



Delft University of Technology

Document Version

Final published version

Citation (APA)

Sibeijn, M. W. (2026). *Scalable Predictive Control for District Heating Networks*. [Dissertation (TU Delft), Delft University of Technology]. <https://doi.org/10.4233/uuid:797abc88-d9ae-4151-9730-a11a46e643ac>

Important note

To cite this publication, please use the final published version (if applicable). Please check the document version above.

Copyright

In case the licence states "Dutch Copyright Act (Article 25fa)", this publication was made available Green Open Access via the TU Delft Institutional Repository pursuant to Dutch Copyright Act (Article 25fa, the Taverne amendment). This provision does not affect copyright ownership. Unless copyright is transferred by contract or statute, it remains with the copyright holder.

Sharing and reuse

Other than for strictly personal use, it is not permitted to download, forward or distribute the text or part of it, without the consent of the author(s) and/or copyright holder(s), unless the work is under an open content license such as Creative Commons.

Takedown policy

Please contact us and provide details if you believe this document breaches copyrights. We will remove access to the work immediately and investigate your claim.

This work is downloaded from Delft University of Technology.

SCALABLE PREDICTIVE CONTROL
for **DISTRICT HEATING NETWORKS**



MAX SIBEIJN

SCALABLE PREDICTIVE CONTROL FOR DISTRICT HEATING NETWORKS

SCALABLE PREDICTIVE CONTROL FOR DISTRICT HEATING NETWORKS

Dissertation

for the purpose of obtaining the degree of doctor
at Delft University of Technology
by the authority of the Rector Magnificus,
Prof. dr. ir. H. Bijl,
chair of the Board for Doctorates
to be defended publicly on
Tuesday 12 May 2026 at 12:30

Max Willem SIBEIJN

This dissertation has been approved by the (co)promoters.

Composition of the doctoral committee:

Rector Magnificus,	chairperson
Prof. dr. T. Keviczky,	Delft University of Technology, <i>promotor</i>
Dr. S.D. Gonçalves Melo Pequito,	Instituto Superior Técnico, <i>promotor</i>
Dr. M. Khosravi,	Delft University of Technology, <i>copromotor</i>

Independent members:

Prof. dr. T. Faulwasser	Hamburg University of Technology, Germany
Prof. dr. K. Hooman	Delft University of Technology
Dr. A. La Bella	Politecnico di Milano, Italy
Dr. ir. I.W.M. Pothof,	Delft University of Technology
Prof. dr. ir. J.W. van Wingerden,	Delft University of Technology, <i>reserve member</i>

This research was funded by the Local Inclusive Future Energy (LIFE) City Project (MOOI32019), supported by the Ministry of Economic Affairs and Climate and the Ministry of the Interior and Kingdom Relations of the Netherlands.



Keywords: District heating, model predictive control, operational optimization, sensor placement, dissipativity theory, dual dynamic programming, Koopman operators.

Printed by: Ridderprint.nl

Cover design: © Sila Akçakoca, 2026

Copyright © 2026 by M.W. Sibeijn

ISBN 978-94-6518-293-3

An electronic copy of this dissertation is available at
<https://repository.tudelft.nl/>.

Contents

Summary	xi
Samenvatting	xiii
Terminology & Notation	xv
Acronyms	xv
Variables	xv
Notation	xvi
1. Introduction	1
1.1. Operational Management of District Heating Networks	3
1.2. Challenges in Predictive Control of District Heating Networks	6
1.3. Problem Formulation: Predictive Control of District Heating Networks	7
1.4. Research Contributions and Dissertation Outline	8
2. Modeling for Control of District Heating Networks	11
2.1. District Heating Network Model	12
2.1.1. Graph Model	12
2.1.2. Edge Dynamics	13
2.1.3. Nodal Constraints	13
2.1.4. AROMA: A Benchmark for District Heating Networks	14
2.2. DHN Model: A Tractable Reformulation for the Hydraulics	15
2.2.1. Fixing the Flow Direction	15
2.2.2. Independent Flows	16
2.3. DHN Model: A Tractable Reformulation for the Thermal Dynamics	18
2.3.1. Thermal Node Model	18
2.3.2. Interconnected DHN Model	18
2.A. Appendix: Derivation of Adjacency Matrix	22
3. Sensor Placement in District Heating Networks	23
3.1. Introduction	24
3.2. Optimal Sensor Placement Problem for District Heating Networks	24
3.3. Time-Averaged Frequential Gramian	25
3.3.1. Approximate Gramian Computation	27
3.3.2. Properties of Frequential Gramians	29
3.4. Greedy Sensor Selection	30
3.4.1. Supermodular and Non-Supermodular Set Functions	30
3.4.2. The Greedy Algorithm	31
3.5. Numerical Experiments and Results	32

3.6.	Conclusion	34
4.	Economic Nonlinear Model Predictive Control for District Heating Networks	35
4.1.	Introduction	36
4.2.	Treatment of the Momentum Equation	37
4.3.	Convex Reformulation of the Kirchhoff Loop Constraints	38
4.4.	Time Discretization of the Thermal Dynamics	43
4.5.	Economic MPC Formulation for District Heating Networks	44
4.5.1.	Problem Formulation	44
4.5.2.	Objective Function	45
4.5.3.	Improving Numerical Performance	46
4.6.	Numerical Experiments and Results	46
4.6.1.	High-fidelity Model	47
4.6.2.	Parameters and Data	47
4.6.3.	Economic Performance and Comparisons	48
4.6.4.	Added Value of Storage	51
4.6.5.	Added Value of Multiple Producers	53
4.6.6.	Computational Study	54
4.6.7.	Limitations and Future Work	55
4.7.	Conclusion	55
5.	Stability Analysis of Economic MPC for DHNs using Dissipativity Theory	57
5.1.	Introduction	58
5.2.	Strict Dissipativity for Thermal Networks	58
5.2.1.	Dissipativity of Closed Thermal Networks	59
5.2.2.	Dissipativity of Open Thermal Networks	60
5.2.3.	Numerical Verification of Strict Dissipativity	61
5.3.	Economic Nonlinear Model Predictive Control	63
5.4.	Numerical Experiments and Results	64
5.5.	Discussion and Conclusion	65
6.	Scalable Predictive Control for DHNs: A Dual Dynamic Programming Approach	69
6.1.	Introduction	70
6.2.	Short-Sightedness of MPC for DHN	71
6.3.	Dual Dynamic Programming	72
6.3.1.	Prediction Form MPC	74
6.3.2.	Approximate Lower Bound on Value Function	74
6.3.3.	Infeasible Subproblems	75
6.4.	Numerical Examples	77
6.5.	Conclusion	81
6.A.	Appendix: Lower Bound on Value Function	81
7.	Scalable Predictive Control for DHNs: A Physics-Guided Koopman Approach	83
7.1.	Introduction	84
7.2.	Thermo-Hydraulic Dynamics	85

- 7.3. Koopman Operator Theory 86
 - 7.3.1. Koopman Operator Framework 87
 - 7.3.2. Data-Driven Model Approximation for DHNs 88
 - 7.3.3. Numerical Identification Algorithm 88
- 7.4. Observable Selection: A Physics-Guided Approach 89
- 7.5. Koopman Predictive Control for District Heating Networks 91
 - 7.5.1. Koopman Predictive Control for Unidirectional DHNs 91
 - 7.5.2. Extension to Bidirectional DHNs 92
- 7.6. Numerical Experiments and Results 93
 - 7.6.1. Closed-Loop Experiments 94
 - 7.6.2. Data Acquisition, Learning, and Implementation Details 95
 - 7.6.3. Predictive Performance of the Koopman Operator 98
 - 7.6.4. Predictive Control for Unidirectional DHNs 100
 - 7.6.5. Predictive Control for Bidirectional DHNs 102
- 7.7. Conclusion 104

- 8. Conclusions and Recommendations 105**
 - 8.1. Conclusions 106
 - 8.2. Recommendations 108

- Bibliography 111**

- Curriculum Vitæ 121**

- List of Publications 123**

Summary

THE transition to sustainable energy requires cleaner and more efficient ways to consume heat. We are used to our homes and offices being warm, yet behind this expectation lies infrastructure that must continuously balance supply and demand: heat must be produced, transported, and delivered exactly when and where it is needed. District heating networks accomplish this by circulating hot water through underground pipelines, connecting heat sources to thousands of buildings across a city. Operating such a network efficiently is far from trivial. Decisions made now about how much heat to produce, at what temperature, and through which routes have consequences that unfold over hours as hot water slowly travels through kilometers of pipe. Getting these decisions right can mean significant cost savings; getting them wrong leads to wasted energy or discomfort.

This challenge is fundamentally one of planning under complexity. The operator must anticipate future demand, account for heat losses during transport, respect physical limitations of pumps and pipes, and respond to fluctuating energy prices, all while the state of the network is only partially observable through a limited set of sensors. Traditional approaches rely on simple rules and operator experience, which suffice for conventional high-temperature networks but fall short as renewable sources have lower temperatures, requiring more advanced anticipative control strategies.

This thesis develops computational methods that enable district heating networks to be operated efficiently in real time. The underlying physics lead to mathematical optimization problems that are, in their original form, too complex to solve within practical time limits. This thesis introduces techniques that reformulate these problems into more tractable forms, decompose them across time to reduce computational burden, and learn efficient representations from operational data. We also address the question of where to place sensors to best reconstruct the network state, and establish theoretical conditions that guarantee stable closed-loop operation.

These methods are validated on realistic network models and compared against conventional rule-based strategies. The results show that predictive control can meaningfully reduce operating costs while improving constraint satisfaction, that our sensor placement method provides a practical tool for selecting informative measurement locations under tight instrumentation budgets, and that computational decomposition and data-driven surrogate models each bring solve times within real-time limits without sacrificing performance. We further find evidence of inherent stability properties of the closed-loop system, even in the absence of terminal constraints. Together, these findings demonstrate that the fundamental barriers to deploying advanced control in district heating networks can be systematically addressed.

Samenvatting

DE energietransitie vraagt om verduurzaming van de warmtevoorziening. Hoewel we verwarming in huizen en kantoren als vanzelfsprekend beschouwen, schuilt hierachter een infrastructuur die continu vraag en aanbod moet balanceren. Warmtenetten spelen hierin een sleutelrol door warm water via ondergrondse leidingen te transporteren en zo warmtebronnen te verbinden met duizenden stedelijke gebouwen.

Het efficiënt beheren van zo'n netwerk is echter een complex vraagstuk. Operationele beslissingen over productiehoeveelheden, temperaturen en distributieroutes hebben effecten die door de trage stroming van water door kilometerslange leidingen pas uren later merkbaar zijn. Optimale besluitvorming levert hierbij aanzienlijke kostenbesparingen op, terwijl suboptimale keuzes leiden tot energieverstopping of comfortverlies.

In de kern is dit een complex planningsprobleem. De operator moet anticiperen op de toekomstige warmtevraag, rekening houden met warmteverliezen en fysieke netwerkbeperkingen, en inspelen op fluctuerende energieprijzen. Dit alles moet gebeuren terwijl de status van het netwerk slechts beperkt zichtbaar is via een beperkt aantal sensoren. Waar traditionele, op ervaring gebaseerde aansturing volstaat voor conventionele netten, schiet deze tekort voor moderne systemen met meerdere bronnen, opslagmogelijkheden en tweerichtingsverkeer.

