, and W. Birk (2024a). “A Graph Theoretical Approach to Modeling of District Energy Networks”. *IEEE Transactions on Control Systems Technology* **32**:5, pp. 1616–1630. ISSN: 1558-0865. DOI: 10.1109/TCST.2023.3345213.
- Simonsson, J., K. T. Atta, and W. Birk (2024b). “Semi-decentralized temperature control in district heating systems”. *Journal of Process Control* **140**, p. 103251. ISSN: 0959-1524. DOI: 10.1016/j.jprocont.2024.103251.
- Strehle, F., J. E. Machado, M. Cucuzzella, A. J. Malan, J. M. Scherpen, and S. Hohmann (2022). “Port-Hamiltonian Modeling of Hydraulics in 4th Generation District Heating Networks”. In: *2022 IEEE 61st Conference on Decision and Control (CDC)*. ISSN: 2576-2370, pp. 1182–1189. DOI: 10.1109/CDC51059.2022.9992887.
- Strehle, F., J. E. Machado, M. Cucuzzella, A. J. Malan, S. Hohmann, and J. M. A. Scherpen (2024). “A Unifying Passivity-Based Framework for Pressure and Volume Flow Rate Control in District Heating Networks”. *IEEE Transactions on Control Systems Technology* **32**:4, pp. 1323–1340. ISSN: 1558-0865. DOI: 10.1109/TCST.2024.3365250.
- Strehle, F., J. Vieth, M. Pfeifer, and S. Hohmann (2021). “Passivity-Based Stability Analysis of Hydraulic Equilibria in 4th Generation District Heating Networks”. *IFAC-PapersOnLine*. 7th IFAC Workshop on Lagrangian and Hamiltonian Methods for Nonlinear Control LHMNC 2021 **54**:19, pp. 261–266. ISSN: 2405-8963. DOI: 10.1016/j.ifacol.2021.11.088.
- Sweetnam, T., C. Spataru, M. Barrett, and E. Carter (2019). “Domestic demand-side response on district heating networks”. en. *Building Research & Information* **47**:4, pp. 330–343. ISSN: 0961-3218, 1466-4321. DOI: 10.1080/09613218.2018.1426314.
- Tan, Y., W. Moase, C. Manzie, D. Nei, and I. Mareels (2010). “Extremum seeking from 1922 to 2010”. In: *Proceedings of the 29th Chinese Control Conference*. ISSN: 2161-2927, pp. 14–26.

- Trip, S., T. Scholten, and C. D. Persis (2019). “Optimal regulation of flow networks with transient constraints”. *Automatica* **104**, pp. 141–153. ISSN: 0005-1098. DOI: 10.1016/j.automatica.2019.02.046.
- van der Heijde, B., M. Fuchs, C. Ribas Tugores, G. Schweiger, K. Sartor, D. Basciotti, D. Müller, C. Nytsch-Geusen, M. Wetter, and L. Helsen (2017a). “Dynamic equation-based thermo-hydraulic pipe model for district heating and cooling systems”. *Energy Conversion and Management* **151**, pp. 158–169. ISSN: 0196-8904. DOI: <https://doi.org/10.1016/j.enconman.2017.08.072>.
- van der Heijde, B., M. Sourbron, F. Vega Arance, R. Salenbien, and L. Helsen (2017b). “Unlocking flexibility by exploiting the thermal capacity of concrete core activation”. *Energy Procedia* **135**. 11th International Renewable Energy Storage Conference, IRES 2017, 14-16 March 2017, Düsseldorf, Germany, pp. 92–104. ISSN: 1876-6102. DOI: <https://doi.org/10.1016/j.egypro.2017.09.490>.
- Vandermeulen, A., B. van der Heijde, and L. Helsen (2018). “Controlling district heating and cooling networks to unlock flexibility: A review”. en. *Energy* **151**, pp. 103–115. ISSN: 03605442. DOI: 10.1016/j.energy.2018.03.034.
- Wang, N., S. You, Y. Wang, H. Zhang, Q. Miao, X. Zheng, and L. Mi (2018). “Hydraulic resistance identification and optimal pressure control of district heating network”. *Energy and Buildings* **170**, pp. 83–94. ISSN: 0378-7788. DOI: <https://doi.org/10.1016/j.enbuild.2018.04.003>.
- Wang, Y., S. You, H. Zhang, W. Zheng, X. Zheng, and Q. Miao (2017). “Hydraulic performance optimization of meshed district heating network with multiple heat sources”. *Energy* **126**, pp. 603–621. DOI: 10.1016/j.energy.2017.03.044.
- Zecevic, A. and D. Siljak (2005). “Global low-rank enhancement of decentralized control for large-scale systems”. *IEEE Transactions on Automatic Control* **50**:5, pp. 740–744. DOI: 10.1109/TAC.2005.847054.
- Zheng, X., Z. Shi, Y. Wang, H. Zhang, and Z. Tang (2024). “Digital twin modeling for district heating network based on hydraulic resistance identification and heat load prediction”. *Energy* **288**, p. 129726. ISSN: 0360-5442. DOI: <https://doi.org/10.1016/j.energy.2023.129726>.

Paper I

Combating District Heating Bottlenecks Using Load Control

Felix Agner Pauline Kergus Richard Pates Anders Rantzer

Abstract

The 4th generation of district heating systems face a potential problem where lowered water temperatures lead to higher flow rates, which requires higher hydraulic capacity in terms of pipe and pump sizes. This increases the effect of the already existing issue of hydraulic bottlenecks, causing peripheral units (customers) to experience reduced flow rates. A coordinating control strategy is presented in this work aimed at reducing the effect of such bottlenecks on the comfort of customers. This is done by distributing the flow deficit over many units rather than a few. Previous works mainly focus on MPC-structured controllers that depend on complex system models and online optimization techniques. This work proposes a method that requires little information about models for individual units and minimal IT communication between control systems. The proposed method is compared with a traditional control strategy and an optimal baseline in a simulation study. This shows that the proposed method can decrease the worst case indoor temperature deviations.

Reprinted, with permission, from F. Agner et al. (2022). “Combating district heating bottlenecks using load control”. *Smart Energy* **6**. ISSN: 2666-9552. DOI: 10.1016/j.segy.2022.100067.

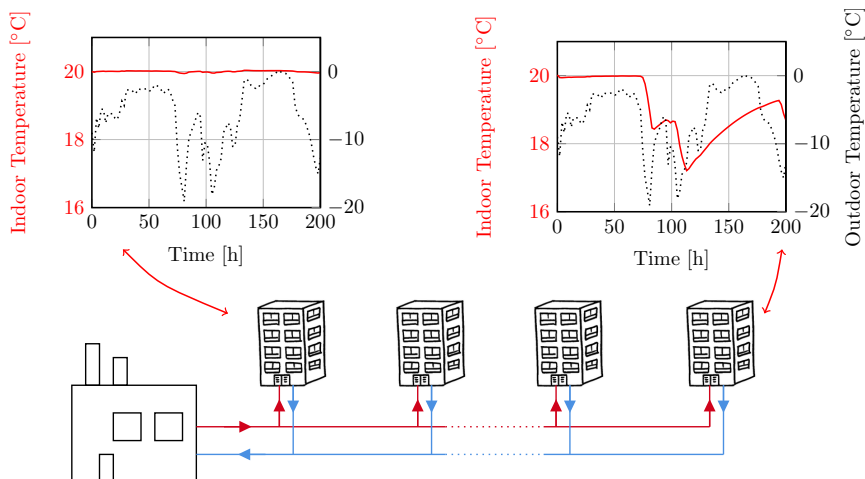


Figure 1. Discrepancy in indoor temperature (red) between units connected to the grid. When the outdoor temperature (dotted) becomes critically cold, units close to the pressure source are able to maintain indoor comfort temperature while units far from the source are not.

1. Introduction

One important puzzle piece of the smart energy system of the future is the integration of a variety of energy sources and distribution methods. This allows harnessing synergies and reducing the impact of stochastic fluctuations in energy supply and demand [Mathiesen et al., 2015]. District heating systems have been shown to be a powerful tool in this energy system, but research indicates that a transformation of district heating from the old 3rd generation to a new 4th generation is needed. An important characteristic of this emerging 4th generation that allows it to be integrated into the overall energy chain is reduced supply water temperatures which would allow using previously untapped heat sources such as renewable sources and industrial waste heat [Lund et al., 2014]. In theory, the reduction in supply temperature should be accompanied by an equal drop in return temperature, leading to an equal temperature difference and thus no alternation in the necessary flow [Lund et al., 2018]. However, lowering the building return temperatures requires an improvement in space heating technology [Lund et al., 2017; Lund et al., 2018] and if such is not the case there may be a reduction in differential temperature. This leads to higher flows needed to distribute the same amount of power, implying that the piping and pumping power of 4th generation district heating systems may have to be dimensioned for higher capacity. This presents an additional cost. It also reduces the potential of retrofitting existing infrastructure for lower grid temperatures, which otherwise may prove a cost-effective solution [Lund et al., 2018; Brange et

al., 2019]. If the grid capacity is not dimensioned for higher flows, it may lead to **bottlenecks** [Brange et al., 2019]. Bottlenecks imply that the pressure losses in part of the distribution grid are too high, leaving an insufficient differential pressure between supply and return pipes in subsequent parts of the network. This happens when the water flow is too great for the dimensions of the piping. Buildings (units) connected in these parts may find it hard to extract sufficient flows to keep indoor temperatures at comfort level. In fact, this is not only a hypothetical problem in future generations of district heating but is already a problem in currently operating networks [Frederiksen and Werner, 2013]. This phenomenon arises under **peak load conditions**, i.e. when the flows in the network are high, coinciding with when the outdoor temperature is low. Figure 1 shows this problem, based on simulation which will be explained later in this article. When the outdoor temperature becomes too low, the indoor temperatures start to differ from each other. Buildings close to the pressure source maintain comfort temperature while it becomes cold in buildings further down the distribution line. Reducing the effect of bottlenecks could increase **robustness** to low outdoor temperatures, in the sense that a drop in outdoor temperature would cause reduced worst-case deviations in indoor temperature. Apart from the possibility of reducing supply temperatures, this could also grant the possibility of extending existing networks, and designing larger new networks with less concern for critical outdoor temperature and the influence this would have on customer comfort. One approach to tackling this issue is through the use of demand side management, as suggested in [Frederiksen and Werner, 2013; Vandermeulen et al., 2018; Brange et al., 2019; Guelpa and Verda, 2021].

Demand side management is an umbrella term for different ways of altering the demand of customers connected to a grid. There is a rich history of demand side management in the power grid literature, but it has also begun making an appearance in the district heating literature [Guelpa and Verda, 2021]. This article focuses specifically on direct load control, i.e. directly altering and deciding the heat load of customers. A common approach to direct load control is to use **centralized** optimization with the objective of minimizing some operational cost for the entire network. For instance an optimization scheme was introduced in [Bhattacharya et al., 2019] to improve fairness of heat distribution in a line network. As the optimization problems tend to grow drastically with the number of connected units, these methods can run into problems of scalability. Another approach is to have **de-centralized** optimal controllers, such as in [Saletti et al., 2020], and then combine their signals to compute desired heat production. However in this scenario there is no coordination between the different units to ensure that they request a load that is feasible.

To ensure that any enacted heat loads are within the system constraints, a good model of the grid is needed. In practice, the considered distribution model depends on the purpose of the model as well as the design of the network, as summarized in [Sarbu et al., 2019]. Some works choose to disregard aspects of the constraints given by the network such as pressure losses or time delays [Bhattacharya et al., 2019].

One common approach to model flows and pressures of the network is to assume that the specific heat load at each unit has to be met, and then using these loads to calculate the flows realized in the network [Wang et al., 2017; Benonysson et al., 1995; Larsen et al., 2002]. However this approach does not hold in the case where the desired heat loads of each building cannot be met due to the corresponding flows being too large for the network to handle. To the authors' knowledge, there is little work on describing the limits on the flows in the system.

This work is focused on understanding how bottlenecks can be combated through direct load control, in such a way that the hydraulic constraints of the network are taken into account and the control structure remains scalable for a large number of connected units. The idea is to combine the increased ease of implementation of distributed controllers with the system-level benefits of a centralized strategy. An important distinction to make is that many works on demand side management try to optimize some operational cost, e.g. cost of energy production units. In this work we assume that the system operates at full capacity, and the objective is simply to distribute the supplied energy fairly between customers.

The contribution of this work is in three parts;

1. formulating the constraints limiting the unit flows in a line-structured district heating grid;
2. introducing a load coordination scheme that builds on the traditional control architecture of district heating such that it should be easy to implement in existing networks, and;
3. comparing two control architectures; traditional control and the aforementioned load coordination architecture, with an optimal baseline reference through a simulation study.

The work is presented as follows: Section 3.1 introduces a mathematical formulation of the problem. The notion of robustness to outdoor temperatures is introduced here. Section 3 presents the mathematical models of the network and the units connected to the network. Section 4 defines two different control architectures and an optimal baseline reference, which are then compared in a simulation study described in Section 5. The results and future work are finally discussed in Section 6.

2. List of Notation

Table 1 lists the notation used in this article. In addition to this table, three other conventions are used:

- Vectors are noted with bold face, e.g. \mathbf{q} denotes a vector of flows.
- Indices are used to reference elements of a vector or matrix. If the vector notation already contains subset notation, these are separated by a comma.

For instance, $T_{c,i}$ refers to element i in the vector of comfort temperatures \mathbf{T}_c . M_{ij} refers to matrix element of row index i , column index j .

Table 1. Table of notation used in this article.

Symbol	Description	Unit
J	Cost function related to discomfort	$^{\circ}\text{C}\text{s}$
T	Temperature	$^{\circ}\text{C}$
T_c	Comfort temperature	$^{\circ}\text{C}$
T_{in}	Indoor temperature	$^{\circ}\text{C}$
e	Difference between indoor and comfort temperature	$^{\circ}\text{C}$
T_{hs}	Temperature of water in heating system	$^{\circ}\text{C}$
T_{ext}	Outdoor temperature	$^{\circ}\text{C}$
T_{sup}	Primary side supply temperature	$^{\circ}\text{C}$
T_{ret}	Primary side return temperature	$^{\circ}\text{C}$
t	Time	s
t_s	Sampling time	s
q	Water flow	kg/s
\mathcal{Q}	The set of admissible system water flows	-
C_{in}	Heat capacity of indoor area	J/ $^{\circ}\text{C}$
C_{hs}	Heat capacity of water in radiator systems	J/ $^{\circ}\text{C}$
C_w	Specific heat capacity of water	J/kg/ $^{\circ}\text{C}$
R_{ext}	Heat resistance between building interior and exterior	$^{\circ}\text{C}/\text{W}$
R_{hs}	Heat resistance between radiator system and building interior	$^{\circ}\text{C}/\text{W}$
P	Furnished heat power	-
A, B_g, B_{ext}	Matrices defining the dynamics of simulated buildings	-
Δp	Differential pressure	Pa
\mathcal{L}	Network loop	-
F	Network incidence matrix	-
a	Hydraulic resistance	Pa/(kg/s) ²
c	Pump curve parameters	-
r	Pump frequency ratio	-
α_0, α_1	Heating system temperature set-point parameters	-
k	Building P-controller gain	kg/s/ $^{\circ}\text{C}$
δ	Flow set-point deviation	kg/s
γ	Coordination weight factor	s/ $^{\circ}\text{C}/\text{kg}$
λ	Coordination price factor	-

3. Problem and System Formulation

This section formalizes the problem of this work. Part 3.1 puts the problem to be solved in mathematical form. Part 3.2 presents the model of building temperature dynamics and the union between buildings and the district heating grid. Part 3.3 explains the hydraulic model of the distribution network, dictating the constraints on hot water flow in the system. Formulating these constraints in closed form constitutes the first contribution of the article.

3.1 Problem Formulation

The control problem of this work is to maintain comfortable indoor temperatures in all buildings connected to a district heating network even under extreme disturbances in the form of low outdoor temperatures. The control signal deciding the amount of heat furnished to each building i is the flow of hot water q_i through their substation. Mathematically, the problem is formulated as:

$$\min_{\mathbf{q}(t_k)} J(\mathbf{T}) \quad (1a)$$

$$\text{subject to} \quad \mathbf{T}(t_{k+1}) = f(\mathbf{T}(t_k), \mathbf{q}(t_k), T_{\text{ext}}(t_k)) \quad (1b)$$

$$\mathbf{q}(t_k) \in \mathcal{Q}. \quad (1c)$$

What this means is that we want to minimize some discomfort J related to the indoor temperatures \mathbf{T} in the connected buildings. These temperatures \mathbf{T} evolve according to dynamics f , which depend on the furnished flows \mathbf{q} and the outdoor temperature T_{ext} . This relationship f will be detailed in the next section, 3.2 and is in this work modelled linearly. Lastly, the furnished flow \mathbf{q} is limited by the capacity of the distribution system. The set \mathcal{Q} of flows that can be realized in the system is the subject of Section 3.3.

The cost function J should capture the discomfort experienced by each customer. To define this discomfort, consider the temperature deviation $e_i(t_k)$ for each unit i connected to the grid at each point in time t_k . $e_i(t_k)$ is the deviation between the desired comfort temperature $T_{c,i}$ and the actual indoor temperature $T_{\text{in},i}(t_k)$.

$$e_i(t_k) = T_{c,i} - T_{\text{in},i}(t_k) \quad (2)$$

The discomfort J_i experienced by a unit during a time period $t = t_1, t_2, \dots, t_K$ can then be defined as

$$J_i = \sum_{k=1}^K |e_i(t_k)t_s|, \quad (3)$$

where t_s is the time in between times t_k and t_{k+1} . Note also three candidates for measuring the system-level discomfort, J_1 , J_2 and J_∞ :

$$J_1 = \sum_{k=0}^K \frac{1}{N} \sum_{i=1}^N |e_i(t_k)t_s| \quad (4)$$

$$J_2 = \sum_{k=0}^K \sqrt{\frac{1}{N^2} \sum_{i=1}^N |e_i(t_k)t_s|^2} \quad (5)$$

$$J_\infty = \sum_{k=0}^K \max_i (|e_i(t_k)t_s|) \quad (6)$$

Here J_1 is a metric for the sum of discomfort experienced by all units, J_2 is a metric for the total discomfort where larger units discomfort are penalized more, and J_∞ is a metric for the worst discomfort experienced in the grid. The scenario we want to avoid is for extreme discomfort levels to arise in any unit, and for this reason the J_∞ -cost is the cost that will be used in the controller design of Section 4. The two remaining costs, J_1 and J_2 will be used for evaluation as a point of reference.

REMARK 1

Some works also consider optimizing over the power required to actuate the flows and temperatures in the grid and thus minimize the cost of running the system. For instance [Saletti et al., 2020] consider the utilized pumping power and [Bacher and Madsen, 2011] consider the electrical heating power in an adjacent problem considering an electrically heated unit. In this work we don't consider the cost of running the system. As we are interested in fair distribution under extremely cold situations, it is assumed that the heat and pumping power supplied to the system will have to be at maximum capacity. The interest is rather in understanding how to distribute this supplied power between connected units.

3.2 Buildings

Here we investigate the dynamics dictating the temperatures in each building, i.e. the function f of (1b). With each building i , we associate two states $T_{in,i}$ and $T_{hs,i}$, representing the mean indoor temperature and mean temperature of heating system circulating water respectively. This allows the construction of the following state space representation:

$$C_{in,i}\dot{T}_{in,i} = -\left(\frac{1}{R_{ext,i}} + \frac{1}{R_{hs,i}}\right)T_{in,i} + \frac{1}{R_{hs,i}}T_{hs,i} + \frac{1}{R_{ext,i}}T_{ext} \quad (7)$$

$$C_{hs,i}\dot{T}_{hs,i} = \frac{1}{R_{hs,i}}T_{in,i} - \frac{1}{R_{hs,i}}T_{hs,i} + P_i, \quad (8)$$

where $C_{in,i}$ and $C_{hs,i}$ is heat capacity of the building interior and heating system respectively. P_i is the heat power extracted from the primary side of the district heating system. The heat energy flow between interior and exterior as well as between heating and system interior are proportional to the inverse of the heat resistances $R_{ext,i}$ and $R_{hs,i}$ respectively. These types of models of varying complexity have been used extensively in literature on modeling building temperature dynamics, [Bacher and Madsen, 2011; Bhattacharya et al., 2019; Saletti et al., 2020], and can be augmented to capture different levels of complexity. In this work, a simple model of buildings is used, motivated by the interest in understanding the general distribution of temperatures in a large set of buildings, rather than the details of one individual building. The presented continuous time state space representation can then be transformed into a discrete time representation of the system if a standard zero-order-hold assumption is made for the inputs T_{ext} and P_i .

The heat energy, P_i , extracted from the network is here assumed to be proportional to the water flow through the primary side pipes of the building substation and the temperature difference between supply and return pipes in the network, $(T_{\text{sup}} - T_{\text{ret}})$:

$$P_i = C_w(T_{\text{sup}} - T_{\text{ret}})q_i \quad (9)$$

where C_w is the specific heat capacity of water. In the simulations and analysis in this work, the supply and return temperatures in the network are considered constant. This simplification is made to simplify simulations and analysis. While these temperatures are not constant in a real system, they are measured in building substations. As such, they could be included in the control strategy, where the now constant values would simply be exchanged for measured values.

To simplify the equations above, we can gather the indoor temperatures of all buildings into one vector \mathbf{T} , and the dynamics can then be put on the following linear form:

$$\mathbf{T}(t_{k+1}) = \mathbf{A}\mathbf{T}(t_k) + \mathbf{B}_q\mathbf{q}(t_k) + \mathbf{B}_{\text{ext}}T_{\text{ext}}(t_k) \quad (10)$$

REMARK 2

In this model we do not take domestic hot water use into account, much due to the difficulty of including a realistic model of this usage. In a typical scenario, the flow q_i through each building would consist of two parts, one for space heating and one for hot water usage.

3.3 Distribution Model

This part formulates the constraints on hot water flows in the distribution network, the first contribution of this article. This corresponds to the set \mathcal{Q} of equation (1c). This work considers primarily a simple network architecture corresponding to a line of N units, as seen in Figure 2. A central pump circulates the water through the pipes, and each substation, with index i , has a control valve that it can use to regulate the water flow q_i through their substation locally. Associate with each pipe and valve a hydraulic resistance $a_i(t)$. $a_i(t)$ is constant for pipes and variable for valves, but bounded below by $a_i(t) \geq a_i^{\min}$ corresponding to a completely open valve. Note that the hydraulic resistance of supply pipes will be denoted a_i^{sup} and for return pipes a_i^{ret} .

The following equations dictate the relation between flows and pressure head in the network. For a pipe or valve i ,

$$\Delta p_i = a_i(t)q_i^2, \quad (11)$$

where Δp_i is the pressure difference between the entrance and exit points of the component, caused by pressure losses due to friction [Frederiksen and Werner,




Figure 2. Simple network structure with only one heat source. Here equation (13) has already been used to calculate the flows in the supply and return pipes as a function of the substation flows q_i . Loop $N - 1$ is illustrated with blue arrows.

2013; Wang et al., 2017]. For a pump j ,

$$\Delta p_j = c_{1,j} q_j^2 + c_{2,j} r_j(t) + c_{3,j} r_j(t)^2 \quad (12)$$

where $c_{1,j}$, $c_{2,j}$ and $c_{3,j}$ are pump parameters that denote the characteristics of a specific pump, and $r_j(t) \leq 1$ is the pump frequency ratio indicating the capacity at which the pump is operating at.

Two laws apply to the flows and pressures in the network [Wang et al., 2017]:

- (i) The sum of directed flows entering a node is 0, so that the volume of water in a specific node does not change.
- (ii) Traversing a loop of pipes in the network results in a zero net change in pressure.

(i) can be expressed as

$$F \mathbf{q}(t) = 0, \quad (13)$$

where F is the incidence matrix of the network. The incidence matrix defines how all the pipes in the grid are connected to nodes (connection points) in the network and is defined as

$$F_{ij} = \begin{cases} 1, & \text{pipe } j \text{ leads to node } i \\ -1, & \text{pipe } j \text{ leads away from node } i \\ 0, & \text{pipe } j \text{ is not connected to node } i \end{cases} \quad (14)$$

When applied to the network in Figure 2, we see that the flows in the supply-and-return pipes can be expressed as sums of the substation flows q_i . The second constraint (ii) can be expressed as

$$\sum_{i \in \mathcal{L}_l} \Delta p_i = 0, \quad (15)$$

where \mathcal{L}_l denotes the l th loop in the network, and Δp_i is the pressure difference along each edge that constitutes that loop [Sulzer Pumps, 2010; Wang et al., 2017].

In Figure 2 we can identify N loops. Loop l starts in the central pump, goes through the supply pipes with resistances a_i^{sup} , then through the valve of substation l with resistance $a_l(t)$, and then back through the return pipes with resistances a_i^{ret} . The net pressure difference along this loop is then

$$\begin{aligned} & q_l(t)^2 a_l(t) + \sum_{j=1}^l \left(a_j^{\text{sup}} + a_j^{\text{ret}} \right) \left(\sum_{i=j}^N q_i(t) \right)^2 \\ &= c_1 \left(\sum_{i=1}^N q_i(t) \right)^2 + c_2 r(t) + c_3 r(t)^2, \end{aligned} \quad (16)$$

where the left expression is the pressure losses in the pipes and the right expression is the pressure head generated by the pump. There are N constraints on this form, one for each loop l for $l = 1 \dots N$. Any flow q that satisfies the inequality

$$\begin{aligned} & q_l(t)^2 a_l^{\min} + \sum_{j=1}^l \left(a_j^{\text{sup}} + a_j^{\text{ret}} \right) \left(\sum_{i=j}^N q_i(t) \right)^2 \\ & \leq c_1 \left(\sum_{i=1}^N q_i(t) \right)^2 + c_2 + c_3 \end{aligned} \quad (17)$$

can also be made to satisfy the equation (16), by choosing a significantly large $r(t)$ and a sufficiently large $a_l(t)$. Therefore any flow that satisfies (17) can be actuated with sufficiently high pumping power and local regulation of the valves and any flow \mathbf{q} that satisfies all of the N equations on the form (17) is feasible, i.e. $\mathbf{q} \in \mathcal{Q}$. Note that the equation (17) is convex in \mathbf{q} (assuming c_1 negative). Therefore the set \mathcal{Q} , as the union of N convex sets on the form (17) and the constraints $\mathbf{q} \geq 0$, is also convex.

$$\mathcal{Q} = \{ \mathbf{q} \mid \mathbf{q} \geq 0, \mathbf{q}^\top M_l \mathbf{q} - c_2 - c_3 \leq 0, l = 1 \dots N \} \quad (18)$$

$$M_l = D^\top A D + E_{l,l} a_l^{\min} \quad (19)$$

Here D is an N by N upper triangular matrix of ones. A is an N by N diagonal matrix with entries $A_{1,1} = a_1^{\text{sup}} + a_1^{\text{ret}} - c_1$, $A_{i,i} = a_i^{\text{sup}} + a_i^{\text{ret}}$ for $2 \leq i \leq l$ and 0 otherwise. $E_{l,l}$ is an N by N matrix of all zeros, except for the element l,l , which is a one. This formulation makes M_l an N by N positive semidefinite matrix since A and $E_{l,l}$ have only positive diagonal entries, and thus \mathcal{Q} is a union of quadratic and linear constraints, making it a convex set.

REMARK 3

The convexity of the set \mathcal{Q} is connected to this specific grid structure, as the direction of the flow in this network is obvious. Indeed a meshed network is not guaranteed to enjoy this convexity of \mathcal{Q} , making optimization over the constraints on \mathbf{q} harder to handle.

REMARK 4

Further operational constraints could also be introduced to restrict \mathbf{q} . For instance, too large flows may cause damage to pipes, or generate noise. An upper flow limitation could easily be added.

4. Control Strategies

This section investigates two potential control strategies, the *traditional* strategy and the *load coordination* strategy. An *optimal* baseline comparison is also introduced. The traditional architecture is where units are not connected through any sort of IT communication, and are simply attempting to maintain their own indoor temperature. In the load coordination architecture, the units calculate their desired loads locally through the exact same method as the traditional architecture, but these loads are then processed in a central computation and altered if they are not feasible. The optimal baseline is an upper bound on performance given the cost defined for the system. In this baseline it is assumed that a central unit has access to a perfect model of the entire system, as well as a posteriori measurements of the disturbance.

REMARK 5

Night set-back is an additional part of control strategies common in for instance Southern Europe. [Frederiksen and Werner, 2013] However, this work considers primarily the Northern European situation where this practice is less common and therefore it will not be considered.

4.1 Traditional Architecture

In traditional DH systems there is no IT communication between units in the network. Each unit will greedily evaluate their own desired flow q_i and actuate it through their control valve. The central pump then ensures that the pressure difference between supply and return pipes in the network is high enough to allow these control valves to actuate any desired flow. Traditionally, the control for each individual building has been done through the following control loop: A temperature curve is calibrated for the building, where a reference temperature $T_{hs,i}^r$ is set for the water circulating in the heating system, $T_{hs,i}$. A controller then tracks this reference through the control signal q_i , i.e. the flow through the substation heat exchanger primary side. This is actuated through altering the control valve opening $a_i(t)$. In this work we assume a simple proportional controller with gain k_i

$$T_{hs,i}^r = \alpha_{0,i} + \alpha_{1,i}T_{ext} \quad (20)$$

$$\tilde{q}_i = k_i(T_{hs,i}^r - T_{hs,i}) \quad (21)$$

Here $\alpha_{0,i}$ and $\alpha_{1,i}$ are calibration parameters for the temperature curve. \tilde{q}_i is the desired flow. When the distribution system is at maximum capacity, the differential pressure at unit i may be too low, and in that case the actual flow q_i will be lower

than \tilde{q}_i . The tuning of the parameters would be done by hand by a technician, based on experience and knowledge of suitable parameters for similar buildings. When a unit is not constrained in the flow $q_i(t)$, the unit should be able to reject the influence of outdoor temperatures such that a stationary outdoor temperature should not cause a stationary deviation in indoor temperature. When investigating the model of each building (7), (8) and (21), we can find that this is fulfilled when

$$1 + R_{hs,i}\beta_i k_i + R_{ext,i}\beta_i k_i \alpha_{1,i} = 0, \quad (22)$$

$$\frac{1}{1 - \alpha_{1,i}} \alpha_{0,i} = T_{c,i}. \quad (23)$$

The details of these relations are covered in Appendix 1. Parameters chosen in this way yield that the building will be able to reject the influence of outdoor temperature and maintain indoor temperature at comfort level. For simulation purposes, the parameters were chosen as

$$k_i = \frac{T_c}{\alpha_{0,i} R_{ext,i} - T_{c,i} R_{hs,i} - T_{c,i} R_{ext,i}} \quad (24)$$

$$\alpha_{1,i} = -\frac{1 + k_i R_{hs,i}}{k_i R_{ext,i}}. \quad (25)$$

$\alpha_{0,i}$ is simply chosen large enough that the denominator of (24) does not become negative.

REMARK 6

In practice, the actuator in the building substation is the control valve, and current implementations of control systems may use this actuator directly to control the secondary side heating system temperature. In this case, the flow q becomes an output of the system rather than an input. This problem is readily overcome through standard cascade control. In this setup, the flow q will be the input that dictates the temperature of heating system water. This flow level will be the set-point for a secondary control loop where the valve position is used to actuate the desired flow. This adds the complexity of including the measurement of the flow into the control process. [Skogestad and Postlethwaite, 2005]

4.2 Load Coordination Architecture

The main contribution of this work is the proposition of the following control strategy: Each unit calculates their desired flow \tilde{q}_i as per the traditional strategy of Section 4.1, equations (20) and (21). However, a central device ensures feasibility and fairness by providing each unit with an adjustment δ_i so that the actuated flow will be $q_i = \tilde{q}_i + \delta_i$. In terms of IT communication and computational complexity, this method would be found between the traditional architecture and other optimization-based approaches. Depending on how δ_i is calculated, the central unit does not need access to internal building measurements, only their desired flow \tilde{q}_i . The explicit

models of building dynamics i.e. equations (7) and (8) are not needed in the central computation. Instead only the tuning parameters of the controllers can be utilized. The tuning for the controllers in each building can be done in a distributed fashion, so that a technician working on one individual unit does not affect the control of the whole system.

The aim of the coordination is that the temperature deviations in each building should be distributed more fairly than in the non-coordinated traditional case. In Appendix 1, we show that given

- the models of the buildings presented in Section 3.2, equations (7), (8) and (21)
- and the local unit controllers from Section 4.1, equations (22) and (23),

then given a constant temperature disturbance, each unit will converge to the following stationary indoor temperature deviation from comfort e_i :

$$e_i = \frac{1}{k_i(1 - \alpha_{1,i})} \delta_i \quad (26)$$

While this stationary deviation fails to capture the time dynamics of the system, it is still a valuable metric. Should the system be subject to a constant outdoor temperature lower than the system is able to reject due to flow constraints, then the indoor temperature deviations will align with this distribution. This motivates the following coordination strategy:

Define the parameters γ_i :

$$\gamma_i = \frac{1}{k_i(1 - \alpha_{1,i})} \quad (27)$$

The interpretation of this parameter is a weight provided to each building, indicating how much the deviation δ_i will affect them. Units with large controller gain parameters (k_i and $\alpha_{1,i}$) will not be as impacted by the deviation term. The coordination then wants to minimize the weighted indoor temperature deviations, which can be formulated as the following optimization problem:

$$\underset{\delta}{\text{minimize}} \quad \max_i |\lambda_i \gamma_i \delta_i| \quad (28)$$

$$\text{subject to} \quad \tilde{\mathbf{q}} - \delta \in \mathcal{Q} \quad (29)$$

\mathcal{Q} is a union of quadratic constraints, and the objective function can be reformulated as a linear program. Therefore this becomes a quadratic program where the number of constraints and decision variables grows linearly with the number of connected units, making the problem readily solvable by standard quadratic program solvers. The actual cost to minimize is J_∞ (6). This is a simplified problem where instead the central coordinator minimizes the weighted stationary temperature that would arise

from the coordination terms δ_i . The weights λ_i are design parameters that could be used to capture the quality of service requirements of specific units. For instance a hospital with harsh climate requirements may have a larger λ_i than for example a residential building. In this work the influence of λ_i will not be investigated, and thus we will from now on assume $\lambda_i = 1$.

REMARK 7

Note that according to the current assumptions of individual unit controllers, this central coordination can be designed without explicit knowledge of the building parameters $R_{\text{ext},i}$, $R_{\text{hs},i}$, $C_{\text{in},i}$ or $C_{\text{hs},i}$. The modelling effort is left to each individual unit in the form of controller tuning.

4.3 Optimal Baseline

While we are interested in comparing the load coordination strategy to the traditional strategy, it is also interesting to see what the upper limit of optimality is. We consider the following problem

$$\min_{\mathbf{q}(t_k)} \sum_{k=0}^K \max_i (|e_i(t_k)t_s|) \quad (30a)$$

$$\text{subject to } \mathbf{T}(t_{k+1}) = A\mathbf{T}(t_k) + B_q\mathbf{q}(t_k) + B_{\text{ext}}T_{\text{ext}}(t_k), \quad (30b)$$

$$\mathbf{q}(t_k) \in \mathcal{Q}. \quad (30c)$$

$$\mathbf{T}(t_0) = \mathbf{T}_0. \quad (30d)$$

which can directly be solved by optimization solvers, as the problem is convex. The problem implies minimizing the cost J_∞ of equation (6), subject to the dynamical constraints of the system. For larger networks and longer time-horizons, it will no longer be feasible to solve the entire problem at once as we have done here without adding computational power.

It should be clarified that this optimal baseline as explored in this paper is only presented as a point of reference for comparison with the other methods. In reality it would be completely unfeasible to have exact knowledge of all system parameters, system states, and knowledge of future disturbances. This comparison serves to give an indication about how much possible improvement a given strategy could theoretically have, given our current cost-evaluation.

REMARK 8

It should be noted that this is distinct from online optimization-and-prediction based strategies such as MPC. Such methods rely on online measurements and predictions of disturbances and state evolutions. The optimal strategy in this work is an a posteriori optimization given full knowledge of disturbances and system models.

5. Simulation and Results

The first part of this section details the setup for the simulation experiments, followed by a part detailing the results.

5.1 Simulation Description

This work was simulated in Matlab, with optimization performed using Yalmip [Löfberg, 2004] with a Mosek optimizer. A network of $N = 25$ buildings, consisting of state space models as per Section 3.2 was generated randomly. Controller parameters k_i , $\alpha_{0,i}$ and $\alpha_{1,i}$ were generated for each building in accordance with Section 4.1. Random parameters were generated for pipes connecting these buildings in a line as per Figure 2, as well as parameters that describe the limits of customer substations. The random generation of parameters was done by setting a nominal value for parameters based on parameters from similar models in other works, and then uniformly generating the parameters in a range from these nominal values. The resulting parameters are listed in Table 2

The distribution pump curve (12) was generated the following way: The pump is dimensioned to handle a peak load that occurs at -15°C outdoor temperature. For each building connected to the grid, the flow required to keep the unit at comfort temperature given an outdoor temperature of -15°C , denoted q_i^{peak} was calculated, given equations (7) and (8). Using these flows in the left side of (16) with $a_l(t) = a_l^{\text{max}}$, the pressure generated by the pump p^{peak} can be calculated.

$$p^{\text{peak}} = \max_l \left(q_l^{\text{peak}^2} a_l^{\text{min}} + \sum_{j=1}^l \left(a_j^{\text{sup}} + a_j^{\text{ret}} \right) \left(\sum_{i=j}^N q_i^{\text{peak}} \right)^2 \right) \quad (31)$$

It is then assumed that at this peak flow rate, the pump is running at full capacity, $r(t) = 1$. The parameters c_i are then found by solving the equation

$$c_1 \left(\sum_{i=1}^N q_i^{\text{peak}} \right)^2 + c_2 + c_3 = p^{\text{peak}} \quad (32)$$

such that c_i are proportional to the corresponding parameters in other literature [Wang et al., 2017].

The system was then simulated subject to an outdoor temperature curve generated from real data. The data was gathered from [Swedish Meteorological & Hydrological Institute, n.d.], from a region in Sweden, chosen to represent a time period of drastically dropping temperature. The readings are hourly measurements and were therefore linearly interpolated to 15 minute intervals in the simulation. The resulting temperature curve is visible in Figure 3. The simulation was done for each of the above listed architectures.

5.2 Results

Figures 4a, 4b and 4c show the evolution of indoor temperatures using the traditional strategy, load coordination strategy and optimal baseline respectively. Recall

Table 2. Model parameters used for simulation.

Index	a^{ext}	a^{sup}	a^{min}	R_{hs}	R_{ext}	C_{hs}	C_{in}	T_c	α_0	α_1	k
Unit	mPa/(kg/s) ²	mPa/(kg/s) ²	mPa/(kg/s) ²	°C/KW	°C/KW	kJ/°C	kJ/°C	°C	°C	-	kg/s°C
1	8.22	8.22	160	275	272	22.6	1110	20	48.2	-1.41	9.15
2	3.14	3.14	64.7	281	301	22.4	1020	20	46.4	-1.32	8.59
3	5.18	5.18	286	225	299	18.9	1050	20	42.1	-1.1	9.54
4	4.28	4.28	297	282	264	21.6	1270	20	49.6	-1.48	9.17
5	7.78	7.78	179	262	261	16.7	1180	20	48.1	-1.4	9.58
6	5.39	5.39	179	223	288	22.1	1110	20	42.6	-1.13	9.79
7	8.49	8.49	139	236	321	15.3	1400	20	41.6	-1.08	8.97
8	3.77	3.77	285	255	277	17.8	1250	20	46.2	-1.31	9.4
9	4.3	4.3	148	285	294	15.5	1230	20	47.2	-1.36	8.63
10	3.54	3.54	80.7	285	268	16	1450	20	49.6	-1.48	9.03
11	3.47	3.47	254	227	306	23.2	1070	20	41.8	-1.09	9.37
12	8.23	8.23	153	286	270	21.9	1350	20	49.4	-1.47	8.99
13	6.35	6.35	114	285	288	18.2	1350	20	47.7	-1.39	8.72
14	6.16	6.16	157	251	302	24.5	1130	20	43.9	-1.2	9.04
15	3.53	3.53	76.8	274	316	15.3	1240	20	44.8	-1.24	8.48
16	8.12	8.12	86	226	321	19.4	946	20	40.9	-1.05	9.14
17	6.62	6.62	296	246	291	18.8	932	20	44.3	-1.21	9.3
18	4.87	4.87	300	273	262	22.7	1220	20	49.8	-1.49	9.19
19	5.92	5.92	201	273	263	23	1370	20	48.9	-1.45	9.33
20	5.2	5.2	67.3	285	271	16.9	1460	20	49.3	-1.46	9
21	3.08	3.08	113	263	313	19.9	978	20	44.2	-1.21	8.68
22	4.15	4.15	143	219	270	19.5	1240	20	43.4	-1.17	10.2
23	3.39	3.39	265	277	311	21.5	1180	20	45.4	-1.27	8.51
24	3.78	3.78	55.8	283	270	22.1	907	20	49.2	-1.46	9.05
25	4.15	4.15	63	265	319	22.5	1100	20	43.9	-1.2	8.56






Figure 3. Outdoor temperature curve used for simulation.


from the problem formulation of Section 3.1 that no unit should experience heavy temperature deviations from the comfort temperature of 20°C. The clear distinction between the strategies is that using the traditional architecture results in a few units deviating greatly from their desired indoor temperature. Using the load coordination strategy, the units are much more aligned, leading to all units experiencing deviations but on a much lower magnitude. Finally, in the optimal baseline the results are even better. The units hardly deviate at all from their desired temperatures, and



(a) Indoor temperature, traditional strategy.



(b) Indoor temperature, load coordination strategy.



(c) Indoor temperature, optimal baseline.

Figure 4. Indoor temperatures registered during the simulation.

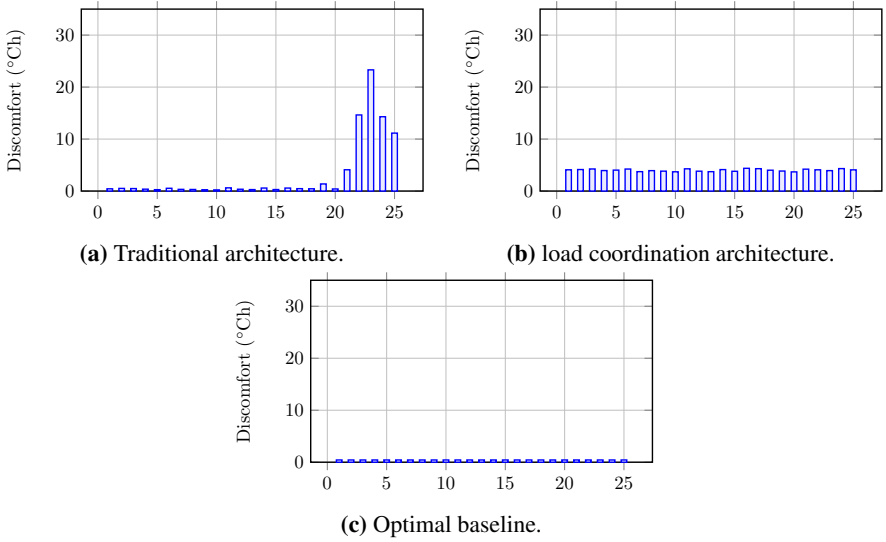


Figure 5. Discomfort experienced by each individual building, indexed 1-25 by their distance from the central distribution pump where 25 is the furthest.

temperatures are deviating equally between all units. In this baseline the units are also pre-heated before the severe drop in temperature, which is not incorporated in the other strategies as they do not include any predictive behavior. The plots of Figures 4 give a hint of what the effect of the different strategies are. However they are also supported by Figure 5. Here the discomfort metric of equation (3) are shown, evaluated on each strategy and unit. Figure 5a shows the inequality generated by the traditional strategy, as units located further from the heat source experience higher discomfort. Meanwhile, Figure 5b shows a much more equal distribution of discomfort. Lastly, Figure 5c shows that there is still a discrepancy between the coordinated strategy and the theoretical lower bound on discomfort. Figure 6 shows the different discomfort metrics of (4), (5) and (6) evaluated through each coordination strategy, corresponding to J_1 , J_2 and J_∞ respectively. We see that the sum of discomfort experienced in units, corresponding to J_1 , is actually improved using traditional architecture than the load coordination architecture. This is quite reasonable, since providing higher flow to units further down the network incurs a higher pressure loss. Thus the total flow provided in the traditional strategy is higher. However, when measured through J_2 and J_∞ , the load coordination strategy outperforms the traditional strategy. This is because the worst-case experience for any unit is much lower with this setup. The optimal baseline shows that there is still potential improvements to be made.




Figure 6. Discomfort metrics defined in equations (4), (5) and (6) (J_1 , J_2 and J_∞ respectively) evaluated through each coordination strategy.

6. Summary

This section concludes the work with some final remarks, followed by potential future outlooks.

6.1 Conclusions

In this work, we investigated the influence of two different architectures for coordinating the flows in a line-structured district heating network. It was shown that utilizing traditional control strategies in each unit can be augmented with a coordination mechanism which reduces the worst-case discomfort experienced by any unit under peak load conditions, at the cost of increasing the mean discomfort, see Figure 6. This coordination can be achieved without explicit models or temperature readings accessed by the central unit. This proof of concept shows how augmenting future district heating systems with smarter controllers can increase the systems' robustness to peak load conditions. The design requirements for future district heating grids can therefore be lowered, allowing for lower grid temperature without as much additional grid capacity in terms of extended piping and pumping power.

However, further improvements can be made to the control strategy when utilizing an optimization-based architecture that allows utilizing information on temperature forecasts to pre-heat units ahead of peak loads. This requires even further complexity, where the central computation unit would have access to individual unit measurements, unit building parameters, and accurate weather forecasts.

The fact that the coordination strategy does not rely on building temperature measurements, and that controllers can be tuned individually for units without affecting the tuning of other units, makes the strategy scalable to growing networks as well as a more privacy-compliant option than a full optimization-based scheme.

6.2 Future Work

The proposed coordination strategy currently does not include the intelligent behavior of the optimal baseline, where the unit indoor temperatures can be utilized for pre-heating before load peaks, often referred to as peak-shaving and valley-filling. The main interest here would be to see if the architecture could maintain the autonomy of unit controllers, while simultaneously including predictive behavior based on an outdoor temperature forecast.

Both the optimal strategy and the proposed coordination strategy currently rely on understanding the set \mathcal{Q} that describes the set of possible flows. This may in practice be harder to estimate than proposed in this work, as specific and accurate parameters for all network parameters may not be known, or degrade and change over time. Therefore it would be interesting to see how these methods hold to uncertainties in network models, as well as data driven methods for estimating the parameters that dictate \mathcal{Q} . While the building model parameters are technically not necessary in the controller coordination, it is reasonable to believe that building controllers will not be as perfectly tuned as proposed in this work. Therefore a study should be conducted to investigate the sensitivity to poorly tuned individual building controllers.

To further simplify the tuning of individual unit controllers, it is likely that more sophisticated unit controllers should be utilized. For instance, a simple PI-controller would allow the elimination of stationary errors when tracking the reference heating system temperature. Therefore including more advanced individual controllers in the analysis would be a valuable extension.

7. Acknowledgements

This work is funded by the European Research Council (ERC) under the European Union's Horizon 2020 research and innovation program under grant agreement No 834142 (ScalableControl).

APPENDIX A - Individual Unit Controller Tuning

This appendix presents the motivating equations behind the tuning of traditional unit controllers and the choice of weights γ_i used in the load coordinating architecture.

To investigate the effects of the control parameters k_i , $\alpha_{0,i}$, $\alpha_{1,i}$ as well as the coordination signal δ_i , we combine equations (7), (8), (9) and (21). This grants the following system description of the unit:

$$\begin{pmatrix} \dot{T}_{\text{in},i} \\ \dot{T}_{\text{hs},i} \end{pmatrix} = A \begin{pmatrix} T_{\text{in},i} \\ T_{\text{hs},i} \end{pmatrix} + B \begin{pmatrix} T_{\text{ext}} \\ \delta_i \end{pmatrix} + C \quad (\text{A.1})$$

where

$$A = \begin{pmatrix} -\frac{1}{R_{\text{ext},i}} - \frac{1}{R_{\text{hs},i}} & \frac{1}{R_{\text{hs},i}} \\ \frac{1}{R_{\text{hs},i}} & -\frac{1}{R_{\text{hs},i}} - C_w(T_{\text{sup}} - T_{\text{ret}})k_i \end{pmatrix} \quad (\text{A.2})$$

$$B = \begin{pmatrix} \frac{1}{R_{\text{ext},i}} & 0 \\ C_w(T_{\text{sup}} - T_{\text{ret}})k_i\alpha_{1,i} & C_w(T_{\text{sup}} - T_{\text{ret}}) \end{pmatrix} \quad (\text{A.3})$$

and

$$C = \begin{pmatrix} 0 \\ C_w(T_{\text{sup}} - T_{\text{ret}})k_i\alpha_{0,i} \end{pmatrix} \quad (\text{A.4})$$

Note that these matrices A and B are not the same matrices as in equation (10). A feasible target for the design of the control parameters k_i , $\alpha_{0,i}$ and $\alpha_{1,i}$ is that when there is no coordination signal δ_i , the building should, given a constant outdoor temperature T_{ext}^0 , be able to reach a given comfort temperature $T_{c,i}$ indoors. We therefore investigate the stationary case where $\dot{T}_{\text{in},i} = \dot{T}_{\text{hs},i} = 0$, $T_{\text{ext}} = T_{\text{ext}}^0$ and $\delta_i = \delta_i^0$. We can find the resulting indoor and heating system temperatures as

$$\begin{pmatrix} T_{\text{in},i}^0 \\ T_{\text{hs},i}^0 \end{pmatrix} = -A^{-1}B \begin{pmatrix} T_{\text{ext}}^0 \\ \delta_i^0 \end{pmatrix} - A^{-1}C. \quad (\text{A.5})$$

Introducing $\beta_i = C_w(T_{\text{sup}} - T_{\text{ret}})$ for brevity, this yields the following stationary indoor temperature:

$$T_{\text{in},i}^0 = \frac{1 + R_{\text{hs},i}\beta_i k_i + R_{\text{ext},i}\beta_i k_i \alpha_{1,i}}{1 + R_{\text{ext},i}\beta_i k_i + R_{\text{hs},i}\beta_i k_i} T_{\text{ext}}^0 \quad (\text{A.6})$$

$$+ \frac{R_{\text{ext},i}\beta_i}{1 + R_{\text{ext},i}\beta_i k_i + R_{\text{hs},i}\beta_i k_i} \delta^0 \quad (\text{A.7})$$

$$+ \frac{R_{\text{ext},i}\beta_i k_i}{1 + R_{\text{ext},i}\beta_i k_i + R_{\text{hs},i}\beta_i k_i} \alpha_{0,i} \quad (\text{A.8})$$

The temperature deviation caused by the external temperature is captured in the term (A.6). To ensure that the outdoor temperature does not cause systematic temperature deviations, the controller gains will have to be chosen so that

$$1 + R_{\text{hs},i}\beta_i k_i + R_{\text{ext},i}\beta_i k_i \alpha_{1,i} = 0 \quad (\text{A.9})$$

Substituting equation (A.9) into the terms (A.6), (A.7) and (A.8), we receive the following resulting indoor temperature:

$$T_{\text{in},i}^0 = \frac{1}{k_i(1 - \alpha_{1,i})} \delta^0 \quad (\text{A.10})$$

$$+ \frac{1}{1 - \alpha_{1,i}} \alpha_{0,i} \quad (\text{A.11})$$

From here, we see that a suitable choice of $\alpha_{0,i}$ is so that the relation

$$\frac{1}{1 - \alpha_{1,i}} \alpha_{0,i} = T_{c,i} \quad (\text{A.12})$$

is fulfilled, i.e. given no coordination term δ_i , the unit should experience comfort temperature.

The remaining deviation caused by the coordination term δ_i is demonstrated in equation (A.10), motivating the weights chosen in section 4.2.

References

- Agner, F., P. Kergus, R. Pates, and A. Rantzer (2022). “Combating district heating bottlenecks using load control”. *Smart Energy* **6**. ISSN: 2666-9552. DOI: 10.1016/j.segy.2022.100067.
- Bacher, P. and H. Madsen (2011). “Identifying suitable models for the heat dynamics of buildings”. en. *Energy and Buildings* **43**:7, pp. 1511–1522. ISSN: 03787788. DOI: 10.1016/j.enbuild.2011.02.005.
- Benonysson, A., B. Bøhm, and H. Ravn (1995). “Operational optimization in a district heating system”. en. *Energy Conversion and Management* **36**, pp. 297–314.
- Bhattacharya, S., C. V., A. V., and K. K. (2019). “Demand response for thermal fairness in district heating networks”. en. *IEEE Transactions on Sustainable Energy* **10**:2, pp. 865–875. ISSN: 1949-3029, 1949-3037. DOI: 10.1109/TSTE.2018.2852629.
- Brange, L., K. Sernhed, and M. Thern (2019). “Decision-making process for addressing bottleneck problems in district heating networks”. en. *International Journal of Sustainable Energy Planning and Management* **20**, pp. 37–50. DOI: <https://doi.org/10.5278/ijsepm.2020.19.4>.
- Frederiksen, S. and S. Werner (2013). *District heating and cooling*. Studentlitteratur.
- Guelpa, E. and V. Verda (2021). “Demand response and other demand side management techniques for district heating: A review”. en. *Energy* **219**. ISSN: 03605442. DOI: 10.1016/j.energy.2020.119440.
- Larsen, H., H. Pálsson, B. Bøhm, and H. Ravn (2002). “Aggregated dynamic simulation model of district heating networks”. en. *Energy Conversion and Management* **43**:8, pp. 995–1019. ISSN: 01968904. DOI: 10.1016/S0196-8904(01)00093-0.
- Löfberg, J. (2004). “Yalmip : a toolbox for modeling and optimization in matlab”. In: *In Proceedings of the CACSD Conference*. Taipei, Taiwan.

- Lund, H., S. Werner, R. Wiltshire, S. Svendsen, J. Thorsen, F. Hvelplund, and B. Mathiesen (2014). “4th generation district heating (4GDH)”. *Energy* **68**, pp. 1–11. ISSN: 03605442. DOI: 10.1016/j.energy.2014.02.089.
- Lund, H., P. A. Østergaard, M. Chang, S. Werner, S. Svendsen, P. Sorknæs, J. E. Thorsen, F. Hvelplund, B. O. G. Mortensen, B. V. Mathiesen, C. Bojesen, N. Duic, X. Zhang, and B. Möller (2018). “The status of 4th generation district heating: research and results”. *Energy* **164**, pp. 147–159. ISSN: 0360-5442. DOI: <https://doi.org/10.1016/j.energy.2018.08.206>.
- Lund, R., D. Østergaard, X. Yang, and B. Mathiesen (2017). “Comparison of low-temperature district heating concepts in a long-term energy system perspective”. *International Journal of Sustainable Energy Planning and Management* **12**, pp. 5–18. DOI: <https://doi.org/10.5278/ijsepm.2017.12.2>.
- Mathiesen, B., H. Lund, D. Connolly, H. Wenzel, P. Østergaard, B. Möller, S. Nielsen, I. Ridjan, P. Karnøe, K. Sperling, and F. Hvelplund (2015). “Smart energy systems for coherent 100% renewable energy and transport solutions”. *Applied Energy* **145**, pp. 139–154. ISSN: 0306-2619. DOI: <https://doi.org/10.1016/j.apenergy.2015.01.075>.
- Saletti, C., A. Gambarotta, and M. Morini (2020). “Development, analysis and application of a predictive controller to a small-scale district heating system”. en. *Applied Thermal Engineering* **165**. ISSN: 13594311. DOI: 10.1016/j.applthermaleng.2019.114558.
- Sarbu, I., M. Mirza, and E. Crasmareanu (2019). “A review of modelling and optimisation techniques for district heating systems”. *International Journal of Energy Research*, pp. 6572–6598. ISSN: 0363-907X, 1099-114X. DOI: 10.1002/er.4600.
- Skogestad, S. and I. Postlethwaite (2005). *Multivariable Feedback Control: Analysis and Design*. 2nd. John Wiley & Sons Ltd, Chichester, p. 423.
- Sulzer Pumps (2010). *Centrifugal pump handbook*. Ed. by Sulzer Pumps. 3rd. Butterworth-Heinemann, Oxford, pp. 27–68. ISBN: 978-0-7506-8612-9. DOI: <https://doi.org/10.1016/B978-0-7506-8612-9.00002-4>.
- Swedish Meteorological & Hydrological Institute (n.d.). <https://www.smhi.se/en/services/open-data/search-smhi-s-open-data-1.81004>.
- Vandermeulen, A., B. van der Heijde, and L. Helsen (2018). “Controlling district heating and cooling networks to unlock flexibility: A review”. en. *Energy* **151**, pp. 103–115. ISSN: 03605442. DOI: 10.1016/j.energy.2018.03.034.
- Wang, Y., S. You, H. Zhang, W. Zheng, X. Zheng, and Q. Miao (2017). “Hydraulic performance optimization of meshed district heating network with multiple heat sources”. *Energy* **126**, pp. 603–621. DOI: 10.1016/j.energy.2017.03.044.

Paper II

Anti-windup Coordination Strategy Around a Fair Equilibrium in Resource Sharing Networks

Felix Agner Pauline Kergus Anders Rantzer
Sophie Tarbouriech Luca Zaccarian

Abstract

We coordinate interconnected agents where the control input of each agent is limited by the control input of others. In that sense, the systems have to share a limited resource over a network. Such problems can arise in different areas and it is here motivated by a district heating example. When the shared resource is insufficient for the combined need of all systems, the resource will have to be shared in an optimal fashion. In this scenario, we want the systems to automatically converge to an optimal equilibrium. The contribution of this paper is the proposal of a control architecture where each separate system is controlled by a local PI controller. The controllers are then coordinated through a global rank-one anti-windup signal. It is shown that the equilibrium of the proposed closed-loop system minimizes the infinity-norm of stationary state deviations. A proof of linear-domain passivity is given, and a numerical example highlights the benefits of the proposed method with respect to the state-of-the-art.

©2023 IEEE. Reprinted, with permission, from F. Agner et al. (2023). “Anti-Windup Coordination Strategy Around a Fair Equilibrium in Resource Sharing Networks”. *IEEE Control Systems Letters* **7**, pp. 2521–2526. ISSN: 2475-1456. DOI: 10.1109/LCSYS.2023.3287252.

1. Introduction

In this paper we consider the problem of asymptotically coordinating a large number of agents that share a central, limited resource towards an optimal equilibrium. Such problems arise in many applications, e.g. optimal power flow [DallAnese and Simonetto, 2018], the TCP protocol [Low et al., 2002; Kelly, 2003] and buffer networks [Bauso et al., 2013; Blanchini et al., 2019; Blanchini et al., 2022]. We consider the motivating example of a district heating network where unfair situations can arise in peak load conditions; buildings close to heat sources stay warm but peripheral buildings become cold. Coordinating central buildings to reduce their heat load in these scenarios would yield a more fair heat distribution [Agner et al., 2022]. For a more detailed view on district heating systems and challenges in district heating control, see e.g. [Frederiksen and Werner, 2013; Vandermeulen et al., 2018]. We consider a representation of such systems given by a linear system with saturating control:

$$\dot{x} = -x + B \text{sat}(u) + w. \quad (1)$$

Here $x \in (x_1, \dots, x_n) \in \mathbb{R}^n$ represents deviations from reference levels for the agents, $w \in \mathbb{R}^n$ is a constant disturbance acting on the system, $u \in \mathbb{R}^n$ represents the control actions of the agents and $B \in \mathbb{R}^{n \times n}$ represents the interconnection among the agents. The saturation function $\text{sat}(\cdot)$ represents the limited nature of the resource in the system. A more detailed description of the system will be given later. Problems of this form are addressed in [Bauso et al., 2013], which shows that feedback control on the form $u = -B^T x$ asymptotically minimizes the cost $x^T x + v^T v$ for (1) where $v = \text{sat}(u)$. Furthermore [Blanchini et al., 2022] designs a controller that asymptotically minimizes varying norms of u in the non-saturated formulation of (1). We extend the asymptotically optimal control design of these previous authors in three ways. First, we consider minimizing the cost function $\|x\|_\infty = \max_i |x_i|$, associated with worst-case fairness. In the district heating example, this objective captures the deviation in the coldest building, which for the specific application is more important than minimizing u . Secondly, we approach scalability of the control strategy in another fashion. Indeed, [Bauso et al., 2013; Blanchini et al., 2022] approach scalability by considering systems where B has a sparse structure, so that with $u = -B^T x$, each agent acts on a few measurements. We consider another scalable control approach of rank-one coordination as utilized in [Madjidian et al., 2017; Lidström and Rantzer, 2016]. In this scenario, the signals from all the agents are combined into one scalar value, and then redistributed to the agents. The advantage of rank-one schemes is that they allow for implementations that scale well and maintain privacy among the agents, as well as scalability, even when B is not sparse. We do this for a specific set of systems (1) where B is an M -matrix, a property that was not previously exploited in this context. Finally, we propose a control law where the results rely only on the structure of B , thus making the implementation robust to modeling errors. To solve the problem under consideration we propose a controller where each agent maintains a local proportional-integral-controller, and coordina-

tion is performed through a global rank-one anti-windup correction. Anti-windup techniques have a long-standing tradition of effective use in combination with integral controllers to improve performance[Galeani et al., 2009]. However, recent results show a strong connection between anti-windup schemes and optimization [Hauswirth et al., 2020a; Hauswirth et al., 2020b], opening the possibility of considering anti-windup loops for optimal equilibrium coordination, in line with what we propose here.

In this article we provide several contributions: We propose the aforementioned controller for driving the system to an optimal equilibrium. We show that under certain conditions on the disturbance w , such an equilibrium exists, is unique, and is in fact uniquely optimal. Analytical proofs of convergence are left outside the scope of this work as we have not yet been able to demonstrate them. However, we provide sufficient conditions for stability in the linear domain $\text{sat}(u) = u$. This leads us to formulate a stability conjecture, subject of future work. A numerical experiment is included to show the effectiveness of the proposed method.

The paper is organized as follows. The system and problem under consideration are formally introduced in Section II, along with the proposed control scheme. The existence of an equilibrium for this system is considered in Section III, and the optimality of the system equilibrium is treated in Section IV. Section V introduces a conjecture on the convergence properties of the proposed closed loop, based on a proof of linear-domain stability under suitable conditions on the PI gains. A numerical example is shown in Section VI. Finally conclusions and future work are covered in Section VII.

Notation: If A is a matrix then denote A_i to be row i of A and $A_{i,j}$ be the element of A at row i and column j . Let $\mathbf{1}$ be a column vector of all ones with dimensions taken from context, and thus $\mathbf{1}\mathbf{1}^T$ is a matrix of all 1's. Denote $\text{sat}(\cdot)$ to be the saturation function $\text{sat}(u) = \max(\min(u, 1), -1)$ and denote the dead-zone function $\text{dz}(u) = u - \text{sat}(u)$. With a slight abuse of notation, the dead-zone and saturation functions applied to vectors operate element-wise. Let the superscript x^0 denote the state x in an equilibrium point, and the superscript x^* to be the value of x which solves an optimization problem. Let the infinite norm $\|\cdot\|_\infty$ of a vector v be defined as the maximum magnitude element $\max_i |v_i|$.

2. System Description and Problem Formulation

2.1 System Description

Consider system (1) where $x_i \in \mathbb{R}$ is the state of each agent $i = 1, \dots, n$ and $u_i \in \mathbb{R}$ is the controller output of each agent. $B \in \mathbb{R}^{n \times n}$ is an M -matrix [Horn and Johnson, 1991; Rantzer and Valcher, 2018]. Such matrices have non-positive off-diagonal entries, thus capturing the fact that each agent negatively impacts the others. Thus, if agent i increases its control input it receives more resources and the other agents receive less resources. Denote $M = B^{-1}$. M -matrices have non-negative inverses

in general and we will assume that M is strictly positive, thus $M_i \mathbf{1} > 0$ for all i . Input w denotes a constant, unknown disturbance affecting the system. In practice the disturbance does not need to be constant, but sufficiently slowly varying. Let us also introduce the index k as a maximizing argument of the following expression

$$k \in \mathcal{K} = \arg \max_i \left| \frac{dz(M_i w)}{M_i \mathbf{1}} \right|, \quad (2)$$

which is, in general, nonunique and characterizes the agent that is most affected by the disturbance w .

2.2 Problem Formulation

We address here the unfair allocation of resources when the agents try to reject a constant disturbance that is too large to drive the system to the origin in view of input saturation. We consider a notion of fairness as described in [Mazumdar et al., 1991] where "*no individual can improve its performance without affecting at least one user adversely*" with regards to the deviations x_i , which we formally describe below.

DEFINITION 1

An equilibrium pair (x^0, u^0) is fair if there is no other equilibrium pair (x^\dagger, u^\dagger) where $\|x^\dagger\|_\infty < \|x^0\|_\infty$, or $\|x^\dagger\|_\infty = \|x^0\|_\infty$ and $|x_i^\dagger| < |x_i^0|$ for some i .

We therefore consider the following problem formulation.

PROBLEM 1

Design a feedback controller driving system (1) from any suitable initial condition to an equilibrium pair (x^0, u^0) , such that $x^ = x^0$ and $v^* = \text{sat}(u^0)$ solves the optimization problem*

$$\underset{x, v}{\text{minimize}} \quad \|x\|_\infty \quad (3a)$$

$$\text{subject to} \quad -x + Bv + w = 0, \quad (3b)$$

$$-1 \leq v \leq 1 \quad (3c)$$

uniquely.

As $\|x_0\|_\infty$ is minimized uniquely, this equilibrium must be fair by Definition 1.

2.3 Proposed Feedback Controller

We propose an individual PI controller for each agent. Each controller has an integral state z_i , and strictly positive gains p_i (proportional gain) and r_i (integral gain). These gains can be tuned locally by the agents. We then introduce a scalar communication signal exchanged among the agents, resulting in a rank-one anti-windup

correction term such that each controller adds the sum of all agents' dead-zones to their integrator input. The full closed-loop system can be written as

$$\dot{x} = -x + B\text{sat}(u) + w \quad (4a)$$

$$\dot{z} = x + \beta \mathbf{1}\mathbf{1}^T \text{dz}(u) \quad (4b)$$

$$u = -Px - Rz, \quad (4c)$$

where P and R are diagonal, positive matrices gathering the controller gains p_i , r_i , $\mathbf{1}\mathbf{1}^T \text{dz}(u)$ is the rank-one anti-windup signal and β is a positive, scalar anti-windup gain. One advantage of the proposed structure is that, under normal circumstances, each PI controller is completely disconnected from the other ones and acts based on local information only. If saturation occurs, the central signal activated and a fairness-oriented coupling emerges from the anti-windup term. Another advantage of the architecture (4) is that when coupling occurs, the coupling signal is merely the sum of the dead-zones for each agent which can be computed efficiently. The summation hides the individual signals, so that when this central signal is redistributed to the agents, each agent does not know the dead-zone values for any of the other individual agents. As such, this global signal lends itself well to scalable and privacy-compliant implementations.

3. Closed-Loop Equilibria

In this section we characterize the equilibria of the proposed closed-loop system (4). From (4), any equilibrium (x^0, z^0) solves the equations

$$0 = -x^0 + B\text{sat}(u^0) + w \quad (5a)$$

$$0 = x^0 + \beta \mathbf{1}\mathbf{1}^T \text{dz}(u^0) \quad (5b)$$

$$u^0 = -Px^0 - Rz^0. \quad (5c)$$

It is not trivial to show whether a solution to (5) exists. This section studies conditions for the existence and uniqueness of such solutions. Note that it is sufficient to study pairs (x^0, u^0) satisfying (5a) and (5b), because the positive definiteness of R implies its invertibility. Hence for any such state-control pair (x^0, u^0) satisfying (5a) and (5b), z^0 can be uniquely determined from (5c).

3.1 Existence of an Equilibrium Point

Recall that $M = B^{-1}$. We provide below a necessary and sufficient condition for (4) to admit an equilibrium.

LEMMA 1

The closed-loop system (5) admits an equilibrium point (x^0, z^0) , if and only if

$$\max_i \frac{M_i w - 1}{M_i \mathbf{1}} \leq \min_j \frac{M_j w + 1}{M_j \mathbf{1}}. \quad (6)$$

Proof. Let us begin with showing that (6) is necessary for the existence of an equilibrium point. (5a) and (5b) can be combined to

$$Mw + \text{sat}(u^0) = -\beta M \mathbf{1}^T \text{dz}(u^0). \quad (7)$$

Thus

$$\frac{M_i w + \text{sat}(u_i^0)}{M_i \mathbf{1}} = -\beta \mathbf{1}^T \text{dz}(u^0), \quad \forall i = 1, \dots, n. \quad (8)$$

If (6) does not hold, then there exist i and j such that

$$\frac{M_i w - 1}{M_i \mathbf{1}} > \frac{M_j w + 1}{M_j \mathbf{1}}. \quad (9)$$

However, (8) implies that

$$\frac{M_i w + \text{sat}(u_i^0)}{M_i \mathbf{1}} = \frac{M_j w + \text{sat}(u_j^0)}{M_j \mathbf{1}}. \quad (10)$$

As $\text{sat}(u_i^0) \geq -1$ and $\text{sat}(u_j^0) \leq 1$, (9) and (10) cannot simultaneously hold, which establishes a contradiction thus proving that there is no equilibrium. This proves that (6) is necessary for the existence of an equilibrium. For the sufficiency, first recall the definition of k , given by (2). Then consider the candidate equilibrium x^0 , u^0 given by

$$x^0 = \mathbf{1} \frac{\text{dz}(M_k w)}{M_k \mathbf{1}} \quad (11a)$$

$$u_k^0 = -\text{sat}(M_k w) - \frac{\text{dz}(M_k w)}{\beta M_k \mathbf{1}} \quad (11b)$$

$$u_i^0 = -M_i w + \frac{M_i \mathbf{1}}{M_k \mathbf{1}} \text{dz}(M_k w), \quad \forall i \neq k. \quad (11c)$$

We show below that when (6) holds, the candidate equilibrium (11) solves (5). Consider 3 scenarios. (i): $\text{dz}(M_k w) = 0$, (ii): $\text{dz}(M_k w) > 0$ and (iii): $\text{dz}(M_k w) < 0$. In scenario (i), $x^0 = 0$, and $u^0 = -Mw$. As $\text{dz}(M_k w) = 0$ in this scenario, $\text{dz}(M_i w) = 0$ for all i . Otherwise (2) would not be maximized by k . This implies that $\text{sat}(u^0) = u^0 = -Mw$ and $\text{dz}(u^0) = 0$. It is thus easy to verify that (5) holds. In scenario (ii), note that the left side of (6) is maximized by index k and can be reformulated as

$$\frac{\text{dz}(M_k w)}{M_k \mathbf{1}} \leq \frac{M_i w + 1}{M_i \mathbf{1}} \quad \forall i = 1, \dots, n. \quad (12)$$

Returning to the candidate equilibrium and (11c) for $i \neq k$,

$$u_i^0 = -M_i w + \frac{M_i \mathbf{1}}{M_k \mathbf{1}} \text{dz}(M_k w) \leq 1 \quad (13)$$

where the inequality is derived from (12). Thus $u_i^0 \leq 1$. In addition,

$$\begin{aligned}
 u_i^0 &= -M_i w + \frac{M_i \mathbf{1}}{M_k \mathbf{1}} dz(M_k w) \\
 &= -\text{sat}(M_i w) - dz(M_i w) + \frac{M_i \mathbf{1}}{M_k \mathbf{1}} dz(M_k w) \\
 &= -\text{sat}(M_i w) + M_i \mathbf{1} \left(\frac{dz(M_k w)}{M_k \mathbf{1}} - \frac{dz(M_i w)}{M_i \mathbf{1}} \right) \\
 &\geq -1,
 \end{aligned} \tag{14}$$

where the last inequality holds because k maximizes (2). This means that $-1 \leq u_i^0 \leq 1$ for all $i \neq k$. Thus

$$\text{sat}(u_i^0) = u_i^0 = -M_i w + \frac{M_i \mathbf{1}}{M_k \mathbf{1}} dz(M_k w), \quad \forall i \neq k, \tag{15}$$

and

$$dz(u_i) = 0, \quad \forall i \neq k. \tag{16}$$

For index k , (11b) provides

$$\text{sat}(u_k^0) = -\text{sat}(M_k w) = -M_k w + dz(M_k w) \tag{17}$$

and

$$dz(u_k^0) = -\frac{dz(M_k w)}{\beta M_k \mathbf{1}}. \tag{18}$$

Combining (15), (16), (17) and (18) yields

$$\text{sat}(u^0) = -M w + M \mathbf{1} \frac{dz(M_k w)}{M_k \mathbf{1}} \tag{19}$$

and

$$\mathbf{1}^T dz(u^0) = -\frac{dz(M_k w)}{\beta M_k \mathbf{1}}. \tag{20}$$

which allows us to easily verify that (x^0, u^0) from (11) solves (5) in scenario (ii). An equal argument can be made for scenario (iii), which we omit for brevity. This shows that given any scenario for $dz(M_k w)$, the candidate equilibrium (11) is valid when (6) holds. Thus (6) is both necessary and sufficient for the existence of an equilibrium. \square

To interpret (6), note that it is satisfied when all entries w_i are similar to each other. For instance, $w = s \mathbf{1}$ for any scalar s trivially satisfies the condition. This makes it a sensible assumption when the disturbance w affects all agents in a similar way. This is for instance the case in the district heating example, where the outdoor temperature is likely to be quite similar for all the buildings located in a specific area. To simplify our follow-up definitions, we will assume that (6) holds with a strict inequality, as formulated below.