Dit proefschrift draagt bij aan de ontwikkeling van rekenmethoden voor de optimale, real-time aansturing van warmtenetten. De onderliggende fysica leidt tot wiskundige modellen die in hun oorspronkelijke vorm te complex zijn om binnen praktische tijdslimieten op te lossen. Wij introduceren technieken om deze problemen te herformuleren en in tijdstappen op te knippen (decompositie) om de rekenlast te verlagen, en gebruiken operationele data om efficiënte modelrepresentaties af te leiden. Tot slot behandelen we de optimale sensorplaatsing voor nauwkeurige staatsreconstructie en formuleren we theoretische voorwaarden die een stabiele regeling garanderen.

Deze methoden zijn gevalideerd op realistische netwerkmodellen en vergeleken met conventionele regelgebaseerde aansturing. De resultaten laten zien dat voorspellende regeling de operationele kosten aanzienlijk kan verlagen en tegelijkertijd beter voldoet aan operationele beperkingen. Onze sensorplaatsingsmethode biedt operators een praktisch hulpmiddel om met een beperkt instrumentatiebudget toch de juiste meetlocaties te selecteren. Zowel decompositie als datagedreven surrogaatmodellen brengen de rekentijd binnen real-time limieten, zonder noemenswaardige concessies aan de regelprestaties. Daarnaast vinden we aanwijzingen voor inherente stabiliteitseigenschappen van het gesloten-lussysteem. Gezamenlijk tonen deze bevindingen aan dat de fundamentele drempels voor geavanceerde aansturing van warmtenetten systematisch geadresseerd kunnen worden.

Terminology & Notation

Acronyms

acronym	description
DHN	district heating network
MPC	model predictive control
EMPC	economic model predictive control
NMPC	nonlinear model predictive control
KPC	Koopman predictive control
OCP	optimal control problem
RBC	rule-based control
ODE	ordinary differential equation
PDE	partial differential equation
DAE	differential algebraic equation
LP	linear program
QP	quadratic program
MILP	mixed-integer linear program
NLP	nonlinear program
SOS	sum of squares

Variables

symbol	units	description
T	K	temperature
T_a	K	ambient soil temperature
\tilde{T}	K	shifted temperature $T - T_a$
v	m s^{-1}	velocity
q	$\text{m}^3 \text{s}^{-1}$	flow rate
p	Pa	pressure
t	s	time
T_p	s	system period
g	m s^{-2}	gravitational acceleration
ρ	kg m^{-3}	density of water
V	m^3	Volume
d	m	diameter
K	-	friction coefficient
U	$\text{J m}^{-2} \text{K}^{-1}$	heat transfer coefficient

symbol	units	description
c_p	$\text{J kg}^{-1} \text{K}^{-1}$	specific heat coefficient

Notation

symbol	description
\mathbb{R}	set of real numbers
\mathbb{R}_+	set of nonnegative real numbers
\mathbb{Z}	set of integers
\mathbb{N}	set of natural numbers
\mathbb{N}_m	set of natural numbers up to finite positive integer m , i.e., $\mathbb{N}_m = \{1, 2, \dots, m\}$
∂_x	partial derivative with respect to x
$[x_i]_{i \in \mathcal{S}}$	(column) vector of all elements x_i , where i in set \mathcal{S}
$ \mathcal{S} $	cardinality of set \mathcal{S}
$\text{tr}(M)$	trace of matrix M
$\log \det(M)$	log-determinant of matrix M
M^\dagger	Moore-Penrose pseudo-inverse of matrix M
$ M _o$	element-wise absolute value operator for matrix M
\odot	element-wise multiplication operator
$\mathbf{1}$	the vector of all ones, i.e., $\mathbf{1} = [1, \dots, 1]^\top$

1

Introduction

THE energy transition requires a major shift from fossil fuel-based generation to renewable energy sources. The heating and cooling sector accounts for half of the total final energy consumption in the EU [1], yet fossil fuels still dominate heat production at nearly 70% as of 2022 [2]. This heavy reliance on fossil fuels makes the transition to sustainable heating necessary. At the same time, heat production is undergoing rapid electrification as the energy sector shifts away from fossil fuels; see, e.g., [3]. This transition, driven by growing electricity demand from heat pumps and electric vehicles alongside the intermittent nature of renewable sources like wind and solar, is placing unprecedented strain on power grids worldwide.



Figure 1.1: District heating network consisting of supply (red) and return (blue) pipelines transporting heated water from producer to consumer and back.

While expanding power grid infrastructure is necessary to accommodate growing electricity demand, such expansion is typically costly, time-consuming, and held back by regulatory challenges [4]. Moreover, the scale of heating demand means that full electrification through individual heat pumps could double electricity consumption [5], requiring grid expansion that is neither economically viable nor achievable within the required timeframes. Therefore, a diversified portfolio of heating technologies is essential for the energy transition.

District heating networks (DHNs) represent a crucial component of this diversified portfolio. DHNs are networks of pipelines that distribute heated water throughout a district or city, as illustrated in Figure 1.1. The core advantage of district heating lies in its ability to capture secondary heat resources that would otherwise be wasted—such as industrial waste heat, excess energy from waste incineration and combined heat and power plants, combustible renewables, and geothermal sources [6]. By leveraging these secondary energy sources, the primary energy input required is significantly lower than the heat delivered to consumers, making district heating a highly efficient technology. Despite supplying only around 12% of heat consumption in the EU as of 2022 [2], DHNs

have the potential to become a key contributor to achieving climate goals if their market share increases [1].

District heating networks offer another significant advantage contributing to the energy transition: *flexibility*. The networks' substantial thermal inertia allows them to decouple heat production from consumption, storing energy when electricity is abundant or cheap and releasing it during peak demand or high-price periods. This flexibility enables DHNs to respond dynamically to electricity market conditions, helping to balance power grid loads and integrate variable renewable energy sources. However, unlocking this flexibility potential requires moving beyond traditional operational practices. Modern DHNs, with their numerous distributed controllable assets and increasingly diverse heat sources, demand advanced control strategies [7–9] that can optimize performance in real-time while managing the complex interactions between production, storage, and consumption.

1.1. Operational Management of District Heating Networks

BEFORE addressing modern district heating networks, we first describe the status quo of scheduling and control in conventional systems. District heating networks typically consist of hydraulically decoupled subnetworks connected through heat exchangers, enabling independent pressure management while facilitating heat transfer between network sections, as illustrated in Figure 1.2. Within these networks, control actions operate at different temporal scales, which we distinguish into two hierarchical levels:

1. *Scheduling*, also called high-level control. This task involves the determination of operational parameters and setpoints for network assets and devices, typically spanning hours to days;
2. *Regulation*, or low-level control. This constitutes the lowest level of actuation, operating at frequencies from seconds to minutes. Regulation is the process of tracking reference setpoints provided by the scheduling layer using physical actuators such as valve positions, pump speeds, and generator power settings.

The regulation layer of conventional DHNs involves multiple control systems acting at different physical locations. From an operational perspective, these systems can be understood through their relationship to operator control authority: heat demand control (consumer-controlled), flow regulation (substation-level, tracking local conditions), and supply temperature and pressure control (plant-level, tracking setpoints) [6]. Understanding this distribution of control authority is essential for positioning the scheduling problem that constitutes our focus.

Heat demand is determined by consumers adjusting thermostatic radiator valves or manual tap water openings, regulating secondary network flows to achieve desired heat transfer. Operators lack control authority over these devices, so demand acts as an exogenous disturbance. To meet this demand, flow control systems at substations regulate valve positions on the primary side, opening valves to increase heat exchanger throughput when needed. However, this mechanism requires both sufficient pressure differential and adequate supply temperature in the primary network.

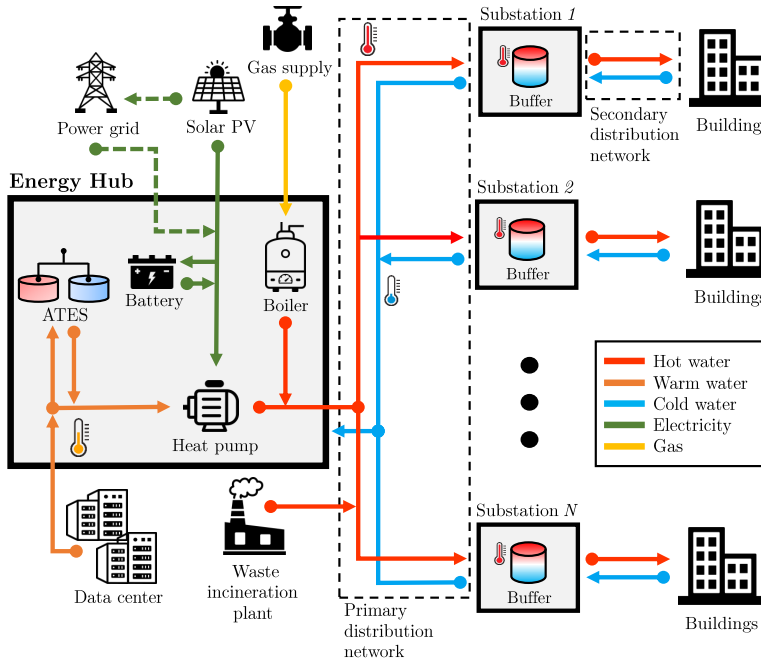


Figure 1.2: A schematic diagram of a district heating network with multiple substations (consumers) and production points.

These conditions—adequate pressure and temperature—are maintained through regulators that control the output of pressure pumps and heat generators within the network (often at the production plant). These regulators are provided with pre-specified setpoints to track, for example, supply temperature setpoints that are classically determined by outdoor temperature using simple rules. The determination of these setpoints represents the scheduling problem: deciding how to coordinate production, storage, and network operation over time horizons of hours to days. This scheduling layer constitutes the primary focus of this thesis. We do not propose replacing the lower-level regulation mechanisms, but rather develop advanced methods for generating the setpoint strategies that govern high-level asset coordination across the network.

Modern district heating networks, designated as *4th and 5th generation district heating systems*, are characterized by their capacity to operate at substantially reduced temperatures with correspondingly reduced thermal losses, their ability to integrate renewable energy sources, their integration within multi-carrier energy systems, and their accommodation of increasingly uncertain supply and demand patterns [8, 10, 11]. As a result, conventional rule-based scheduling methodologies, often based on ad hoc, heuristic approaches [9, 12], prove increasingly inadequate in addressing the inherent complexity of advanced network infrastructures.

The current trend in the literature supports the adoption of so-called *intelligent* control systems [7–9, 13], defined by their ability to adapt their control strategies based

on changing conditions and historical data. More specifically, this approach requires advanced computational algorithms capable of several important functions: predicting system behavior over different time scales, coordinating multiple distributed energy resources, satisfying system constraints, and optimizing control decisions in real-time. These capabilities are expected to improve district heating network operations through systematic temperature reduction, provision of demand flexibility, better management of supply-demand mismatches, and overall efficiency improvements.

Among the intelligent control approaches gaining traction for district heating networks, *model predictive control* (MPC) emerges as a particularly promising solution for operational optimization and scheduling [12–15]. Figure 1.3 provides a schematic overview of MPC applied to DHNs. MPC employs a receding horizon optimization framework that systematically coordinates high-level assets and operational decisions across the network. This approach is especially relevant for modern DHNs, where managing the interplay between production units, storage systems, and demand-side flexibility requires sophisticated decision-making over varying time horizons [13, 15]. By combining predictive capabilities with feedback mechanisms—effectively managing demand forecast errors and model uncertainties—MPC provides a robust optimization-based strategy for operational planning. These characteristics make MPC well-suited to address the scheduling and coordination challenges inherent in 4th and 5th generation district heating networks [7].

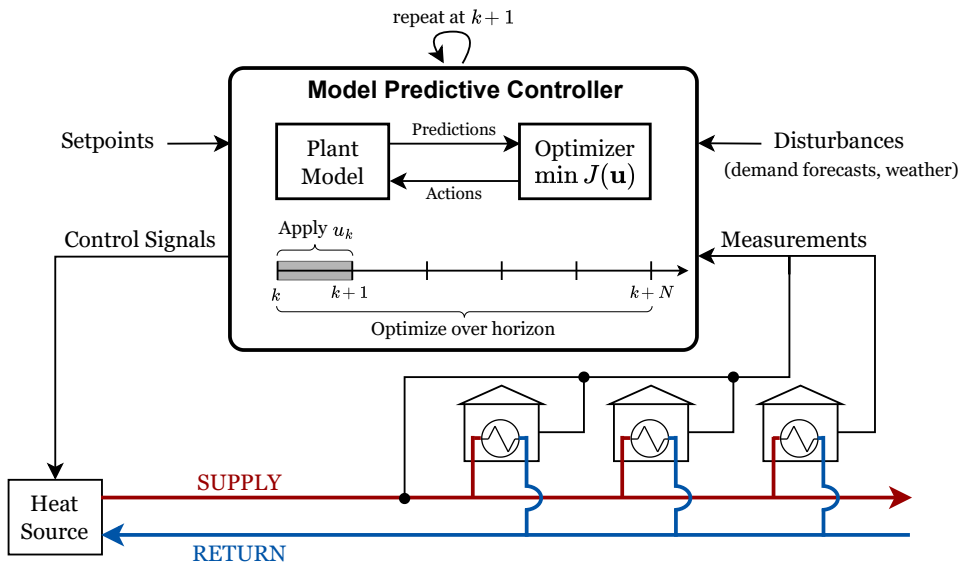


Figure 1.3: Overview of MPC for DHNs. The controller computes an optimal input sequence $\mathbf{u} = (u_k, \dots, u_{k+N-1})$ over a prediction horizon N . Following the receding horizon principle, only u_k is applied, after which new measurements are collected and the optimization is repeated.

1.2. Challenges in Predictive Control of District Heating Networks

WHILE model predictive control offers a systematic framework for DHN optimization, its practical implementation faces three interconnected challenges that must be addressed to enable real-time operation. First, the spatial extent and governing physics of DHNs create modeling challenges that complicate straightforward representation for control. Second, the resulting problem structure creates computational challenges that impact real-time tractability. Third, the cost of comprehensive instrumentation creates observability challenges that limit available state information. Each challenge demands specific methodological solutions to bridge the gap between MPC theory and practical DHN operation.

Modeling Complexity. DHN dynamics are governed by multiscale coupled partial differential equations describing advection-dominated heat transport and nonlinear hydraulic behavior across networks spanning multiple kilometers [16–19]. Accurate prediction requires capturing transport delays that can exceed ten hours [20], preventing the steady-state assumptions common in power system optimization [21] and necessitating spatial discretization schemes that expand model dimensionality into hundreds or thousands of states [22, 23]. The resulting discrete-time models inherit strong nonconvexities from three sources: bilinear coupling between temperature and mass flow in thermal power relationships, quadratic dependence of pressure losses on flow rates, and combinatorial complexity from bidirectional flow modeling. These characteristics place DHN optimal control in the class of large-scale nonlinear programs, raising the question:

RQ1 — How can we develop models for control of DHNs that capture essential dynamics while remaining suitable for embedding within optimization frameworks?

Computational Tractability. Even with careful modeling, the nonconvex optimization problems arising from DHN control are generally solvable only to local optimality, making solution quality highly sensitive to initialization [24, 25]. This is compounded by the need for sufficiently long prediction horizons, often 24 to 48 hours. Longer horizons substantially increase problem size, creating a practical dilemma: operational control requires solutions within minutes, yet problem complexity can render optimizations intractable within these time constraints. For large networks with multiple production units, storage facilities, and consumers, the computational burden becomes prohibitive for real-time implementation. This motivates two related questions:

RQ2 — Can we exploit the mathematical structure of DHN optimal control problems to improve computational efficiency despite inherent nonconvexity?

RQ3 — How can MPC algorithms be scaled to large networks while maintaining real-time capability?

Observability and State Estimation. Effective MPC implementation requires accurate state information, yet full state observation would require temperature and flow sensors at every network node, which is economically infeasible [26]. Network operators must therefore make strategic decisions about sensor placement to ensure observability of critical states within budget constraints, combined with state estimation methods to reconstruct unmeasured variables from available measurements [26, 27]. Standard LTI techniques for sensor placement cannot be applied since DHNs are nonlinear. This raises the question:

RQ4 — Where should sensors be strategically placed to ensure observability of the network state with a limited instrumentation budget?

To address these research questions, we first establish the mathematical background and problem formulation, before presenting the contributions of this dissertation.

1.3. Problem Formulation: Predictive Control of District Heating Networks

WE consider the district heating network dynamics [18, 28, 29] described as

$$M\dot{\zeta}(t) = \vartheta(\zeta(t), u(t), \eta(t)), \quad \zeta(0) = \zeta_0, \quad (1.1)$$

where $\zeta(t)$, $u(t)$, and $\eta(t)$ are the plant states, controls, and disturbances, respectively, defined at time instance $t \in [0, \infty)$, M is a diagonal positive (semi)definite matrix, and ζ_0 refers to the initial conditions of the system.

In DHN control applications, economic nonlinear model predictive control has proven effective for determining optimal control policies for systems described by (1.1), combining real-time operational optimization with trajectory tracking [30, 31] within a receding horizon framework. The practical implementation of such receding horizon controllers for DHNs requires a discrete-time finite-dimensional approximation of the continuous-time system (1.1), expressed as

$$\xi(k+1) = f(\xi(k), u(k), \eta(k)), \quad (1.2)$$

where $\xi(k) \in \mathbb{R}^n$, $u(k) \in \mathbb{R}^m$, and $\eta(k) \in \mathbb{R}^p$ are respectively the state, control input, and disturbance after spatiotemporal discretizations, $k \in \mathbb{Z}$ indicates the discrete-time index, and $f: \mathbb{R}^n \times \mathbb{R}^m \times \mathbb{R}^p \rightarrow \mathbb{R}^n$ is the state-transition map, commonly assumed to be Lipschitz continuous in its arguments. In practice, the disturbance $\eta(k)$ is replaced by its forecasted value $w(k)$, which has bounded error with respect to the true disturbance, i.e., we have

$$\eta(k) = w(k) + \epsilon(k), \quad (1.3)$$

where $\epsilon(k) \in [\underline{\epsilon}, \bar{\epsilon}]$. In many applications, notably energy systems, disturbances often represent consumer energy demand that can be reliably predicted over extended periods using historical consumption data and weather forecasts. Therefore, the disturbance sequence is henceforth assumed to be known over the prediction horizon. While

slight forecast errors inevitably arise in practice, these are naturally compensated for by the receding horizon implementation, which provides feedback through repeated re-optimization with updated forecasts. Accordingly, at each time $k \in \mathbb{Z}$, the problem of finding an optimal trajectory for (1.1) is defined in the form of a finite-horizon optimal control problem as

$$\begin{aligned}
 J_N(\xi(k)) = \min_{\substack{\{x(i)\}_{i=0}^N, \\ \{u(i)\}_{i=0}^{N-1}}} & \sum_{i=0}^{N-1} \ell(x(i), u(i)) + \ell_f(x(N)) \\
 \text{s.t.} & \quad x(i+1) = f(x(i), u(i), w(i)), \\
 & \quad i = 0, \dots, N-1, \\
 & \quad x(i) \in \mathbb{X}, \quad u(i) \in \mathbb{U}, \\
 & \quad x(0) = \xi(k), \quad x(N) \in \mathbb{X}_f,
 \end{aligned} \tag{1.4}$$

where $\xi(k)$ is the given initial state, N is the prediction horizon, $\ell : \mathbb{R}^{n+m} \rightarrow \mathbb{R}$ represents the stage cost, and $\ell_f : \mathbb{R}^n \rightarrow \mathbb{R}$ denotes the terminal cost. Furthermore, the sets \mathbb{X} and \mathbb{U} represent polytopic constraints on the system, and \mathbb{X}_f is the terminal region. In general, the stage cost is a weighted combination of tracking and economic functions, promoting a trade-off between the corresponding objectives. For example, ℓ might penalize temperature deviations from setpoints, minimize heat production, or maximize system-wide profit.

While terminal costs and sets are conventionally employed to guarantee stability and recursive feasibility [31], their systematic determination is intractable for the complex DHN systems (1.1). Therefore, sufficiently long prediction horizons are required to guarantee both practical stability and to promote recursive feasibility [32, 33], compensating for the absence of stabilizing terminal ingredients, and meanwhile, increasing the complexity of the predictive control problem (1.4).

The mathematical structure of problem (1.4) inherits the physical challenges described in Section 1.2. The nonconvex nonlinear dynamics f in (1.2) arise from the bilinear temperature-flow coupling, quadratic friction, and flow directionality requirements discussed previously. The high dimensionality of $\xi \in \mathbb{R}^n$ (with n often in the hundreds or thousands) reflects the spatial discretization needed to capture transport delays across large networks. Moreover, the limited observability of $\xi(k)$ necessitates state estimation methods that can operate with partial measurements.

Solving (1.4) is subject to a hard real-time constraint: the optimization must be completed within the sampling interval to deliver updated control actions before the next decision moment. For the large-scale, nonconvex problems arising from DHN applications, this requires methodological innovations beyond standard MPC algorithms.

1.4. Research Contributions and Dissertation Outline

THIS dissertation makes several principal contributions to the theory and practice of predictive control for district heating networks, addressing the modeling, computational, and observability challenges identified in Section 1.2 in the context

of the formulation presented in Section 1.3. The contributions are stated below—by chapter—in the order they appear in the thesis.¹

2 — Control-oriented modeling framework for DHN dynamics

We develop a systematic modeling approach for district heating networks that balances physical fidelity with computational tractability for use within optimization-based control. The framework employs graph-theoretic representations combined with discretized thermo-hydraulic dynamics. This provides the mathematical foundation for all subsequent control algorithms while maintaining prediction accuracy sufficient for operational decision-making.

3 — Frequential Gramian-based optimal sensor placement

We formulate the sensor placement problem for DHN observability using frequential Gramians, providing a computationally efficient method for determining minimal sensor sets that ensure the required observability. Our approach extends submodular optimization techniques from LTI systems to linear periodic systems, a critical extension for DHNs whose dynamics vary periodically with daily and weekly demand patterns. The resulting sensor configurations enable accurate state estimation for feedback control while respecting practical instrumentation budgets, addressing a key implementation barrier for MPC deployment in large-scale networks.