ASSUMPTION 1

The disturbance w satisfies (6) strictly, namely

$$\max_i \frac{M_i w - 1}{M_i \mathbf{1}} < \min_j \frac{M_j w + 1}{M_j \mathbf{1}}. \quad (21)$$

We assume the strict inequality to enforce uniqueness of the equilibrium, which is studied in the next section.

3.2 Uniqueness of the Equilibrium

Lemma 1 shows that under Assumption 1, there is an equilibrium for the closed-loop system. We study here conditions for this equilibrium to be unique. To enforce the uniqueness of this equilibrium, we assume the following.

ASSUMPTION 2

Either $\mathrm{d}z(Mw) = 0$, or the maximizing argument k given by (2) is unique.

If k is non-unique, an arbitrarily small perturbation of B or w would make it so. In practical applications, w is expected to vary slowly over time. This makes it unlikely that k would be non-unique for an extended period of time, but may also cause k to shift between agents. The analysis of such scenarios requires to study the transient behavior of the system, which is outside the scope of the paper, but will be the subject of future work.

LEMMA 2

If Assumptions 1 and 2 hold, then (11) is the unique equilibrium of the closed-loop system (4).

Proof. Recall from the proof of Lemma 1 that for any equilibrium inducing input u^0 , identity (7) must hold. Now denote

$$t = \beta \mathbf{1}^T \mathrm{d}z(u^0), \quad (22)$$

which allows (7) to be rewritten as

$$\mathrm{sat}(u_i^0) = -M_i w - M_i \mathbf{1} t, \quad \forall i = 1, \dots, n. \quad (23)$$

Note that if $t > 0$, there must exist an $i \in \{1, \dots, n\}$ such that $\mathrm{sat}(u_i^0) = 1$. Similarly, if $t < 0$, there exists an $i \in \{1, \dots, n\}$ such that $\mathrm{sat}(u_i^0) = -1$. Also note that (21) implies that there cannot exist i and j such that $M_i w \geq 1$ and $M_j w \leq -1$. This in turn implies that either $\mathrm{d}z(u^0) \geq 0$ or $\mathrm{d}z(u^0) \leq 0$, where the inequality should be understood componentwise. Now, recalling that k in (2) is unique by assumption, consider 3 scenarios; (i): $\mathrm{d}z(M_k w) = 0$, (ii): $\mathrm{d}z(M_k w) > 0$ and (iii): $\mathrm{d}z(M_k w) < 0$. In scenario (i), we see that $\mathrm{d}z(M_i w) = 0$ for all $i = 1, \dots, n$, as otherwise $|\mathrm{d}z(M_i w)| > 0$ for some i , implying that (2) would be maximized by this i . Thus $|\mathrm{d}z(M_i w)| \leq 1$ for all

i. Through (23), we prove next that this implies $t = 0$. Indeed, assume by an absurd argument that $t > 0$. Then (23) yields $\text{sat}(u_i^0) = -M_i w - M_i \mathbf{1} t < -M_i w \leq 1$ for all $i = 1, \dots, n$. But if $\text{sat}(u_i^0) < 1$ for all i , then $\text{dz}(u_i^0) \leq 0$ for all i and thus t cannot be positive. A parallel contradiction can be built for $t < 0$. Thus we conclude that $t = 0$. In turn, $t = 0$ implies that $\text{dz}(u^0) = 0$, because (22) shows that t is the sum of the entries of $\text{dz}(u^0) = 0$, multiplied by the positive scalar β . As the entries of $\text{dz}(u^0)$ are either all positive or all negative, t can only be 0 if all of the entries of $\text{dz}(u^0)$ are 0. This uniquely fixes $u^0 = -Mw$, which is the same as the candidate solution (11). This in turn uniquely fixes x^0 through (5b), and uniquely fixes z^0 through (5c).

In scenario (ii), (23) implies $t \leq -\frac{\text{dz}(M_k w)}{M_k \mathbf{1}}$, because otherwise $\text{sat}(u_k^0) < -1$. If $t = \frac{\text{dz}(M_k w)}{M_k \mathbf{1}}$ then $\text{sat}(u_k^0) = -1$. For $i \neq k$,

$$\begin{aligned} \text{sat}(u_i^0) &= -M_i w + \frac{M_i \mathbf{1}}{M_k \mathbf{1}} \text{dz}(M_k w) \\ &= -\text{sat}(M_i w) - \text{dz}(M_i w) + \frac{M_i \mathbf{1}}{M_k \mathbf{1}} \text{dz}(M_k w) \\ &= -\text{sat}(M_i w) + M_i \mathbf{1} \left(\frac{\text{dz}(M_k w)}{M_k \mathbf{1}} - \frac{\text{dz}(M_i w)}{M_i \mathbf{1}} \right) \\ &> -1. \end{aligned} \tag{24}$$

The last inequality holds because k uniquely maximizes (2) and $M_i \mathbf{1} > 0$ due to non-negativity and invertibility of M . Thus we conclude that $t = -\frac{\text{dz}(M_k w)}{M_k \mathbf{1}}$, because otherwise $\text{sat}(u_i^0) > -1$ for all i , contradicting the fact that t is negative. For this scenario (ii), (21) can be written as

$$\frac{\text{dz}(M_k w)}{M_k \mathbf{1}} < \frac{M_i w + 1}{M_i \mathbf{1}}, \quad \forall i \neq k, \tag{25}$$

Which can be combined with $t = -\frac{\text{dz}(M_k w)}{M_k \mathbf{1}}$ to show that, for $i \neq k$,

$$\text{sat}(u_i^0) = -M_i w + \frac{M_i \mathbf{1}}{M_k \mathbf{1}} \text{dz}(M_k w) < 1. \tag{26}$$

Inequality (26) implies $|\text{sat}(u_i^0)| < 1$ for all $i \neq k$, and thus $\text{dz}(u_i^0) = 0$ for all $i \neq k$. This implies

$$t = \beta \mathbf{1}^T \text{dz}(u^0) = \beta \text{dz}(u_k^0) \tag{27}$$

and thus

$$\text{dz}(u_k^0) = -\frac{\text{dz}(M_k w)}{\beta M_k \mathbf{1}}. \tag{28}$$

Equations (26) and (28) uniquely determine u^0 , and, together with (11), x^0 and z^0 are uniquely determined. For scenario (iii), a symmetric argument can be followed, which is omitted for brevity, thus completing the proof. \square

4. Optimality

We proved in the previous section that under Assumptions 1 and 2, the proposed closed-loop system has a unique equilibrium, given by (11). In this section, we will prove that this equilibrium is also the unique, optimal solution to (3).

THEOREM 1

If Assumptions 1 and 2 hold, then $x^ = x^0$ and $v^* = \text{sat}(u^0)$ is the unique solution to (3) where x^0 and u^0 are given by (11).*

Proof. Lemma 2 proves that under Assumptions 1 and 2, (x^0, u^0) is a state-input equilibrium pair. This means that $(x^*, v^*) = (x^0, \text{sat}(u^0))$ satisfies the constraints (3b) and (3c) and is therefore feasible. What remains is only to show that it is not only feasible but also uniquely optimal. Consider for establishing a contradiction that there exists $\xi \neq 0$, such that $x^\dagger = x^* + \xi$, along with $v^\dagger = v^* + M\xi$ is also feasible and provides a lower or equal cost than (x^*, v^*) . An equivalent rewriting of (3) using ξ is

$$\underset{\xi}{\text{minimize}} \quad \left\| \mathbf{1} \frac{dz(M_k w)}{M_k \mathbf{1}} + \xi \right\|_\infty \quad (29a)$$

$$\text{subject to} \quad \left\| M(\xi - w + \mathbf{1} \frac{dz(M_k w)}{M_k \mathbf{1}}) \right\|_\infty \leq 1. \quad (29b)$$

First note that if $dz(M_k w) = 0$, then $\xi = 0$ is trivially optimal as any $\xi \neq 0$ would yield a higher cost and thus not be an optimizer. Then consider the case where $dz(M_k w) > 0$. For ξ to provide a lower or equal cost, it must hold that $\xi_i \leq 0$ for all i . However, analyzing constraint (29b) for index k yields

$$|-M_k w + dz(M_k w) + M_k \xi| \leq 1. \quad (30)$$

Since we are focusing on the case $dz(M_k w) > 0$, (30) reduces to

$$|-1 + M_k \xi| \leq 1. \quad (31)$$

Since $\xi \leq 0$, inequality (31) can only hold for $x_i = 0$ as M_k has strictly positive entries. Therefore $\xi = 0$ is uniquely optimal when $dz(M_k w) > 0$. A parallel reasoning can be performed for the case $dz(M_k w) < 0$. Thus (x^*, v^*) is the optimal solution to (3). \square

5. Stability Properties

The results of Sections III and IV established that under Assumptions 1 and 2, the unique equilibrium of the closed-loop system (4) solves the optimization problem (3). In this section we formulate the following conjecture regarding its stability properties.

CONJECTURE 1

Under Assumptions 1 and 2, if $p_i > r_i$ for all $i = 1, \dots, n$, then the proposed controller (4) globally solves Problem 1.

Conjecture 1, subject to its proof, would provide strong properties for the proposed control law, granting stability and optimality for a large family of systems subject to a simple control tuning constraint. The proof however is non-trivial and requires results for saturated systems operating deeply in the saturated regime, which is why it is left outside the scope of this work. Our confidence in Conjecture 1 arises from numerous simulations of randomized systems. In addition, the specific choice of $p_i > r_i$ provides notions of stability for our problem through the following lemma.

LEMMA 3

Assume that $p_i > r_i$ for all i and $w \in \mathcal{L}_2$. Then system (4) is asymptotically stable in the region of linearity where $\text{sat}(u) = u$.

Proof. Define $y = -Bu$. When $\text{sat}(u) = u$ and thus $\text{dz}(u) = 0$, the closed loop system (4) can be reformulated in the frequency domain as

$$sX = -X - Y + W \quad (32a)$$

$$sU = (P - R)X + PY - PW. \quad (32b)$$

These equations are fully diagonal, and can for each agent i be combined to form

$$U_i = \frac{r_i + p_i s}{s(s+1)}(Y_i - W_i) = G_i(s)(Y_i - W_i). \quad (33)$$

This feedback interconnection is represented in Figure 1. When $p_i > r_i$, the transfer function (33) is positive real, making it a passive component [Khalil, 2002]. In addition, due to B being an M-matrix, we know that there exists a positive, diagonal matrix D such that $-DB - B^T D < 0$. This means that the combined upper block of Figure 1 is strictly passive. The multiplication by the positive, diagonal matrix D^{-1}

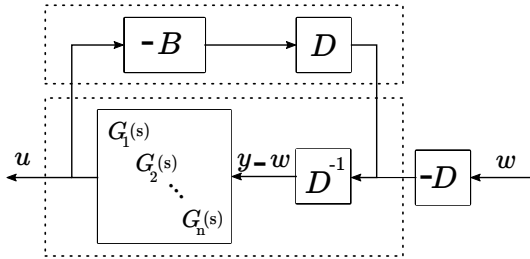


Figure 1. Block diagram, showing the interconnection used in the proof of Lemma 3. (©2023 IEEE)

does not affect the passivity properties of $G_1(s) \dots G_n(s)$. The feedback interconnection between the strictly passive upper block and the passive lower block means that for any $w \in \mathcal{L}_2$, we have $u \in \mathcal{L}_2$ [Desoer and Vidyasagar, 2009]. This means that $\lim_{t \rightarrow \infty} w(t) = 0$, $\lim_{t \rightarrow \infty} u(t) = 0$ and thus clearly $\lim_{t \rightarrow \infty} x(t) = 0$ by (4a). \square

To prove or refute Conjecture 1 in future work, we believe that these passivity properties may be a useful tool. While it can be shown that the condition $P > R$ is conservative, we have also found examples of sufficiently large integral gains causing instability, thereby suggesting that our conjecture is reasonable.

6. Numerical Example

To demonstrate the usefulness of the proposed controller, we investigate $n = 250$ agents interconnected through the matrix $B = D(1.2nI - \mathbf{1}\mathbf{1}^T)$ where $D = \text{diag}(d_1, d_2, \dots, d_n)$ and d_1, \dots, d_n are distributed at even intervals between 0.5 and 1.5. $w(t) = \mathbf{1} \frac{n \sin(t/2\pi)}{2}$. We compare three strategies: First the *coordinated* strategy, consisting in the controller proposed in this paper using the gains $p_i = 1$, $r_i = 1.5$ for all i and $\beta = 1$. Secondly the *uncoordinated* strategy, namely the same PI-controllers as those of the coordinated case, only equipped with a local anti-windup action: $\dot{z}_i = x_i + \beta dz(u_i)$. Finally the *linear saturated decentralized* (lsd) controller $u = -B^T x$ as proposed by [Bauso et al., 2013]. The systems are simulated using the DifferentialEquations toolbox [Rackauckas and Nie, 2017] in Julia. Figure 2 shows the envelopes of the time series over the simulation. In the coordinated case (green), all of the states x are nearly completely synchronized. Under both the uncoordinated (blue) and the lsd (red) strategy, there is a large discrepancy between the maximum and minimum states. Furthermore, the lsd strategy is optimal with regards to a tradeoff between states x and control action u , and therefore no states are driven to the origin with large disturbances w . Figure 3 shows histograms of




Figure 2. Envelopes of the states x for each strategy. The dashed lines constitute the minimum $\min_i x_i(t)$ and the solid lines the maximum $\max_i x_i(t)$. (©2023 IEEE)




Figure 3. Histogram of maximum absolute deviations $\max_t |x_i(t)|$ experienced under each strategy. From deviation 20 to 80, the red and blue bars overlap. (©2023 IEEE)

the worst magnitude deviations in each strategy. We see that both the uncoordinated (blue) and lsd (red) strategies have several agents with larger deviations than any of the agents in the coordinated case. However, both the uncoordinated and lsd strategies also have many agents with lower deviations than that of the coordinated case.

7. Conclusion

In this paper we have presented a controller for coordinating the control actions of agents that share a central resource. We proved that the only equilibrium of this closed-loop system is optimally fair. This optimality concerns the states x , an important extension of the literature which has mainly focused on properties of the control input u . A conjecture was proposed giving conditions for stability of this optimal equilibrium, motivated by passivity of the closed-loop system in the linear domain.

Subject to the proof of Conjecture 1, the proposed method has many advantages. Each agent could tune the gains of a PI-controller locally while maintaining global guarantees of stability. These guarantees are only dependent on the structure of the system and not the model itself (i.e. the B -matrix does not have to be known, only that it has certain properties). The rank-one communication scheme ensures scalability of the implementation which does not require sparsity of B .

Extensions of the work include exploiting the proven passivity property to prove stability with regards to the optimal equilibrium. Further system structures could be considered, for instance more general A -matrices, output feedback, or non-linear interconnections $B(u)$ which maintain similar properties to the current B -structure. Finally, one can consider analyzing and improving transient performance.

References

- Agner, F., P. Kergus, R. Pates, and A. Rantzer (2022). “Combating district heating bottlenecks using load control”. *Smart Energy* **6**. ISSN: 2666-9552. DOI: <https://doi.org/10.1016/j.segy.2022.100067>.
- Agner, F., P. Kergus, A. Rantzer, S. Tarbouriech, and L. Zaccarian (2023). “Anti-Windup Coordination Strategy Around a Fair Equilibrium in Resource Sharing Networks”. *IEEE Control Systems Letters* **7**, pp. 2521–2526. ISSN: 2475-1456. DOI: 10.1109/LCSYS.2023.3287252.
- Bauso, D., F. Blanchini, L. Giarré, and R. Pesenti (2013). “The linear saturated decentralized strategy for constrained flow control is asymptotically optimal”. *Automatica* **49**:7, pp. 2206–2212. ISSN: 0005-1098. DOI: <https://doi.org/10.1016/j.automatica.2013.03.029>.
- Blanchini, F., D. Casagrande, F. Fabiani, G. Giordano, and R. Pesenti (2019). “Network-decentralised optimisation and control: an explicit saturated solution”. *Automatica* **103**, pp. 379–389. ISSN: 0005-1098. DOI: <https://doi.org/10.1016/j.automatica.2019.02.009>.
- Blanchini, F., C. A. Devia, G. Giordano, R. Pesenti, and F. Rosset (2022). “Fair and sparse solutions in network-decentralized flow control”. *IEEE Control Systems Letters* **6**, pp. 2984–2989. DOI: 10.1109/LCSYS.2022.3181341.
- Dall’Anese, E. and A. Simonetto (2018). “Optimal power flow pursuit”. *IEEE Transactions on Smart Grid* **9**:2, pp. 942–952. DOI: 10.1109/TSG.2016.2571982.
- Desoer, C. A. and M. Vidyasagar (2009). *Feedback systems. input-output properties*. Classics in applied mathematics: 55. Society for Industrial and Applied Mathematics. ISBN: 9780898716702.
- Frederiksen, S. and S. Werner (2013). *District heating and cooling*. Studentlitteratur.
- Galeani, S., S. Tarbouriech, M. Turner, and L. Zaccarian (2009). “A tutorial on modern anti-windup design”. *European Journal of Control* **15**:3, pp. 418–440. ISSN: 0947-3580. DOI: <https://doi.org/10.3166/ejc.15.418-440>.
- Hauswirth, A., F. Dörfler, and A. Teel (2020a). “On the differentiability of projected trajectories and the robust convergence of non-convex anti-windup gradient flows”. *IEEE Control Systems Letters* **4**:3, pp. 620–625. DOI: 10.1109/LCSYS.2020.2988515.
- Hauswirth, A., F. Dörfler, and A. Teel (2020b). “On the robust implementation of projected dynamical systems with anti-windup controllers”. In: *2020 American Control Conference (ACC)*, pp. 1286–1291. DOI: 10.23919/ACC45564.2020.9147378.
- Horn, R. A. and C. R. Johnson (1991). *Topics in Matrix Analysis*. Cambridge University Press. DOI: 10.1017/CB09780511840371.

- Kelly, F. (2003). “Fairness and stability of end-to-end congestion control”. *European Journal of Control* **9**:2, pp. 159–176. ISSN: 0947-3580. DOI: <https://doi.org/10.3166/ejc.9.159-176>.
- Khalil, H. K. (2002). *Nonlinear systems*. Prentice Hall. ISBN: 0130673897.
- Lidström, C. and A. Rantzer (2016). “Optimal H-infinity state feedback for systems with symmetric and Hurwitz state matrix”. eng. In: *American Control Conference (ACC), 2016*. IEEE - Institute of Electrical and Electronics Engineers Inc., 3366–3371. ISBN: 978-1-4673-8683-8. DOI: {10.1109/ACC.2016.7525437}.
- Low, S., F. Paganini, and J. Doyle (2002). “Internet congestion control”. *IEEE Control Systems Magazine* **22**:1, pp. 28–43. DOI: 10.1109/37.980245.
- Madjidian, D., L. Mirkin, and A. Rantzer (2017). “ H_2 Optimal coordination of homogeneous agents subject to limited information exchange”. *IEEE Transactions on Automatic Control* **62**:3, pp. 1424–1430. DOI: 10.1109/TAC.2016.2571782.
- Mazumdar, R., L. Mason, and C. Douligeris (1991). “Fairness in network optimal flow control: optimality of product forms”. *Communications, IEEE Transactions on* **39**, pp. 775–782. DOI: 10.1109/26.87140.
- Rackauckas, C. and Q. Nie (2017). “DifferentialEquations.jl—a performant and feature-rich ecosystem for solving differential equations in julia”. *Journal of Open Research Software* **5**:1.
- Rantzer, A. and M. E. Valcher (2018). “A tutorial on positive systems and large scale control”. In: *2018 IEEE Conference on Decision and Control (CDC)*, pp. 3686–3697. DOI: 10.1109/CDC.2018.8618689.
- Vandermeulen, A., B. van der Heijde, and L. Helsen (2018). “Controlling district heating and cooling networks to unlock flexibility: A review”. en. *Energy* **151**, pp. 103–115. ISSN: 03605442. DOI: 10.1016/j.energy.2018.03.034.

Paper III

Decentralized PI-control and Anti-windup in Resource Sharing Networks

Felix Agner Jonas Hansson Pauline Kergus Anders Rantzer
Sophie Tarbouriech Luca Zaccarian

Abstract

We consider control of multiple stable first-order agents which have a control coupling described by an M-matrix. These agents are subject to incremental sector-bounded input nonlinearities. We show that such plants can be globally asymptotically stabilized to a unique equilibrium using fully decentralized proportional-integral controllers equipped with anti-windup and subject to local tuning rules. In addition, we show that when the nonlinearities correspond to the saturation function, the closed loop asymptotically minimizes a weighted 1-norm of the agents state mismatch. The control strategy is finally compared to other state-of-the-art controllers on a numerical district heating example.

Reprinted, with permission, from F. Agner et al. (2024). “Decentralized PI-control and anti-windup in resource sharing networks”. *European Journal of Control*, p. 101049. ISSN: 0947-3580. DOI: 10.1016/j.ejcon.2024.101049.

1. Introduction

In this paper we consider the control of agents sharing a central distribution system with limited capacity. These are systems where the positive action of one agent negatively impacts others. This type of competitive structure can arise in many domains, for instance internet congestion control [Low et al., 2002; Kelly, 2003] and district heating systems [Agner et al., 2022]. In the district heating scenario, the structure arises because of the hydraulic constraints of the grid. If one agent (building) locally decides to increase their heat demand by opening their control valves, this will lead to higher flow rates and greater frictional pressure losses. These losses make it so that other agents now receive lower flow rates [Agner et al., 2022]. We consider a simple description of such systems:

$$\dot{x} = -Ax + Bf(u) + w. \quad (1)$$

Here each agent i is associated with a state x_i , and these states are gathered in the vector x . The agents are subject to an external disturbance w and interconnected via the matrix B . The nonlinear function $f(\cdot)$ can for instance represent the common phenomenon of input saturation, which motivates this work. A is assumed diagonal. We will more formally describe the plant in Section 2.

In multi-agent systems such as (1), decentralized controllers are desirable. Semi-decentralized control strategies for multi-agent systems subject to input saturation have been considered in the following works. In [Dal Col et al., 2019], each networked agent is equipped with a local controller that receives the control input of its neighbors. In [Ofodile and Turner, 2016], semi-decentralized anti-windup was considered for stable SISO plants that are decentralized in the linear domain, but become coupled during saturation. This is demonstrated on unmanned aerial vehicles. These works focus on stabilization when the disturbance w in plant (1) is energy bounded. In this work we focus instead on the asymptotic properties of plant (1), which become important when w is expected to vary slower than the dynamics of the plant and can be approximated as constant. Previous works considering asymptotic optimality for plants of the form (1) are [Bauso et al., 2013] and [Agner et al., 2023]. In [Agner et al., 2023], it was shown that, when B is an M-matrix, decentralized PI-controllers with a rank-one coordinating anti-windup scheme can minimize the cost $\max_i |x_i|$. In [Bauso et al., 2013], it was shown that the static controller $u = -B^\top x$ asymptotically minimizes the cost $x^\top Ax + v^\top v$ where $v = \text{sat}(u)$. This result also extends to the case when B is not an M-matrix. Both of these control strategies maintain certain scalability properties: With $u = -B^\top x$ [Bauso et al., 2013], any sparsity structure in the B -matrix is maintained and the rank-one coordination scheme of [Agner et al., 2023] admits scalable implementations. However, the most scalable control solution is one that is fully decentralized. In this work, we analyze (1) under a fully decentralized PI (proportional-integral) control strategy. In general, it is non-trivial that decentralized PI-controllers are stabilizing, let alone fulfill any optimality criterion. However, in this work we show not only that

our strategy minimizes asymptotic costs of the form $\sum_{i=1}^n \gamma_i |x_i|$ but also that the resulting equilibrium is globally asymptotically stable under decentralized controller tuning rules.

The paper is organized as follows: Section 2 presents the considered plant and control strategy. Section 3 presents the main results of the paper, namely equilibrium existence and uniqueness, global asymptotic stability, and equilibrium optimality for our considered closed loop. A motivating numerical example consisting in the flow control of a simplified district-heating network is subsequently given in section 4. The proofs of the main results are presented in sections 5, 6, and 7 respectively. Conclusions and future work are covered in section 8.

Notation: v_i denotes element i of vector $v \in \mathbb{R}^n$, A_i denotes row i of matrix $A \in \mathbb{R}^{n \times m}$, and $A_{i,j}$ denotes its (i, j) -th element. A matrix A is strictly diagonally row-dominant if $|A_{i,i}| > \sum_{j \neq i} |A_{i,j}|$ for all i . A is strictly diagonally column-dominant if A^\top , denoting the transpose of A , is strictly diagonally row-dominant. Matrix $B \in \mathbb{R}^{n \times n}$ is called positive stable if all of its eigenvalues have positive real part. We denote $S \succ 0$ ($S \succeq 0$) if $S \in \mathbb{R}^{n \times n}$ is symmetric and positive definite (semi-definite). Similarly, for two symmetric matrices $S_1, S_2 \in \mathbb{R}^{n \times n}$ we denote $S_1 \succ S_2$ ($S_1 \succeq S_2$) if $S_1 - S_2 \succ 0$ ($S_1 - S_2 \succeq 0$). Let the 2-norm of a vector $x \in \mathbb{R}^n$ be given by $\|x\|_2 = (\sum_{i=1}^n x_i^2)^{1/2}$. Let the 1-and-infinity-norms of a vector $x \in \mathbb{R}^n$ be given by $\|x\|_1 = \sum_{i=1}^n |x_i|$ and $\|x\|_\infty = \max_i |x_i|$ respectively. Let the norm $\|A\|_2$ of a matrix A be the induced 2-norm. Let $\mathbf{1} \in \mathbb{R}^n$ be a vector of all ones, where n is taken in context. We say that a function $f: \mathbb{R} \rightarrow \mathbb{R}$ is increasing (non-decreasing) if $y > x$ implies that $f(y) > f(x)$ ($f(y) \geq f(x)$).

2. Problem Data and Proposed Controller

We consider control of plants of the form (1) where vector $x \in \mathbb{R}^n$ gathers the states x_i of each agent, $A \in \mathbb{R}^{n \times n}$, and $w \in \mathbb{R}^n$ is a constant disturbance acting on the plant. $B \in \mathbb{R}^{n \times n}$ couples the control-inputs of the agents. The input nonlinearity $f: \mathbb{R}^n \rightarrow \mathbb{R}^n$ is characterized by the following assumption.

ASSUMPTION 1

$f(x) = [f_1(x_1), f_2(x_2), \dots, f_n(x_n)]^\top$ has components f_i satisfying $f_i(0) = 0$ and incrementally sector-bounded in the sector $[0, 1]$, namely satisfying $0 \leq (f_i(y) - f_i(x)) / (y - x) \leq 1$ for all $x \in \mathbb{R}$, $y \in \mathbb{R}$, $x \neq y$.

Note that Assumption 1 implies that f is non-decreasing and Lipschitz with Lipschitz constant 1. Since $f(0) = 0$, f also enjoys a sector $[0, 1]$ condition.

Stability properties for feedback with incrementally sector-bounded nonlinearities has long been considered in the literature. As far back as [Zames, 1966] it was used for input-output stability analysis. Both [Zhang et al., 2013] and [DeLellis and Bernardo, 2012] consider the type of diagonally partitioned incrementally sector-bounded functions that we consider here, whereas [Giaccagli et al., 2023; Zhang et

al., 2014; Giaccagli et al., 2022] consider a richer class of incremental sector-bound constraints $(f(x) - f(y) - S_1(x - y))^\top (f(x) - f(y) - S_2(x - y)) \leq 0$ for all $x \in \mathbb{R}^n$, $y \in \mathbb{R}^n$. Here S_1 and S_2 are real symmetric matrices with $0 \preceq S_1 \prec S_2$.

We will consider function pairs $f(\cdot)$, $h(\cdot)$ where $f(x) + h(x) = x$. These pairs fulfill the following property, the proof of which is in the appendix.

LEMMA 1

Let $f : \mathbb{R}^n \rightarrow \mathbb{R}^n$ satisfy Assumption 1. Then $h(u) = u - f(u)$ also satisfies Assumption 1.

The considered class of function pairs is well motivated by the common case $f(x) = \text{sat}(x)$ where $\text{sat}(x) = \max(\min(x, \mathbf{1}), -\mathbf{1})$ and $h(x) = \text{dz}(x) = x - \text{sat}(x)$.

We propose controlling the plant (1) with fully decentralized PI controllers having decentralized anti-windup for each agent $i = 1, \dots, n$.

$$\dot{z}_i = x_i + s_i h_i(u_i) \quad (2)$$

$$u_i = -p_i x_i - r_i z_i \quad (3)$$

where z_i is the integral state, u_i is the controller output, $p_i > 0$ and $r_i > 0$ are proportional and integral controller gains respectively, $s_i > 0$ is an anti-windup gain, and $h(u) = u - f(u)$ is an anti-windup signal. Note that while the notation h is not needed (indeed we could equivalently replace $h(u)$ with $u - f(u)$), we will use the pair f, h both to simplify the exposition and to highlight that f is the nonlinearity acting on the plant while h is the nonlinearity acting on the controller. We assume that the closed loop system satisfies the following assumption.

ASSUMPTION 2

A is a diagonal positive definite matrix, B is an M-matrix, and w is a constant disturbance. The controller parameters p_i , r_i , and s_i , for $i = 1, \dots, n$, are all positive.

The M-matrix property which we consider for B has the following standard definition [Horn and Johnson, 1991, p. 113].

DEFINITION 1

A matrix $B \in \mathbb{R}^{n \times n}$ is called an M-matrix if B is positive stable and all off-diagonal elements of B are non-positive.

M-matrices hold certain exploitable properties as listed in Theorem 2.5.3 of [Horn and Johnson, 1991, pp. 114-115]. We summarize the ones we employ in this paper in the following proposition.

PROPOSITION 1

If $B \in \mathbb{R}^{n \times n}$ has only non-positive off-diagonal elements, then the following statements are equivalent:

- (i) B is positive stable, that is, B is an M -matrix.
- (ii) DB is an M -matrix for every positive definite diagonal matrix D .
- (iii) There exists a diagonal positive definite matrix U such that UB and UBU^{-1} are strictly column-diagonally dominant.
- (iv) There exists a diagonal positive definite matrix Q such that $QB + B^\top Q \succ 0$.

3. Main Results

In this section we will cover the main results of this paper. In particular, we will consider the proposed control law (2)–(3) for the plant (1). We will show that this closed loop system admits an equilibrium for any constant disturbance w . We will additionally show that this equilibrium is globally asymptotically stable and enjoys a notion of optimality. We will leave the proofs for Sections 5 to 7.

Let us first consider the existence of an equilibrium, which corresponds to well-posedness of the equations (1)–(3) with $\dot{x} = \dot{z} = 0$.

THEOREM 1—EQUILIBRIUM EXISTENCE AND UNIQUENESS

Let f satisfy Assumption 1 and let Assumption 2 hold. Then for each constant $w \in \mathbb{R}^n$, the closed loop (1)–(3) has a unique equilibrium (x^0, z^0) , inducing input u^0 from (3), which satisfies (1)–(3) with $\dot{x} = \dot{z} = 0$.

In addition to the existence of the unique equilibrium (x^0, z^0) , we can also show that it is globally asymptotically stable under the following assumption on the control parameters.

ASSUMPTION 3

Assume that $a_i p_i > r_i$ and $p_i s_i < 1$ for all i , where a_i are the diagonal elements of A in (1) and p_i , r_i , and s_i are the controller gains in (2)–(3).

THEOREM 2—GLOBAL ASYMPTOTIC STABILITY

Let f satisfy Assumption 1 and let $f(u) + h(u) = u$. Let Assumptions 2 and 3 hold. Then there is a globally asymptotically stable equilibrium for the closed loop (1)–(3).

REMARK 1

The tuning rules of Assumption 3 are fully decentralized. Each agent i can tune their own controller gains to satisfy $r_i < a_i p_i$ and $s_i < 1/p_i$.

Let us now focus on the case where the function pair $f(\cdot)$ and $h(\cdot)$ are given by the pair $\text{sat}(\cdot)$ and $\text{dz}(\cdot)$ respectively, motivated by classical anti-windup for saturating controllers. Let γ_i be positive scalar weights, and consider the problem of minimizing the weighted sum of all state errors $\sum_{i=1}^n \gamma_i |x_i|$. We can define this problem through the optimization problem

$$\underset{x, v}{\text{minimize}} \quad \sum_{i=1}^n \gamma_i |x_i| = \|\Gamma x\|_1 \quad (4a)$$

$$\text{subject to} \quad -Ax + Bv + w = 0, \quad (4b)$$

$$-\mathbf{1} \leq v \leq \mathbf{1}. \quad (4c)$$

where $\Gamma = \text{diag}\{\gamma_1, \dots, \gamma_n\}$. The inequalities (4c) are considered componentwise. This problem can be motivated by a district heating example. Let w be the outdoor temperature, x_i be the deviation from the comfort temperature for each agent i , and let Bv denote the heat provided to the agents, limited by (4c). Then if $\Gamma = I$, this corresponds to minimizing the total discomfort experienced by all agents. One could consider γ_i to be a cost describing the severity of agent i deviating from the comfort temperature, where γ_i would be high for e.g. a hospital. Note that this cost does not capture the notion of *fairness* as considered in [Agner et al., 2023]. For instance, with $\Gamma = I$, $x = [n, 0, \dots, 0]^\top$, and $y = [1, 1, \dots, 1]^\top$ we achieve the same costs $\|\Gamma x\|_1 = \|\Gamma y\|_1$. With the problem (4) defined, the following holds.

THEOREM 3—EQUILIBRIUM OPTIMALITY

Let Assumption 2 hold and let $\Gamma A^{-1}B$ be a strictly diagonally column-dominant M -matrix. Let $f(u) = \text{sat}(u)$ and $h(u) = \text{dz}(u) = u - \text{sat}(u)$. Let (x^0, z^0) be an equilibrium for the closed loop system in (1)–(3), associated with input u^0 . Then $x^* = x^0$ and $v^* = f(u^0)$ solves (4).

REMARK 2

For an arbitrary choice of γ , it does not necessarily hold that $\Gamma A^{-1}B$ is strictly diagonally column-dominant. Thus Theorem 3 cannot be used as a design-method where we first fix the weights γ according to some performance criterion and then calculate a correspondingly optimal controller. However, the set of weights γ_i such that the required condition is satisfied will always be non-empty. We can see this because A is a positive definite diagonal matrix by Assumption 2 and thus $A^{-1}B$ is also an M -matrix by Proposition 1 (ii). Proposition 1 (iii) then shows that a positive definite diagonal matrix Γ such that $\Gamma A^{-1}B$ is strictly diagonally column-dominant must exist. Thus Theorem 3 yields a non-empty set of performance criteria for which the fully decentralized control strategy cannot be outperformed by a more complex control architecture.

4. Numerical Example

This motivating example compares three different control strategies on a simplified, linear model of 10 buildings connected in a district heating grid. The compared strategies are the same as the ones considered in [Agner et al., 2023]. Each building i has identical thermodynamics on the form

$$\dot{x}_i = -\frac{a_i}{C_i}(x_c + x_i - T_{\text{ext}}(t)) + \frac{1}{C_i}\dot{Q}_i(u), \quad (5)$$

where x_i denotes agent i 's indoor temperature deviation from the comfort temperature x_c , C_i is the heat capacity of each building and T_{ext} is the outdoor temperature. \dot{Q}_i is the heat supplied to building i . This heat supply is given by

$$\dot{Q} = B \text{sat}(u), \quad (6)$$

where B represents the network interconnection. The simulation was conducted with $a_i = 0.167 [\text{kW}/\text{C}^\circ]$, $C_i = 2.0 [\text{kWh}/\text{C}^\circ]$, $p_i = 2.5 [1/\text{C}^\circ]$, $r_i = 0.2 [1/\text{C}^\circ\text{h}]$, and $s_i = 2.0 [\text{C}^\circ]$ for all i . The parameters a_i , C_i are chosen close to the values found in [Bacher and Madsen, 2011] which discusses parameter estimation for a single-family building. The Matrix B is selected as $B_{i,i} = 12 \forall i$, $B_{i,j} = -0.15 \min(i, j) \forall i \neq j$ in units $[\text{kW}]$. Matrix B is constructed such that fully opened control valves ($\text{sat}(u) = \mathbf{1}$) gives \dot{Q} representing a reasonable peak heat demand for small houses. In this scenario, \dot{Q}_i is high for buildings with i small (close to the production facility). We simulate the system using the `DifferentialEquations` toolbox in Julia [Rackauckas and Nie, 2017], for an outdoor temperature scenario given by data from the city of Gävle, Sweden in October 2022 during which the temperature periodically drops to almost -20°C . The data is gathered from the Swedish Meteorological and Hydrological Institute (SMHI). This slowly time-varying disturbance also brings insight into how our proposed controller handles a signal w which is not constant. We compare three different controllers and three different cost functions. The first controller is the fully *decentralized* PI-controller considered in this paper. Secondly the *coordinating* controller consists of the same PI-controllers as the decentralized case, but with the coordinating rank-1 anti-windup signal $\dot{z}_i = x_i + \beta \mathbf{1}^\top \text{dz}(u)$ considered in [Agner et al., 2023]. Finally, the *static* controller is given by $u = -B^\top C^{-1}x$ as considered in [Bauso et al., 2013], where C is the diagonal matrix of all heat capacities C_i .