4 — Convexified hydraulic economic nonlinear MPC formulation

We introduce an economic MPC formulation that exploits the convex structure of DHN thermal dynamics through convexification of the hydraulic constraints related to Kirchhoff's laws. This approach reduces operating costs by 9% compared to standard rule-based control while achieving fewer constraint violations.

5 — Dissipativity analysis for economic MPC of DHNs

We establish dissipativity properties for the economic MPC formulation developed in Chapter 4, providing theoretical guarantees for closed-loop stability and performance without terminal constraints or terminal costs when the stage cost is linear. For nonlinear cost functions, we use sum of squares programming to computationally verify strict dissipativity by constructing valid storage functions for several example costs, thereby extending the stability guarantees beyond the linear cost case. This analysis bridges the gap between practical MPC implementations that omit terminal constraints and the theoretical literature requiring them for stability guarantees.

6 — Computational acceleration through dual dynamic programming

We develop a horizon decomposition strategy using dual dynamic programming that resolves the computational trade-off between MPC short-sightedness with brief horizons and intractability with long horizons. The method approximates the tail-end value function through nested Benders' decomposition, enabling shorter optimization horizons while preserving long-term economic optimality. Numerical

¹Related publications per chapter are as follows: 2: [34], 3: [35], 4: [34], 5: [36], 6: [37], 7: [38].

results demonstrate significant computational speedup compared to the full nonlinear optimization while maintaining comparable economic performance.

7 — Koopman operator-based predictive control

We propose a Koopman operator framework for learning DHN dynamics that preserves the linear MPC structure while capturing essential nonlinearities from the DHN dynamics. We develop a physics-guided systematic methodology to construct meaningful Koopman observables by integrating the graph topology and fundamental thermodynamic conservation laws governing the DHN. This data-driven approach significantly reduces computational burden compared to physics-based nonlinear MPC while maintaining control performance, offering a promising alternative when physics-based models are unavailable or computationally prohibitive.

Collectively, these contributions provide a comprehensive methodology for implementing predictive control in modern district heating networks, advancing both the theoretical understanding and practical deployment of advanced control strategies for thermal energy systems.

An overview of the dissertation structure and the recommended reading sequence is presented in Figure 1.4.

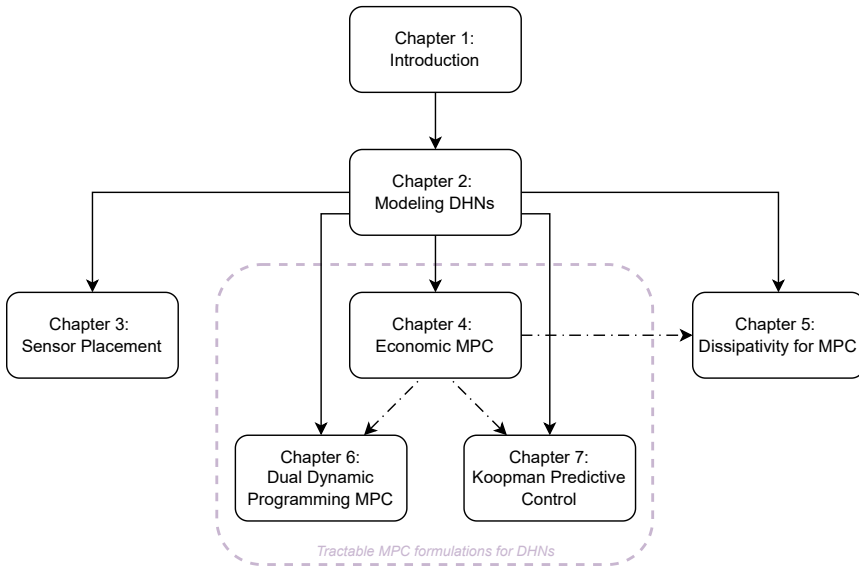


Figure 1.4: Structural overview of the dissertation. Solid arrows indicate prerequisite chapters that should be read first; dash-dotted lines indicate optional reading sequences left to the reader’s discretion.

2

Modeling for Control of District Heating Networks

This chapter develops an adaptable modeling framework suitable for real-time optimization and multi-purpose modeling of district heating networks. The framework integrates physical governing equations with graph-theoretic principles and offers adjustable spatial resolution, establishing the foundation for subsequent chapters.

This chapter is based on [34].

2.1. District Heating Network Model

IN this section, we detail the fundamental modeling choices and conventions. Firstly, we define the graph-theoretic notions used to model the DHN in Section 2.1.1. Subsequently, we discuss the transient thermal dynamics of pipe flow in Section 2.1.2, and conservation constraints that are enforced through nodes in Section 2.1.3. Finally, we introduce an example DHN called the AROMA network in Section 2.1.4.

The modeling framework builds upon a fundamental characteristic of DHNs: their symmetric structure consisting of parallel supply and return pipelines. The supply network delivers hot water from heat sources to consumers, while the mirrored return network carries the cooled water back for reheating. This configuration enables cost-effective installation through shared infrastructure while supporting various pipe arrangements for efficient heat distribution [39, 40].

2.1.1. Graph Model

The DHN consists of hydraulic and thermal components such as pipes, junctions, pumps, valves, heat exchangers, and buffers. We model this network as a strongly connected directed graph $\mathcal{G} = (\mathcal{N}, \mathcal{E})$ with a set of nodes \mathcal{N} that represent junctions in the DHN, which are connected by edges $\mathcal{E} \subseteq \mathcal{N} \times \mathcal{N}$, representing pipelines that may be fitted with pumps, valves, or heat exchangers. Note that strong connectivity of \mathcal{G} ensures no mass exits the system.

For the ease of discussion, let \mathcal{N} be characterized as $\mathcal{N} := \{1, 2, \dots, |\mathcal{N}|\}$. The *adjacency matrix* $\mathbf{D} \in \{0, 1\}^{|\mathcal{N}| \times |\mathcal{N}|}$ of \mathcal{G} describes node-to-node connections, i.e., for any $(i, j) \in \mathcal{N} \times \mathcal{N}$, we have

$$\mathbf{D}_{ij} = \begin{cases} 1, & \text{if } (i, j) \in \mathcal{E}, \\ 0, & \text{otherwise.} \end{cases} \quad (2.1)$$

Given an enumeration for the edges of the graph modeling our DHN, the *incidence matrix* $\mathbf{E} \in \{-1, 0, 1\}^{|\mathcal{N}| \times |\mathcal{E}|}$ of \mathcal{G} describes the edge-to-node relationship, i.e., for any $(i, k) \in \mathcal{N} \times \{1, 2, \dots, |\mathcal{E}|\}$, we have

$$\mathbf{E}_{ik} = \begin{cases} -1, & \text{if the } k^{\text{th}} \text{ edge exits node } i, \\ 1, & \text{if the } k^{\text{th}} \text{ edge enters node } i, \\ 0, & \text{otherwise.} \end{cases} \quad (2.2)$$

Moreover, let \mathcal{S}_f be the set of fundamental cycles in \mathcal{G}^1 , such that we have the fundamental loop matrix $\mathbf{F} \in \{0, 1\}^{|\mathcal{S}_f| \times |\mathcal{E}|}$ defined for any $(\ell, k) \in \mathcal{S}_f \times \mathcal{E}$ by

$$\mathbf{F}_{\ell k} = \begin{cases} 1, & \text{if edge } k \in \mathcal{S}_{f,\ell} \text{ and cycle } \mathcal{S}_{f,\ell} \in \mathcal{S}_f, \\ 0, & \text{otherwise,} \end{cases}$$

where the rows of \mathbf{F} are independent.

In our model of the DHN, nodes correspond to volume-less junctions and edges correspond to pipes that may be equipped with heat exchangers, valves, or pumps [29]. Storage systems are modeled by a fixed volume set of pipe segments, similar to [41].

¹see Section 2.2.2 for details on the derivation of \mathcal{S}_f

2.1.2. Edge Dynamics

We model the dynamics of an edge representing a pipe, potentially equipped with a heat exchanger, pump, or a valve, through an approximation of the 1-dimensional compressible Euler equations and the thermal energy equation for cylindrical pipes [18, 23]. Hence, the dynamics of edge $e \in \mathcal{E}$ are described through the following PDEs:

$$\partial_t \rho_e + \partial_x (\rho_e v_e) = 0, \quad (2.3)$$

$$\partial_t (\rho_e v_e) + \partial_x (\rho_e v_e^2) + \partial_x p_e + \rho_e g \hat{z}_e + K_e \frac{\rho_e}{2d_e} |v_e| v_e = 0, \quad (2.4)$$

$$\partial_t T_e + v_e \partial_x T_e + \frac{p_e}{\rho_e c_p} \partial_x v_e - \frac{K_e}{2c_p d_e} |v_e| v_e^2 + \frac{4U_e}{\rho_e c_p d_e} (T_e - T_a) = 0, \quad (2.5)$$

which are essentially characterizing the temporal and spatial evolution of three central variables; *temperature* $T_e(t, x)$ [K], *flow velocity* $v_e(t, x)$ [m s^{-1}], and *pressure* $p_e(t, x)$ [Pa] of water. The other parameters are the density of water ρ_e [kg m^{-3}], gravity g [m s^{-2}], slope of pipe \hat{z}_e [-], friction coefficient K_e [-], diameter of pipe d_e [m], heat transfer coefficient U_e [$\text{J m}^{-2} \text{K}^{-1}$], specific heat capacity of water c_p [$\text{J kg}^{-1} \text{K}^{-1}$], and ambient temperature T_a [K].

For the purpose of scheduling and control, several simplifying assumptions are in place. First, liquid water is effectively incompressible, and its density varies by less than 5% over the typical DHN operating range of 40-120 °C. This variation is negligible compared to other sources of uncertainty such as demand forecast errors and heat loss estimation. Hence, it is commonly assumed [18, 23] that the water inside the system is incompressible and has constant density, i.e., we have $\partial_t \rho_e = 0$, and $\partial_x \rho_e = 0$, for each $e \in \mathcal{E}$. It then follows from the continuity equation that $\partial_x v_e = 0$ for each $e \in \mathcal{E}$. Moreover, heat generated through friction is negligible in practice [23], particularly compared to other terms in (2.5). Therefore, we omit the term $\frac{K_e}{2c_p d_e} |v_e| v_e^2$ from (2.5). Furthermore, since pipelines are typically laid underground at constant depth [40], we assume, without loss of generality, that there are no elevation differences throughout the network, i.e., $\hat{z}_e = 0$, $\forall e \in \mathcal{E}$. Additionally, considering the significant separation in time scales between thermal and hydraulic dynamics, and since the frictional term in (2.4) dominates the inertial term [42], one can neglect dynamics on the flow rate, i.e., $\partial_t v_e = 0$, for any $e \in \mathcal{E}$. Combined with $\partial_x v_e = 0$, this eliminates all inertial terms from equation (2.4), resulting in quasi-steady hydraulics. Thus, for the dynamics of edge $e \in \mathcal{E}$, we have the following equations:

$$\partial_x v_e = 0 \quad (2.6)$$

$$\partial_x p_e + K_e \frac{\rho_e}{2d_e} |v_e| v_e = 0, \quad (2.7)$$

$$\partial_t T_e + v_e \partial_x T_e + \frac{4U_e}{\rho_e c_p d_e} (T_e - T_a) = 0. \quad (2.8)$$

2.1.3. Nodal Constraints

The nodal constraints follow from conservation laws. At each network node, the conservation of mass principle combined with the incompressibility assumption requires

that the total volumetric flow rate entering the node equals the total volumetric flow rate exiting the node. More precisely, for any node $n \in \mathcal{N}$, we define the edge sets $\mathcal{E}_{\rightarrow n} = \{e \in \mathcal{E} : e \text{ enters } n\}$ and $\mathcal{E}_{n \rightarrow} = \{e \in \mathcal{E} : e \text{ exits } n\}$ as the set of entering and exiting edges, respectively. Subsequently, we obtain

$$\sum_{e \in \mathcal{E}_{\rightarrow n}} q_e(t) = \sum_{e \in \mathcal{E}_{n \rightarrow}} q_e(t), \quad (2.9)$$

where $q_e(t) = \Phi_e v_e(t)$ is the flow rate in pipe e with Φ_e the cross-section of pipe. Additionally, energy balance should also be considered for any node, which can be described as a mixing rule determining the relationship between the exit temperature of a node as a function of temperatures of entering flows. Accordingly, one can employ a mixing rule that takes a flow-weighted average of the incoming temperatures, as in [29] and [18], to obtain the temperature of a node. More precisely, for the temperature of any node $n \in \mathcal{N}$, we have

$$T_n(t) = \frac{\sum_{e \in \mathcal{E}_{\rightarrow n}} q_e(t) T_e(t)}{\sum_{e \in \mathcal{E}_{\rightarrow n}} q_e(t)}. \quad (2.10)$$

Also, the temperature of any edge exiting a node is expected to be the same as the node temperature, i.e., we have

$$T_e(t) = T_n(t), \quad \forall n \in \mathcal{N}, \forall e \in \mathcal{E}_{n \rightarrow}. \quad (2.11)$$

The equations presented so far give rise to a graph theoretical model for DHNs. Thus, the system is described through a set of partial differential algebraic equations (PDAEs), namely (2.7) – (2.11), which is used as the base formulation for our future derivation in this chapter and thesis. It is worth noting the DHN model in the current form is not applicable for optimization purposes mainly due to the complex infinite-dimensional nature of the mentioned PDAEs. Accordingly, in the following sections, we develop more tractable and yet realistic models, suitable for optimization-based control strategies.

2.1.4. AROMA: A Benchmark for District Heating Networks

We introduce the AROMA district heating network, which is widely used in DHN literature as a benchmark example to evaluate the performance of numerical algorithms, see, e.g., [18, 41]. Therefore, throughout this thesis, we employ this network to demonstrate the developed methods. The network, depicted in Figure 2.1, features multiple consumers, producers and a storage unit. Originally, the AROMA network consisted of only a single producer without any storage included [18]. Recently [41], storage was featured in the network by augmenting the AROMA network model. Specifications and parameters of the AROMA network are provided in Section 4.6.2. Considering that we are focused on the general case of DHNs, namely with multiple producers/prosumers, for the demonstration of the proposed modeling scheme and the numerical validation of the developed methodology in this chapter, we additionally augment the AROMA DHN to accommodate the mentioned extensions and features. Nevertheless, one should note that our methodologies are applicable beyond the AROMA network.




Figure 2.1: The AROMA district heating network with multiple consumers, a prosumer, a storage buffer, and a loop.

2.2. DHN Model: A Tractable Reformulation for the Hydraulics

IN this section, we improve the tractability of the DHN model introduced in Section 2.1, with a particular focus on the hydraulic dynamics described by the equations (2.7) and (2.9). To this end, we first introduce a new graph to accommodate bidirectional flows. Secondly, we describe the method used to compute a set of independent cycles that fully describe the mass flowing through the network. Thirdly, we address the treatment of the momentum equation and introduce a convex reformulation for these constraints.

2.2.1. Fixing the Flow Direction

Considering the graph abstraction of a district heating network as in Figure 2.1, each pipeline is represented by a string of edges connected in series. Thus, in the coarsest representation of our system, we have exactly one edge for each pipeline until it reaches a junction. Let us denote this minimal directed graph as $\mathcal{G}_0 = (\mathcal{N}_0, \mathcal{E}_0)$. Hence, the network has $|\mathcal{E}_0|$ pipelines, where with respect to each $e \in \mathcal{E}_0$, we define flow variables q_e that belong to the interval $\mathcal{Q}_e := \{q_e : \underline{q}_e \leq q_e \leq \bar{q}_e\}$, where $\bar{q}_e > 0$.

While seemingly a relatively minor detail, allowing the direction of the flow to switch during the operational phase has significant implications for the modeling procedure and complicates the design of a tractable controller. Therefore, for the edges where directional switching is allowed, i.e., for $e \in \mathcal{E}_0$ with $\underline{q}_e < 0$ and $\bar{q}_e > 0$, we decompose the corresponding flow variable as

$$q_e = q_e^+ - q_e^- \text{ with } q_e^+, q_e^- \geq 0 \text{ and } q_e^+ q_e^- = 0. \quad (2.12)$$

In terms of the graph, we obtain a new graph $\mathcal{G}_0^+ = (\mathcal{N}_0, \mathcal{E}_0^+)$ with additional set of edges $\mathcal{E}_0^+ = \mathcal{E}_0 \cup \Delta_{\mathcal{E}_0}$, where $\Delta_{\mathcal{E}_0} := \{(v, u) \notin \mathcal{E}_0 \mid e = (u, v) \in \mathcal{E}_0 \text{ and } \underline{q}_e < 0\}$. To illustrate, we depict \mathcal{G}_0^+ for the AROMA network in Figure 2.2. Here, \mathcal{G}_0^+ is directed, and double-sided arrows indicate pipelines that allow bidirectional flow. The direction matching the original orientation in \mathcal{G}_0 will have the flow q_e^+ associated with it, while the newly added reverse edge has q_e^- associated with it.

Remark 1 In typical DHNs, bidirectional flow is neither necessary nor practical for all pipeline segments. However, bidirectional flow capability is essential for enabling




Figure 2.2: A directed graph abstraction of the AROMA district heating network with fixed flow directions. Double-sided arrows indicate two opposite facing edges.

prosumer and storage tank operation in the network. The determination of flow directionality is made during the design phase through systematic analysis of network architecture. Therefore, not all edges in Figure 2.2 are bidirectional. \diamond

2.2.2. Independent Flows

We denote the vector of edge flow rates as $q = [q_e]_{e \in \mathcal{E}_0^+}$. The flow on an edge depends linearly on the flow on all other edges in the network due to conservation of mass. As a result, we can reduce the number of free flow variables that we need to optimize for. To this end, we introduce the *reduced loop matrix* \mathbf{F}_r which maps the reduced flow vector $q_r \in \mathbb{R}_+^{m_r}$ to q through

$$q = \mathbf{F}_r^\top q_r, \quad (2.13)$$

with $m_r < |\mathcal{E}_0^+|$. Intuitively, the vector q_r represents the flow which circulates the network, meaning that it passes through a supply section and its mirrored return section reaching back to its starting point. In the process, these flows pass through producers, consumers, or storages. Additionally, we introduce the *fundamental loop matrix* \mathbf{F} which maps all fundamental flows $q_f \in \mathbb{R}^{m_f}$ to q , i.e.,

$$q = \mathbf{F}^\top q_f, \quad (2.14)$$

with $m_f \leq m_r$. The vector q_f depends linearly on the elements of q_r . The matrices \mathbf{F}_r and \mathbf{F} consist of rows that describe directed cycles within \mathcal{G}_0^+ . The difference is that \mathbf{F} is full row rank while \mathbf{F}_r does not need to be (if \mathbf{F}_r is full rank then \mathbf{F} and \mathbf{F}_r coincide). Similarly, q_r has nonnegative elements only, whereas the entries of q_f are not required to be nonnegative. The reason for introducing both \mathbf{F}_r and \mathbf{F} is that the nonnegativity of q_r in (2.13) significantly improves the numerical results compared to using (2.14), while F is primarily employed for theoretical purposes, as further detailed in Section 4.3.

In [43], a method is presented to compute the fundamental loop matrix \mathbf{F} by setting the free flow variables as the flows through the chords of the spanning tree of \mathcal{G}_0 . Consequently, a fundamental loop is defined as the loop that is formed whenever a chord is re-connected to the spanning tree. Then, the fundamental loop matrix \mathbf{F} has elements $\mathbf{F}_{ij} \in \{-1, 0, 1\}$, for all i and j , depending on the orientation of the chord and whether an edge is part of a fundamental loop.

To preserve nonnegativity of all elements in \mathbf{F} , we develop a different approach to compute \mathbf{F} . To this end, we need the notion of a *directed cycle* [44].

Definition 1 (Directed cycle) A directed cycle \mathcal{C} in \mathcal{G}_0 is a sequence of nodes and edges as $(n_0, e_1, n_1, e_2, \dots, e_k, n_k)$ such that

1. n_0, n_1, \dots, n_k are different nodes;
2. $e_i = (n_{i-1}, n_i)$ for $i = 1, \dots, k$;
3. $n_0 = n_k$,

where two sequences are an equivalent cycle if one can be obtained from the other by a cyclic permutation.

Let $\mathcal{S}(\mathcal{G}_0) := \{\mathcal{C}_i(\mathcal{G}_0)\}$ be the set of all directed cycles within \mathcal{G}_0 . The matrix \mathbf{F}_r and \mathbf{F} can be computed through the following procedure:

1. Compute all directed cycles $\mathcal{S}(\mathcal{G}_0^+)$ in \mathcal{G}_0^+ ;
2. Remove any cycles that consist of at most two nodes;
3. Remove any asymmetric cycles, i.e., cycles for which their path is not mirrored between the supply and return network. We call this reduced set $\mathcal{S}_r(\mathcal{G}_0^+)$;
4. Define reduced loop matrix \mathbf{F}_r with elements

$$\mathbf{F}_{r,ij} = \begin{cases} 1, & \text{if edge } j \in \mathcal{C}_i \text{ and cycle } \mathcal{C}_i \in \mathcal{S}_r(\mathcal{G}_0^+), \\ 0, & \text{otherwise,} \end{cases}$$

where \mathbf{F}_r has $m_r = |\mathcal{S}_r(\mathcal{G}_0^+)|$ rows and $|\mathcal{E}_0^+|$ columns;

5. Compute \mathbf{F} by performing a pivoted QR factorization on \mathbf{F}_r^\top as

$$\mathbf{F}_r^\top P = QR, \quad (2.15)$$

where $R \in \mathbb{R}^{|\mathcal{E}_0^+| \times m_r}$ is an upper triangular matrix, $Q \in \mathbb{R}^{|\mathcal{E}_0^+| \times |\mathcal{E}_0^+|}$ is an orthonormal matrix, $P \in \mathbb{R}^{m_r \times m_r}$ is a permutation matrix, and the columns of $\mathbf{F}_r^\top P$ are projected onto an orthonormal basis spanned by the columns of Q . If $\text{rank}(\mathbf{F}_r) = m_f < m_r$, all elements R_{ii} with $i > m_f$ are zero, indicating that corresponding columns of $\mathbf{F}_r^\top P$ lie in the subspace of all prior columns. Hence, we obtain the fundamental loop matrix by taking only the first m_f columns of $\mathbf{F}_r^\top P$ as follows:

$$\mathbf{F} = \left(\mathbf{F}_r^\top P \begin{bmatrix} I_{m_f} \\ \mathbf{0} \end{bmatrix} \right)^\top. \quad (2.16)$$

The set of remaining cycles is denoted as $\mathcal{S}_f(\mathcal{G}_0^+)$.