Figure 1 shows the resulting deviations x during the simulations. At around hour 100, the outdoor temperature is critically low. At this time, the buildings do not receive sufficient heat, regardless of the control strategy. Figure 1a shows that with the decentralized strategy, the worst deviations become larger than with the coordinating strategy (Figure 1b). However, not all buildings experience temperature deviations, whereas with the coordinating strategy, all the buildings share the discomfort. Lastly, the static controller has large deviations experienced by many

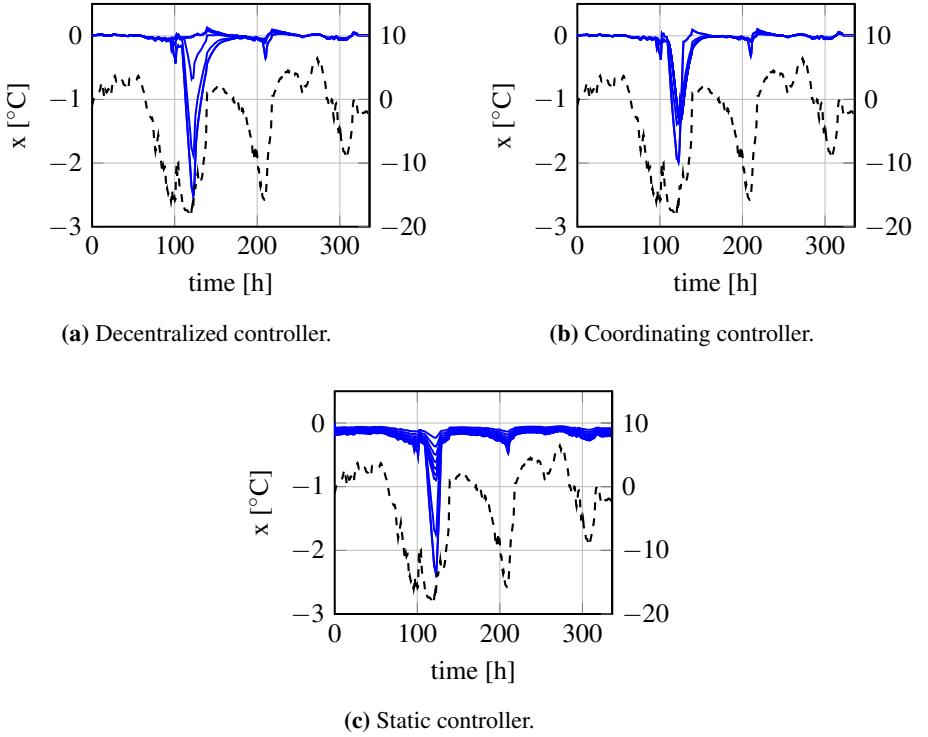


Figure 1. Temperature deviations x (blue, left axis) for each strategy and the outdoor temperature w (black, dotted, right axis). Around hour 100, w becomes critically low and the indoor temperatures drop as the controllers saturate.

buildings. Even when the outdoor temperature is manageable, the static controller has a constant offset from the comfort temperature, highlighting the usefulness of the integral action. We evaluate the performance through the cost functions

$$J_1 = \frac{1}{T} \int_0^T \|x(t)\|_1 dt, \quad (7)$$

$$J_\infty = \frac{1}{T} \int_0^T \|x(t)\|_\infty dt, \quad (8)$$

$$J_2 = \frac{1}{T} \int_0^T x(t)^\top L x(t) + \text{sat}(u(t))^\top \text{sat}(u(t)) dt. \quad (9)$$

where T is the simulation time and L is a diagonal matrix where each element is given by $l_i = \frac{q_i}{C_i}$. The cost J_1 mimics the optimality notion considered in this paper, J_∞ mimics the optimality notion considered in [Agner et al., 2023], and J_2 mimics the optimality considered in [Bauso et al., 2013]. Table 1 shows the resulting evaluations. The coordinating controller gives minimal worst-case deviations J_∞ , but J_1

Table 1. Costs (7)–(9) evaluated over the simulation.

	Decentralized	Coordinating	Static
J_∞	0.17	0.13	0.28
J_1	0.67	0.9	1.96
J_2	3.52	3.52	3.49

is minimized in the decentralized strategy. This result, i.e. that the total discomfort is minimized by decentralized control but the worst-case discomfort is minimized by coordination, is found also in [Agner et al., 2022] where a nonlinear model of the grid hydraulics and a 2-state model of building dynamics is employed. On the weighted cost J_2 , all controllers provide similar performance. The static controller slightly outperforms the other two in this scenario, but it is outperformed in every other measure.

5. Proof of Equilibrium Existence and Uniqueness

We will now prove Theorem 1 through the use of Banach’s fixed-point theorem [Agarwal et al., 2018]. This proof requires the following two lemmas, the proofs of which are found in the appendix.

LEMMA 2

Let $f : \mathbb{R}^n \rightarrow \mathbb{R}^n$ and $h : \mathbb{R}^n \rightarrow \mathbb{R}^n$ where $h(x) = x - f(x)$ satisfy Assumption 1. Then $\tilde{f} : \mathbb{R}^n \rightarrow \mathbb{R}^n$ and $\tilde{h} : \mathbb{R}^n \rightarrow \mathbb{R}^n$ given by $\tilde{f}(x) = f(x + x^0) - f(x^0)$ and $\tilde{h}(x) = h(x + x^0) - h(x^0)$ for some $x^0 \in \mathbb{R}^n$ also satisfy Assumption 1 and $\tilde{h}(x) + \tilde{f}(x) = x$.

LEMMA 3

Let $f : \mathbb{R}^n \rightarrow \mathbb{R}^n$ and $h : \mathbb{R}^n \rightarrow \mathbb{R}^n$ where $h(x) = x - f(x)$ satisfy Assumption 1. Then $\tilde{f} : \mathbb{R}^n \rightarrow \mathbb{R}^n$ and $\tilde{h} : \mathbb{R}^n \rightarrow \mathbb{R}^n$ given by $\tilde{f}(x) = Df(D^{-1}x)$ and $\tilde{h}(x) = Dh(D^{-1}x)$ where D is a diagonal positive definite matrix also satisfy Assumption 1 and $\tilde{h}(x) + \tilde{f}(x) = x$.

Proof of Theorem 1. Denote by S a diagonal matrix gathering the positive anti-windup gains s_i , $i = 1, \dots, n$. We can rearrange (1)–(3) by imposing $\dot{x} = \dot{z} = 0$, which yields

$$0 = h(u^0) + S^{-1}A^{-1}Bf(u^0) + S^{-1}A^{-1}w. \quad (10)$$

If there is a unique u^0 solving (10) then $x^0 = A^{-1}(Bf(u^0) + w)$ and $z^0 = R^{-1}(-Px^0 - u^0)$ are uniquely determined by (1) and (3) respectively, where $R = \text{diag}\{r_1, \dots, r_n\}$ is invertible by Assumption 3. Hence we need only show that there is a unique u^0 solving (10) for the proof to be complete. Let D be a diagonal positive definite matrix such that $DS^{-1}A^{-1}BD^{-1}$ is strictly diagonally column-dominant. Note that such a D always exists by Proposition 1 (iv) because A

and S are diagonal positive definite and B is an M-matrix. Left-multiply (10) by D and insert multiplication by $I = D^{-1}D$ before $f(u^0)$ to obtain

$$0 = Dh(u^0) + DS^{-1}A^{-1}BD^{-1}Df(u^0) + DS^{-1}A^{-1}w. \quad (11)$$

Introduce the change of variables $\hat{B} = DS^{-1}A^{-1}BD^{-1}$, $\zeta = Du^0$, and $\hat{w} = DS^{-1}A^{-1}w$. Then (11) yields

$$0 = Dh(D^{-1}\zeta) + \hat{B}Df(D^{-1}\zeta) + \hat{w}. \quad (12)$$

Here we can use Lemma 3 to replace $f(\zeta)$, $h(\zeta)$ with $\hat{f}(\zeta) = Df(D^{-1}\zeta)$, $\hat{h}(\zeta) = Dh(D^{-1}\zeta)$. Note that $\hat{f}(\cdot)$, $\hat{h}(\cdot)$ satisfy Assumption 1 and $\hat{f}(\zeta) + \hat{h}(\zeta) = \zeta$. Introduce a scalar k satisfying $k > \max(1, 2\max_i \hat{B}_{i,i})$. Divide (12) by $-k$, add ζ to the left-hand side, and $\zeta = \hat{f}(\zeta) + \hat{h}(\zeta)$ to the right-hand side of (12) to obtain

$$\zeta = -\frac{1}{k} \left((1-k)\hat{h}(\zeta) + (\hat{B} - kI)\hat{f}(\zeta) + \hat{w} \right). \quad (13)$$

We define the right-hand side of this expression as $T_w(\zeta)$, defined for a specific w . By showing that T_w is a contractive mapping for any \hat{w} , we can use Banach's fixed point theorem [Agarwal et al., 2018] to show that there is a unique solution $\zeta = T_w(\zeta)$ (and thus a unique $u^0 = D^{-1}\zeta$) for any \hat{w} (and thus any $w = ASD^{-1}\hat{w}$). Consider any $\alpha \in \mathbb{R}^n$, $\beta \in \mathbb{R}^n$. Then

$$T_w(\alpha) - T_w(\beta) = \frac{-1+k}{k} (\hat{h}(\alpha) - \hat{h}(\beta)) + \frac{-\hat{B} + kI}{k} (\hat{f}(\alpha) - \hat{f}(\beta)). \quad (14)$$

Here we use Lemma 2 to introduce $\tilde{h}(\alpha - \beta) = \hat{h}(\alpha) - \hat{h}(\beta)$ and $\tilde{f}(\alpha - \beta) = \hat{f}(\alpha) - \hat{f}(\beta)$. Denote $\Delta = \alpha - \beta$ and $\Delta^+ = T_w(\alpha) - T_w(\beta)$. Then

$$|\Delta_i^+| \leq \frac{k-1}{k} |\tilde{h}_i(\Delta_i)| + \frac{k-\hat{B}_{i,i}}{k} |\tilde{f}_i(\Delta_i)| + \sum_{j \neq i} \frac{|\hat{B}_{i,j}|}{k} |\tilde{f}_j(\Delta_j)|. \quad (15)$$

Therefore

$$\begin{aligned} \|\Delta^+\|_1 &= \sum_{i=1}^n |\Delta_i^+| \\ &\leq \sum_{i=1}^n \left(\frac{k-1}{k} |\tilde{h}_i(\Delta_i)| + \frac{k-\hat{B}_{i,i}}{k} |\tilde{f}_i(\Delta_i)| + \sum_{j \neq i} \frac{|\hat{B}_{j,i}|}{k} |\tilde{f}_j(\Delta_j)| \right). \end{aligned} \quad (16)$$

Due to the diagonal column-dominance of \hat{B} and the definition of k , it holds that $k > \hat{B}_{i,i} > \sum_{j \neq i} |\hat{B}_{j,i}|$. Thus, selecting $\lambda = \frac{k-1}{k} < 1$, $\mu_i = \frac{k-(\hat{B}_{i,i} - \sum_{j \neq i} |\hat{B}_{j,i}|)}{k} < 1$, $\gamma_i =$

$\max(\lambda, \mu_i) < 1$, and $\bar{\gamma} = \max_i \gamma_i < 1$, we obtain

$$\begin{aligned} \|\Delta^+\|_1 &\leq \sum_{i=1}^n \lambda |\tilde{h}_i(\Delta_i)| + \mu_i |\tilde{f}_i(\Delta_i)| \\ &\leq \sum_{i=1}^n \gamma_i (|\tilde{h}_i(\Delta_i)| + |\tilde{f}_i(\Delta_i)|) \\ &\leq \sum_{i=1}^n \bar{\gamma} |\Delta_i| = \bar{\gamma} \|\Delta\|_1. \end{aligned} \quad (17)$$

Note that $|\tilde{h}_i(\Delta_i)| + |\tilde{f}_i(\Delta_i)| = |\Delta_i|$ since $\tilde{f}_i(\Delta_i)$ and $\tilde{h}_i(\Delta_i)$ always have the same sign by Assumption 1, and sum to Δ_i . This proves that T_w is a contraction mapping with respect to the metric $\|\cdot\|_1$. Thus, by Banach's fixed point theorem, for each w and the ensuing $\hat{w} = DS^{-1}A^{-1}w$ there is a unique ζ such that (13) holds, and thus a $u^0 = D^{-1}\zeta$ such that (10) holds, which completes the proof. \square

6. Proof of Global Asymptotic Stability

Given the existence of an equilibrium (x^0, z^0) and the associated input u^0 , consider the change of variables $\tilde{z} = -R(z - z^0)$, $\tilde{u} = u - u^0$, $\tilde{f}(\tilde{u}) = f(u^0 + \tilde{u}) - f(u^0)$, and $\tilde{h}(\tilde{u}) = h(u^0 + \tilde{u}) - h(u^0)$. Due to Lemma 2, $\tilde{f}(\cdot)$, $\tilde{h}(\cdot)$ satisfy Assumption 1, and $\tilde{f}(\tilde{u}) + \tilde{h}(\tilde{u}) = \tilde{u}$. This allows rewriting the (1)–(3) as

$$\begin{bmatrix} \dot{\tilde{z}} \\ \dot{\tilde{u}} \end{bmatrix} = \begin{bmatrix} -RP^{-1} & RP^{-1} \\ A - RP^{-1} & -A + RP^{-1} \end{bmatrix} \begin{bmatrix} \tilde{z} \\ \tilde{u} \end{bmatrix} - \begin{bmatrix} 0 \\ PB \end{bmatrix} \tilde{f}(\tilde{u}) - \begin{bmatrix} RS \\ RS \end{bmatrix} \tilde{h}(\tilde{u}) \quad (18)$$

where P , R , and S are diagonal matrices gathering the controller parameters p_i , r_i , and s_i . Stabilizing this system to $\tilde{z} = \tilde{u} = 0$ is equivalent to stabilizing the original system to the equilibrium $x = x^0$, $z = z^0$, and $u = u^0$. We will therefore now prove Theorem 2 with a Lyapunov-based argument considering system (18).

Proof of Theorem 2. Consider the Lyapunov function candidate

$$\begin{aligned} V(\tilde{z}, \tilde{u}) &= \sum_{i=1}^n \int_0^{\tilde{z}_i} q_i \left(a_i \frac{p_i}{r_i} - 1 \right) (\tilde{f}_i(\zeta) + \varepsilon \zeta) d\zeta \\ &\quad + \sum_{i=1}^n \int_0^{\tilde{u}_i} q_i (\tilde{f}_i(\zeta) + \varepsilon \zeta) d\zeta \end{aligned} \quad (19)$$

where scalars $q_i > 0$ and $\varepsilon > 0$ are parameters to be fixed later. For any such choice of parameters, V is positive definite and radially unbounded because $\tilde{f}_i(\zeta) + \varepsilon \zeta$ is increasing in ζ and zero at zero. Also $a_i \frac{p_i}{r_i} - 1 > 0$ due to Assumption 3. The time

derivative of V along the trajectories of system (18) is given by

$$\dot{V}(\tilde{z}, \tilde{u}) = -(\tilde{f}(\tilde{z}) + \varepsilon \tilde{z} - \tilde{f}(\tilde{u}) - \varepsilon \tilde{u})^\top \tilde{D}(\tilde{z} - \tilde{u}) \quad (20a)$$

$$- (\tilde{f}(\tilde{z}) + \varepsilon \tilde{z})^\top \tilde{D}PS\tilde{h}(\tilde{u}) \quad (20b)$$

$$- (\tilde{f}(\tilde{u}) + \varepsilon \tilde{u})^\top QRS\tilde{h}(\tilde{u}) \quad (20c)$$

$$- (\tilde{f}(\tilde{u}) + \varepsilon \tilde{u})^\top QPB\tilde{f}(\tilde{u}) \quad (20d)$$

where \tilde{D} is a diagonal matrix gathering the positive elements $q_i(a_i - r_i/p_i)$ and Q is a diagonal matrix gathering the positive elements q_i . To simplify this expression, we split it into

$$\dot{V}(\tilde{z}, \tilde{u}) = \dot{V}_1(\tilde{z}, \tilde{u}) + \dot{V}_2(\tilde{z}, \tilde{u}) \quad (21)$$

where $\dot{V}_1(\tilde{z}, \tilde{u})$ corresponds to the terms (20a)–(20b) and $\dot{V}_2(\tilde{z}, \tilde{u})$ corresponds to the terms (20c)–(20d). Since \tilde{D} and $\tilde{D}PS$ are diagonal, \dot{V}_1 can be analyzed for each i individually. $\tilde{f}_i(\zeta_i) + \varepsilon \zeta_i$ is increasing in ζ_i , therefore $\text{sign}(\tilde{f}_i(\tilde{z}_i) + \varepsilon \tilde{z}_i - \tilde{f}_i(\tilde{u}_i) - \varepsilon \tilde{u}_i) = \text{sign}(\tilde{z}_i - \tilde{u}_i)$ and thus (20a) is negative semi-definite. If \tilde{z}_i and \tilde{u}_i have the same sign, (20b) contributes negatively to \dot{V}_1 . If they have opposite signs the contribution is positive, but then (20a) only comprises negative terms as $(\tilde{f}_i(\tilde{z}_i) + \varepsilon \tilde{z}_i - \tilde{f}_i(\tilde{u}_i) - \varepsilon \tilde{u}_i) \tilde{D}_{i,i}(\tilde{z}_i - \tilde{u}_i) = (|\tilde{f}_i(\tilde{z}_i) + \varepsilon \tilde{z}_i| + |\tilde{f}_i(\tilde{u}_i) - \varepsilon \tilde{u}_i|) \tilde{D}_{i,i}(|\tilde{z}_i| + |\tilde{u}_i|)$. Indeed, since $p_i s_i < 1$ from Assumption 3 and $|\tilde{h}_i(\tilde{u}_i)| \leq |\tilde{u}_i|$ from Assumption 1, then (20a) as developed above dominates (20b) which is upper bounded by $|\tilde{f}_i(\tilde{z}_i) + \varepsilon \tilde{z}_i| \tilde{D}_{i,i} |\tilde{h}_i(\tilde{u}_i)|$. Thus \dot{V}_1 is negative semi-definite. We now turn our attention to \dot{V}_2 . Note that \tilde{u} , $\tilde{f}(\tilde{u})$, and $\tilde{h}(\tilde{u})$ elementwise have the same sign and QRS is diagonal, positive definite. Thus

$$\begin{aligned} & (\tilde{f}(\tilde{u}) + \varepsilon \tilde{u})^\top QRS\tilde{h}(\tilde{u}) \\ &= (\tilde{f}(\tilde{u}) + \varepsilon \tilde{f}(\tilde{u}) + \varepsilon \tilde{h}(\tilde{u}))^\top QRS\tilde{h}(\tilde{u}) \\ &= (1 + \varepsilon) \tilde{f}(\tilde{u})^\top QRS\tilde{h}(\tilde{u}) + \varepsilon \tilde{h}(\tilde{u})^\top QRS\tilde{h}(\tilde{u}) \geq \varepsilon \beta \|\tilde{h}(\tilde{u})\|_2^2 \end{aligned} \quad (22)$$

where β is the minimum diagonal element of QRS . Note also that

$$\begin{aligned} & (\tilde{f}(\tilde{u}) + \varepsilon \tilde{u})^\top QPB\tilde{f}(\tilde{u}) = (1 + \varepsilon) \tilde{f}(\tilde{u})^\top QPB\tilde{f}(\tilde{u}) \\ & \quad + \varepsilon \tilde{h}(\tilde{u})^\top QPB\tilde{f}(\tilde{u}). \end{aligned} \quad (23)$$

Fix now the weights q_i in such a way that $QPB + B^\top PQ$ is positive definite. This is possible by Proposition 1 (iv) because B is an M-matrix according to Assumption 2. Therefore $\exists \alpha > 0$ such that $QPB + B^\top PQ \succ 2\alpha I$. Thus the first term of (23) satisfies

$$(1 + \varepsilon) \tilde{f}(\tilde{u})^\top QPB\tilde{f}(\tilde{u}) \geq (1 + \varepsilon) \alpha \|\tilde{f}(\tilde{u})\|_2^2. \quad (24)$$

We also note that the second term in (23) satisfies

$$\varepsilon \tilde{h}(\tilde{u})^\top QPB\tilde{f}(\tilde{u}) \geq -\varepsilon \gamma \|\tilde{f}(\tilde{u})\|_2 \|\tilde{h}(\tilde{u})\|_2 \quad (25)$$

where $\gamma = \|QPB\|_2$. Thus, combining the bounds in (22), (24) and (25) within (20c)–(20d), we obtain

$$\begin{aligned} \dot{V}_2(\tilde{z}, \tilde{u}) &\leq -(1 + \varepsilon)\alpha \|\tilde{f}(\tilde{u})\|_2^2 - \varepsilon\beta \|\tilde{h}(\tilde{u})\|_2^2 + \varepsilon\gamma \|\tilde{f}(\tilde{u})\|_2 \|\tilde{h}(\tilde{u})\|_2 \\ &= \begin{pmatrix} \|\tilde{f}(\tilde{u})\|_2 \\ \|\tilde{h}(\tilde{u})\|_2 \end{pmatrix}^\top \begin{pmatrix} -(1 + \varepsilon)\alpha & \frac{1}{2}\varepsilon\gamma \\ \frac{1}{2}\varepsilon\gamma & -\varepsilon\beta \end{pmatrix} \begin{pmatrix} \|\tilde{f}(\tilde{u})\|_2 \\ \|\tilde{h}(\tilde{u})\|_2 \end{pmatrix}. \end{aligned} \quad (26)$$

We may now select the Lyapunov function parameter ε sufficiently small such that $\left(\alpha + \varepsilon\alpha - \frac{\varepsilon\gamma^2}{4\beta}\right) > 0$. This makes the quadratic form (26) negative definite. Thus $\dot{V}_2(\tilde{z}, \tilde{u}) = 0$ if and only if $\tilde{f}(\tilde{u}) = \tilde{h}(\tilde{u}) = 0$, i.e. if and only if $\tilde{u} = 0$. In this case, $\dot{V}_1(\tilde{z}, \tilde{u})$ is clearly negative definite in \tilde{z} . Thus $\dot{V}(\tilde{z}, \tilde{u})$ is negative definite, which implies that the origin is globally asymptotically stable for system (18). Equivalently, the equilibrium (x^0, z^0) , with input u^0 , is therefore globally asymptotically stable for the original system (1)–(3). \square

7. Proof of Equilibrium Optimality

Here we prove Theorem 3.

Proof of Theorem 3. Firstly, it is clear that $v^* = \text{sat}(u^0)$ and $x_i^* = x_i^0 = -s_i \text{dz}(u_i^0)$ for all i satisfies (4b) due to x^0, z^0 being an equilibrium, and satisfies (4c) because $\text{sat}(\cdot)$ is bounded in the range $[-1, 1]$. Consider, for establishing a contradiction, that there exists $\mu \neq 0$ such that $v^\dagger = v^* + \mu$ and $x^\dagger = A^{-1}Bv^\dagger + A^{-1}w = x^* + A^{-1}B\mu$ is the optimal solution to (4) with a smaller cost (4a) than the one obtained by x^*, v^* . Then μ solves the optimization problem

$$\underset{\mu}{\text{minimize}} \quad \sum_{i=1}^n |\gamma_i x_i^* + \tilde{B}_i \mu| \quad (27a)$$

$$\text{subject to} \quad -\mathbf{1} \leq v^* + \mu \leq \mathbf{1}. \quad (27b)$$

where \tilde{B}_i is row i of the matrix $\tilde{B} = \Gamma A^{-1}B$. The equilibrium of (2) implies $x_i^* = -s_i \text{dz}(u_i^0)$. Therefore we can leverage (27b) to see that $x_i^* > 0 \implies u_i^0 < -1 \implies v_i = -1 \implies \mu_i \geq 0$ and conversely $x_i^* < 0 \implies u_i^0 > 1 \implies v_i = 1 \implies \mu_i \leq 0$. Combining this with Γ and A both being diagonal, positive definite and the fact that B is an M-matrix which implies that $\tilde{B}_{i,i} > 0$, we obtain $|\gamma_i x_i + \tilde{B}_{i,i} \mu_i| = |\gamma_i x_i| +$

$|\tilde{B}_{i,i}\mu_i|$ for all i . Thus (27a) can be expanded as follows

$$\begin{aligned}
 \sum_{i=1}^n |\gamma_i x_i + \tilde{B}_i \mu| &\geq \sum_{i \neq j} \left(|\gamma_i x_i + \tilde{B}_{i,i}\mu_i| - \left| \sum_{i \neq j} \tilde{B}_{i,j}\mu_j \right| \right) \\
 &\geq \sum_{i=1}^n (|\gamma_i x_i| + |\tilde{B}_{i,i}||\mu_i|) - \sum_{i=1}^n \sum_{j \neq i} |\tilde{B}_{i,j}||\mu_j| \quad (28) \\
 &= \sum_{i=1}^n |\gamma_i x_i| + \sum_{k=1}^n \left(|\tilde{B}_{k,k}| - \sum_{j \neq k} |\tilde{B}_{j,k}| \right) |\mu_k|.
 \end{aligned}$$

Since \tilde{B} is diagonally column-dominant, then $|\tilde{B}_{k,k}| - \sum_{j \neq k} |\tilde{B}_{j,k}|$ is positive for all k . Thus this expression is minimized by $\mu = 0$, which completes the proof. \square

8. Conclusions

In this paper we considered fully decentralized PI-control for a class of interconnected systems subject to incrementally sector-bounded nonlinearities. We showed that for systems where the input matrix is an M-matrix, fully decentralized PI-controllers globally asymptotically stabilize a specific equilibrium. Furthermore, this equilibrium is optimal in that it minimizes costs of the form $\sum_{i=1}^n \gamma_i |x_i|$. The proposed control strategy was employed in a numerical example of a simplified district heating system model. The example showed that, with our decentralized strategy, the total discomfort in the system is minimized, at the cost of higher worst-case discomforts when compared with a alternative coordinated control strategies. We have thus demonstrated that a fully decentralized and easily tuned control law constitutes a relevant design for a large class of systems.

Open questions include analysis of the transient response, and finding controller tuning rules accordingly. This could encompass the case when w is not constant but slowly time-varying, such as in the simulation study in Section 4. Furthermore, to better capture the district heating application, a richer class of systems should be considered: Multi-state models for each building, as well as more complex, nonlinear models of the interconnection B can be considered.

Declaration of Competing Interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Acknowledgements

Felix Agner, Jonas Hansson and Anders Rantzer are members of the ELLIIT Strategic Research Area at Lund University.

This work is funded by the European Research Council (ERC) under the European Union's Horizon 2020 research and innovation program under grant agreement No 834142 (ScalableControl).

This work was funded by Wallenberg AI, Autonomous Systems and Software Program (WASP) funded by the Knut and Alice Wallenberg Foundation.

Work supported in part by the MUR via grant DOCEAT, CUP E63C22000410001, number 2020RTWES4, and by the ANR via grant OLYMPIA, number ANR-23-CE48-0006.

Appendix

We prove here suitable properties of the function class characterized by Assumption 1, as stated in Lemmas 1, 2 and 3. To simplify the exposition, we drop the subscript i .

Proof of Lemma 1. Clearly, $h(0) = 0 - f(0) = 0$. Additionally,

$$\frac{h(y) - h(x)}{y - x} = \frac{y - f(y) - x + f(x)}{y - x} = 1 - \frac{f(y) - f(x)}{y - x} \in [0, 1] \quad (\text{A.1})$$

which shows that $0 \leq (h(y) - h(x))/(y - x) \leq 1$ if $x \neq y$, concluding the proof. \square

Proof of Lemma 2. Clearly, $\tilde{f}(0) = f(x^0) - f(x^0) = 0$. In addition,

$$\frac{\tilde{f}(y) - \tilde{f}(x)}{y - x} = \frac{f(y + x^0) - f(x + x^0)}{(y + x^0) - (x + x^0)} \in [0, 1] \quad (\text{A.2})$$

which shows that $0 \leq (\tilde{f}(y) - \tilde{f}(x))/(y - x) \leq 1$ if $x \neq y$. Finally $\tilde{f}(x) + \tilde{h}(x) = f(x + x^0) - f(x^0) + h(x + x^0) - h(x^0) = x + x^0 - x^0 = x$, concluding the proof. \square

Proof of Lemma 3. $\tilde{f}(0) = D^{-1}f(0) = 0$. Additionally,

$$\frac{\tilde{f}(y) - \tilde{f}(x)}{y - x} = \frac{df(y/d) - df(x/d)}{y - x} = \frac{f(y/d) - f(x/d)}{y/d - x/d} \in [0, 1]. \quad (\text{A.3})$$

Thus $0 \leq (\tilde{f}(y) - \tilde{f}(x))/(y - x) \leq 1$ if $x \neq y$. Finally $\tilde{f}(x) + \tilde{h}(x) = Df(D^{-1}x) + Dh(D^{-1}x) = D(f(D^{-1}x) + h(D^{-1}x)) = DD^{-1}x = x$, concluding the proof. \square

References

- Agarwal, P., M. Jleli, and B. Samet (2018). *Fixed Point Theory in Metric Spaces: Recent Advances and Applications*. 1st ed. See p. 5 for reference. Springer Singapore, Singapore. ISBN: 978-981-13-2913-5. DOI: <https://doi.org/10.1007/978-981-13-2913-5>.
- Agner, F., J. Hansson, P. Kergus, A. Rantzer, S. Tarbouriech, and L. Zaccarian (2024). “Decentralized PI-control and anti-windup in resource sharing networks”. *European Journal of Control*, p. 101049. ISSN: 0947-3580. DOI: [10.1016/j.ejcon.2024.101049](https://doi.org/10.1016/j.ejcon.2024.101049).
- Agner, F., P. Kergus, R. Pates, and A. Rantzer (2022). “Combating district heating bottlenecks using load control”. *Smart Energy* **6**. ISSN: 2666-9552. DOI: <https://doi.org/10.1016/j.segy.2022.100067>.
- Agner, F., P. Kergus, A. Rantzer, S. Tarbouriech, and L. Zaccarian (2023). “Anti-windup coordination strategy around a fair equilibrium in resource sharing networks”. *IEEE Control Systems Letters*.
- Bacher, P. and H. Madsen (2011). “Identifying suitable models for the heat dynamics of buildings”. en. *Energy and Buildings* **43**:7, pp. 1511–1522. ISSN: 03787788. DOI: [10.1016/j.enbuild.2011.02.005](https://doi.org/10.1016/j.enbuild.2011.02.005).
- Bauso, D., F. Blanchini, L. Giarre, and R. Pesenti (2013). “The linear saturated decentralized strategy for constrained flow control is asymptotically optimal”. *Automatica* **49**:7, pp. 2206–2212. ISSN: 0005-1098. DOI: <https://doi.org/10.1016/j.automatica.2013.03.029>.
- Dal Col, L., I. Queinnec, S. Tarbouriech, and L. Zaccarian (2019). “Regional H_∞ synchronization of identical linear multiagent systems under input saturation”. *IEEE Transactions on Control of Network Systems* **6**:2, pp. 789–799. DOI: [10.1109/TCNS.2018.2877742](https://doi.org/10.1109/TCNS.2018.2877742).
- DeLellis, P. and M. di Bernardo (2012). “Adaptive pinning control of complex networks of lur’e systems”. In: *2012 IEEE 51st IEEE Conference on Decision and Control (CDC)*, pp. 6060–6064. DOI: [10.1109/CDC.2012.6426548](https://doi.org/10.1109/CDC.2012.6426548).
- Giaccagli, M., V. Andrieu, S. Tarbouriech, and D. Astolfi (2022). “Infinite gain margin, contraction and optimality: an LMI-based design”. *European Journal of Control* **68**. 2022 European Control Conference Special Issue, p. 100685. ISSN: 0947-3580. DOI: <https://doi.org/10.1016/j.ejcon.2022.100685>.
- Giaccagli, M., V. Andrieu, S. Tarbouriech, and D. Astolfi (2023). “LMI conditions for contraction, integral action, and output feedback stabilization for a class of nonlinear systems”. *Automatica* **154**, p. 111106. ISSN: 0005-1098. DOI: <https://doi.org/10.1016/j.automatica.2023.111106>.
- Horn, R. A. and C. R. Johnson (1991). *Topics in Matrix Analysis*. Cambridge University Press. DOI: [10.1017/CB09780511840371](https://doi.org/10.1017/CB09780511840371).

- Kelly, F. (2003). “Fairness and stability of end-to-end congestion control”. *European Journal of Control* **9**:2, pp. 159–176. ISSN: 0947-3580. DOI: <https://doi.org/10.3166/ejc.9.159-176>.
- Low, S., F. Paganini, and J. Doyle (2002). “Internet congestion control”. *IEEE Control Systems Magazine* **22**:1, pp. 28–43. DOI: 10.1109/37.980245.
- Ofodile, N. A. and M. C. Turner (2016). “Decentralized approaches to antiwindup design with application to quadrotor unmanned aerial vehicles”. *IEEE Transactions on Control Systems Technology* **24**:6, pp. 1980–1992. DOI: 10.1109/TCST.2016.2521799.
- Rackauckas, C. and Q. Nie (2017). “DifferentialEquations.jl—a performant and feature-rich ecosystem for solving differential equations in julia”. *Journal of Open Research Software* **5**:1.
- Zames, G. (1966). “On the input-output stability of time-varying nonlinear feedback systems part one: conditions derived using concepts of loop gain, conicity, and positivity”. *IEEE Transactions on Automatic Control* **11**:2, pp. 228–238. DOI: 10.1109/TAC.1966.1098316.
- Zhang, F., H. L. Trentelman, and J. M. Scherpen (2013). “Robust synchronization of lur’e networks with incremental nonlinearities”. In: *52nd IEEE Conference on Decision and Control*, pp. 3439–3444. DOI: 10.1109/CDC.2013.6760410.
- Zhang, F., H. L. Trentelman, and J. M. Scherpen (2014). “Fully distributed robust synchronization of networked lure systems with incremental nonlinearities”. *Automatica* **50**:10, pp. 2515–2526. ISSN: 0005-1098. DOI: <https://doi.org/10.1016/j.automatica.2014.08.033>.

Paper IV

On PI-control in Capacity-Limited Networks

Felix Agner Anders Rantzer

Abstract

This paper concerns control of a class of systems where multiple dynamically stable agents share a nonlinear and bounded control-interconnection. The agents are subject to a disturbance which is too large to reject with the available control action, making it impossible to stabilize all agents in their desired states. In this nonlinear setting, we consider two different anti-windup equipped proportional-integral control strategies and analyze their properties. We show that a fully decentralized strategy will globally, asymptotically stabilize a unique equilibrium. This equilibrium also minimizes a weighted sum of the tracking errors. We also consider a light addition to the fully decentralized strategy, where rank-1 coordination between the agents is introduced via the anti-windup action. We show that any equilibrium to this closed-loop system minimizes the maximum tracking error for any agent. A remarkable property of these results is that they rely on extremely few assumptions on the interconnection between the agents. Finally we illustrate how the considered model can be applied in a district heating setting, and demonstrate the two considered controllers in a simulation.

Submitted to Automatica November 2024. Preprint available: F. Agner and A. Rantzer (2024). *On PI-control in capacity-limited networks*. arXiv: 2411.14077 [eess.SY]

1. Introduction

In this paper we consider control systems where a large number of interconnected agents share a limited resource, with the goal of utilizing this resource in an optimal fashion. This form of problem arises in many real-world domains: Communication networks [Kelly et al., 1998; Low and Lapsley, 1999; Low et al., 2002; Kelly, 2003], power systems [DallAnese and Simonetto, 2018; Molzahn et al., 2017; Ortmann et al., 2023], building cooling systems [Kallesøe et al., 2019; Kallesøe et al., 2020], district heating and cooling networks [Agner et al., 2022], and distributed camera systems [Martins et al., 2020; Martins and Årzén, 2021].