Matrix \mathbf{F} , constructed via QR factorization with column pivoting on \mathbf{F}_r^\top , inherits nonnegativity from \mathbf{F}_r and possesses full column rank, ensuring mass conservation within the network.

Remark 2 We note that the rows of the computed fundamental loop matrix span the same basis as the row space of the matrices introduced in [43] and [29]. The difference is that, in this context, we exclude asymmetric cycles without loss of generality. \diamond

2.3. DHN Model: A Tractable Reformulation for the Thermal Dynamics

2.3.1. Thermal Node Model

To obtain a tractable formulation for the thermal dynamics in the network, a suitable discrete spatial approximation of (2.8) is required, for which we apply an upwind scheme. The upwind finite volume discretization is standard practice in the district heating literature, see, e.g., [18, 22, 23, 42], and provides well-known stability guarantees for hyperbolic conservation laws [45]. The well-posedness of the underlying PDE and the discretized system has been established in [28]. Therefore, the dynamics of the i^{th} finite volume cell of water can be described by the following scalar continuous-time ordinary differential equation (ODE):

$$\begin{aligned} V_i \dot{T}_i &= -q_i(T_i - T_{i-1}) - \alpha_i(T_i - T_a) + w_i, \\ y_i &= q_i(T_i - T_a), \end{aligned} \quad (2.17)$$

where $V_i = \Phi_i \Delta x_i$ denotes the cell volume for some spatial discretization step Δx_i , q_i denotes the mass flow of water, T_i denotes the temperature of water in the cell, T_{i-1} is the temperature of the inflow into the cell, $\alpha_i = 4U_i V_i / \rho c_p d_i$ is the heat loss coefficient, T_a is the ambient temperature, w_i is a variable denoting the transfer of heat from or to the environment, and y_i is the output which is proportional to the exergy. Note that the choice of spatial discretization step Δx_i has implications for the stability and accuracy of the resulting discrete-time models, particularly in relation to the Courant-Friedrichs-Lewy (CFL) condition. This is discussed further in Section 4.4. Additionally, note that, similarly to [29], when the cell represents a heat exchanger, w_i indicates the transfer of heat from one side to the other side, and otherwise, the term w_i can be dropped from (2.17). Hence, w_i is a control variable if it corresponds to a controllable producer, and a disturbance if it corresponds to an uncontrollable producer or consumer.

We define the state variable \tilde{T}_i as $\tilde{T}_i = T_i - T_a$, assuming that the ambient temperature is equal and constant for all cells, and consider the corresponding dynamics described as

$$\begin{aligned} V_i \dot{\tilde{T}}_i &= -(q_i + \alpha_i) \tilde{T}_i + u_i + w_i, \\ y_i &= q_i \tilde{T}_i. \end{aligned} \quad (2.18)$$

Note that (2.17) is equivalent to (2.18) when $u_i = q_i \tilde{T}_{i-1}$. From here on, we consider any finite volume cell as a node, or more precisely, a thermal node (TN), in the graph. The thermal node has a compartmental structure as shown in Figure 2.5a.

2.3.2. Interconnected DHN Model

To construct the model of the thermal system, we approximate the spatial evolution of temperature along pipelines, described by (2.8), through a finite sequence of partitions. We adopt an approach similar to [22], where nodes represent volumes of water and edges represent flow rates.

Closed- and open-loop networks

We introduce two types of network topologies that may be used to represent a DHN.

Definition 2 (Closed thermal network) A closed thermal network is a network for which \mathcal{G}_0 is strongly connected. See Figure 2.3 for an example of a closed thermal network. In practice, district heating networks are always closed circuits, ensuring mass is conserved within the system.




Figure 2.3: Schematic representation of a **closed thermal network** with five thermal nodes.

Definition 3 (Open thermal network) An open thermal network is a network for which \mathcal{G}_0 is weakly connected. Additionally, to ensure mass conservation, all source nodes have an external input acting on them (and all sink nodes have an output flow associated to them). See Figure 2.4 for an example of an open thermal network.




Figure 2.4: Schematic representation of a **open thermal network** with four thermal nodes.

Remark 3 The rationale for introducing open thermal network formulations is that decoupling the supply and return networks can prove advantageous in certain instances,

for example, to facilitate model reduction. Fundamentally, both open and closed network models are governed by identical graph-theoretic principles, and the subsequent modeling procedures apply equally to both configurations. It should be noted, however, that open network formulations require the introduction of dummy nodes at each consumer to ensure that the incidence matrix properly accounts for outgoing edges connecting to the return network. Furthermore, open thermal networks have an additional input at producer nodes, representing the total energy of the incoming water from the return network. \diamond

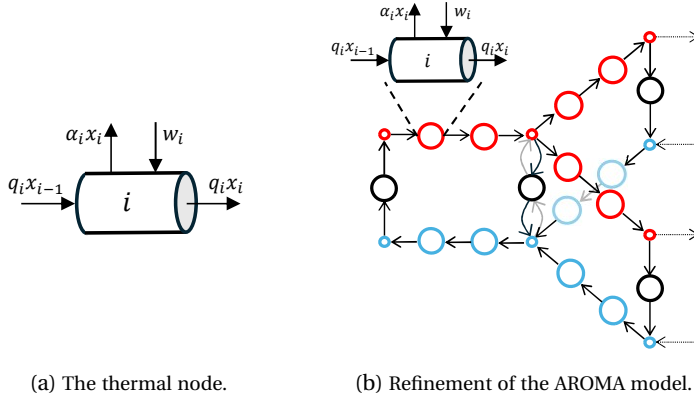


Figure 2.5: Illustration of thermal node model and mesh refinement of the DHN graph. (a) shows the conceptual thermal node model with its dynamics defined by (2.18), while (b) demonstrates the injection of thermal nodes in a section of the AROMA network.

Updating the graph

We introduce intermediate nodes on all edges of the original graph to increase the granularity of the model and improve the approximation of (2.8). These nodes follow the compartmental dynamics of (2.18), i.e., thermal nodes. Given the original graph $\mathcal{G}_0^+ = (\mathcal{N}_0, \mathcal{E}_0^+)$ with adjacency matrix \mathbf{D}_0^+ , we add $l_{x,i}$ thermal nodes to each edge i , creating $m_c = \sum_{i=1}^{|\mathcal{E}_0^+|} l_{x,i}$ total new nodes. The resulting augmented graph has adjacency matrix \mathbf{D} of dimension $|\mathcal{N}_0| + m_c$, as partially shown in Figure 2.5b. Further details and explanation on how to obtain \mathbf{D} are provided in Appendix 2.A.

Remark 4 The extension of the graph does not change the independent flow distribution in the network presented in Section 2.2. Without loss of generality and for the ease of notation, we use q to denote the extended mass flow variable. Nonetheless, considering the introduced graph augmentation, the precise definition of the flow variable in the new setting is $q = [q_i \otimes \mathbf{I}_{p_i}]_{i \in \mathcal{E}_0^+}$. The same is true for the loop matrices \mathbf{F} and \mathbf{F}_r , which will have additional columns due to the added edges but remain structurally the same. Therefore, we henceforth keep the same notation for these matrices when referring to the extended graph. \diamond

From graph to state-space

Let $\mathcal{G} = (\mathcal{N}, \mathcal{E})$ be the resulting graph with $|\mathcal{N}| = |\mathcal{N}_0| + m_c$ nodes and incidence matrix \mathbf{E} , where $|\mathcal{N}_0|$ and m_c denote the number of junctions and thermal nodes, respectively. To model energy flow rates between nodes, we define edge variable $\varphi_e(q_e, \tilde{T}_e) = q_e \tilde{T}_e$, which requires the root node temperature for each edge since temperatures are node-specific quantities. We obtain the network-wide energy flows by

$$\varphi = \frac{1}{2} \mathbf{Q} (|\mathbf{E}| - \mathbf{E})^\top \tilde{T}, \quad (2.19)$$

where $|\mathbf{E}|$ denotes the element-wise absolute value operator of \mathbf{E} , i.e., $|\mathbf{E}| = [|\mathbf{E}_{ik}|]$, \tilde{T} is the vector of state variables defined as $\tilde{T} = [\tilde{T}_i]_{i \in \mathcal{N}}$, and \mathbf{Q} is a diagonal matrix defined as $\mathbf{Q} = \text{diag}(q_i)_{i \in \mathcal{E}}$. Subsequently, by multiplying φ with the full incidence matrix, we obtain the energy balance on each node. More precisely, we have

$$\phi = \mathbf{E}\varphi, \quad (2.20)$$

where $\phi \in \mathbb{R}^{|\mathcal{N}|}$ represents the nodal rate of change in energy due to in-flows and out-flows. By including external effects from heat exchangers and environment, we can introduce the complete state-space description through the thermal dynamics of the network. To this end, given flow vector q , we define matrix $A(q)$ as

$$A(q) = \frac{1}{2} \mathbf{E}\mathbf{Q} (|\mathbf{E}| - \mathbf{E})^\top - D_\alpha, \quad (2.21)$$

where D_α is the diagonal matrix of heat loss coefficients characterized as $D_\alpha = \text{diag}(\alpha_i)_{i \in \mathcal{N}}$. Also, let \mathcal{W} denote the set of nodes corresponding to heat exchangers in the network, i.e., $\mathcal{W} := \{i \in \mathcal{N} : w_i \neq 0\}$. Furthermore, we define a $|\mathcal{N}| \times |\mathcal{W}|$ matrix, denoted by B , with the entry in the i^{th} row and j^{th} column given by

$$B_{ij} = \begin{cases} \frac{1}{\rho c_p}, & \text{if } j = 1, \dots, |\mathcal{W}| \text{ and } i = i_j, \\ 0, & \text{otherwise,} \end{cases} \quad (2.22)$$

where $i_1, i_2, \dots, i_{|\mathcal{W}|}$ are the elements of \mathcal{W} sorted in an increasing order, i.e., $i_1 < i_2 < \dots < i_{|\mathcal{W}|}$. Accordingly, we describe the thermal dynamics of a *closed or open* thermal network by

$$\begin{aligned} V \dot{\tilde{T}} &= A(q) \tilde{T} + Bw, \\ z &= C \tilde{T}, \end{aligned} \quad (2.23)$$

where $V = \text{diag}(V_i)_{i \in \mathcal{N}}$ is the volume matrix, $w = [w_i]_{i \in \mathcal{W}}$ is the vector of external inputs and disturbances. Note that \mathcal{W} can be partitioned into disjoint subsets \mathcal{W}_C and \mathcal{W}_P to distinguish between consumer and producer interactions, respectively.