From a control-theoretic perspective, this family of problems poses several interesting challenges. Firstly, the multi-agent setting calls for control solutions which are distributed or decentralized to maintain scalability in large networks. Secondly, the nonlinearity imposed by the resource constraint means that a fully linear systems perspective will be insufficient. Thirdly, it is often the case that a detailed system model is difficult to obtain. Hence an explicit system model may be unavailable for control design. Finally, due to the constrained resource of the system, it is often impossible to drive the system to a preferable state for all agents. Hence it becomes interesting to analyze the optimality of any equilibrium stabilized by the closed-loop system.

Early works in this direction concerned with congestion in communication networks [Kelly et al., 1998; Low and Lapsley, 1999]. Since then, a larger body of literature has grown. An often-considered approach is to design the closed loop system to act as a gradient-descent algorithm [Krishnamoorthy and Skogestad, 2022; Hauswirth et al., 2024], in order to ensure optimality of the resulting equilibrium. This approach faces the challenge that the gradient of the steady-state map from input to equilibrium states needs to be known. Additionally, the resulting controller inherits the structure of this gradient, which may in general be dense. While works have been published in the directions of data-driven estimation of this gradient [He et al., 2024], there are still major challenges in multi-agent and continuous-time settings. For specific problem-instances, asymptotically optimal control solutions with structural sparsity have been shown. For network flow-control, distributed solutions have been found which yield asymptotic optimality [Bauso et al., 2013; Blanchini et al., 2022]. For agents connected via a saturated, linear map, where the linear part corresponds to an M-matrix, fully decentralized and rank-1 coordinated control has been considered [Agner et al., 2024a; Agner et al., 2023]. These two works consider anti-windup-equipped proportional-integral control. Anti-windup has a long history of use in dynamic controllers for plants with input saturations, typically with the purpose of ensuring that the behavior of the controller in the saturated region does not drastically differ from the unsaturated behavior [Galeani et al., 2009]. However, recent works have also shown that anti-windup has a useful application in real-time optimization [Hauswirth et al., 2020] as it holds an interpretation of projection onto the feasibility set of the system. In [Martins et al., 2020; Martins and Årzén, 2021],

an anti-windup-based controller is heuristically proposed and used to coordinate the allocation of a limited volume of disk space within a distributed camera system, informing the different cameras in the network of the current resource availability and thus improving the resource usage.

In this paper we present the following contributions. We study an extension of the model of capacity-constrained systems considered in [Agner et al., 2024a; Agner et al., 2023] to a fully nonlinear setting. We show that this extension to the nonlinear domain is crucial for modeling real world systems by explicitly demonstrating how the model can capture a district heating network. For the considered model, we consider the same two forms of controllers based on anti-windup-equipped PI control as considered in [Agner et al., 2024a; Agner et al., 2023]. Firstly a fully decentralized control structure, and secondly a structure which introduces light rank-1 coordination between the agents. We show that the results presented in [Agner et al., 2024a; Agner et al., 2023] still hold in a fully nonlinear setting. In particular, the fully decentralized controller globally, asymptotically stabilizes the system, and both of the considered controllers admit closed-loop equilibria which are optimal in the following ways: The fully decentralized controller minimizes a cost on the form $\sum a_i v_i |x_i|$, and the coordinated controller minimizes the largest control error $\|x\|_\infty$.

We formally introduce the considered plant and problem formulation in Section 2. We present the two considered control strategies, along with their associated theoretical results on stability and optimality in Section 3. We demonstrate the applicability of the considered controllers in a motivating example based on district heating in Section 4, along with a simulation. In Section 5 we prove the main results of the paper and we finally conclude the paper in Section 6.

1.1 Notation

For a vector $v \in \mathbb{R}^n$, we denote v_i to be element i of v . We denote $\text{diag}(v)$ to be a diagonal matrix with the elements of the vector v along its' diagonal. We denote $\mathbb{R}_{\geq 0}$ ($\mathbb{R}_{> 0}$) to be the set of non-negative (positive) numbers. If v and u are two vectors in \mathbb{R}^n , we say that $v \geq u$ ($v > u$) if $v - u \in \mathbb{R}_{\geq 0}^n$ ($v - u \in \mathbb{R}_{> 0}^n$). We denote $[v]_+$ to be the element-wise non-negative parts of the elements of v , such that if $u = [v]_+$, then $u_i = \max(v_i, 0)$. Conversely, $[v]_- = v - [v]_+$. We define the saturation function as $\text{sat}(u)_i = \max(\underline{l}_i, \min(\bar{l}_i, u_i))$. $\text{sat}(\cdot)$ maps from \mathbb{R}^n to a set $\mathcal{S} = \{v \in \mathbb{R}^n \mid \underline{l}_i \leq v_i \leq \bar{l}_i, \forall i = 1, \dots, n\}$ defined by the bounds \bar{l}, \underline{l} . When $\text{sat}(\cdot)$ is applied to an element of a vector, e.g., $\text{sat}(u_i)$, the bounds $\bar{l}_i, \underline{l}_i$ are implicitly used. We define the dead-zone nonlinearity $\text{dz}(u) = u - \text{sat}(u)$. We denote the sign-function $\text{sign}(x) = x/|x|$ when $x \neq 0$ and $\text{sign}(x) = 0$ for $x = 0$. When we apply $\text{sign}(\cdot)$ to a vector, the operation is performed element-wise. For a vector $x \in \mathbb{R}^n$ we use the l_1 -and- l_∞ -norms $\|x\|_1 = \sum_{i=1}^n |x_i|$ and $\|x\|_\infty = \max_i |x_i|$ respectively.

2. Problem Formulation

In this section we first introduce the considered plant, and subsequently the associated control problem.

2.1 Plant Description

We consider the control of multi-agent systems where the dynamics of agent $i \in 1, \dots, n$ can be described by the following dynamics.

$$\dot{x}_i = -a_i x_i + b_i(\text{sat}(u)) + w_i \quad (1)$$

Here x_i denotes the scalar state of agent i which should be maintained close to 0. $a_i \in \mathbb{R}_{>0}$ models a stable internal behavior of agent i . $w_i \in \mathbb{R}$ is a disturbance acting on agent i , assumed to be constant. b_i is the i 'th component of a nonlinear interconnection $b : \mathcal{S} \rightarrow \mathcal{B}$ between the agents. Here \mathcal{S} is the range of the saturation function. We consider the case where b is not explicitly known and hence cannot be used in control design and actuation. However, we assume that b holds certain exploitable properties:

ASSUMPTION 1

(Input-output properties of b) $b : \mathcal{S} \rightarrow \mathcal{B}$ is a continuous function. There exists $\eta \in \mathbb{R}_{>0}^n$ such that for any pair $\bar{v}, \underline{v} \in \mathcal{S}$ where $\bar{v} \geq \underline{v}$ and $\bar{v} \neq \underline{v}$,

(i) $b_i(\bar{v}) - b_i(\underline{v}) < 0$ if $\bar{v}_i = \underline{v}_i$, and

(ii) $\eta^\top (b(\bar{v}) - b(\underline{v})) > 0$.

Assumption (i) encodes competition between the agents: if other agents *increase* their control action while agent i maintains their control input ($\bar{v}_i = \underline{v}_i$ and $\bar{v} \geq \underline{v}$), the resource granted to agent i decreases ($b_i(\bar{v}) - b_i(\underline{v}) < 0$). Assumption (ii) encodes that if all agents increase their system input ($\bar{v} \geq \underline{v}$), the output of the system increases ($\eta^\top (b(\bar{v}) - b(\underline{v})) > 0$). This increase concerns a weighted output, governed by a weight η . If b satisfies (ii) for many different vectors $\eta > 0$, then the results of this paper hold for any such choice of η .

REMARK 1

Assumption 1 is satisfied in the linear case when $b(v) = Bv$ and $B \in \mathbb{R}^{n \times n}$ is an M -matrix. (i) then corresponds to the non-positivity of B 's off-diagonal elements. B also has a positive left eigenvector η with associated positive eigenvalue λ such that $\eta^\top B = \lambda \eta^\top$, which implies (ii). This is the case investigated in [Agner et al., 2024a; Agner et al., 2023]. We refer to [Horn and Johnson, 1991, pp. 113-115] for a more detailed definition of M -matrices and a list of their properties.

REMARK 2

Note that \mathcal{S} is an n -dimensional box and thus compact. As b is continuous, \mathcal{B} is therefore also compact due to the extreme value theorem.

2.2 Problem Description

In an ideal scenario, a controller should drive the system (1) to the origin ($x = 0$), which means that there are no control errors. This is unfortunately not always possible. The dynamics (1) dictate that any equilibrium state-input pair (x^0, u^0) yielding $\dot{x} = 0$ must satisfy $a_i x_i^0 = b_i(\text{sat}(u^0)) + w_i$ for all $i = 1, \dots, n$. But when the disturbance w is large, we may find that $-w \notin \mathcal{B}$ as the image \mathcal{B} of b is compact. Thus it becomes impossible to stabilize the origin. In this scenario, our aim is to design controllers which stabilize an equilibrium close to the origin, where we will consider two such notions of "close". The multi-agent setting also provides the complication that the controllers should require little to no communication. Furthermore, we have no explicit model of b , and can therefore not use it for control design or actuation.

3. Considered Controllers and Main Results

In this section we will define two proportional-integral control strategies. In the unsaturated region, both controllers are equivalent and fully decentralized. In the saturated region they are equipped with different anti-windup compensators. One of these anti-windup compensators is fully *decentralized* and the other is *coordinating* using rank-1 communication. We will show how the closed-loop equilibria of these two strategies minimize the distance to the origin by two different metrics.

3.1 Decentralized Control

The first control strategy we investigate is also the simplest, namely the fully decentralized strategy. Each agent $i = 1, \dots, n$, is equipped with an integral error z_i , proportional and integral gains $k_{P,i} \in \mathbb{R}_{>0}$ and $k_{I,i} \in \mathbb{R}_{>0}$, and an anti-windup gain $k_{A,i} \in \mathbb{R}_{>0}$. Their closed loop system is therefore described by

$$\dot{x}_i = -a_i x_i + b_i(\text{sat}(u)) + w_i \quad (2a)$$

$$\dot{z}_i = x_i + k_{A,i} \text{dz}_i(u) \quad (2b)$$

$$u_i = -k_{P,i} x_i - k_{I,i} z_i. \quad (2c)$$

We assume that the controller gains of each agent are tuned according to the following rule.

ASSUMPTION 2

For all agents $i = 1, \dots, n$, it holds that $k_{P,i} a_i > k_{I,i}$ (the proportional gain dominates the integral gain) and $k_{P,i} k_{A,i} < 1$ (the proportional gain and the anti-windup gain are limited).

Note that this control strategy is fully decentralized not only in terms of actuation, but also in terms of Assumption 2. The controller tuning also requires no explicit model of the interconnection b . For this closed-loop system, we present the following qualities, which we will later prove in Section 5.

THEOREM 1—GLOBAL ASYMPTOTIC STABILITY

Let Assumptions 1 and 2 hold. Then the closed loop system (2) formed by the decentralized controller has a unique, globally asymptotically stable equilibrium.

This theorem is proven in Section 5.2. By an equilibrium in this context, we mean a pair (x^0, z^0) with associated control input u^0 which solves (2) with $\dot{x} = \dot{z} = 0$. We can show that this equilibrium is optimal in the following sense.

THEOREM 2—EQUILIBRIUM OPTIMALITY

Let Assumptions 1 and 2 hold, and recall the vector η from Assumption 1. Let (x^0, u^0) be the equilibrium state-input pair stabilized by the decentralized controller (2). Consider any other pair $x^\dagger \in \mathbb{R}^n$, $u^\dagger \in \mathbb{R}^n$ which forms an equilibrium for the open-loop system, i.e., which solves (1) with $\dot{x} = 0$. If $\text{sat}(u^\dagger) \neq \text{sat}(u^0)$, then

$$\sum_{i=1}^n \eta_i a_i |x_i^0| < \sum_{i=1}^n \eta_i a_i |x_i^\dagger|. \quad (3)$$

This theorem is proven in Section 5.3. This optimality guarantee is given for the objective $\sum_{i=1}^n \eta_i a_i |x_i^0|$, characterized by a and η . Hence this result does not yield a control design method for minimizing general costs on the form $\|Wx\|_1$, where W is an arbitrary weight. Rather, it highlights that for an interesting class of problems, this fully decentralized controller which is designed without explicit parameterization of b can still provide a notion of optimality. In fact, if Assumption 1 is satisfied for a whole set of vectors η , the optimality of Theorem 2 holds for all such vectors.

3.2 Coordinating Control

The second control strategy we consider introduces a coordinating anti-windup signal. The proposed closed-loop system is given by

$$\dot{x}_i = -a_i x_i + b_i (\text{sat}(u)) + w_i \quad (4a)$$

$$\dot{z}_i = x_i + k^C \mathbf{1}^\top \text{dz}(u) \quad (4b)$$

$$u_i = -k_{P,i} x_i - k_{I,i} z_i. \quad (4c)$$

The only difference from the decentralized strategy (2) is the coordinating anti-windup term $k^C \mathbf{1}^\top \text{dz}(u)$, where $k^C \in \mathbb{R}_{>0}$ is an anti-windup gain. This coordinating controller is also fully decentralized in the unsaturated domain $\text{dz}(u) = 0$. When saturation occurs, the communication is rank-1, hence it can be implemented simply though one shared point of communication, or via scalable consensus-protocols [Olfati-Saber and Murray, 2004]. The coordinating term $k^C \mathbf{1}^\top \text{dz}(u)$ heuristically embeds the following idea: If the current disturbance on the system is large, and an agent requires more control action than the saturation allows ($\text{dz}(u_k)$ large for some $k \in 1, \dots, n$), then this will enter into the coordinating term $k^C \mathbf{1}^\top \text{dz}(u)$ and make all other agents reduce their control action, freeing more of the shared resource.

We assume that the coordinating controller is designed according to the following rules.

ASSUMPTION 3

For all agents $i = 1, \dots, n$, it holds that $a_i k_{P,i} = (1 + \alpha) k_{I,i}$, where $\alpha \in \mathbb{R}_{>0}$ is a tuning gain known to all agents. Additionally, the anti-windup gain $k^C \in \mathbb{R}_{>0}$ is chosen sufficiently small, such that $\frac{k^C}{2} \mathbf{1}^\top k_P \leq 1$.

Equation (4b) imposes the equilibrium condition $-x^0 = k^C \mathbf{1} \mathbf{1}^\top \text{dz}(u^0)$, i.e., x^0 is parallel to $\mathbf{1}$. This means that in any closed-loop equilibrium, the imposed control error is shared equally between all agents. In a sense, this means that the resource is being shared in a fair fashion between the agents. In fact, the imposed equilibrium will be optimally fair in the following sense.

THEOREM 3—EQUILIBRIUM OPTIMALITY

Let Assumptions 1 and 3 hold. Assume that (x^0, u^0) is an equilibrium stabilized by the coordinating controller (4). Consider any other pair $x^\dagger \in \mathbb{R}^n$, $u^\dagger \in \mathbb{R}^n$ forming an equilibrium for the open-loop system, i.e., they solve (1) with $\dot{x} = 0$. If $\text{sat}(u^\dagger) \neq \text{sat}(u^0)$ then

$$\|x^0\|_\infty < \|x^\dagger\|_\infty. \quad (5)$$

This theorem is proven in Section 5.3. While this is a strong result, it is only interesting if two implicit assumptions are satisfied: That such an equilibrium exists, and that it is globally (or at least locally) asymptotically stable. However, this is not always the case. As stated previously, (4b) imposes that any equilibrium x^0 is parallel to $\mathbf{1}$. At the same time, (4a) imposes that any equilibrium satisfies $x^0 = A^{-1}(b(\text{sat}(u)) + w)$ where $A = \text{diag}(a)$. As b is bounded, these two relations can only hold if $A^{-1}w$ is approximately parallel to $\mathbf{1}$, i.e., if w has a similar effect on each agent. An exact characterization of such a condition on w is outside the scope of this work. We refer to [Agner et al., 2023] for the case where b is linear. Furthermore, even when an equilibrium exists, it is non-trivial to show that the equilibrium will be stable. Such an exercise is outside the scope of most regular stability analysis for saturating systems, where it is often assumed that the stabilized equilibrium lies in the unsaturated region. We can however show the following result, which applies when the disturbance is small enough to be rejected. This theorem is proven in Section 5.2.

THEOREM 4—GLOBAL ASYMPTOTIC STABILITY

Let Assumptions 1 and 3 hold. Additionally, assume that $b(\bar{l}) + w > 0$ and $b(\underline{l}) + w < 0$. Then the closed loop system (4) formed by the coordinating controller has a unique, globally asymptotically stable equilibrium.




Figure 1. A district heating network. A heating plant (node 23) heats up water and pumps it out to consumers (nodes 1-22) through the supply-side network (red edges). The water subsequently returns through the return-side network (blue edges).

4. Motivating Example - District Heating

To illustrate the usefulness of the theoretical results, we consider district heating networks as a motivating example. Figure 1 shows a schematic example of such a system. Typically in existing networks of traditional design, one or a few large heating plants produce hot water which is pumped out to consumers via a network of pipelines (red edges in Figure 1). Each consumer is equipped with a valve to regulate the amount of hot water they receive. This water runs through a heat exchanger in which heat is transferred to the internal heating system of the building. The water subsequently returns through another network of pipes (blue edges in Figure 1) which is symmetric to the supply-side network. The primary aim in the network is to supply enough hot water to the consumers, such that they can maintain comfort temperatures within their buildings. We consider a simple dynamical model for the temperature T_i in each building $i = 1, \dots, n$, which consumer i would like to maintain at a reference temperature T_i^r . Hence the tracking error is $x_i = T_i - T_i^r$. The dynamics guiding the tracking error x_i is thus given by

$$c_i \dot{x}_i = -\hat{a}_i (T_i^r + x_i - T_o) + c_{p,w} \rho_w \delta_i q_i(v). \quad (6)$$

Here c_i is the heat capacity of building i . \hat{a}_i is the thermal conductance and T_o is the outdoor temperature, acting as a disturbance. Hence the first term $-\hat{a}_i (T_i^r + x_i - T_o)$ corresponds to diffusion of heat between the interior and exterior of the building. $c_{p,w}$ and ρ_w are the specific heat capacity and density of water respectively (both assumed constant) and q_i is the volume flow rate going through the building heat exchanger. δ_i is the difference in supply-and-return temperature before and after the heat exchanger. Hence the second term $c_{p,w} \rho_w \delta_i q_i(v)$ corresponds to the heat provided to the building through the heat exchanger. The volume flow rate q_i is

regulated by the valve positions $v = \text{sat}(u)$. Note that the flow rate q_i provided to consumer i is influenced by the valve positions v of all consumers in the network, not only v_i . This is because the pressure distribution in the network is affected by all of the flow rates in the network. Herein lies the main connection to our theoretical results. As the central pump is limited in its maximum capacity, and the valves themselves are saturated, the volume flow rate and hence the heat that can be supplied to the consumers is limited. Hence when the disturbance T_o is sufficiently low, the available capacity becomes insufficient.

To complete the connection between the temperature model (6) and the agent dynamics (1) as we have considered in this paper, we can identify $a_i = \frac{\hat{a}_i}{c_i}$, $b_i(\text{sat}(u)) = \frac{c_{p,w}\rho_w\delta_i}{c_i} q_i(\text{sat}(u))$ and $w_i = \frac{\hat{a}_i}{c_i}(T_o = T_i^r)$. Secondly, we make the following simplifying assumption.

ASSUMPTION 4

The delta temperatures δ_i , $c_{p,w}$ and ρ_w are all constant.

In practice this assumption will not hold exactly. The delta temperature changes slightly with several factors, such as the supply-temperature in the network, the activity on the secondary side of the heat exchanger (i.e., the side facing the consumers internal heating system). There are also slight temperature-dependent variations in the density of the water. In general however and over shorter time-spans, these variations are much smaller than the variations in volume flow rate. This assumption means that the final verification to make is that q satisfies Assumption 1. Under the assumption that we use common static models for the valves and pipes in the network such as in [De Persis et al., 2014; Agner et al., 2022; Jeeninga et al., 2023], that the network is tree-structured, and at the pump at the root of the tree operates at constant capacity, we can show the following. Hence q satisfies Assumption 1.

PROPOSITION 1

Given two sets of valve positions $\bar{v} \geq v$,

$$(i) \quad q_i(\bar{v}) - q_i(v) \leq 0 \text{ if } \bar{v}_i = v_i, \text{ and}$$

$$(ii) \quad \mathbf{1}^\top (q(\bar{v}) - q(v)) > 0 \text{ if } \bar{v} \neq v.$$

We omit the proof of this proposition, as it demands a technical description of district heating hydraulic models. However, we can motivate the proposition in the following way. If all agents incrementally open their valves ($\bar{v} \geq v$), this reduces the resistance in the system, which means that the total throughput increases ($\mathbf{1}^\top q(\bar{v}) > \mathbf{1}^\top q(v)$), i.e., (ii). However, as the total throughput increases, so do pressure losses in the pipelines. Hence, if one valve i is unchanged ($\bar{v}_i = v_i$), the flow rate through the valve will decrease due to reduced differential pressure ($q_i(\bar{v}) \leq q_i(v)$), i.e., (i).




Figure 2. Outdoor temperature $T_o(t)$ used in simulation.

4.1 Numerical Example

To investigate the effect of the considered control strategies in a district heating setting, we perform a simulation of a small district heating network. The network is structured as in Figure 1, and each building is subject to the dynamics given in (6). For simplicity, we consider a homogeneous building stock with $c_i = 2.0[\text{kWh/K}]$, $a_i = 1.2[\text{kW/K}]$ and $\delta_i = 50.0[\text{K}] \forall i = 1, \dots, n$. We have $c_w = 1.16 \cdot 10^{-3}[\text{kWh/kgK}]$ and $\rho_w = 10^3[\text{kg/m}^3]$. While we omit a detailed description of district heating hydraulics here, we use the same type of graph-based modeling as is used in [Agner et al., 2024b]. We assume that the heating plant supplies a differential pressure of $0.6 \cdot 10^6[\text{Pa}]$. The difference between the input and output of each pipe e is given by $\Delta p_e = s_e |q_e| q_e$ where s_e corresponds to a hydraulic resistance. We use $s_e = 0.9[\text{Pa} / (\text{m}^3/\text{h})^2]$ for the long edge connecting nodes 23 and 24. We use $s_e = 0.25[\text{Pa} / (\text{m}^3/\text{h})^2]$ for the edges connecting nodes 24, 25, 26, 30 and 33. We use $s_e = 0.05[\text{Pa} / (\text{m}^3/\text{h})^2]$ for the pipes connecting nodes 26-27-28-29, nodes 30-31-32 and nodes 33-34-35-36. Finally we use $s_e = 2.5[\text{Pa} / (\text{m}^3/\text{h})^2]$ for the connection to each consumer. The pressure difference between supply-and-return-side for consumer i is modeled as $\Delta p_i(q_i, v_i) = \left(5 + \frac{30}{(v_i + 1.001)^2}\right) q_i^2$. Here $v_i = \text{sat}(u_i)$ is limited in the $v_i \in [-1, 1]$. The component $5q_i^2$ corresponds to inactive components of the consumer substation, i.e., the heat exchanger and internal piping. The remaining component $30q_i^2/(v_i + 1.001)^2$ corresponds to the pressure loss over the valve.

We subject the buildings to an outdoor temperatures disturbance T_o as seen in Figure 2, acting equally on all buildings. The temperature drops critically to below -25°C around 50 hours into the simulation. The temperature is based on temperature data from Gävle, Sweden on January 18th-21st, 2024. The data is collected from the Swedish Meteorological and Hydrological Institute.

We consider the two control policies analyzed in this paper, namely the *decentralized* and *coordinating* control policies. We employ identical controllers for each consumer with $k_{P,i} = 1.0$, $k_{I,i} = 1.0$, $k_{A,i} = 1.0$ for all $i = 1, \dots, n$ in the decentralized case, and $k_{P,i} = 1.0$, $k_{I,i} = 1.0$, $k_{A,i} = 1.0$, $k^C = 0.5$ in the coordinating case. As

a benchmark, we compare these strategies to optimal counter-parts. In these benchmarks, the volume flow rate $q(t)$ is distributed optimally in each instance of the simulation as the solution to the problem

$$\underset{x, q}{\text{minimize}} \quad J(x) \quad (7a)$$

$$\text{subject to} \quad (6) \text{ with } \dot{x}_i = 0 \text{ for } i = 1, \dots, n, \quad (7b)$$

$$q \in \mathcal{Q} \quad (7c)$$

where we use $J(x) = \|x\|_1$ and $J(x) = \|x\|_\infty$ respectively. This problem corresponds to calculating a flow rate q which is feasible within the hydraulic constraints of the network (i.e., (7b)) which generates an equilibrium (i.e., (7b)) which minimizes $J(x)$. To see how this optimization problem can be cast as a convex problem, we refer to [Agner et al., 2022] in which it is shown that \mathcal{Q} is convex.

We use the `DifferentialEquations` toolbox [Rackauckas and Nie, 2017] in Julia to simulate the system, utilizing the FBDf solver. We use the `NonlinearSolve` [Pal et al., 2024] toolbox to calculate q as a function of the valve positions $v = \text{sat}(u)$. We use the `Convex` toolbox [Udell et al., 2014] with the Mosek optimizer to find the optimal trajectories for the benchmark comparisons.

The results of the four simulations are seen in Figure 3. Under all four policies, the temperatures in the buildings drop at several points during the simulation, and most significantly starting after around 40 hours. This is because of the extremely cold temperature at this time, for which the available pumping capacity is insufficient. We can first compare the results of using the decentralized strategy to the results of using the optimal equilibrium input with regards to $\|x\|_1$, as seen in Figures 3a and 3c respectively. We find that they are effectively the same, except for minor oscillations around the equilibrium in the PI-controller case, caused by the integral action in the controller. The same comparison can be drawn between Figures 3b and 3d, showing the results of using the coordinating strategy and the optimal input with regards to $\|x\|_\infty$. This is in line with Theorems 2 and 3, where we expect our controllers to track the optimal equilibrium. While the interpretation of the optimal cost in Theorem 2 is obscured by the weight η , we see in this example how it can correspond to the combined tracking error of all agents $\|x\|_1$.

When comparing the results of using the decentralized strategy to using the coordinating strategy, we see the following: In the fully decentralized simulation, several of the buildings far away from the heating plant drop below 5°C , whereas buildings close to the heating facility maintain comfort temperature. On the other hand, none of the buildings drop below 10°C . However, none of the buildings maintain comfort temperature either. Which strategy is to be preferred is debatable and perhaps situational. Arguably in the extreme scenario of this simulation, the decentralized strategy might be preferred. Consumers will be severely dissatisfied if their indoor temperatures drop by 10°C , hence it may be better to have a lower number of consumers be very dissatisfied than to have the whole building stock be moderately dissatisfied. However, if we instead consider the temperatures distribution at

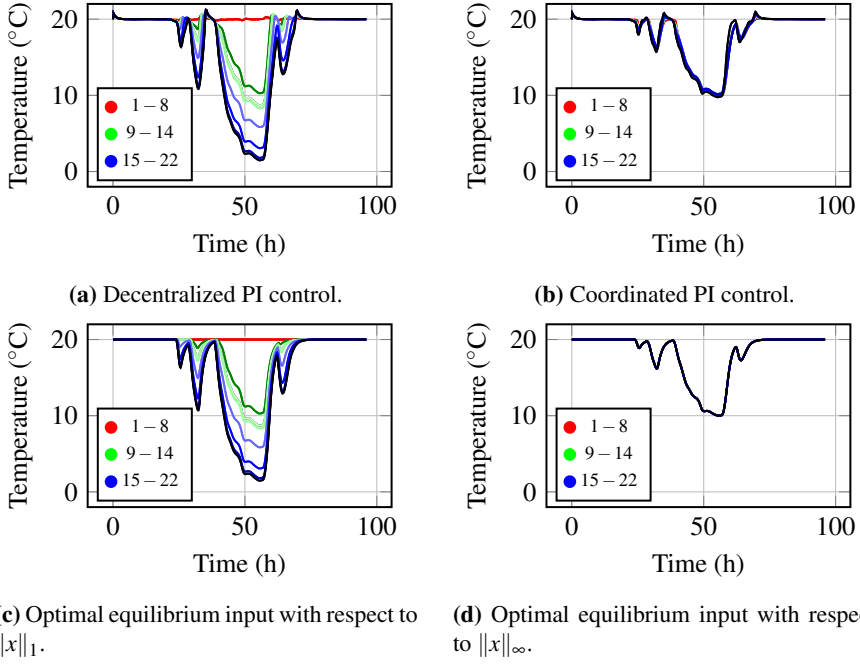


Figure 3. Resulting indoor temperatures. The three clusters of buildings, i.e., nodes 1-8, 9-14 and 15-22 are colored in red, green and blue respectively. The darker shades of each color shows the buildings further down each line.

around 30 hours into the simulations, we see that under the decentralized case there are buildings experiencing reductions in indoor temperature by 10°C , whereas in the coordinated case the worst reduction is approximately 4°C . In this case, it is arguably preferable to coordinate, as a 4°C temperature reduction is acceptable for a shorter period of time, whereas 10°C is too extreme to be tolerated by most consumers. The exact trade-offs and results will depend on the specific system and temperature levels at hand. What is interesting is that both of these behaviors which are optimal under different perspectives are achievable with such simple control techniques.

5. Main Proofs

We will now move on to prove the main theoretical results of this paper as presented in section 3. We will prove the stability results for both proposed controllers, i.e., Theorems 1 and 4, followed by the optimality of their equilibria, i.e., Theorems 2 and 3. First however, we will introduce a few extra properties of the interconnection b which are required for the subsequent proofs.

5.1 Additional Nonlinearity Properties

The proofs of both of the following Lemmas are found in the Appendix.

LEMMA 1

Let b satisfy Assumption 1. Consider any pair $v \in \mathcal{S}$, $\tilde{v} \in \mathcal{S}$ where $v \neq \tilde{v}$. Let $\mathcal{I}^\pm = \{i \in 1, \dots, n \mid v_i \neq \tilde{v}_i\}$ and let $\mathcal{I}^0 = \{i \in 1, \dots, n \mid v_i = \tilde{v}_i\}$. Then

$$\begin{aligned} & \sum_{i \in \mathcal{I}^\pm} \eta_i \text{sign}(\tilde{v}_i - v_i) (b_i(\tilde{v}) - b_i(v)) \\ & > \sum_{k \in \mathcal{I}^0} \eta_k |b_k(\tilde{v}) - b_k(v)|. \end{aligned} \quad (8)$$

This lemma is proven in the appendix. An interpretation of this lemma is that the individual change in output $b_i(v) - b_i(\tilde{v})$ goes mostly along the same sign as the corresponding individual change in input $v_i - \tilde{v}_i$. This value for all agents who have changed their inputs (\mathcal{I}^\pm) dominates the change in output affecting all of the agents who did not change their inputs (\mathcal{I}^0). We continue with the following property of b .

LEMMA 2

Let b satisfy Assumption 1. Then for any pair $\bar{v} \in \mathcal{S}$ and $\underline{v} \in \mathcal{S}$ where $\bar{v} \neq \underline{v}$, if $b(\bar{v}) \geq b(\underline{v})$, then $\bar{v} > \underline{v}$.

This lemma is proven in the appendix. In the linear case where b is an M-matrix ($b(v) = Bv$), this property can be likened with positivity of the inverse of this matrix ($B^{-1} > 0$ element-wise).

5.2 Stability Proofs

We will prove both Theorem 1 and 4 using Lyapunov-based arguments. To do so, we will first introduce a change of coordinates from (x, z) to (ζ, u) , where $u_i = -k_{p,i}x_i - k_{1,i}z_i$ and $\zeta_i = -k_{1,i}z_i$ for $i = 1, \dots, n$. We also introduce the matrices $P = \text{diag}(k_p)$, $R = \text{diag}(k_1)$, $C = RP^{-1}$ and $D = \text{diag}(a) - C$. Note that P , R , C and D are all positive definite, diagonal matrices under either Assumption 2 or Assumption 3. The closed-loop system in these new coordinates is then given by

$$\begin{bmatrix} \dot{\zeta} \\ \dot{u} \end{bmatrix} = \begin{bmatrix} -C & C \\ D & -D \end{bmatrix} \begin{bmatrix} \zeta \\ u \end{bmatrix} - \begin{bmatrix} 0 \\ P \end{bmatrix} b(\text{sat}(u)) - \begin{bmatrix} 0 \\ P \end{bmatrix} w - \begin{bmatrix} R \\ R \end{bmatrix} Sdz(u). \quad (9)$$

Here S denotes the anti-windup compensation, i.e., $S = \text{diag}(k_{A,i})$ for the decentralized controller and $S = k^C \mathbf{1}\mathbf{1}^\top$ for the coordinating controller.

We will first prove stability of the decentralized controller, beginning with ensuring that there exists an equilibrium.

LEMMA 3

Let Assumptions 1 and 2 hold. Then the decentralized closed loop system (2) has at least one equilibrium.

This can be easily proven with Brouwer's fixed-point theorem. We provide an outline of the proof in the Appendix, omitting the details for brevity. We can now prove stability of the decentralized closed-loop system.

Proof of Theorem 1. Under the taken assumptions, Lemma 3 provides that the original system (2), and thus also (9), has at least one equilibrium which we denote (ζ^0, u^0) . We introduce the shifted variables $\tilde{\zeta} = \zeta - \zeta^0$ and $\tilde{u} = u - u^0$, and the shifted notation $\text{s}\tilde{\text{a}}\text{t}(u) = \text{sat}(u^0 + \tilde{u}) - \text{sat}(u^0)$, $\tilde{b}(\text{s}\tilde{\text{a}}\text{t}(\tilde{u})) = b(\text{s}\tilde{\text{a}}\text{t}(u) + \text{sat}(u^0)) - b(\text{sat}(u^0))$ and $\tilde{\text{d}}\tilde{z}(\tilde{u}) = \text{dz}(u^0 + \tilde{u}) - \text{dz}(u^0) = \tilde{u} - \text{s}\tilde{\text{a}}\text{t}(u)$. In this coordinate frame, the closed-loop dynamics are

$$\begin{bmatrix} \dot{\tilde{\zeta}} \\ \dot{\tilde{u}} \end{bmatrix} = \begin{bmatrix} -C & C \\ D & -D \end{bmatrix} \begin{bmatrix} \tilde{\zeta} \\ \tilde{u} \end{bmatrix} - \begin{bmatrix} 0 \\ P \end{bmatrix} \tilde{b}(\text{s}\tilde{\text{a}}\text{t}(\tilde{u})) - \begin{bmatrix} R \\ R \end{bmatrix} S\tilde{\text{d}}\tilde{z}(\tilde{u}). \quad (10)$$

The aim is now to show that (10) is globally, asymptotically stable with regards to the origin, which is equivalent with Theorem 1. Now recall the vector η of Assumption 1, which we use to define the following Lyapunov function candidate.

$$V(\tilde{\zeta}, \tilde{u}) = \sum_{i=1}^n \frac{\eta_i d_i}{p_i c_i} |\tilde{\zeta}_i| + \frac{\eta_i}{p_i} |\tilde{u}_i| \quad (11)$$

While V is not strictly continuously differentiable, we note this as a technicality. The function $|\cdot|$ can be exchanged with an arbitrarily close approximation which is continuously differentiable in the origin. For this proof, we will maintain the convention that

$$\frac{d}{dt}|x| = \text{sign}(x)\dot{x}. \quad (12)$$

Denote $H = \text{diag}(\eta)$. We then find that

$$\begin{aligned} & \dot{V}(\tilde{\zeta}, \tilde{u}) \\ &= \text{sign}(\tilde{\zeta})^\top H D P^{-1} C^{-1} \dot{\tilde{\zeta}} + \text{sign}(\tilde{u})^\top H P^{-1} \dot{\tilde{u}} \\ &= \text{sign}(\tilde{\zeta})^\top H D P^{-1} C^{-1} \left(C(\tilde{u} - \tilde{\zeta}) - R S \tilde{\text{d}}\tilde{z}(\tilde{u}) \right) \\ & \quad + \text{sign}(\tilde{u})^\top H P^{-1} \left(D(\tilde{\zeta} - \tilde{u}) \right. \\ & \quad \left. - P \tilde{b}(\text{s}\tilde{\text{a}}\text{t}(\tilde{u})) - R S \tilde{\text{d}}\tilde{z}(\tilde{u}) \right) \\ &= - \left(\text{sign}(\tilde{u}) - \text{sign}(\tilde{\zeta}) \right)^\top H D P^{-1} (\tilde{u} - \tilde{\zeta}) \end{aligned} \quad (13a)$$

$$- \text{sign}(\tilde{\zeta})^\top H D P^{-1} C^{-1} R S \tilde{\text{d}}\tilde{z}(\tilde{u}) \quad (13b)$$

$$- \text{sign}(\tilde{u})^\top H P^{-1} R S \tilde{\text{d}}\tilde{z}(\tilde{u}) \quad (13c)$$

$$- \text{sign}(\tilde{u})^\top H \tilde{b}(\text{s}\tilde{\text{a}}\text{t}(\tilde{u})). \quad (13d)$$

The terms (13a)–(13b) act fully diagonally, hence we can analyze their sign contribution for each $i \in 1, \dots, n$ individually. If either $\tilde{u}_i = 0$ and $\tilde{\zeta}_i \neq 0$ or $\tilde{u}_i \neq 0$ and $\tilde{\zeta}_i = 0$, then clearly (13b) contributes with 0, and (13a) becomes strictly negative in the non-zero variable. If $\text{sign}(\tilde{u}_i) = \text{sign}(\tilde{\zeta}_i)$, then (13a) contributes with 0, but (13b) contributes with a negative semidefinite term, which is strictly negative if $\tilde{d}\tilde{z}(\tilde{u})_i \neq 0$. If $\text{sign}(\tilde{u}_i) = -\text{sign}(\tilde{\zeta}_i) \neq 0$, (13b) contributes with a positive semidefinite term $\frac{\eta_i d_i r_i s_i}{p_i c_i} |\tilde{d}\tilde{z}(\tilde{u})_i|$. However, in this case (13a) contributes with a strictly negative term $-\frac{\eta_i d_i}{p_i} (|\tilde{\zeta}_i| + |\tilde{u}_i|)$. This negative term dominates the positive semidefinite term because of Assumption 2 where the anti-windup gain is bounded, and because $|\tilde{d}\tilde{z}(\tilde{u})_i| \leq |\tilde{u}_i|$. Hence the contribution of (13a)–(13b) is negative semidefinite, and strictly negative when $\tilde{d}\tilde{z}(\tilde{u}) \neq 0$. The term (13c) is trivially negative semidefinite, and strictly negative when $\tilde{d}\tilde{z}(\tilde{u}) \neq 0$. Finally, let $\mathcal{I}^0 = \{i \in 1, \dots, n \mid \text{s}\tilde{\text{a}}\tilde{\text{t}}(\tilde{u})_i = 0\}$. Note that $\text{sign}(\tilde{u}_i) = 0 \implies \text{sign}(\text{s}\tilde{\text{a}}\tilde{\text{t}}(\tilde{u})_i) = 0$, and if $\text{sign}(\tilde{u}_i) \neq 0$ then either $i \in \mathcal{I}^0$ or $\text{sign}(\text{s}\tilde{\text{a}}\tilde{\text{t}}(\tilde{u})_i) = \text{sign}(\tilde{u}_i)$. Hence the term (13d) can be bounded by

$$-\text{sign}(\tilde{u})^\top H\tilde{b}(\text{s}\tilde{\text{a}}\tilde{\text{t}}(\tilde{u})) \quad (14)$$

$$\geq \text{sign}(\text{s}\tilde{\text{a}}\tilde{\text{t}}(\tilde{u}))^\top H\tilde{b}(\text{s}\tilde{\text{a}}\tilde{\text{t}}(\tilde{u})) \quad (15)$$

$$-\sum_{j \in \mathcal{I}^0} \eta_j |\tilde{b}_j(\text{s}\tilde{\text{a}}\tilde{\text{t}}(\tilde{u}))|. \quad (16)$$

As Assumption 1 holds, we can employ Lemma 1 to show that this expression is strictly negative when $\text{s}\tilde{\text{a}}\tilde{\text{t}}(\tilde{u}) \neq 0$. All together, if $\tilde{u} \neq 0$, we must have $\text{s}\tilde{\text{a}}\tilde{\text{t}}(\tilde{u}) \neq 0$, $\tilde{d}\tilde{z}(\tilde{u}) \neq 0$ or both, in which case $\dot{V}(\tilde{\zeta}, \tilde{u})$ will be negative due to the above arguments. If $\tilde{u} = 0$, the term (13a) is negative definite in $\tilde{\zeta}$. Hence $\dot{V}(\tilde{\zeta}, \tilde{u})$ is negative definite. As $V(\tilde{\zeta}, \tilde{u}) > 0$ and $\dot{V}(\tilde{\zeta}, \tilde{u}) < 0$ for any non-zero pair $(\tilde{\zeta}, \tilde{u})$ the equilibrium (ζ^0, u^0) is globally asymptotically stable for (9) and must therefore be unique. This can be translated to a unique equilibrium (x^0, z^0) in the original coordinate frame, thus concluding the proof. \square

To prove Theorem 3 regarding stability of the coordinating closed-loop system, we utilize the following result.

LEMMA 4

Assume that Assumptions 1 and 3 hold. Furthermore, assume that $b(\bar{l}) + w > 0$ and $b(\underline{l}) + w < 0$. Then, from any initial condition, the coordinating closed-loop system (4) will converge to a forward invariant set in which the control signal is unsaturated, i.e., $\text{d}z(u) = 0$.

Proof. For this proof, we will employ the coordinate frame (ζ, u) as in (9), now with $S = k^C \mathbf{1}\mathbf{1}^\top$. Consider the Lyapunov function candidate

$$V(\zeta, u) = \frac{1}{2} \text{dz}(\zeta)^\top D R^{-1} C^{-1} \text{dz}(\zeta) + \frac{1}{2} \text{dz}(u)^\top R^{-1} \text{dz}(u). \quad (17)$$

Here $\text{dz}(\zeta)$ is to be understood to element-wise use the same bounds \bar{l} and \underline{l} as $\text{dz}(u)$, i.e., if $\zeta = u$ then $\text{dz}(\zeta) = \text{dz}(u)$. Note that

$$\frac{1}{2} \cdot \frac{d}{dt} \text{dz}(x)^2 = \text{dz}(x) \frac{\partial \text{dz}(x)}{\partial x} \dot{x} = \text{dz}(x) \dot{x}, \quad (18)$$

as $\frac{\partial \text{dz}(x)}{\partial x} = 0$ when $\text{dz}(x) = 0$, and $\frac{\partial \text{dz}(x)}{\partial x} = 1$ when $\text{dz}(x) \neq 0$. Hence we find that

$$\begin{aligned} \dot{V}(\zeta, u) &= \text{dz}(\zeta)^\top D R^{-1} C^{-1} \dot{\zeta} + \text{dz}(u)^\top R^{-1} \dot{x} \\ &= \text{dz}(\zeta)^\top D R^{-1} C^{-1} \left(C(u - \zeta) - k^C R \mathbf{1}\mathbf{1}^\top \right) \\ &\quad + \text{dz}(u)^\top R^{-1} (D(\zeta - u) \\ &\quad - P(b(\text{sat}(u)) + w) - k^C R \mathbf{1}\mathbf{1}^\top) \\ &= -(\text{dz}(u) - \text{dz}(\zeta))^\top R^{-1} D(u - \zeta) \end{aligned} \quad (19)$$

$$- k^C \text{dz}(\zeta)^\top C^{-1} D \mathbf{1}\mathbf{1}^\top \text{dz}(u) \quad (20)$$

$$- k^C \text{dz}(u)^\top \mathbf{1}\mathbf{1}^\top \text{dz}(u) \quad (21)$$

$$- \text{dz}(u)^\top R^{-1} P(b(\text{sat}(u)) + w) \quad (22)$$

We will begin with the last term (22), which is strictly negative when $\text{dz}(u) \neq 0$. We can see this by noting that $R^{-1}P$ is a positive diagonal matrix, and then identifying that for any $i \in 1, \dots, n$, if $\text{dz}_i(u) > 0$, $\text{sat}_i(u) = \bar{l}_i$, and thus Assumption 1 (i) yields that

$$b_i(\text{sat}(u)) + w_i \geq b_i(\bar{l}) + w_i > 0.$$

The opposite relation can be shown when $\text{dz}_i(u) < 0$. For the remaining terms, we first note that

$$\begin{aligned} &(\text{dz}(u) - \text{dz}(\zeta))^\top R^{-1} D(u - \zeta) \\ &\geq (\text{dz}(u) - \text{dz}(\zeta))^\top R^{-1} D(\text{dz}(u) - \text{dz}(\zeta)), \end{aligned} \quad (23)$$

which can trivially be shown by utilizing the definitions of $\text{sat}(\cdot)$ and $\text{dz}(\cdot)$. Furthermore, Assumption 3 yields that $C^{-1}D = \alpha I$ and $R^{-1}D = \alpha P^{-1}$. Hence (19)–(21)

can be combined and upper bounded by

$$\begin{aligned}
& -\alpha (\mathrm{d}z(u) - \mathrm{d}z(\zeta))^\top P^{-1} (\mathrm{d}z(u) - \mathrm{d}z(\zeta)) \\
& -\alpha k^C \mathrm{d}z(\zeta)^\top \mathbf{1} \mathbf{1}^\top \mathrm{d}z(u) \\
& -k^C \mathrm{d}z(u)^\top \mathbf{1} \mathbf{1}^\top \mathrm{d}z(u) \\
& = -\alpha (\mathrm{d}z(u) - \mathrm{d}z(\zeta))^\top (P^{-1} - \frac{k^C}{2} \mathbf{1} \mathbf{1}^\top) (\mathrm{d}z(u) - \mathrm{d}z(\zeta)) \\
& \quad - \frac{k^C \alpha}{2} \mathrm{d}z(v)^\top \mathbf{1} \mathbf{1}^\top \mathrm{d}z(v) \\
& \quad - k^C \left(1 + \frac{\alpha}{2}\right) \mathrm{d}z(u)^\top \mathbf{1} \mathbf{1}^\top \mathrm{d}z(u).
\end{aligned}$$

This expression is negative semi-definite under the condition that $(P^{-1} - \frac{k^C}{2} \mathbf{1} \mathbf{1}^\top) \succeq 0$. This is equivalent to

$$\begin{bmatrix} P^{-1} & \mathbf{1} \\ \mathbf{1}^\top & \frac{2}{k^C} \end{bmatrix} \succeq 0, \quad (24)$$

as $(P^{-1} - \frac{k^C}{2} \mathbf{1} \mathbf{1}^\top)$ is the Schur complement of this matrix. The condition (24) is once again equivalent to

$$\frac{2}{k^C} - \mathbf{1}^\top P \mathbf{1} \geq 0 \iff \frac{k^C}{2} \mathbf{1}^\top P \mathbf{1} = \frac{k^C}{2} \mathbf{1}^\top k_P \geq 0. \quad (25)$$

This holds, due to Assumption 3. Hence $(P^{-1} - \frac{k^C}{2} \mathbf{1} \mathbf{1}^\top) \succeq 0$ and thus (19)–(21) is negative semi-definite. Therefore we have shown that $V(\zeta, u) > 0$ when $\mathrm{d}z(u) \neq 0$ and $\dot{V}(\zeta, u) < 0$ when $\mathrm{d}z(u) \neq 0$. Therefore the region where $\mathrm{d}z(u) = 0$ is globally attracting and forward invariant. \square

Proof of Theorem 4. Lemma 4 proves that all trajectories of the closed-loop system will converge to, and remain in, the region where $\mathrm{d}z(u) = 0$. Here, the closed-loop systems of the coordinating and decentralized controllers are equivalent, and Assumption 3 implies that also Assumption 2 holds (disregarding the statements about the anti-windup gains, as they are inactive). Hence, we can invoke Theorem 1 which applies to the decentralized closed-loop system. \square

5.3 Optimality Proofs

We continue now to prove Theorems 2 and 3.

Proof of Theorem 2. To simplify notation, we will introduce $v^0 = \text{sat}(u^0)$ and $v^\dagger = \text{sat}(u^\dagger)$. Note that $v^\dagger \neq v^0$ by assumption. As both pairs (x^0, u^0) and (x^\dagger, u^\dagger) satisfy (1) with $\dot{x} = 0$, we can conclude that

$$\eta_i a_i x_i^\dagger = \eta_i (b_i(v^\dagger) + w_i) = \eta_i a_i x_i^0 + \eta_i (b_i(v^\dagger) - b_i(v^0))$$

for all $i = 1, \dots, n$. Therefore

$$\begin{aligned}
 \sum_{i=1}^n \eta_i a_i |x_i^\dagger| &= \sum_{i=1}^n \eta_i |a_i x_i^0 + b_i(v^\dagger) - b_i(v^0)| \\
 &\geq \sum_{i \in \mathcal{I}^\pm} \eta_i a_i |x_i^0| + \eta_i \text{sign}(x_i^0) (b_i(v^\dagger) - b_i(v^0)) \\
 &\quad + \sum_{j \in \mathcal{I}^0} \eta_j |b_j(v^\dagger) - b_j(v^0)|
 \end{aligned} \tag{26}$$

where $\mathcal{I}^\pm = \{i \in 1, \dots, n \mid x_i^0 \neq 0\}$ and $\mathcal{I}^0 = \{i \in 1, \dots, n \mid x_i^0 = 0\}$. As (x^0, u^0) satisfies (2b) with $\dot{z} = 0$, and $k_A > 0$, we know that $\text{sign}(x^0) = -\text{sign}(\text{d}z(u^0))$. Hence we can continue to expand (26) to

$$\begin{aligned}
 &\sum_{i=1}^n \eta_i a_i |x_i^0| \\
 &+ \sum_{i \in \mathcal{I}^\pm} \eta_i \text{sign}(\text{d}z_i(u^0)) (b_i(v^0) - b_i(v^\dagger)) \\
 &+ \sum_{j \in \mathcal{I}^0} \eta_j |b_j(v^0) - b_j(v^\dagger)|.
 \end{aligned} \tag{27}$$

We would here like to apply Lemma 1, but this requires the sets $\mathcal{J}^\pm = \{i \in 1, \dots, n \mid v_i^0 \neq v_i^\dagger\}$ and $\mathcal{J}^0 = \{i \in 1, \dots, n \mid v_i^0 = v_i^\dagger\}$. To continue, we note the following. For all $i \in \mathcal{I}^\pm$, $\text{d}z_i(u^0) \neq 0$, and hence $v_i^0 = \bar{l}_i$ or $v_i^0 = \underline{l}_i$. Therefore the following statements hold.

$$\begin{aligned}
 &\eta_i \text{sign}(\text{d}z_i(u^0)) (b_i(v^0) - b_i(v^\dagger)) \\
 &= \eta_i \text{sign}(v_i^\dagger - v_i^0) (b_i(v^0) - b_i(v^\dagger)), \quad \forall i \in \mathcal{I}^\pm \cap \mathcal{J}^\pm
 \end{aligned} \tag{28}$$

$$\begin{aligned}
 &\eta_i \text{sign}(\text{d}z_i(u^0)) (b_i(v^0) - b_i(v^\dagger)) \\
 &\geq -\eta_i |b_i(v^0) - b_i(v^\dagger)|, \quad \forall i \in \mathcal{I}^\pm \cap \mathcal{J}^0
 \end{aligned} \tag{29}$$

$$\begin{aligned}
 &\eta_i |b_i(v^0) - b_i(v^\dagger)| \\
 &\geq \eta_i \text{sign}(\text{d}z_i(u^0)) (b_i(v^0) - b_i(v^\dagger)), \quad \forall i \in \mathcal{I}^0 \cap \mathcal{J}^\pm
 \end{aligned} \tag{30}$$

$$\begin{aligned}
 &\eta_i |b_i(v^0) - b_i(v^\dagger)| \\
 &\geq -\eta_i |b_i(v^0) - b_i(v^\dagger)|, \quad \forall i \in \mathcal{I}^0 \cap \mathcal{J}^0.
 \end{aligned} \tag{31}$$

We can use these relations to reorganize and upper-bound the sums in (27), and

therefore state that

$$\begin{aligned}
& \sum_{i=1}^n \eta_i a_i |x_i^\dagger| \\
& \geq \sum_{i=1}^n \eta_i a_i |x_i^0| \\
& \quad + \sum_{i \in \mathcal{J}^\pm} \eta_i \text{sign}(v^0 - v_i^\dagger) (b_i(v^0) - b_i(v^\dagger)) \\
& \quad - \sum_{j \in \mathcal{J}^0} \eta_j |b_j(v^0) - b_j(v^\dagger)| \\
& > \sum_{i=1}^n \eta_i a_i |x_i^0|. \tag{32}
\end{aligned}$$

The final inequality derives from Lemma 1, as we assume that Assumption 1 holds and $v^0 \neq v^\dagger$. \square

Proof of Theorem 3. Since both (x^0, u^0) and (x^\dagger, u^\dagger) solve (1) with $\dot{x} = 0$, we know that

$$\begin{aligned}
x^\dagger &= A^{-1} (b(\text{sat}(u^\dagger)) + w) \\
&= x^0 + A^{-1} (b(\text{sat}(u^\dagger)) - b(\text{sat}(u^0))), \tag{33}
\end{aligned}$$

where $A = \text{diag}(a)$. We also know that

$$x^0 = -k^C \mathbf{1} \mathbf{1}^\top \text{dz}(u^0) \tag{34}$$

because (x^0, u^0) is an equilibrium for the closed-loop system (4) and thus satisfy (4b) with $\dot{z} = 0$. We therefore know that x^0 is proportional to the vector $\mathbf{1}$, and thus $\|x^0\|_\infty = \max_i |x_i^0|$ is maximized by all $i = 1, \dots, n$ simultaneously. Consider first the case where $x^0 > 0$. Then the contradictory notion that $\|x^\dagger\|_\infty \leq \|x^0\|_\infty$ would therefore require that $x^\dagger \leq x^0$, and by (33) also $b(\text{sat}(u^\dagger)) - b(\text{sat}(u^0)) \leq 0$. This is however impossible, because under Assumption 1, Lemma 2 states that $b(\text{sat}(u^\dagger)) - b(\text{sat}(u^0)) \leq 0$ requires $\text{sat}(u^\dagger) < \text{sat}(u^0)$. This is incompatible with (34), stating that $\mathbf{1}^\top \text{dz}(u^0) < 0$, i.e., there must exist at least one $i \in 1, \dots, n$ such that $\text{sat}_i(u^0) = L_i$, which means that necessarily $\text{sat}_i(u^\dagger) \geq \text{sat}_i(u^0)$, establishing a contradiction. We can make a symmetric argument to discard the possibility that $x^0 < 0$ and $x^\dagger \geq x^0$. Thus the only remaining option is $x^0 = x^\dagger = 0$, which requires $b(\text{sat}(u^\dagger)) = b(\text{sat}(u^0))$, and thus $u^\dagger = u^0$, contradicting the assumption of the theorem statement. This concludes the proof. \square

6. Conclusion

In this paper we considered a particular class of multi-agent systems, where the agents are connected through a capacity-constrained nonlinearity. For this type of system, we considered two proportional-integral controllers equipped with anti-windup compensation: One which was fully decentralized, and one in which the anti-windup compensator introduces a rank-1 coordinating term. We showed that the equilibria of these two closed-loop system were optimal, in the sense that they minimized the size of the control errors x in terms of the costs $\sum a_i \eta_i |x_i|$ and $\|x\|_\infty$ respectively. Additionally, we showed that the fully decentralized strategy provides guarantees of global, asymptotic stability with regards to a unique equilibrium. For the coordinating controller, we demonstrated global asymptotic stability when the disturbance can be rejected.

To demonstrate the applicability of the considered model, we showed how it can capture the indoor temperatures of buildings connected through a district heating network. In this setting, the capacity-constrained nonlinear interconnection consists of the hydraulics mapping the valve positions of each building to the resulting flow rates in the system. We demonstrated in a numerical example how the two considered controllers could then achieve different design goals - minimizing the average or the worst-case temperature deviations in the system respectively.

There are plenty of outlooks for future work: The internal agent dynamics are currently simple and on the form $-a_i x_i$. This could perhaps be extended to more complex dynamics, where some stability assumptions are placed on the dynamics of each agent. Another outlook is analyzing the transient behavior of these systems. This would include understanding the effect of slowly time-varying $w(t)$ and $b(v, t)$. In the district heating setting which we considered in this paper, this would account for changes in outdoor temperature and changes in the supply-temperature in the network. The cost-functions which are asymptotically minimized by the considered controllers could perhaps be generalized. Developments in this direction include design of controllers which maintain scalability and structure when considering other cost functions, as well as quantifying the suboptimality attained in utilizing one of the considered controllers in this paper for other cost functions. Finally, stronger stability guarantees can likely be established for the coordinating controller, which are applicable even when the stabilized equilibrium lies in the saturated domain.

7. Acknowledgments

This work is funded by the European Research Council (ERC) under the European Union's Horizon 2020 research and innovation program under grant agreement No 834142 (ScalableControl).

The authors are members of the ELLIIT Strategic Research Area at Lund University.

8. Remaining Proofs

Proof of Lemma 1. Introduce the difference between the inputs μ as $\mu = \tilde{v} - v$. We will then split the set \mathcal{I}^\pm into the two sets $\mathcal{I}^+ = \{i \in 1, \dots, n \mid \text{sign}(\mu_i) = 1\}$ and $\mathcal{I}^- = \{i \in 1, \dots, n \mid \text{sign}(\mu_i) = -1\}$. For all $i \in \mathcal{I}^+$, we can invoke Assumption 1 (i) to state that

$$b_i(v + \mu) \geq b_i(v + [\mu]_+), b_i(v) \leq b_i(v + [\mu]_-). \quad (\text{A.1})$$

Conversely, for any $j \in \mathcal{I}^-$,

$$b_j(v + \mu) \leq b_j(v + [\mu]_-), b_j(v) \geq b_j(v + [\mu]_+). \quad (\text{A.2})$$

Finally for all $k \in \mathcal{I}^0$, Assumption 1 (i) implies

$$\begin{aligned} |b_k(v + \mu) - b_k(v)| &\geq b_k(v + \mu) - b_k(v) \\ &\geq b_k(v + [\mu]_+) - b_k(v + [\mu]_-). \end{aligned} \quad (\text{A.3})$$

Hence can expand (8) as

$$\begin{aligned} &\sum_{i \in \mathcal{I}^+} \eta_i (b_i(v + \mu) - b_i(v)) \\ &+ \sum_{j \in \mathcal{I}^-} \eta_j (b_j(v) - b_j(v + \mu)) \end{aligned} \quad (\text{A.4})$$

$$\begin{aligned} &- \sum_{k \in \mathcal{I}^0} \eta_k |b_k(v + \mu) - b_k(v)| \\ &\geq \sum_i \eta_i (b_i(v + [\mu]_+) - b_i(v + [\mu]_-)) \end{aligned} \quad (\text{A.5})$$

$$= \eta^\top (b(v + [\mu]_+) - b(v + [\mu]_-)). \quad (\text{A.6})$$

□

By Assumption 1 (ii), this quantity is strictly positive when $[\mu]_+ - [\mu]_- \geq 0$, which holds for any $\mu \neq 0$, thus concluding the proof.

Proof of Lemma 2. Let $\mu = \bar{v} - \underline{v}$. Consider first the contradictory notion that $[\mu]_- \neq 0$. Then Assumption 1 (ii) yields that

$$\eta^\top (b([\mu]_- + \underline{v}) - b(\underline{v})) < 0. \quad (\text{A.7})$$

Hence $\exists k \in 1, \dots, n$ such that $b_k([\mu]_- + \underline{v}) - b_k(\underline{v})$, and thus due to 1 (i) we can conclude that $\mu_k < 0$. But since $\mu_k < 0$, we can also use Assumption 1 (i) to find

$$b_k(\mu + \underline{v}) - b_k(\underline{v}) \geq b_k([\mu]_- + \underline{v}) - b_k(\underline{v}) < 0. \quad (\text{A.8})$$

□

This contradicts the assumption of the Lemma, and hence we conclude that $[\mu]_- = 0$. Therefore $\mu \geq 0$. Now consider the notion that $\exists j$ such that $\mu_j = 0$. Then if $\mu \neq 0$, Assumption 1 (i) states that $b_j(\mu + \underline{v}) - b_j(\underline{v}) < 0$, which contradicts the lemma assumption. Hence either $\mu > 0$ or $\mu = 0$. Clearly $\mu = 0$ is impossible, as this would contradict the assumption of the lemma. Thus $\mu = \bar{v} - \underline{v} > 0$.

Proof sketch of Lemma 3.. An equilibrium (x^0, z^0) with associated stationary control input u^0 satisfies (2) with $\dot{x} = \dot{z} = 0$. Note that to solve this system of equations, it is sufficient to find u^0 which satisfies

$$0 = b_i(\text{sat}(u_i^0)) + w_i + a_i k_{A,i} dz(u_i^0), \forall i = 1, \dots, n.$$

This u^0 uniquely fixes x^0 through (2a) and z^0 through (2c). Equation (18) equates to solving a fixed-point problem $u^0 = T(u^0)$, where we define the map T as

$$T_i(u^0) = u_i^0 - \alpha (b_i(\text{sat}(u_i^0)) + w_i + a_i k_{A,i} dz(u_i^0))$$

where $\alpha > 0$ is chosen sufficiently small. T is clearly forward-invariant with respect to a sufficiently large box \mathcal{C} of size c , $\mathcal{C} = \{u \in \mathbb{R}^n \mid \|u\|_\infty \leq c\}$, which allows us to invoke Brouwer's fixed-point theorem. \square

References

- Agner, F., J. Hansson, P. Kergus, A. Rantzer, S. Tarbouriech, and L. Zaccarian (2024a). “Decentralized pi-control and anti-windup in resource sharing networks”. *European Journal of Control*, p. 101049. ISSN: 0947-3580. DOI: <https://doi.org/10.1016/j.ejcon.2024.101049>.
- Agner, F., C. M. Jensen, A. Rantzer, C. S. Kallesøe, and R. Wisniewski (2024b). “Hydraulic parameter estimation for district heating based on laboratory experiments”. *Energy* **312**, p. 133462. ISSN: 0360-5442. DOI: <https://doi.org/10.1016/j.energy.2024.133462>.
- Agner, F., P. Kergus, R. Pates, and A. Rantzer (2022). “Combating district heating bottlenecks using load control”. *Smart Energy* **6**. ISSN: 2666-9552. DOI: <https://doi.org/10.1016/j.segy.2022.100067>.
- Agner, F., P. Kergus, A. Rantzer, S. Tarbouriech, and L. Zaccarian (2023). “Anti-windup coordination strategy around a fair equilibrium in resource sharing networks”. *IEEE Control Systems Letters* **7**, pp. 2521–2526. DOI: 10.1109/LCSYS.2023.3287252.
- Agner, F. and A. Rantzer (2024). *On PI-control in capacity-limited networks*. arXiv: 2411.14077 [eess.SY].

- Bauso, D., F. Blanchini, L. Giarré, and R. Pesenti (2013). “The linear saturated decentralized strategy for constrained flow control is asymptotically optimal”. *Automatica* **49**:7, pp. 2206–2212. ISSN: 0005-1098. DOI: <https://doi.org/10.1016/j.automatica.2013.03.029>.
- Blanchini, F., C. A. Devia, G. Giordano, R. Pesenti, and F. Rosset (2022). “Fair and sparse solutions in network-decentralized flow control”. *IEEE Control Systems Letters* **6**, pp. 2984–2989. DOI: 10.1109/LCSYS.2022.3181341.
- DallAnese, E. and A. Simonetto (2018). “Optimal power flow pursuit”. *IEEE Transactions on Smart Grid* **9**:2, pp. 942–952. DOI: 10.1109/TSG.2016.2571982.
- De Persis, C., T. N. Jensen, R. Ortega, and R. Wisniewski (2014). “Output Regulation of Large-Scale Hydraulic Networks”. *IEEE Transactions on Control Systems Technology* **22**:1, pp. 238–245. ISSN: 1558-0865. DOI: 10.1109/TCST.2012.2233477.
- Galeani, S., S. Tarbouriech, M. Turner, and L. Zaccarian (2009). “A tutorial on modern anti-windup design”. *European Journal of Control* **15**:3, pp. 418–440. ISSN: 0947-3580. DOI: <https://doi.org/10.3166/ejc.15.418-440>.
- Hauswirth, A., F. Dörfler, and A. Teel (2020). “On the robust implementation of projected dynamical systems with anti-windup controllers”. In: *2020 American Control Conference (ACC)*, pp. 1286–1291. DOI: 10.23919/ACC45564.2020.9147378.
- Hauswirth, A., Z. He, S. Bolognani, G. Hug, and F. Dörfler (2024). “Optimization algorithms as robust feedback controllers”. en. *Annual Reviews in Control* **57**, p. 100941. ISSN: 13675788. DOI: 10.1016/j.arcontrol.2024.100941.
- He, Z., S. Bolognani, J. He, F. Dörfler, and X. Guan (2024). “Model-free non-linear feedback optimization”. *IEEE Transactions on Automatic Control* **69**:7, pp. 4554–4569. DOI: 10.1109/TAC.2023.3341752.
- Horn, R. A. and C. R. Johnson (1991). *Topics in Matrix Analysis*. Cambridge University Press. DOI: 10.1017/CB09780511840371.
- Jeeninga, M., J. E. Machado, M. Cucuzzella, G. Como, and J. Scherpen (2023). “On the Existence and Uniqueness of Steady State Solutions of a Class of Dynamic Hydraulic Networks via Actuator Placement”. In: *2023 62nd IEEE Conference on Decision and Control (CDC)*. ISSN: 2576-2370, pp. 3652–3657. DOI: 10.1109/CDC49753.2023.10384050.
- Kallesøe, C. S., B. K. Nielsen, A. Overgaard, and E. B. Sørensen (2019). “A distributed algorithm for auto-balancing of hydronic cooling systems”. In: *2019 IEEE Conference on Control Technology and Applications (CCTA)*, pp. 655–660. DOI: 10.1109/CCTA.2019.8920702.
- Kallesøe, C. S., B. K. Nielsen, and A. Tsouvalas (2020). “Heat balancing in cooling systems using distributed pumping”. *IFAC-PapersOnLine* **53**:2, 21st IFAC World Congress, pp. 3292–3297. ISSN: 2405-8963. DOI: <https://doi.org/10.1016/j.ifacol.2020.12.1139>.

- Kelly, F. P., A. K. Maulloo, and D. K. H. Tan (1998). “Rate control for communication networks: shadow prices, proportional fairness and stability”. en. *Journal of the Operational Research Society* **49**:3, pp. 237–252. ISSN: 0160-5682, 1476-9360. DOI: 10.1057/palgrave.jors.2600523.
- Kelly, F. (2003). “Fairness and Stability of End-to-End Congestion Control*”. en. *European Journal of Control* **9**:2-3, pp. 159–176. ISSN: 09473580. DOI: 10.3166/ejc.9.159-176.
- Krishnamoorthy, D. and S. Skogestad (2022). “Real-Time optimization as a feedback control problem A review”. en. *Computers & Chemical Engineering* **161**, p. 107723. ISSN: 00981354. DOI: 10.1016/j.compchemeng.2022.107723.
- Low, S. and D. Lapsley (1999). “Optimization flow control. I. Basic algorithm and convergence”. *IEEE/ACM Transactions on Networking* **7**:6, pp. 861–874. ISSN: 10636692. DOI: 10.1109/90.811451.
- Low, S., F. Paganini, and J. Doyle (2002). “Internet congestion control”. *IEEE Control Systems Magazine* **22**:1, pp. 28–43. ISSN: 1941-000X. DOI: 10.1109/37.980245.
- Martins, A. and K.-E. Årzén (2021). “Dynamic Management of Multiple Resources in Camera Surveillance Systems”. In: *2021 American Control Conference (ACC)*. IEEE, New Orleans, LA, USA, pp. 2061–2068. ISBN: 9781665441971. DOI: 10.23919/ACC50511.2021.9482666.
- Martins, A., M. Lindberg, M. Maggio, and K.-E. Årzén (2020). “Control-Based Resource Management for Storage of Video Streams”. en. *IFAC-PapersOnLine* **53**:2, pp. 5542–5549. ISSN: 24058963. DOI: 10.1016/j.ifacol.2020.12.1564.
- Molzahn, D. K., F. Dörfler, H. Sandberg, S. H. Low, S. Chakrabarti, R. Baldick, and J. Lavaei (2017). “A Survey of Distributed Optimization and Control Algorithms for Electric Power Systems”. *IEEE Transactions on Smart Grid* **8**:6, pp. 2941–2962. ISSN: 1949-3061. DOI: 10.1109/TSG.2017.2720471.
- Olfati-Saber, R. and R. Murray (2004). “Consensus problems in networks of agents with switching topology and time-delays”. *IEEE Transactions on Automatic Control* **49**:9, pp. 1520–1533. DOI: 10.1109/TAC.2004.834113.
- Ortmann, L., C. Rubin, A. Scozzafava, J. Lehmann, S. Bolognani, and F. Dörfler (2023). “Deployment of an Online Feedback Optimization Controller for Reactive Power Flow Optimization in a Distribution Grid”. In: *2023 IEEE PES Innovative Smart Grid Technologies Europe (ISGT EUROPE)*. IEEE, Grenoble, France, pp. 1–6. DOI: 10.1109/ISGTEUROPE56780.2023.10408057.
- Pal, A., F. Holtorf, A. Larsson, T. Loman, F. Schaefer, Q. Qu, A. Edelman, C. Rackauckas, et al. (2024). “Nonlinearsolve.jl: high-performance and robust solvers for systems of nonlinear equations in julia”. *arXiv preprint arXiv:2403.16341*.

- Rackauckas, C. and Q. Nie (2017). “DifferentialEquations.jl—a performant and feature-rich ecosystem for solving differential equations in julia”. *Journal of Open Research Software* **5**:1.
- Udell, M., K. Mohan, D. Zeng, J. Hong, S. Diamond, and S. Boyd (2014). “Convex optimization in Julia”. *SC14 Workshop on High Performance Technical Computing in Dynamic Languages*. arXiv: 1410.4821 [math-oc].

Paper V

Hydraulic Parameter Estimation for District Heating Based on Laboratory Experiments

**Felix Agner Christian Møller Jensen Anders Rantzer
Carsten Skovmose Kallesøe Rafal Wisniewski**

Abstract

In this paper we consider calibration of hydraulic models for district heating networks based on operational data. We extend previous theoretical work on the topic to handle real-world complications, namely unknown valve characteristics and hysteresis. We generate two datasets in the Smart Water Infrastructure Laboratory in Aalborg, Denmark, on which we evaluate the proposed procedure. In the first data set the system is controlled in such a way to excite all operational modes in terms of combinations of valve set-points. Here the best performing model predicted volume flow rates within roughly 5 and 10 % deviation from the mean volume flow rate for the consumer with the highest and lowest mean volume flow rates respectively. This performance was met in the majority of the operational region. In the second data set, the system was controlled in order to mimic real load curves. The model trained on this data set performed similarly well when evaluated on data in the operational range represented in the training data. However, the model performance deteriorated when evaluated on data which was not represented in the training data.

Reprinted, with permission, from F. Agner et al. (2024). “Hydraulic parameter estimation for district heating based on laboratory experiments”. *Energy* **312**, p. 133462. ISSN: 0360-5442. DOI: <https://doi.org/10.1016/j.energy.2024.133462>.

1. Introduction

The energy systems of the near-future face many challenges such as higher penetration of renewable energy production, and large and fluctuating energy demand. To face these challenges, a promising outlook is to introduce a stronger coupling between smart grids, smart thermal networks and smart gas networks. In this proposed architecture, the emerging 4th generation of district heating networks, characterized by decentralized heat production, significant integration with smart power and gas networks, and low distribution temperatures, is likely to take a central role [Lund et al., 2014]. Operating these new systems come with many challenges related to control and operation [Vandermeulen et al., 2018]. To face these new challenges, smart systems and techniques for modeling and simulating the systems will be key. While modeling and simulating the thermodynamics of district heating network for operational optimization now has a now rather long tradition [Benonysson et al., 1995], recent works have also started to delve deeper into the hydraulics of district heating networks. Hydraulic models are useful for simulating the hydraulic state of district heating networks [Vesterlund et al., 2016; Wang et al., 2020; Dénarié et al., 2023]. This information allows system operators to optimally control pumps at maximum energy efficiency [Guelpa et al., 2016; Wang et al., 2017; Wang et al., 2018; Zheng et al., 2023] or distribute heat in a fair fashion when the heat demand exceeds the available network capacity [Agner et al., 2022]. In essence, hydraulic models describe the relationship between pressures and volume flow rates for hydraulic system components such as valves and pipes. Traditionally, such models have consisted of fully white-box models. For pipes this implies models based on friction coefficients and pipe dimensions. For valves, it has implied manufacturer-provided characteristic curves. While the emerging energy systems pose many tough engineering problems, they also open the door to improving these models by tuning them based on operational data. This is due to the expected increase of smart metering in modern systems [Vandermeulen et al., 2018]. This data-driven approach holds many advantages compared to traditional methods. Data driven models can be updated to better fit the observed behavior of a component under operation and they can be adapted over time to compensate for degradation. On the contrary, gathering and maintaining correct information to build traditional white-box models for hydraulic components, e.g. pipe dimensions and friction coefficients, can be tedious and error-prone. While one could consider fully black-box models for specific purposes, there are advantages to maintaining a component-based grey-box model structure. This approach maintains physical interpretability and allows the model to be used in applications other than simulation, such as fault detection [Bahlawan et al., 2022].