Additionally, in the case of open thermal networks, one must distinguish between two components of the thermal power $w_i \in \mathcal{W}_P$: the power supplied by a connected heat exchanger (if present) and the rate of energy provided by the incoming fluid from the return network. Specifically, this can be expressed as $w_i = w_{i,h} + w_{i,f}$ for any $i \in \mathcal{W}_P$, where $w_{i,h}$ denotes the heat exchanger contribution and $w_{i,f}$ represents the fluid power

input which itself is expressed as the product of flow rate with incoming temperature, i.e., $w_{i,f} = q_i \tilde{T}_{i,\text{in}}$.

The modeling structure adopted throughout this thesis varies between open and closed network formulations depending on the specific requirements of each chapter. More specifically, the following configurations are employed: Chapter 3 makes no assumptions regarding network topology and encompasses both open and closed network models; Chapter 4 utilizes a closed network representation; Chapter 5 considers both configurations; Chapter 6 employs an open thermal network representation; and finally, Chapter 7 assumes a closed network representation.

2.A. Appendix: Derivation of Adjacency Matrix

To construct the model of the thermal system, we approximate the spatial evolution of temperature along pipelines, described by (2.8), through a finite sequence of partitions. We adopt an approach similar to [22], where nodes represent volumes of water and edges represent flow rates.

Consider the original graph $\mathcal{G}_0^+ = (\mathcal{N}_0, \mathcal{E}_0^+)$ with adjacency matrix \mathbf{D}_0^+ . We introduce a vector $l_x \in \mathbb{R}^{|\mathcal{E}_0^+|}$ that contains the number of additional nodes being inserted on each edge. The new graph is then augmented on the original graph, i.e., we introduce $m_c = \sum_{i=1}^{|\mathcal{E}_0^+|} l_{x,i}$ new nodes and remove all original edges from \mathcal{G}_0^+ . The new connections are added to the adjacency matrix \mathbf{D} with dimension $|\mathcal{N}_0| + m_c$ as follows:

1. We consider all original nodes to be labeled by their order in the adjacency matrix, i.e., $1, 2, \dots, |\mathcal{N}_0|$, and all newly added nodes labelled as $|\mathcal{N}_0|+1, |\mathcal{N}_0|+2, \dots, |\mathcal{N}_0|+m_c$.
2. For all $e = (i, j) \in \mathcal{E}_0^+$, the newly added path from i to j is $p_{ij} = (i, |\mathcal{N}_0| + k + 1, \dots, |\mathcal{N}_0| + k + l_{x,e}, j)$, where $k = \sum_{i=1}^{e-1} l_{x,i}$.
3. If the path has fixed direction such that $(j, i) \notin \mathcal{E}_0^+$, we define

$$\mathbf{D}_{p_{ij}, p_{ij}} = \begin{bmatrix} 0 & 1 & 0 & \cdots & 0 \\ 0 & 0 & 1 & & 0 \\ \vdots & \vdots & \ddots & \ddots & \vdots \\ 0 & 0 & \cdots & 0 & 1 \\ 0 & 0 & \cdots & 0 & 0 \end{bmatrix}, \quad (2.24)$$

otherwise, if $(j, i) \in \mathcal{E}_0^+$, we define

$$\mathbf{D}_{p_{ij}, p_{ij}} = \begin{bmatrix} 0 & 1 & 0 & \cdots & 0 \\ 1 & 0 & 1 & & 0 \\ \vdots & \ddots & \ddots & \ddots & \vdots \\ 0 & \vdots & \ddots & 0 & 1 \\ 0 & 0 & \cdots & 1 & 0 \end{bmatrix}. \quad (2.25)$$

3

Sensor Placement in District Heating Networks

District heating networks (DHNs) are essential in providing efficient heating services to urban areas through networked pipes. The performance of these systems critically depends on the strategic placement of thermal storage buffers (actuators) and temperature sensors throughout the network. Due to the inherent slow dynamics of thermal transport, these systems exhibit significant delays and periodic behaviors that necessitate time-varying analysis approaches. This chapter presents a frequency-domain framework for optimal sensor placement in DHNs, focusing on metrics derived from frequential Gramians. We provide rigorous analysis of two key metrics, namely the trace and log-determinant of the frequential Gramian, establishing submodularity properties and performance guarantees for greedy selection algorithms. Our theoretical framework naturally handles both the periodic nature of DHNs and their slow transients, outperforming standard approaches in estimation accuracy.

This chapter is based on [35].

3.1. Introduction

A fundamental challenge in district heating network (DHN) operation is the strategic placement of thermal storage buffers (actuators) and temperature sensors throughout the network to control it [29, 34]. The slow dynamics of thermal transport, with delays of hours from generation to consumption [20, 37], complicate this task. While prior work explored optimal storage [46] and sensor placement [47], challenges persist. Poor actuator placement delays responses and wastes energy, while suboptimal sensors hinder thermal state estimation [48].

Traditional actuator and sensor placement methods rely on controllability and observability metrics from linear system theory [49, 50], using Gramian analysis to assess energy requirements for state control and measurement. However, these approaches are insufficiently equipped to handle the slow periodic dynamics present in DHNs, characterized by daily demand cycles [51]. Beyond classic LTI-based techniques, several approaches were developed for nonlinear or time-varying systems, including empirical Gramians [52], data-driven methods [53], and time-domain analysis for time-varying systems [54]. While powerful in specific contexts, these methods are not necessarily suited for DHNs: empirical Gramians and time-domain approaches require intensive computation for large-scale networks, and data-driven approaches raise concerns regarding the collection and quality of training data and the design of basis functions.

Recent advances in frequency-domain analysis, particularly frequential Gramians, provide an efficient framework for studying DHN dynamics. These methods offer both theoretical and computational advantages over time-domain approaches [55], naturally capturing the system's periodicity and slow thermal transients.

This chapter presents a frequency-domain framework for optimal sensor placement in DHNs, focusing on observability metrics derived from frequential Gramians. We provide rigorous analysis of two key metrics: the trace and log-determinant of the frequential observability Gramian, establishing submodularity properties and performance guarantees for greedy selection algorithms. Our theoretical framework naturally handles both the periodic nature of DHN dynamics and their slow transients, offering computational advantages over traditional time-domain approaches. Furthermore, we demonstrate how these metrics can be efficiently computed using truncated Fourier representations.

3.2. Optimal Sensor Placement Problem for District Heating Networks

IN this section, we formalize the minimal sensor placement problem for DHNs. By providing a precise mathematical formulation of the optimization problem. In essence, the sensor and actuator placement problem for DHNs is an offline design problem. Furthermore, the flow rate $q(t)$ typically follows cyclical flow patterns driven by consistent consumer demand cycles [51]. Despite variations in these daily demand profiles due to external factors, such as weather conditions and occupancy rates, the underlying periodic nature of these phenomena provides a sound basis for analysis. Hence, for this instance of the problem we consider $Q(t) := \text{diag}(q_e(t))_{e \in \mathcal{E}}$ from (2.21) to be a known periodic function, leading to the control-autonomous periodic linear

continuous-time DHN dynamics

$$\dot{x}(t) = A(t)x(t), \quad A(t + T_p) = A(t), \quad (3.1)$$

where $x(t)$ is used to represent shifted temperature state \tilde{T} , $A(t)$ comes from (2.21), and $T_p \in \mathbb{R}$ denotes the period of the system.

We now formally present the problem of determining the optimal placement of sensors and actuators within the network. Consider the following T_p -periodic linear system, representing the DHN dynamics described for any $t \in \mathbb{R}$ as

$$\begin{aligned} \dot{x}(t) &= A(t)x(t) + B_{\mathcal{I}}u(t), & A(t + T_p) &= A(t), \\ y(t) &= C_{\mathcal{J}}x(t), \end{aligned} \quad (3.2)$$

where $x(t) \in \mathbb{R}^m$, $u(t) \in \mathbb{R}^{|\mathcal{I}|}$, $y(t) \in \mathbb{R}^{|\mathcal{J}|}$, and $A(\cdot)$ is continuous. Furthermore, $B_{\mathcal{I}}$ and $C_{\mathcal{J}}$ are structured matrices where $B_{\mathcal{I}} = I_m(\mathcal{I})$ and $C_{\mathcal{J}} = I_m(\mathcal{J})$. Here, I_m is the $m \times m$ identity matrix, with \mathcal{I} and \mathcal{J} indicating the positions of columns and rows corresponding to dedicated actuators and sensors that control and measure single nodes, respectively. While we restrict the analysis in this chapter to the placement of sensors, the results readily extend to the placement of actuators by invoking arguments similar to duality between controllability and observability in LTI systems. Furthermore, one should note that the arguments and discussion in this chapter can be generalized to any systems satisfying (3.2).

Our goal is to identify a minimal subset of sensors $\mathcal{J} \subseteq \{1, 2, \dots, m\}$ ensuring sufficient sensing capabilities. To quantify sensing performance, we introduce a performance metric $\rho: 2^{\mathbb{N}_m} \mapsto \mathbb{R}$ based on the *time-averaged frequential observability Gramian* $\Lambda_{\mathcal{J}}$, defined rigorously in Section 3.3, which generalizes the classical observability Gramian to periodically time-varying linear systems as a measure of observability. We employ two established metrics: the *geometric mean energy*, defined as $\rho(\mathcal{J}) = \log \det(\Lambda_{\mathcal{J}}^{-1})$, and the *average energy*, defined as $\rho(\mathcal{J}) = \text{tr}(\Lambda_{\mathcal{J}}^{-1})$, quantifying the effort required for state reconstruction across all state dimensions. Our objective is to solve the following optimization problem:

$$\begin{aligned} \min_{\mathcal{J} \subseteq \{1, 2, \dots, m\}} & \quad |\mathcal{J}| \\ \text{s.t.} & \quad \rho(\mathcal{J}) \leq \kappa \end{aligned} \quad (3.3)$$

where κ is a positive threshold constant. For the metrics under consideration, this formulation represents an instance of well-known NP-hard problems [56].

3.3. Time-Averaged Frequential Gramian

IN this section, we establish frequential Gramian theory, computation, and properties for submodular optimization.

The frequential Gramian derives from Fourier coefficients that encode structural information across the frequency modes of the periodic system. For a T_p -periodic matrix $A(\cdot)$, Floquet's theorem yields a T_p -periodic coordinate change $z(t) = Q(t)x(t)$ in which the uncontrolled dynamics take the time-invariant form $\dot{z}(t) = Fz(t)$, where F is constant. The eigenvalues of F are the so-called Floquet exponents, which determine the stability

of perturbations of periodic orbits. We henceforth assume F has no eigenvalue on the imaginary axis.

For the Floquet transformed system, there exist T_p -periodic Gramians that can be computed at discrete time points as the outer product of T_p -periodic frequency-domain factors [55, Prop. 2]. When these Gramians are mapped from Floquet coordinates back to physical coordinates, the resulting Gramians can be constructed using analogous periodic frequency-domain factors. Specifically, the frequential Gramian in physical coordinates, indicating the frequential energy content with respect to time lag t , is defined as [55, Eqn. 21]:

$$G_f(t, \mathcal{J}) = \frac{1}{2\pi} \int_{-\infty}^{\infty} \Psi(\gamma, t, \mathcal{J}) \Psi(\gamma, t, \mathcal{J})^H d\gamma, \quad (3.4)$$

where G_f is well-defined (see Lemma 4), and $\Psi : \mathbb{R} \times [0, T_p) \times 2^{\mathbb{N}_m} \mapsto \mathbb{C}^{m \times m}$ are the frequency-domain factors given by

$$\Psi(\gamma, t, \mathcal{J}) = \sum_{k \in \mathbb{Z}} \Psi_k(\gamma, \mathcal{J}) e^{ik\omega t}, \quad t \in [0, T_p), \quad (3.5)$$

where $\omega = \frac{2\pi}{T_p}$ is the *fundamental frequency*.