The notion of data-driven model calibration has been studied longer in the field of water distribution networks, as reviewed in [Savic et al., 2009], than in the field of district heating. However, district heating networks typically have a distinct structure from water distribution networks, which can be exploited in modeling efforts:

district heating networks form closed hydraulic circuits, whereas water distribution networks end in open outlets at the points of consumption. Several works aimed particularly at modeling in the district heating domain have been published in recent years. The authors of [Wang et al., 2018] performed pipe resistance estimation and subsequently used the identified model to demonstrate effectiveness of an advanced control strategy in simulation. They estimated the resistance of all pipes in the supply-and-return-lines separately, based on measurement of pressure at the supply-and-return connections of each substation. In [Liu et al., 2020], the authors investigated the same strategy while demonstrating a way to account for meshed structures in the system by successively closing sets of loop-generating pipes. In [Luo et al., 2022], yet a third simulation study with this approach was performed, where the authors now focused on reducing the number of required measurements in the system. In [Agner et al., 2023], an extension of the framework was proposed where estimation of valve characteristics was included through exploiting the closed hydraulic circuit structure of district heating networks. A common trend among these works is that they are based on mathematical analysis of the networks along with simulations. Zheng et. al. on the other hand perform pipe parameter identification and compare their model to real operational data from a real system in Tianjin, China [Zheng et al., 2024]. This is the first and only work to the authors' knowledge that performs a study on real data. However, this study is restricted to models of pipe parameters, excluding valve models. They are also unable to share their data with the scientific community. Studying data from real systems is highly important, as many issues which are not visible in simulation can emerge. Measurement noise causes an obvious issue, but also problems based on unmodeled properties of the system. A clear example of such an aspect which is not covered in previous literature on hydraulic district heating model calibration, but which is a well-known issue to engineers working on physical systems, is valve hysteresis [Hägglund, 2023].

In this work, we perform experiments in a laboratory setup as a step towards bridging the gap between the theoretical literature and real application. The study was performed in the Smart Water Infrastructure Laboratory (SWIL) at Aalborg University [Val Ledesma et al., 2021]. We present the following contributions.

1. We build a test-bed simulating the hydraulic properties of four buildings in a line-structured district heating network.
2. We generate two data-sets. Firstly a data-set designed to excite all operational modes of the system. Secondly a data-set based on real district heating load curves. These data sets are open and available for future work.
3. We extend the modeling approach presented in [Agner et al., 2023] to handle real world complications. Namely, valve characteristics are not known a priori, and valves are subject to hysteresis.

The paper is organized in the following way. We formally introduce the problem and the considered parameter estimation framework in Section 2. We then go into

detail regarding the specific model structures we consider in Section 3. The experimental setup is described in Section 4. We present and discuss the results of our study in Section 5. We conclude the paper in Section 6 and present topics for future work in Section 7.

1.1 Notation

We denote volume flow rates (l/min) by the variable q , pressures (mH₂O) by the variable p and valve set-points by the variable v . We use subscripts to connect the measured values to a specific component, e.g. q_i is the volume flow rate measured in the component of index i . For a matrix M we denote $M_{i,j}$ to be the entry of M at row i and column j . We denote the ramp function as

$$\text{ramp}_a^b(x) = \begin{cases} 0, & \text{if } x \leq a \\ (x-a)/(b-a), & \text{if } a < x \leq b \\ 1, & \text{else} \end{cases} \quad (1)$$

where we assume that $b > a$. We denote a column vector of all 1's of size n as $\mathbf{1}_n$ and an n by m matrix of all zeros as $\mathbf{0}_{n \times m}$.

2. Hydraulic Parameter Estimation

We consider the problem of finding parameters for hydraulic models of district heating networks, based on operational data corresponding to load conditions measured at times $t = 1, \dots, T$. We restrict ourselves to networks that have a tree-shaped structure, an example of which can be seen in Figure 1. We model the supply-side network as a structured graph $\mathcal{G} = (\mathcal{E}, \mathcal{N})$ which we assume to have a tree-structure. \mathcal{E} is the set of edges which represent pipes, and \mathcal{N} is the set of nodes where these pipes connect. This graph-based perspective is a standard way of representing district heating hydraulics [De Persis and Kalløe, 2011]. The supply-and-return networks are assumed to be symmetrical and therefore the return network can be represented by the same graph \mathcal{G} , with two exceptions. Firstly, we assume that the edges have opposite direction in the supply and return networks. Hence as the water flows out through the supply network and returns via the return network, the sign of the flow rates through these edges will be the same. Secondly, we consider one node in the graph to be the root of the tree, which we denote α in the supply network and β in the return network. This root can represent e.g. a production facility, or a connection to a larger network to which the considered network is a small sub-network. A subset $\mathcal{V} \subset \mathcal{N}$ of the nodes correspond to consumers. They connect the supply and return lines via control valves. In the considered example, $\mathcal{V} = \{1, 2, 3, 4\}$. In a real setting, a single consumer may in fact be represented by two (or more) control valves, as domestic hot water and space heating are typically hydraulically separated [Frederiksen and Werner, 2013, p.365]. For each valve $i \in \mathcal{V}$, there is a




Figure 1. A schematic example of a tree-structured district heating network. The nodes $\mathcal{N} = \{\alpha, 1, \dots, 7\}$ are connected by edges $\mathcal{E} = \{1, \dots, 7\}$. The supply-and-return-networks are symmetrical except for their edges having opposite direction. Valves $\mathcal{V} = \{1, \dots, 4\}$ connect the supply and return lines. Root nodes α and β in the supply-and-return networks respectively connect the branch either to a production facility, or a greater part of the network. The path $\mathcal{L}_2 = \{2, 5, 6\}$ highlighted in green contains the edges leading from α to β via valve 2.

set \mathcal{L}_i of edges corresponding to a path from α to β through valve i . For instance, $\mathcal{L}_2 = \{2, 5, 6\}$ in our example as highlighted in Figure 1. For the considered parameter estimation protocol, we assume that we can measure the root pressure difference $\Delta p_0 = p_\alpha - p_\beta$. We additionally assume that we can log the control valve set-points v_i for each consumer and the volume flow rates q_i through each valve $i \in \mathcal{V}$.

The networks are also assumed to have the following property.

ASSUMPTION 1

For all $i \in \mathcal{N}$, either $i \in \mathcal{V}$ or i is connected to at least 3 edges.

In practice, this assumption means that no pipes in the model are placed in direct series. Any physical interconnection of pipes in direct series would be represented by one single, equivalent pipe in the model. For instance, it would be reasonable to expect that a consumer substation has a physical interconnection of several pipe sections in series with the control valve and a heat exchanger. The heat exchanger and the pipes would then be concatenated into one single pipe in the model.

With the considered model in mind, we have two types of hydraulic components; *pipes* and *valves*. For each such component, we aim to parameterize a model which maps the volume flow rate through the component to the pressure difference.

2.1 Pipe Parameterization

For pipes, we consider models stemming from the Darcy-Weisbach equation. If the pipe corresponds to edge i , which leads from node j to node k , the relation between the pressure loss $\Delta p_i = p_j - p_k$ and the volume flow rate q_i along the direction of the pipe is given by

$$\Delta p_i = s_i q_i |q_i|^\gamma. \quad (2)$$

Here s_i is the model parameter which our framework aims to estimate from data, typically referred to as the pipes hydraulic resistance. γ is an exponent which in the literature typically takes the value 1 or slightly lower e.g. 0.87. We limit the scope of this study to consider the assumption $\gamma = 1$. We make the following assumption regarding the parameters s_i .

ASSUMPTION 2

The pipe resistance s_i is equal for the corresponding supply-and-return-lines.

It was shown in [Agner et al., 2023] that this assumption along with Assumption 1 is sufficient for unique identification of model parameters. In practice, it is unlikely in that this assumption holds exactly due to slight differences in the pipes. However, we cannot expect to find unique parameters for these pipes with our methodology, given that we only measure pressure at two points in the network (α and β). However, for many model applications such as estimating volume flow rates (see Appendix 7), only the sum of the parameters for the supply-and-return-lines is necessary.

2.2 Valve Parameterization

The second and final hydraulic model component type we consider is valves. For valve i which connects node i in the supply network to the return network, Δp_i is the pressure difference between these nodes. Δp_i is given by the volume flow rate q_i and the valve set-point v_i which varies from fully closed ($v_i = 0$) to fully open ($v_i = 1$) in the following way.

$$q_i = K_{v,i} k_i(v_i) \sqrt{\Delta p_i} \implies \Delta p_i = \frac{1}{K_{v,i}^2 k_i(v_i)^2} q_i^2 \quad (3)$$

where $K_{v,i}$ represents the hydraulic admittance of the valve when it is fully open. This is often referred to as the flow coefficient of the valve, or the k_{vS} or c_{vS} value in Nordic or English literature respectively [Frederiksen and Werner, 2013, p. 399]. k_i represents the valve characteristics, i.e. the mapping between the valve set-point and the valve admittance. The valve characteristic function may be linear, e.g. $k_i(v_i) = v_i$. It may however also have other shapes. One common valve characteristic is quick opening, where at low set-points, a small increase in valve set-point will lead to a large increase in flow-through. Another common characteristic is equal percentage which is in a sense the opposite of quick opening. Here it is expected that the valve needs to provide a larger change in volume flow rate in response to small changes in set-point at higher levels of opening [Hägglund, 2023, pp. 82-83]. Additionally, one might consider valve characteristics where the valve fully closes at a value $v > 0$ and fully opens at a value $v < 1$, to ensure that the full capacity of the valve is used at high set-points, and that the valve is fully closed at near-0 set-points. Figure 2 shows three examples of $k_i(\cdot)$, corresponding to a linear valve ($k_i(v_i) = v_i$), an equal percentage valve which fully closes at 0.1 and fully opens




Figure 2. Examples of valve characteristics $k_i(v_i)$.

at 0.8 ($k_i(v_i) = \text{ramp}_{0.1}^{0.8}(v_i)^2$), and a quick-opening valve which fully closes at 0.2 and fully opens at 0.9 ($k_i(v_i) = \text{ramp}_{0.2}^{0.9}(v_i)^{0.5}$). As we cannot necessarily know the valve characteristics of all of the valves we aim to parameterize, we propose a linear combination of typical characteristics, aimed to capture (3):

$$\Delta p_i = \left(\frac{\theta_{i,1}}{k_1(v_i)^2} + \frac{\theta_{i,2}}{k_2(v_i)^2} + \cdots + \frac{\theta_{i,K}}{k_K(v_i)^2} \right) q_i^2. \quad (4)$$

Here $k_k(\cdot)$, $k = 1, \dots, K$ correspond to different possible valve characteristics and $\theta_{i,k}$ corresponds to a tunable parameter for valve i in relation to the valve characteristic $k_k(\cdot)$. Somehow $1/k_k(\cdot)^2$ are basis of all possible characteristics of a valve in a K -dimensional space. Hence while the actual valve characteristic may not be perfectly captured by any of the functions $k_k(\cdot)$, a linear combination of them may be able to replicate the actual behavior of the valve.

2.3 Parameter Estimation

To estimate the values of s_i (for pipes) and $\theta_{i,k}$ (for valves), we assume that we have access to data from T different steady state load conditions. We first consider one such load condition at time t . Here we measure the pressure difference at the tree root $\Delta p_0(t) = p_\alpha(t) - p_\beta(t)$, the valve flow rates $q_i(t)$, and valve set-points $v_i(t)$ where $i \in \mathcal{V}$.

We can use the valve volume flow rates $q_i(t)$ to calculate the volume flow rates in all of the pipes. To do so, we first define the basic incidence matrix $B \in \mathbb{R}^{(n_V-1) \times n_E}$ for our considered graph representation of the network \mathcal{G} . Here n_E and n_V are the

number of edges and nodes respectively. The elements of B are given by

$$B_{i,j} = \begin{cases} 1, & \text{if edge } j \text{ leads to node } i, \\ -1, & \text{if edge } j \text{ leads from node } i, \\ 0, & \text{else.} \end{cases} \quad (5)$$

We omit the row corresponding to the root node α or β , which together with the tree-structure of \mathcal{G} ensures the invertibility of B . Hence we can find the vector $q_{\mathcal{E}}$ corresponding to the vector of all flow rates in \mathcal{E} as

$$q_{\mathcal{E}}(t) = B^{-1} q_{\mathcal{N}}(t) \quad (6)$$

where $q_{\mathcal{N}}(t) \in \mathbb{R}^{n_{\mathcal{N}}-1}$ corresponds to the total volume flow rate out of each node $i \in \mathcal{N} \setminus \alpha$. I.e.

$$q_{\mathcal{N},i}(t) = \begin{cases} q_i(t), & \text{if } i \in \mathcal{V}, \\ 0, & \text{else.} \end{cases} \quad (7)$$

With this information, the volume flow rates through all pipes and valves are known. We can now consider one $i \in \mathcal{V}$, and the corresponding path \mathcal{L}_i from α to β through valve i . It must hold that the total pressure difference $\Delta p_0(t)$ between α and β corresponds to the pressure drop along each traversed component in \mathcal{L}_i , i.e.

$$\begin{aligned} \Delta p_0(t) &= \Delta p_i(t) + 2 \sum_{j \in \mathcal{L}_i} \Delta p_j(t) \\ &= \left(\frac{\theta_{i,1}}{k_1(v_i(t))^2} + \frac{\theta_{i,2}}{k_2(v_i(t))^2} + \dots + \frac{\theta_{i,K}}{k_K(v_i(t))^2} \right) q_i(t)^2 + 2 \sum_{j \in \mathcal{L}_i} s_j q_j(t) |q_j(t)|. \end{aligned} \quad (8)$$

Here $\Delta p_j(t)$ is equal for the edges in the supply-and-return networks due to Assumption 2 and symmetry, yielding the factor 2 in front of the sum. We can use all paths \mathcal{L}_i where $i = 1, \dots, n_{\mathcal{V}}$ to formulate a system of equations

$$\mathbf{1}_{n_{\mathcal{V}}} \Delta p_0(t) = \begin{bmatrix} F_1(t) & \mathbf{0}_{1 \times K} & \dots & \mathbf{0}_{1 \times K} \\ \mathbf{0}_{1 \times K} & F_2(t) & \dots & \mathbf{0}_{1 \times K} \\ \vdots & \vdots & \ddots & \vdots \\ \mathbf{0}_{1 \times K} & \mathbf{0}_{1 \times K} & \dots & F_{n_{\mathcal{V}}}(t) \end{bmatrix} \begin{bmatrix} \theta_1 \\ \theta_2 \\ \vdots \\ \theta_{n_{\mathcal{V}}} \end{bmatrix} + G(t)s = F(t)\theta + G(t)s. \quad (9)$$

Here $G(t) \in \mathbb{R}^{n_{\mathcal{L}} \times n_{\mathcal{E}}}$ is a matrix defined as

$$G_{i,j} = \begin{cases} 2q_j(t)|q_j(t)| & \text{if edge } j \in \mathcal{L}_i, \\ 0 & \text{else.} \end{cases} \quad (10)$$

The vectors θ_i for $i = 1, \dots, n_{\mathcal{V}}$ gather the parameters $\theta_{i,k}$ for $k = 1, \dots, K$, and $F_i(t) \in \mathbb{R}^{1 \times K}$ is given by

$$F_i(t) = \left[\frac{q_i(t)^2}{k_1(v_i(t))^2} \quad \frac{q_i(t)^2}{k_2(v_i(t))^2} \quad \dots \quad \frac{q_i(t)^2}{k_K(v_i(t))^2} \right]. \quad (11)$$

We now move from considering only one load condition at time t to concatenating the data from all T load conditions into one large system of equations

$$\Phi \begin{bmatrix} \theta \\ s \end{bmatrix} = y \quad (12)$$

where

$$\Phi = \begin{bmatrix} F(1) & G(1) \\ F(2) & G(2) \\ \vdots & \vdots \\ F(T) & G(T) \end{bmatrix}, \quad y = \begin{bmatrix} \mathbf{1}_{n_v} \Delta p_0(1) \\ \mathbf{1}_{n_v} \Delta p_0(2) \\ \vdots \\ \mathbf{1}_{n_v} \Delta p_0(T) \end{bmatrix}. \quad (13)$$

This way Φ is a data-matrix constructed by measuring flow rates and valve set-points and y consists of measured root pressure measurements. Once Φ and y are constructed, there are many ways to choose s and θ to fit the observed data. We choose to fit the parameters numerically as the solution to

$$\underset{s, \theta}{\text{minimize}} \quad \left\| \Phi \begin{bmatrix} s \\ \theta \end{bmatrix} - y \right\|_1 \quad (14a)$$

$$\text{subject to} \quad s \geq 0, \quad (14b)$$

$$\theta \geq 0. \quad (14c)$$

The choice of using $\|\cdot\|_1$ in (14a) is based on the general robustness to outliers typically provided by this cost function. The positivity constraints (14b)–(14c) enforce physical feasibility of the model. We solved (14) using the python toolbox CVXPY [Diamond and Boyd, 2016; Agrawal et al., 2018] using the SciPy solver. A natural extension of (14) is to introduce parameter regularization to (14a). This would reduce the risk of overfitting and promote sparseness in the model depending on the choice of regularization employed. Our investigation has however showed that the models we found are already sparse, and therefore leave an investigation of the impact of regularization outside the scope of this paper.

3. Model Structures and Data Preprocessing

In this section we detail three different model structures which we subsequently compare on our experimental data:

A: A simple model with a linear valve curve assumption.

B: Parameterized, nonlinear valve curves.

C: The same valve curves as B, but including data preprocessing for hysteresis compensation.

3.1 Model A: Naive Valve Curve Parameterization

Our first and most simple model equates to the model considered in [Agner et al., 2023]. Here we assume that all valves are fully linear and thus can be parameterized by only one function $k(\cdot)$:

$$k(v_i) = v_i. \quad (15)$$

3.2 Model B: Enhanced Valve Curve Parameterization

To account for nonlinear valve characteristics, we consider also a larger number of parameterization functions $k_k(\cdot)$. To make an educated guess about reasonable valve curve shapes, Figure 3 shows a scatter plot of $q_1(t)$ over $v_1(t)$ for valve 1 in the experimental setup to be described later. While this scatter-plot does not account for changes in differential pressure over the valve, we can clearly see that a linear curve from 0 to 1 does not capture this shape. Instead, it appears that the curve should have a sub-linear curve corresponding more to equal percentage characteristics. Additionally, the valve is clearly fully open already around $v_1 \approx 0.9$, and fully closed around $v_1 \approx 0.2$. To handle these nonlinear patterns, we consider valve characteristic functions on the form

$$k_k(v_i) = \text{ramp}_{a_k}^{b_k}(v_i)^{c_k}. \quad (16)$$

for different choices of a_k , b_k and c_k . Based on our intuition from Figure 3, we considered all such combinations corresponding to

$$a_k \in [0.10, 0.15, 0.20, 0.25], \quad (17)$$

$$b_k \in [0.80, 0.85, 0.90, 0.95, 1.0], \quad (18)$$

$$c_k \in [1.0, 1.25, 1.5], \quad (19)$$




Figure 3. Scatter-plot of flow rates $q_1(t)$ against valve set-points $v_1(t)$ for valve 1 in the experiment setup to be described in the subsequent section.

resulting in $K = 60$ different valve parameterization functions $k_k(\cdot)$ for each valve. We chose the values for a_k and b_k as heuristically reasonable values for full opening and closing, and the values of c_k as values corresponding to reasonable curves for equal percentage valves, which was deemed appropriate based on Figure 3.

3.3 Model C: Hysteresis Compensation

Our third and most sophisticated model uses the same valve characterization as model B. Model C however accounts for the effect of hysteresis. This is an effect which typically manifests in control valves, causing an issue where the valve spindle does not move for small changes in the set-point given to the valve [Häggglund, 2023, p. 88]. We account for this issue in the following way. Instead of using the valve set-points $v_i(t)$ for model fitting and validation, we utilize filtered versions $\hat{v}_i(t)$ which are meant to estimate the actual valve spindle position. We first define a hand-tuned parameter $\delta = 0.015$. The value 0.015 was found heuristically by testing the values 0.05, 0.10, 0.15, 0.20 and then choosing the best one. The spindle is assumed to remain still whenever the set-point moves a distance less than δ from the current estimated spindle position. Once the reference moves outside of this range, the estimated spindle position will lag behind by a distance δ . These assumptions are captured in the following filter:

$$\hat{v}_i(t) = \begin{cases} \hat{v}_i(t-1) & \text{if } |v_i(t) - \hat{v}_i(t-1)| \leq \delta \\ v_i(t) - \delta & \text{if } v_i(t) \geq \hat{v}_i(t-1) + \delta \\ v_i(t) + \delta & \text{else.} \end{cases} \quad (20)$$

The filtered time series are initiated with the first measured valve set-point $\hat{v}_i(0) = v_i(0)$.

4. Experiment Description

Our experiment setup represents four domestic buildings connected through a line-structured district heating network as seen in Figure 1. This is the same structure that we used as an example in Section 2. The four consumers are connected to a single production facility. One could also imagine that the root node in the laboratory setup represents a connection between the sub-branch of these four consumer to a larger network. We will now describe the laboratory equipment used to represent the production facility, the pumping, and the consumers.

4.1 Equipment

The experiments were conducted in Aalborg University's Smart Water Infrastructure Lab (SWIL). SWIL is a state-of-the-art water and district heating laboratory that consists of configurable modules which can be used to build a wide variety of experimental setups [Val Ledesma et al., 2021]. The available module types are:

- Pumping stations
- Consumer stations
- Piping stations
- Heating stations

Examples of consumer and piping stations are seen in Figure 5. Pumping and consumer stations are particularly configurable and may be used to emulate e.g. gravity sewers or elevated reservoirs. The experimental network corresponding to the schematic Figure 1 consists of a pumping station, two piping stations, and two consumer stations as seen on Figure 4. The consumer stations represent two consumers (buildings) each. This corresponds to internal piping, a control valve, and a heat exchanger in series connection.

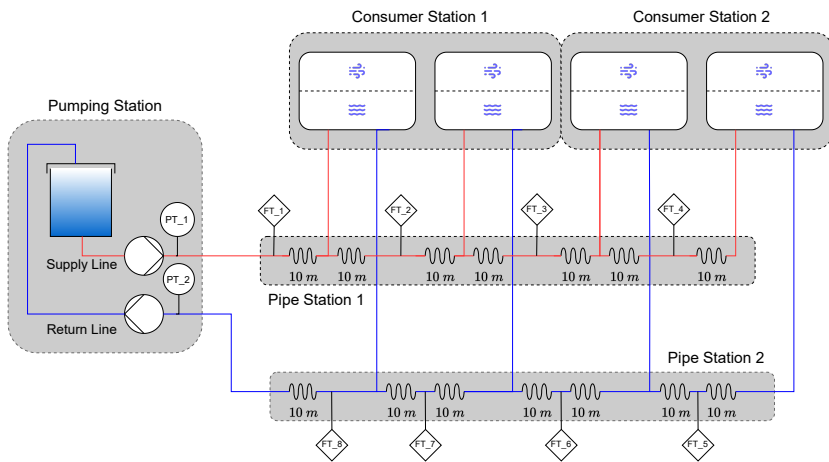


Figure 4. Diagram of the experimental setup used to recreate Figure 1. PT and FT denote respectively flow and pressure transmitters, while the two blocks in each consumer station are water-to-air heat exchangers. Each heat exchanger is connected in series with a control valve used to regulate the volume flow rate.

4.2 Measurements

Flows and pressure were measured using the piping station sensors, respectively an Endress & Hauser Proline Promag 10 and Grundfos Direct Sensor RPI+T 0 – 1.6, connected to Beckhoff I/O modules. The data was collected at 1 Hz and transmitted from the laboratory modules to a central control unit (CCU). The laboratory modules are controlled by soft PLCs comprising CodeSys software running on a Raspberry Pi. The CCU communicates with the modules over Modbus TCP through a



(a) On the left are two pipe stations. These stations measure volume flow rates and pressures, and can be configured for pipe sections of varying lengths. On the right is a pumping station, where water is pumped from a water reservoir into the system, until finally returning to the reservoir.

(b) Two consumer stations, which in total simulate four consumers. The bottom sections are connected to the pipe units. Water then flows up and through the water-to-air heat exchangers at the top, via a control valve.

Figure 5. Images from the laboratory setup.

Simulink simulation using the MATLAB Industrial Communications Toolbox and the Simulink Real-Time Pacer [Vallabha, 2016], and raw measurements are converted to respectively l/min and mH₂O at this stage. The consumer volume flow rates q_1, \dots, q_4 were subsequently calculated as

$$q_1 = q_{FT_1} - q_{FT_2}, \quad (21)$$

$$q_2 = q_{FT_2} - q_{FT_3}, \quad (22)$$

$$q_3 = q_{FT_3} - q_{FT_4}, \quad (23)$$

$$q_4 = q_{FT_4} \quad (24)$$

where e.g. q_{FT_1} is the volume flow rate measured in sensor FT_1. The pressure measurements PT_1 and PT_2 were used to represent the network root pressures p_α and p_β respectively.

4.3 Data Sets

We generated two experimental datasets, henceforth referred to as the *exciting* and *realistic* data sets.

In the *exciting* data set, we kept the pump working at full capacity and set the valve set-points randomly to individually and uniformly sampled values in the interval $[0.3, 1.0]$ every 40 seconds. The lower bound of 0.3 corresponds to valves which are in practice almost fully closed. This bound was included to avoid that a large




Figure 6. Illustration of the method used to generate the *exciting* data set. Valve set-points are randomly generated every 40 seconds. The subsequent 10 seconds of data is discarded to remove transient effects, and the mean of the remaining 30 seconds of data is logged as a data point. Valve set-points are lower-bounded by 0.3.

part of the data corresponds to closed valves. During the full experiment, the pump was operating constantly at full capacity. In post-processing, the first 10 seconds in each 40 second interval was removed in order to ignore the transient phase of each step. The mean of remaining 30 seconds was used to generate a data point for later model fitting and validation. This process is shown in Figure 6.

The *exciting* data set was produced to give us the best possible chance of accurately estimating the system parameters, but would be hard to replicate in a real district heating network. In the *realistic* data set, we instead designed the system to emulate the recorded volume flow rates in four residential buildings in Nuremberg, Germany, from 2022. We used two weeks of hourly mean values. We scaled the magnitude of the recorded flow rates such that they would fit in the operational range of our experimental setup. In the experiment, this data served as the reference volume flow rate for each consumer, which we tracked using the control valves of the consumer units. Each control valve was equipped with hand-tuned PID controller for this purpose. The pump was operating constantly at full capacity during the experiment to ensure sufficient system pressurization. For the purpose of experiment tractability, we scaled down the time such that 1 hour of reference data corresponded to 2 minutes in the experiment. The resulting *exciting* volume flow rates are shown in Figure 7. We can see that the consumption of these consumers follows a rather typical pattern. There are slow dynamics, governed by the demand for space heating driven by the outdoor temperature, along with faster intra-day peaks driven by domestic hot water consumption. We used mean values of 30 seconds of experimental data for the subsequent model fitting and validation which corresponds to 15 minute mean value measurements in the real district heating setting. The reference volume flow rates correspond to the combined load for domestic hot water and space heating. In a real system, these loads are often separated into two different hydraulic loops controlled by separate valves [Frederiksen and Werner, 2013, p.365], but such a separation would be infeasible in our small laboratory setup.




Figure 7. Flow rates $q_i(t)$ in the realistic data set, recorded during the experiment. 48 minutes corresponds to 24 hours of real world time, in which we can see typical daily consumption patterns.

4.4 Model Evaluation

We choose to evaluate the predictive power of the model from a given set of control inputs (valve set-points v_i and differential pressure Δp_0 provided by the pump) to system outputs (flow rates through the four control valves representing consumers). Appendix A describes how to use the parameterized models to calculate estimated volume flow rates \hat{q}_i in this way. We then investigate the prediction errors $e_i = q_i - \hat{q}_i$ to judge the model performance. This is a metric which is relevant when using the model to design flow rate controllers. We can note that this is not the objective for which the parameters are tuned. The training objective (14a) rather reflects the ability of the model to predict pressure drops given the valve set-points and flow rates in the system. The component-wise modeling of the network is what allows this flexibility in model application. This multi-purpose capacity is testimony to the value of considering grey-box component-models of this type over black-box models such as neural networks or regression models without structure considerations. There are other relevant choices of evaluation that could also be considered. In [Agner et al., 2023], ground-truth parameter value comparison was used. This is however clearly infeasible in the real setting where no ground truth is available.

5. Results

The following section details the results of the experiments. Firstly the model structures A, B and C trained on the exciting data set are investigated. Afterwards we evaluate the best performing model structure C trained on the realistic data set.

5.1 Exciting training data

Table 1 shows the resulting pipe parameters s_i using the different models. Pipes 1-4 lead to only one of the four consumers, whereas pipes 5-7 are shared between the

consumers in the sense that the e.g. the volume flow rate for all four consumers comes through pipe 5. Of interest is that the resistances s_1 , s_2 and s_3 were estimated as 0 for all three models trained on the exciting data set. One explanation for this is that our data does not represent the section of the valve curve where the valve is fully closed. To see why this matters, consider Figure 8a. Here we see a hydraulic schematic representing one consumer in the network, i.e. two symmetric pipes (here represented by a total resistance s) and a valve (with curve $K_v k(v)$). The relationship between $\Delta p = p_i - p_j$, i.e. the pressure difference between supply-and-return-lines where the consumer is connected, and the flow rate q for the consumer will be

$$\Delta p = \left(s + \frac{1}{K_v^2 k(v)^2} \right) q^2 \quad (25)$$

and therefore


$$q = \frac{K_v k(v)}{\sqrt{s K_v^2 k(v)^2 + 1}} \sqrt{\Delta p}. \quad (26)$$

We can denote $\hat{k}(v) = \frac{K_v k(v)}{\sqrt{s K_v^2 k(v)^2 + 1}}$. Consider two such choices of $\hat{k}(v)$: Firstly, $\hat{k}_1(v)$ where $s = 1.5$, $k(v) = \text{ramp}_{0.1}^{0.9}(v)^{1.5}$ and $K_v = 1$. Secondly, $\hat{k}_2(v)$ with $s = 0$ (no pipe), $k(v) = \text{ramp}_{0.22}^{0.9}(v)^{0.75}$, $K_v = 0.63$. Figure 8b shows the plots of $\hat{k}_1(v)$ and $\hat{k}_2(v)$. Here we see that if we consider only valve-positions above $v = 0.3$, the curves are almost equivalent even though one of them has no pipe resistance. In the same way, due to the fact that we chose to omit valve openings below 0.3 in the training data, the model cannot clearly discriminate between the valve curve parameters and the pipe parameters. This is an important realization, demonstrating that effective data sets should include also valve set-points in the lower range of operation. There are a couple of possible explanations as to why $s_4 \neq 0$ for models B and C while the other three are 0. Firstly pipes 1-3 should be equivalent due to the experiment configuration, with pipe 4 being longer than the others and thus more resistive. Secondly, since consumer 4 is the furthest from the pump, it also experiences the lowest differential pressure, which causes the volume flow rate to occasionally drop to 0. This could influence the estimation procedure towards $s_4 \neq 0$. Pipes 5, 6 and 7 do not experience the same issue as discussed with regards to Figure 8, most likely because they are not in strict series with a valve. Pipes 6 and 7 should be similar in size. However, in models B and C, pipe 7 appears roughly a factor 5 longer than pipe 6, indicating that either the model does not capture the contribution of the pipes fully, or there are other reasons such as small twists and turns in the piping causing a difference between pipes 5 and 6.

The resulting valve curve models become very sparse, an effect likely caused by the positivity constraint (14c) which prevents overfitting. Even though there are in total $K = 60$ parameters $\theta_{i,k}$ for each valve i , resulting in 240 total valve curve parameters, only one of the valve curves has more than 5 non-zero parameters $\theta_{i,k}$. In total, model B has 18 parameters $\theta_{i,k} > 0$ and model C has 16. This model sparsity

Model	Data	s_1	s_2	s_3	s_4	s_5	s_6	s_7
A	E	0.0	0.0	0.0	0.0	0.0089	0.00082	0.021
B	E	0.0	0.0	0.0	0.00067	0.0039	0.0046	0.029
C	E	0.0	0.0	0.0	0.015	0.0038	0.0045	0.029
C	R	0.0	0.0	0.0	0.0	0.0044	0.0	0.049

Table 1. Pipe parameters estimated for each model structure on either the exciting (E) or realistic (R) dataset.



(a) Hydraulic schematic of one consumer in the district heating network, represented by two symmetrical pipes with combined resistance s and valve curve $K_v k(v)$.

(b) Plots of $\hat{k}_1(v)$ and $\hat{k}_2(v)$ which relate the flow rate q and the valve position v when $p_i - p_j = 1$.

Figure 8. An illustration of a consumer connected to the network, represented by two symmetrical pipes and a valve. For a large portion of the valve curve, two different models of the hydraulic admittance of this consumer could have very similar properties. Here $\hat{k}_1(v)$ represents a pipe and a valve, but $\hat{k}_2(v)$ represents only a valve with no pipe. The difference between the two models becomes evident only when the valve is almost closed.

motivates the decision to omit parameter regularization from the scope of this paper. The full equation for each valve curve and each model as per (3) is provided in Appendix 7. We can also investigate the valve characteristics visually, which yields Figure 9. As all four consumers are constructed with equal equipment, we would expect the four valve characteristic curves to be equivalent. However, this is not the case. Rather, valves 1 and 2 have similar characteristics, whereas valves 3 and 4 have much lower curves. This means that given equal differential pressure and valve set-points, the model predicts that valve 4 would yield a lower volume flow rate than the other valves. One possible explanation for this behavior is that the exponent γ in (2) and q_i^2 in (3) do not perfectly capture the real behavior of the system. If these exponents are too high, it means that the expected pressure drops at high flow rates will be over-estimated. Clearly, the volume flow rates and valve openings are strongly correlated (see Figure 3), hence the model can counter-act




Figure 9. Valve curves $K_{v,i}k_i(v_i)$ for each valve i under each model structure A, B and C trained on the exciting data set.

this effect by increasing the admittance at high valve set-point values. The reason this could make valve 4 stand out from the others is that the average volume flow rates over valve 4 are much lower than the other valves (the mean volume flow rates are $\bar{q}_1 = 4.60$ l/min, $\bar{q}_2 = 4.34$ l/min, $\bar{q}_3 = 2.91$ l/min, and $\bar{q}_4 = 2.11$ l/min respectively) and thus the model does not need to over-estimate the admittance at high flow rates to the same extent.

Figure 10 shows the prediction errors $e_i(t) = q_i(t) - \hat{q}_i(t)$ plotted over the corresponding valve set-points $v_i(t)$. Each column represents one model structure A, B or C, and each row represents one valve index $i = 1, \dots, 4$. Figure 10 shows that A has a strong correlation between prediction errors and valve set-points, which means that this naive parameterization of the valve curves is clearly insufficient. Both models B and C have a much lower correlation between valve set-points and prediction errors, with a few exceptions. Valves 1-3 display strong correlation between e_i and v_i around $v_i \approx 0.85$ which is the inflection point where the valve enters the fully closed position. It is possible that further tuning of the chosen parameterization $k_k(\cdot)$ could remove this issue. Valve 4 shows a cluster of outlier errors around $v_4 \approx 0.3$. We can find an explanation in Figure 11 where we see the flow rate $q_4(t)$ scattered over the valve set-point $v_4(t)$ in the exciting data set. We can

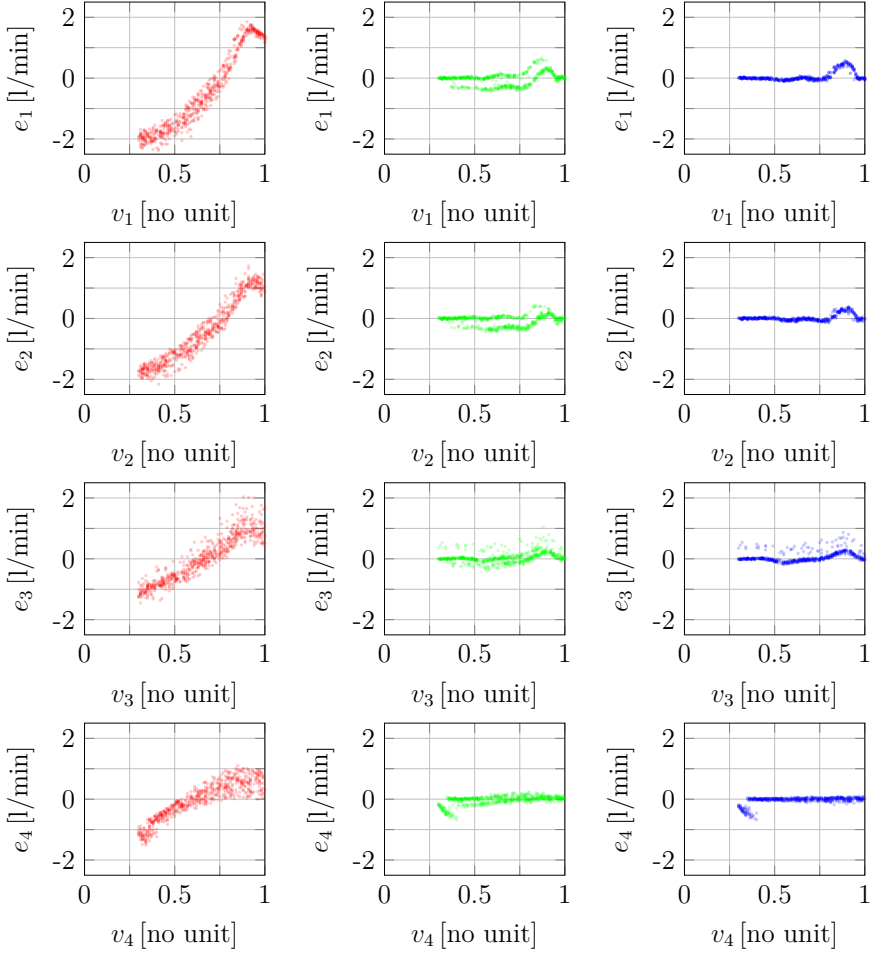


Figure 10. Prediction errors $e_i = q_i - \hat{q}_i$ plotted over the valve set-point v_i for each valve i and each model structure A, B and C trained on the exciting data set. The plots are arranged as valve 1-4 from top to bottom and model A (red), B (green), and C (blue) from left to right.




Figure 11. Scatter-plot of flow rates $q_4(t)$ against valve set-points $v_4(t)$ in the exciting data set.

note that there are barely any recorded volume flow rates between 0 and 0.3 l/min. This means that when the valve is almost closed and the differential pressure over the valve is low, there is a nonlinear behavior where the volume flow rate sharply cuts from 0.3 to 0, which is clearly not captured in the model. Valve 4 is the valve which is the furthest from the pump and hence experiences the lowest differential pressure. Hence the volume flow rate in the other valves never becomes sufficiently low to exhibit this behavior. Finally there are outliers in the prediction errors e_3 . These outliers are likely caused by measurement errors which can occur due to e.g. air bubbles forming in the system.

One difference between models B and C is that the prediction errors for model B appear segmented in one upper and one lower cluster. This is most clearly visible in valves 1 and 2. In Figure 12 we show an enlarged view of these plots. Here we highlight the data points where the change in valve set-point is positive in red ($v_i(t) - v_i(t-1) > 0$), and negative in blue. The two clusters are clearly separated by the movement direction of the valve set-point, indicating the effect of hysteresis. In model C, this effect is no longer visible, demonstrating the strength of the hysteresis compensation.

In total, model C is the strongest model when trained and tested on the exciting data set. Histograms of the prediction errors $e_i(t)$ for this model are shown in Figure 13. All four error distributions have large tails. In Figure 10 we saw that these tails arise due to not fully capturing the valve characteristics for valves 1 and 2 and the data outliers of valves 3 and 4. Outside of these tails, a vast majority of the errors lie in the range $[-0.2, 0.2]$ l/min, which can be compared with the mean flow rates \bar{q}_i exhibited in the system, given by $\bar{q}_1 = 4.60$ l/min, $\bar{q}_2 = 4.34$ l/min, $\bar{q}_3 = 2.91$ l/min, and $\bar{q}_4 = 2.11$ l/min respectively in the exciting data set. Hence in the well-tuned regions of the valve curve, errors remain within 5 or 10 % of the mean flow rates for valves 1-2 and 3-4 respectively.

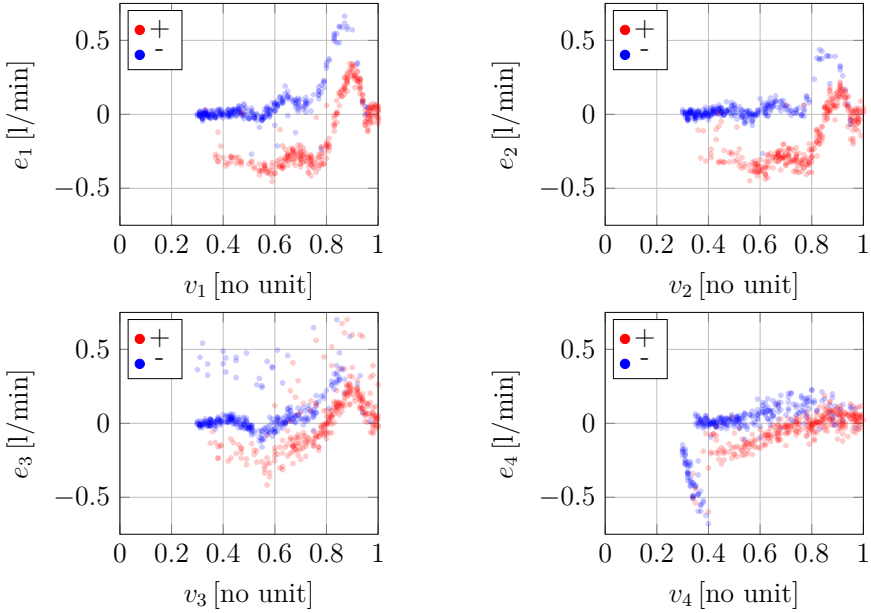


Figure 12. Prediction errors $e_i = q_i - \hat{q}_i$ for each valve i using model B trained on the exciting data set, scattered over the valve set-points v_i . The points where the valve opened since the last data point ($v_i(t) - v_i(t-1) > 0$) are shown in red, and the points where the valve closed are shown in blue.

5.2 Realistic training data

Figure 14 shows the evaluated valve curves gained from training model structure C on the exciting or realistic data set respectively. We can note that in the lower regions of the valve curves, there is a significant overlap between the two curves. However, above a valve set-point of roughly 0.6, we see the curves starting to deviate. This is likely due to the fact that the model trained on the realistic data set has not seen any data within this region. This is also reflected in Figure 15 which shows the prediction errors of using model C trained on the realistic data set. The model behaves rather well when evaluated on the realistic data set and maintains the errors within about 0.3 l/min. However, the realistic data only represents valve set-points in limited intervals. The green bar in the figures highlights the interval between the 5th and 95th quantiles of the valve set-points in the training data. The model is unable to extrapolate beyond the limited intervals of valve set-points represented in the training data which is reflected in the plots on the right of Figure 15. Here the model is evaluated on the exciting data set, which covers a larger range of valve operation. In the region of the valve set-points where the model was trained, prediction errors are relatively small. Outside of these intervals however, the models behave extremely poorly. This is not surprising, as no such data was used in train-

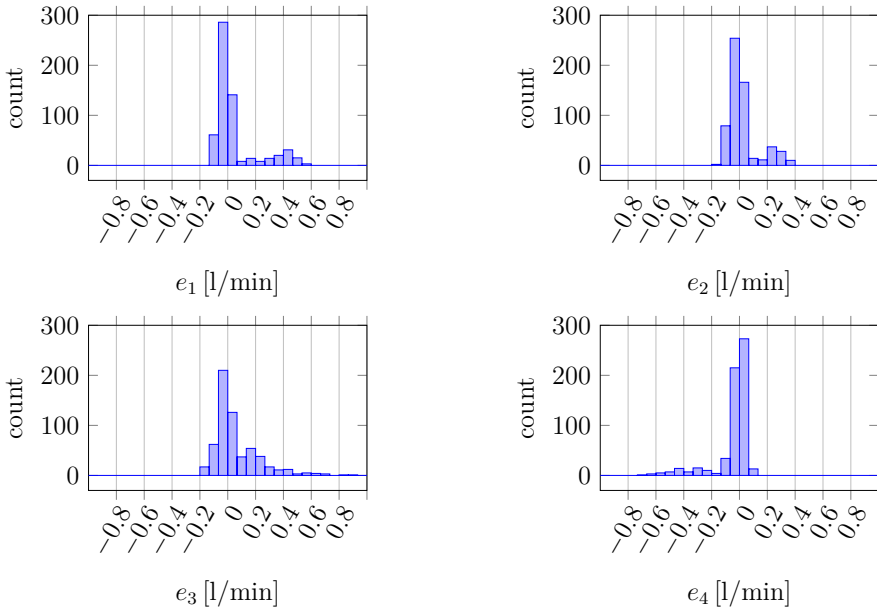


Figure 13. Histograms of prediction errors $e_i = q_i - \hat{q}_i$ using model structure C trained on the exciting data set.

ing. While it may be difficult to guarantee that the full operational range of each valve is represented in data from a real system, it may also be unlikely that the operational range present in the data is quite so restrictive as in our example. Firstly, if such a small range of a valves operational ranges is used, this would imply that the valve may be poorly dimensioned for the building in which it is installed. Secondly, we are only using two weeks worth of representational data. We can see already in Figure 15 that the last 30 % of the data has been shifted in comparison to the training data. Hence if a longer period of time is considered, it becomes increasingly likely that a larger portion of the valves operational range is explored. Thirdly, the reference values for this data are based on hourly means of recorded volume flow rates. The mean operation naturally hides the peaks and valleys of the operational modes and hence one could expect larger variance in a real data set.

6. Conclusion

In this paper we modeled the hydraulic properties of a district heating network, and tuned the parameters of said model based on operational data. This method was evaluated on laboratory data from the Smart Water Infrastructure laboratory in Aalborg. We found that in a realistic setting, the characteristics of valves are hugely




Figure 14. Valve curves $K_{v,i}k_i(v_i)$ for each valve i for model structure C trained on the exciting (E, red) and realistic (R, blue) datasets respectively.

important for accurately describing the system properties. A simple assumption of linear valve characteristics is therefore a naive approach. To remedy this issue, we proposed modeling the valve characteristic with a linear combination of many possible valve characteristics, which improved model performance. Additionally, we found that valve hysteresis played a significant role in the setup, which we remedied through the use of a preprocessing filter before doing model calibration. We compared the modeling procedure on two data set, one designed to test the system in all of its' possible operational mode and one designed to mimic an operational district heating network. We found that in the more realistic data set, the valves only operated in a limited region. This led to the model extrapolating poorly to operational modes outside of this region. However, in the first data set where a larger portion of the systems operational range was explored, the best performing model was able to predict volume flow rates within 5 and 10 % deviation from the mean values, excluding the portion of the valve curve where the fit was the most poor. The performance in this region could likely be improved by further tailoring the parameterization of the valve curve.

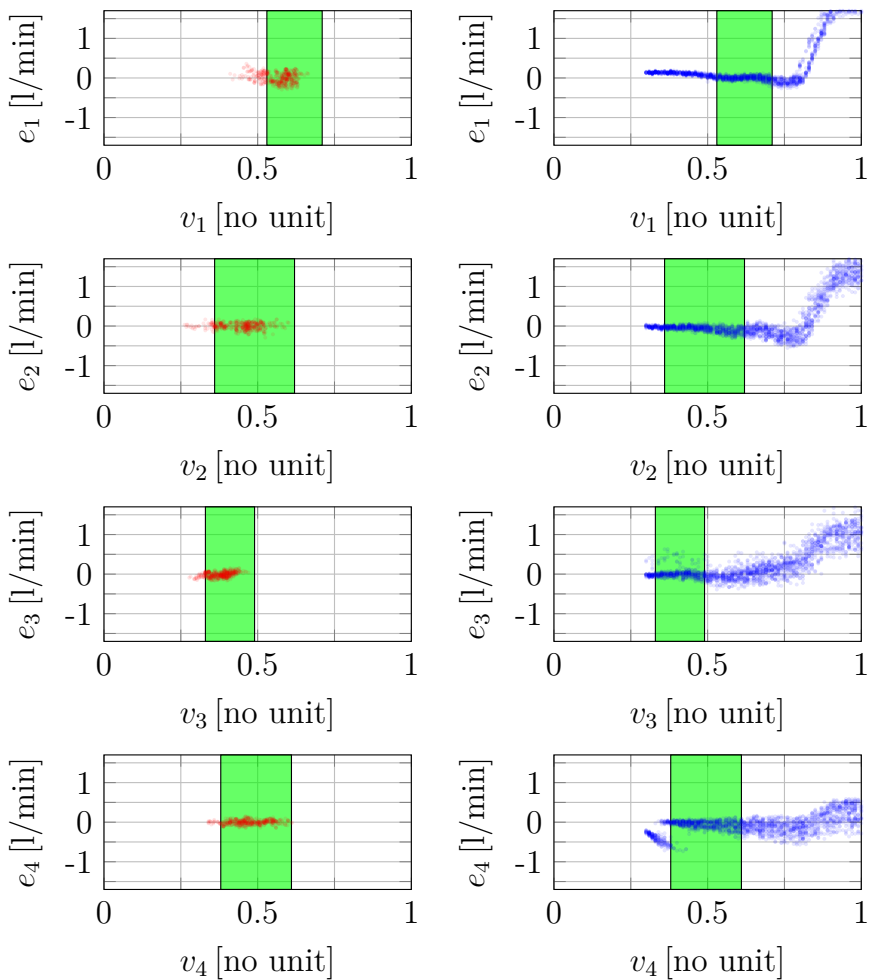


Figure 15. Prediction errors $e_i = q_i - \hat{q}_i$ for each valve i using model C trained on the realistic data set and evaluated on both the realistic data set (red, left) and the exciting data set (blue, right). The figures are arranged as valve 1 to 4 from top to bottom. The green bar outlines the interval between the 5th and 95th quantiles for v_i in the realistic training data.

7. Future work

There are several open questions and directions for extending the results of this paper. Firstly, it is not common in current district heating networks to have access to measurements of valve positions. While this data may become available with future smart installations, a current approach could be to jointly estimate the model parameters and the valve positions. Such an approach would present many interesting challenges, as the quality of the parameter estimation depends on the accuracy of the valve position estimation, and the valve position estimation conversely depends on the accuracy of the model calibration.

A further investigation can be conducted with regards to choosing the filtering parameter δ for hysteresis compensation. In this experimental setup we have only four valves which can be assumed to have similar characteristics. Thus a shared, heuristically chosen δ is viable. For a larger and more heterogeneous network, this approach is no longer valid and thus a reasonable search approach is needed.

It appears that a further investigation into the exponent γ relating volume flow rates and differential pressures over pipes and valves could be an avenue to model improvement. However, such an investigation can quickly become complex. E.g. should this exponent be shared between valves and pipes, and should each component have its own exponent? Tuning this parameter based on data is a non-trivial task with regards to maintaining computational tractability.

Finally one can consider an extension to meshed distribution networks. This causes the issue that the volume flow rates in pipes in the supply-and-return networks are not explicitly known by measuring the valve volume flow rates. Previous works have solved for these values in tandem with model parameters using non-convex optimizers such as particle swarm solvers [Zheng et al., 2023]. Such previous works have however been limited to considering only pipe resistances and not valve characteristics. Therefore they could be less concerned with the computational complexity and convergence issues which may differentiate non-convex from convex problems.

Acknowledgements

Felix Agner and Anders Rantzer are members of the ELLIIT Strategic Research Area at Lund University.

This work is funded by the European Research Council (ERC) under the European Union's Horizon 2020 research and innovation program under grant agreement No 834142 (ScalableControl).

The work has been partially supported by the Poul Due Jensen Foundation through the Smart Water Infrastructure Laboratory project and the Swift project.

Christian Møller Jensen is funded by Danish Innovation Fund grant no. 3129-00019B.

Data Availability

The data sets prepared and used in this paper are openly available via GitHub under a CC-BY 4.0 license [Agner, 2024].

Appendix A - Forward estimation

In a line-structured network such as the one we investigate, our parameterized model can be used for forward estimation. This implies calculating the resulting flow rates q_i given the valve set-points v_i and the differential pressure Δp_0 . To accomplish this we can realize that given a parameterized valve model on the form (4) and given a valve set-point v_i , (4) can be replaced by

$$\Delta p_i = r_i q_i^2 \quad (\text{A.1})$$

where $r_i = \left(\frac{\theta_{i,1}}{k_1(v_i)^2} + \frac{\theta_{i,2}}{k_2(v_i)^2} + \dots + \frac{\theta_{i,K}}{k_K(v_i)^2} \right)$. In this way once the spindle position is fixed, a valve functions like a pipe in the sense of how it relates pressure drops and flow rates. To calculate the total amount of water flowing from α to β , we can reduce our system in a series of steps to one single, equivalent resistance (see Figure A.1). First we replace s_4 and r_4 with an equivalent resistance $\hat{s}_4 = 2s_4 + r_4$ (see Figure A.1a). We can now find a relation between q_3 and q_4 by establishing that

$$\hat{s}_4 q_4^2 = (2s_3 + r_3) q_3^2$$

and hence

$$q_3 = \sqrt{\frac{\hat{s}_4}{2s_3 + r_3}} q_4.$$


We will now lump the part of the network right of valve 2 into one equivalent resistance \hat{s}_3 (see Figure A.1b). The flow rate going through \hat{s}_3 is $q_3 + q_4$, and hence it must hold that

$$\hat{s}_3 (q_3 + q_4)^2 = 2s_6 (q_3 + q_4)^2 + \hat{s}_4 q_4^2 \quad (\text{A.2})$$


$$\begin{aligned} \implies \hat{s}_3 &= 2s_6 + \frac{q_4^2}{(q_3 + q_4)^2} \hat{s}_4 \\ &= 2s_6 + \frac{q_4^2}{q_4^2 \left(\sqrt{\hat{s}_4 / (2s_3 + r_3)} + 1 \right)^2} \hat{s}_4 \\ &= 2s_6 + \frac{(2s_3 + r_3) \hat{s}_4}{\sqrt{\hat{s}_4} + \sqrt{2s_3 + r_3}}. \end{aligned} \quad (\text{A.3})$$

We can iteratively calculate \hat{s}_i in this way according to


$$\hat{s}_i = 2s_{i+3} + \frac{(2s_i + r_i) \hat{s}_{i+1}}{\sqrt{\hat{s}_{i+1}} + \sqrt{2s_i + r_i}} \quad (\text{A.4})$$



(a) Valve resistance r_4 and pipe resistances s_4 are replaced by one equivalent resistance \hat{s}_4 .



(b) The parallel connection of valves 3 and 4 are reduced to one equivalent resistance \hat{s}_3 .



(c) The network is reduced to one equivalent resistance \hat{s}_1 .

Figure A.1. Step-by-step illustration for reducing the full hydraulic representation of the network into one equivalent resistance with given valve set-points.

until reaching \hat{s}_1 which directly connects α to β (see Figure A.1c). Now clearly the total flow rate $\hat{q}_1 = q_1 + q_2 + q_3 + q_4$ must satisfy

$$\hat{q}_1 = \sqrt{\frac{\Delta p_0}{\hat{s}_1}}. \quad (\text{A.5})$$

Now we can calculate the individual flow rates q_1 , q_2 , q_3 and q_4 by expanding the network again in the other direction. We here note that it must hold that

$$\begin{aligned} \hat{s}_1 \hat{q}_1^2 &= \hat{s}_2 (\hat{q}_1 - q_1)^2 \\ \Rightarrow q_1 &= \left(1 - \sqrt{\frac{\hat{s}_1}{\hat{s}_2}}\right) \hat{q}_1 \end{aligned}$$

and also that $\hat{q}_2 = q_2 + q_3 + q_4 = \hat{q}_1 - q_1$. We can thus calculate all of the flow rates recursively according to the rules

$$q_i = \left(1 - \sqrt{\frac{\hat{s}_i}{\hat{s}_{i+1}}}\right) \hat{q}_i \quad (\text{A.6})$$

$$\hat{q}_{i+1} = \hat{q}_i - q_i. \quad (\text{A.7})$$

Appendix B - Valve Characteristic Equations

The expressions relating flow rates, valve set-points and differential pressure for the valves in each model can be written down on closed form. For model A trained on the exciting data set, this relation is given by

$$\Delta p_1 = \frac{0.047}{v_1^2} q_1^2 \quad (\text{B.1a})$$

$$\Delta p_2 = \frac{0.054}{v_2^2} q_2^2 \quad (\text{B.1b})$$

$$\Delta p_3 = \frac{0.083}{v_3^2} q_3^2 \quad (\text{B.1c})$$

$$\Delta p_4 = \frac{0.16}{v_4^2} q_4^2. \quad (\text{B.1d})$$

For model B trained on the exciting data set, the expression is given by

$$\Delta p_1 = \left(\frac{0.011}{\text{ramp}_{0.15}^{0.9}(v_1)^{3.0}} + \frac{0.049}{\text{ramp}_{0.15}^{0.95}(v_1)^{3.0}} + \frac{0.0041}{\text{ramp}_{0.2}^{0.95}(v_1)^{3.0}} + \frac{0.00015}{\text{ramp}_{0.25}^{0.95}(v_1)^{3.0}} \right) q_1^2 \quad (\text{B.2a})$$

$$\Delta p_2 = \left(\frac{0.009}{\text{ramp}_{0.1}^{0.9}(v_2)^{3.0}} + \frac{0.0068}{\text{ramp}_{0.1}^{0.95}(v_2)^{3.0}} + \frac{0.039}{\text{ramp}_{0.15}^{0.95}(v_2)^{3.0}} + \frac{0.0068}{\text{ramp}_{0.2}^{0.95}(v_2)^{3.0}} + \frac{0.0011}{\text{ramp}_{0.25}^{0.95}(v_2)^{3.0}} \right) q_2^2 \quad (\text{B.2b})$$

$$\Delta p_3 = \left(\frac{0.039}{\text{ramp}_{0.15}^{0.95}(v_3)^{3.0}} + \frac{0.029}{\text{ramp}_{0.15}^{1.0}(v_3)^{3.0}} + \frac{0.0044}{\text{ramp}_{0.2}^{1.0}(v_3)^{3.0}} + \frac{0.00072}{\text{ramp}_{0.25}^{1.0}(v_3)^{3.0}} \right) q_3^2 \quad (\text{B.2c})$$

$$\Delta p_4 = \left(\frac{0.11}{\text{ramp}_{0.1}^{0.9}(v_4)^{2.0}} + \frac{0.031}{\text{ramp}_{0.1}^{1.0}(v_4)^{2.0}} + \frac{0.0043}{\text{ramp}_{0.2}^{0.9}(v_4)^{3.0}} + \frac{0.0078}{\text{ramp}_{0.25}^{0.9}(v_4)^{3.0}} + \frac{0.033}{\text{ramp}_{0.25}^{0.95}(v_4)^{3.0}} \right) q_4^2 \quad (\text{B.2d})$$

For model C trained on the exciting data set, the expressions are

$$\Delta p_1 = \left(\frac{0.0093}{\text{ramp}_{0.15}^{0.9}(v_1)^{3.0}} + \frac{0.03}{\text{ramp}_{0.15}^{0.95}(v_1)^{3.0}} + \frac{0.025}{\text{ramp}_{0.2}^{0.95}(v_1)^{3.0}} \right) q_1^2 \quad (\text{B.3a})$$

$$\Delta p_2 = \left(\frac{0.0065}{\text{ramp}_{0.15}^{0.9}(v_2)^{3.0}} + \frac{0.038}{\text{ramp}_{0.15}^{0.95}(v_2)^{3.0}} + \frac{0.018}{\text{ramp}_{0.2}^{0.95}(v_2)^{3.0}} + \frac{0.0022}{\text{ramp}_{0.25}^{0.95}(v_2)^{3.0}} \right) q_2^2 \quad (\text{B.3b})$$

$$\Delta p_3 = \left(\frac{0.04}{\text{ramp}_{0.15}^{0.95}(v_3)^{3.0}} + \frac{0.0089}{\text{ramp}_{0.2}^{0.95}(v_3)^{3.0}} + \frac{0.025}{\text{ramp}_{0.2}^{1.0}(v_3)^{3.0}} \right) q_3^2 \quad (\text{B.3c})$$

$$\Delta p_4 = \left(\frac{0.023}{\text{ramp}_{0.1}^{0.85}(v_4)^{2.0}} + \frac{0.069}{\text{ramp}_{0.1}^{0.9}(v_4)^{2.0}} + \frac{0.0085}{\text{ramp}_{0.2}^{0.9}(v_4)^{3.0}} + \frac{0.036}{\text{ramp}_{0.25}^{0.9}(v_4)^{3.0}} + \frac{0.001}{\text{ramp}_{0.25}^{0.95}(v_4)^{3.0}} + \frac{0.023}{\text{ramp}_{0.25}^{1.0}(v_4)^{3.0}} \right) q_4^2 \quad (\text{B.3d})$$

Finally, when trained on the realistic data set, model C yields the following expressions.

$$\Delta p_1 = \left(\frac{0.11}{\text{ramp}_{0.2}^{0.8}(v_1)^{3.0}} + \frac{0.00014}{\text{ramp}_{0.25}^{0.8}(v_1)^{3.0}} \right) q_1^2 \quad (\text{B.4a})$$

$$\Delta p_2 = \left(\frac{0.11}{\text{ramp}_{0.15}^{0.8}(v_2)^{3.0}} + \frac{0.024}{\text{ramp}_{0.25}^{0.8}(v_2)^{2.5}} + \frac{0.0029}{\text{ramp}_{0.25}^{0.8}(v_2)^{3.0}} \right) q_2^2 \quad (\text{B.4b})$$

$$\Delta p_3 = \left(\frac{0.085}{\text{ramp}_{0.15}^{0.8}(v_3)^{3.0}} + \frac{0.065}{\text{ramp}_{0.2}^{0.8}(v_3)^{3.0}} \right) q_3^2 \quad (\text{B.4c})$$

$$\Delta p_4 = \left(\frac{0.077}{\text{ramp}_{0.1}^{0.8}(v_4)^{2.0}} + \frac{0.026}{\text{ramp}_{0.1}^{0.8}(v_4)^{2.5}} + \frac{0.12}{\text{ramp}_{0.25}^{0.8}(v_4)^{3.0}} \right) q_4^2. \quad (\text{B.4d})$$

References

- Agner, F. (2024). *DH hydraulic parameter estimation*. Version 0.1.0. DOI: 10.5281/zenodo.11946409.
- Agner, F., C. M. Jensen, A. Rantzer, C. S. Kallesøe, and R. Wisniewski (2024). “Hydraulic parameter estimation for district heating based on laboratory experiments”. *Energy* **312**, p. 133462. ISSN: 0360-5442. DOI: <https://doi.org/10.1016/j.energy.2024.133462>.

- Agner, F., P. Kergus, R. Pates, and A. Rantzer (2022). “Combating district heating bottlenecks using load control”. *Smart Energy* **6**, p. 100067. ISSN: 2666-9552. DOI: 10.1016/j.segy.2022.100067.
- Agner, F., P. Kergus, R. Pates, and A. Rantzer (2023). “Hydraulic parameter estimation in district heating networks”. *IFAC-PapersOnLine* **56**:2. 22nd IFAC World Congress, pp. 5438–5443. ISSN: 2405-8963. DOI: <https://doi.org/10.1016/j.ifacol.2023.10.194>.
- Agrawal, A., R. Verschueren, S. Diamond, and S. Boyd (2018). “A rewriting system for convex optimization problems”. *Journal of Control and Decision* **5**:1, pp. 42–60.
- Bahlawan, H., N. Ferraro, A. Gambarotta, E. Losi, L. Manservigi, M. Morini, C. Saletti, P. Ruggero Spina, and M. Venturini (2022). “Detection and identification of faults in a District Heating Network”. *Energy Conversion and Management* **266**, p. 115837. ISSN: 0196-8904. DOI: 10.1016/j.enconman.2022.115837.
- Benonysson, A., B. Bøhm, and H. F. Ravn (1995). “Operational optimization in a district heating system”. en. *Energy Conversion and Management* **36**:5, pp. 297–314. ISSN: 01968904. DOI: 10.1016/0196-8904(95)98895-T.
- De Persis, C. and C. S. Kallesoe (2011). “Pressure regulation in nonlinear hydraulic networks by positive and quantized controls”. *IEEE Transactions on Control Systems Technology* **19**:6, pp. 1371–1383. DOI: 10.1109/TCST.2010.2094619.
- Dénarié, A., M. Aprile, and M. Motta (2023). “Dynamical modelling and experimental validation of a fast and accurate district heating thermo-hydraulic modular simulation tool”. *Energy* **282**, p. 128397. ISSN: 0360-5442. DOI: 10.1016/j.energy.2023.128397.
- Diamond, S. and S. Boyd (2016). “CVXPY: A Python-embedded modeling language for convex optimization”. *Journal of Machine Learning Research* **17**:83, pp. 1–5.
- Frederiksen, S. and S. Werner (2013). *District heating and cooling*. Studentlitteratur.
- Guelpa, E., C. Toro, A. Sciacovelli, R. Melli, E. Sciubba, and V. Verda (2016). “Optimal operation of large district heating networks through fast fluid-dynamic simulation”. *Energy* **102**, pp. 586–595. ISSN: 0360-5442. DOI: 10.1016/j.energy.2016.02.058.
- Hägglund, T. (2023). *Process Control in Practice*. De Gruyter, Berlin, Boston. ISBN: 9783111104959. DOI: [doi:10.1515/9783111104959](https://doi.org/10.1515/9783111104959).
- Liu, Y., P. Wang, and P. Luo (2020). “Pipe hydraulic resistances identification of district heating networks based on matrix analysis”. *Energies* **13**:11. ISSN: 1996-1073. DOI: 10.3390/en13113007.

- Lund, H., S. Werner, R. Wiltshire, S. Svendsen, J. Thorsen, F. Hvelplund, and B. Mathiesen (2014). “4th generation district heating (4GDH)”. *Energy* **68**, pp. 1–11. ISSN: 03605442. DOI: 10.1016/j.energy.2014.02.089.
- Luo, P., H. Wang, Y. Liu, Q. Du, and Z. Zhang (2022). “Resistance characteristic parameters estimation of hydraulic model in heating networks based on real-time operation data”. *Buildings* **12**:6. ISSN: 2075-5309. DOI: 10.3390/buildings12060743.
- Savic, D. A., Z. S. Kapelan, and P. M. Jonkergouw (2009). “Quo vadis water distribution model calibration?” *Urban Water Journal* **6**:1, pp. 3–22. DOI: 10.1080/15730620802613380. eprint: <https://doi.org/10.1080/15730620802613380>.
- Val Ledesma, J., R. Wisniewski, and C. S. Kallesøe (2021). “Smart water infrastructures laboratory: reconfigurable test-beds for research in water infrastructures management”. *Water* **13**:13. ISSN: 2073-4441. DOI: 10.3390/w13131875.
- Vallabha, G. (2016). “Real-time pacer for simulink”. *Published in Matlab Central repository*.
- Vandermeulen, A., B. van der Heijde, and L. Helsen (2018). “Controlling district heating and cooling networks to unlock flexibility: A review”. en. *Energy* **151**, pp. 103–115. ISSN: 03605442. DOI: 10.1016/j.energy.2018.03.034.
- Vesterlund, M., A. Toffolo, and J. Dahl (2016). “Simulation and analysis of a meshed district heating network”. en. *Energy Conversion and Management* **122**, pp. 63–73. ISSN: 01968904. DOI: 10.1016/j.enconman.2016.05.060.
- Wang, N., S. You, Y. Wang, H. Zhang, Q. Miao, X. Zheng, and L. Mi (2018). “Hydraulic resistance identification and optimal pressure control of district heating network”. *Energy and Buildings* **170**, pp. 83–94. ISSN: 0378-7788. DOI: <https://doi.org/10.1016/j.enbuild.2018.04.003>.
- Wang, Y., K. Shi, X. Zheng, S. You, H. Zhang, C. Zhu, L. Li, S. Wei, C. Ding, and N. Wang (2020). “Thermo-hydraulic coupled analysis of meshed district heating networks based on improved breadth first search method”. *Energy* **205**, p. 117950. ISSN: 0360-5442. DOI: <https://doi.org/10.1016/j.energy.2020.117950>.
- Wang, Y., S. You, H. Zhang, W. Zheng, X. Zheng, and Q. Miao (2017). “Hydraulic performance optimization of meshed district heating network with multiple heat sources”. en. *Energy* **126**, pp. 603–621. ISSN: 03605442. DOI: 10.1016/j.energy.2017.03.044.
- Zheng, X., Z. Shi, Y. Wang, H. Zhang, and H. Liu (2023). “Thermo-hydraulic condition optimization of large-scale complex district heating network: A case study of Tianjin”. en. *Energy* **266**, p. 126406. ISSN: 03605442. DOI: 10.1016/j.energy.2022.126406.

Zheng, X., Z. Shi, Y. Wang, H. Zhang, and Z. Tang (2024). “Digital twin modeling for district heating network based on hydraulic resistance identification and heat load prediction”. *Energy* **288**, p. 129726. ISSN: 0360-5442. DOI: <https://doi.org/10.1016/j.energy.2023.129726>.

Reglering av nätverk med begränsad kapacitet

Felix Agner


Institutionen för Reglerteknik

Populärvetenskaplig sammanfattning av doktorsavhandlingen *Control of Capacity-Constrained Networks in a District Heating Setting*, januari 2025. Avhandlingen kan laddas ner från: <http://www.control.lth.se/publications>

En riktigt kall vinterdag drar in över en liten stad som i vanliga fall hålls varm av ett fjärrvärmenät. Det kalla vädret gör att de automatiska styrsystemen i stadens byggnader klickar igång. Styrventiler i varje byggnad öppnas för att öka flödet av vatten som går genom byggnadernas fjärrvärmecentraler. Genom alla de rör som kopplar ihop byggnaderna i nätet ökar nu flödet av vatten. Detta leder till stora tryckförluster på grund av ökad friktion i rören. Pumparna i det lokala kraftvärmeverket går redan i full fart för att trycksätta systemet och möta det stora värmebehovet, så nu när tryckförlusterna ökar uppstår problem. Längst ut i nätet gör tryckförlusterna att trycket i systemet är för lågt. De byggnader som befinner sig långt borta från kraftvärmeverket får inte det vatten de behöver för att hålla värmen uppe.

Det här är ett exempel på ett system där kapacitetsbrist råder. De tillgängliga resurserna kan inte möta behovet hos alla de som är anslutna till nätet. Liknande exempel återfinns i många tekniska system där flera aktörer delar på samma resurser: I de sydsvenska elnäten diskuteras kapacitetsbristen allt mer flitigt. Bredbandsnäten har i hastig takt byggts ut för att kunna täcka behovet hos de uppkopplade kunderna. Smarta kamerasystem behöver distribuera lagrings- och beräkningsresurser i nätet dynamiskt.

I system likt de här exemplen är många enheter sammankopplade, och



Ett fjärrvärmenät förses med värme från ett kraftvärmeverk. När det blir väldigt kallt ute kan behovet av värme vara för stort för att kunna mötas med nätets begränsade kapacitet. När det händer får byggnader nära värmekällan det varmvatten de behöver, medan byggnader längre bort märker att flödet sinar.

varje enhet har ofta ett eget styrsystem som försöker reglera en last. De här systemen designas ofta för att vara decentraliserade. Med det menas att varje styrsystem bara mäter och reglerar lokalt. Till exempel kan styrsystemet för fjärrvärmecentralen i ett hus mäta temperaturer och flöden lokalt, men inte i grannens hus. Utöver detta kommunicerar inte styrsystemen med varandra. Den här decentraliserade design-strukturen har många fördelar. Det behövs ingen avancerad kommunikationsutrustning, och nya enheter kan enkelt anslutas till nätverket.

Den fullt decentraliserade strukturen är dock inte perfekt. I exemplet med den kalla vinterdagen ser vi en anledning till varför. Om alla styrsystem agerar utan kännedom om tillståndet i resten av nätverket så kan det uppstå obalans i hur resurser fördelas. Den här avhandlingen tacklar frågan om hur styrsystem kan designas så att de fördelar en begränsad resurs på ett optimalt sätt. Avhandlingen fokuserar på fjärrvärmesystem, och hur flödet av varmvatten ska fördelas i "kall vinterdag"- exemplet.

En intressant slutsats från avhandlingen är att den enkla, decentraliserade lösningen som finns i dagens fjärrvärmesystem i viss mån är optimal. Med den designen levereras så stor mängd totalt flöde som möjligt. Men även om den totala mängden flöde maximeras så betyder det inte att den är rättvist fördelad. Avhandlingen visar också hur man med små och enkla utvecklingar av nuvarande system kan se till att flödet fördelas mer rättvist i systemet. De byggnader som tidigare blev kalla får mer värme, genom att övriga styrsystem minskar sitt eget flöde.