In what follows, we provide the Fourier coefficients $\Psi_k(\gamma, \mathcal{J})$ following [55, Prop. 3], where the reader can find further details on their construction. Let $(A_k)_{k \in \mathbb{Z}}$ be the Fourier coefficients of $A(\cdot)$, and define $\mathbf{R}_k = (-ik\omega I_m + A_0)$. Given the corresponding Fourier basis for $A(\cdot)$, we define the infinite-dimensional matrix $\mathbf{T} = [\mathbf{T}_{ij}]_{i,j=-\infty}^{\infty}$, consisting of the block matrices $\mathbf{T}_{ij} \in \mathbb{C}^{m \times m}$ defined as

$$\mathbf{T}_{ij} = \begin{cases} \mathbf{R}_i, & \text{if } i = j, \\ A_{i-j}, & \text{if } i \neq j. \end{cases} \quad (3.6)$$

Subsequently, for any $\gamma \in \mathbb{R}$, define $\mathcal{H}(\gamma) = (i\gamma I_\infty - \mathbf{T})^{-1}$. It is worth noting that $\mathcal{H}(\gamma)$ is called the *harmonic resolvent operator*, mapping harmonic components between inputs and outputs, i.e., $[\mathcal{H}(\gamma)]_{k,\ell}$ represents the m by m block of $\mathcal{H}(\gamma)$ mapping the inputs at frequency $\gamma + \ell\omega$ to the outputs at frequency $\gamma + k\omega$, for any $k, \ell \in \mathbb{Z}$. Accordingly, for any $k \in \mathbb{Z}$, $\gamma \in \mathbb{R}$, and $\mathcal{J} \subseteq \mathbb{N}_m$, the function $\Psi_k : \mathbb{R} \times 2^{\mathbb{N}_m} \rightarrow \mathbb{C}^{m \times m}$ is defined analogously to [55, Eqn. 31]; however, $C_{\mathcal{J}}$ is treated as time-invariant, so only the $j = 0$ component is nonzero:

$$\Psi_k(\gamma, \mathcal{J}) = [\mathcal{H}(\gamma)^H]_{k,0} C_{\mathcal{J}}^H. \quad (3.7)$$

Thus, we have a closed-form expression for the Fourier coefficients, which subsequently enables the computationally efficient approximation of the frequential Gramian in (3.4).

In this chapter, we are interested in optimal sensor placement to maximize the average temporal sensing performance for the DHN. Therefore, we define the *time-averaged frequential Gramian* $\Lambda_{\mathcal{J}} : 2^{\mathbb{N}_m} \rightarrow \mathbb{R}^{m \times m}$ as

$$\Lambda_{\mathcal{J}} = \frac{1}{T_p} \int_0^{T_p} G_f(\tau, \mathcal{J}) d\tau. \quad (3.8)$$

Proposition 1 For a given set of sensors $\mathcal{J} \subseteq \mathbb{N}_m$ and a system with period T_p , the time-averaged frequency Gramian $\Lambda_{\mathcal{J}} \in \mathbb{R}^{m \times m}$ satisfies

$$\Lambda_{\mathcal{J}} = \frac{1}{2\pi} \int_{-\infty}^{\infty} \sum_{k \in \mathbb{Z}} \Psi_k(\gamma, \mathcal{J}) \Psi_k(\gamma, \mathcal{J})^H d\gamma. \quad (3.9)$$

Proof: Using Tonelli's theorem, we can exchange the order of integration to get

$$\Lambda_{\mathcal{J}} = \frac{1}{2\pi T} \int_{-\infty}^{\infty} \int_0^{T_p} \Psi(\gamma, \tau, \mathcal{J}) \Psi(\gamma, \tau, \mathcal{J})^H d\tau d\gamma. \quad (3.10)$$

Subsequently, substituting in the Fourier series representation of $\Psi(\gamma, \tau, \mathcal{J})$ from (3.5), the inner integral becomes

$$\begin{aligned} & \int_0^{T_p} \Psi(\gamma, \tau, \mathcal{J}) \Psi(\gamma, \tau, \mathcal{J})^H d\tau \\ &= \int_0^{T_p} \sum_{k \in \mathbb{Z}} \sum_{l \in \mathbb{Z}} \Psi_k(\gamma, \mathcal{J}) \Psi_l(\gamma, \mathcal{J})^H e^{i(k-l)\omega\tau} d\tau \\ &= T_p \sum_{k \in \mathbb{Z}} \Psi_k(\gamma, \mathcal{J}) \Psi_k(\gamma, \mathcal{J})^H, \end{aligned} \quad (3.11)$$

where the last equality follows from Fubini's theorem and the orthogonality of complex exponentials, i.e., the integral $\int_0^{T_p} e^{i(k-l)\omega\tau}$ equals T_p if $k = l$, and 0 if $k \neq l$. Therefore, the time-averaged frequency Gramian takes the form given in (3.9), which completes the proof. ■

Interpretation: For each γ , the operator $\Psi_k(\gamma, \mathcal{J})$ characterises how the k^{th} Floquet harmonic of the state manifests in the measured output channels indexed by \mathcal{J} . The Hermitian product $\Psi_k(\gamma, \mathcal{J}) \Psi_k(\gamma, \mathcal{J})^H$ therefore quantifies the observable energy of that harmonic. Integrating over γ and summing over $k \in \mathbb{Z}$ collect these contributions into $\Lambda_{\mathcal{J}}$.

Remark 5 The temporal dependency of $G_f(t, \mathcal{J})$ on t stems from the system's convergence to a periodic trajectory. For any t , the Gramian already contains full period information. While $G_f(t, \mathcal{J})$ could be used to determine optimal active sensors, tracking the period for sensor scheduling presents practical implementation challenges difficult to overcome in real-world applications. Hence, we focus on the time-averaged frequency Gramian for sensor placement. ◇

3.3.1. Approximate Gramian Computation

The formulation in (3.9) reveals that the time-averaged frequency Gramian decomposes into a sum of contributions from each frequency mode. Nonetheless, exact computation of $\Lambda_{\mathcal{J}}$ is challenging as it requires inverting an infinite-dimensional matrix and evaluating an improper integral. We introduce an approximation to make this computation tractable and tight, namely by truncation of higher-order Fourier modes and discretization of the frequency domain. In particular, we consider a uniformly sampled set of frequencies $\Gamma := \{\gamma_1, \dots, \gamma_L\}$, and a truncation order r of the Fourier series of $A(\cdot)$, i.e.,

$\widehat{A}_r(t) := \sum_{k=-r}^r A_k e^{ik\omega t}$. This consequently reduces \mathbf{T} to a finite matrix of dimension $(2r+1)m \times (2r+1)m$, with $\widehat{\mathcal{H}}_r(\gamma)$ representing the corresponding truncation of $\mathcal{H}(\gamma)$. We formalize the mentioned results below.

Proposition 2 Given Γ and r , for the time-averaged frequential Gramian $\Lambda_{\mathcal{J}}$, we have

$$\Lambda_{\mathcal{J}} = \frac{1}{\pi} \sum_{\gamma_l \in \Gamma} \xi_l \sum_{k=-r}^r \widehat{\Psi}_k(\gamma_l, \mathcal{J}) \widehat{\Psi}_k(\gamma_l, \mathcal{J})^H + \mathcal{R}_{\mathcal{J}}(r, L), \quad (3.12)$$

where ξ_l are quadrature coefficients,

$$\widehat{\Psi}_k(\gamma_l, \mathcal{J}) = [\widehat{\mathcal{H}}_r(\gamma_l)^H]_{k,0} C_{\mathcal{J}}^H$$

are approximations of the Fourier coefficients $\Psi_k(\gamma_l, \mathcal{J})$ at frequency γ_l , and the residual term $\mathcal{R}_{\mathcal{J}}(r, L)$ goes to zero as $r \rightarrow \infty$ and $L \rightarrow \infty$.

Proof: The convergence of the residual term to zero is directly implied from the uniform convergence of the Fourier series of Ψ over compact subsets of the domain of γ and the integrability of $G_f(t)$. ■

Corollary 3 The time-averaged frequential Gramian can be efficiently approximated and computed by determining $\Psi_k(\gamma_l, \mathcal{J})$ only over a narrow band of positive frequencies.

Proof: For the efficient approximation of the truncated Gramian, i.e., the first term in right-hand side of (3.12), we exploit the fact that for any $\gamma \in \mathbb{R}$, there exists an $\ell \in \mathbb{Z}$ such that $\alpha = \gamma - \ell\omega \in (-\omega/2, \omega/2]$, and

$$\Psi_k(\gamma, \mathcal{J}) = [\mathcal{H}(\alpha)^H]_{k,\ell} C_{\mathcal{J}}^H. \quad (3.13)$$

Additionally, since the dynamics are real-valued, their Fourier coefficients are symmetric and we get from (3.7) and the definition of \mathbf{T} that $\Psi_k(-\gamma, \mathcal{J}) = \text{conj}(\Psi_k(\gamma, \mathcal{J}))$. These properties reduce the frequency evaluation range to $\alpha \in (-\omega/2, \omega/2] \cap [0, \infty) = [0, \omega/2]$. For each shift index $\ell \in \{-r, \dots, 0, \dots, r\}$, define

$$\mathbf{C}_{\mathcal{J}}^{(\ell)} = \mathbf{e}_{\ell+r+1} \otimes C_{\mathcal{J}}^H \quad (3.14)$$

where $\mathbf{e}_{\ell+r+1}$ is the $(\ell+r+1)$ -th unit basis vector in \mathbb{R}^{2r+1} . We find $\Psi_{\mathcal{J},\alpha_l}^{(\ell)} \in \mathbb{R}^{(2r+1)m \times |\mathcal{J}|}$ by solving the linear system

$$(\mathbf{i}\alpha_l I_m - \mathbf{T})^H \Psi_{\mathcal{J},\alpha_l}^{(\ell)} = \mathbf{C}_{\mathcal{J}}^{(\ell)}, \quad (3.15)$$

for $\alpha_l \in [0, \omega/2]$ and $l \in \{1, \dots, L\}$. We can retrieve the approximate value of $\Psi_k(\gamma_l, \mathcal{J})$ at $\gamma_l = \alpha_l + c\omega$ by taking the k^{th} block of $\Psi_{\mathcal{J},\alpha_l}^{(\ell)}$, i.e., we have

$$\widehat{\Psi}_k(\gamma_l, \mathcal{J}) = \left[\Psi_{\mathcal{J},\alpha_l}^{(\ell)} \right]_k = [\widehat{\mathcal{H}}_r(\alpha_l)^H]_{k,\ell} C_{\mathcal{J}}^H. \quad (3.16)$$

Applying numerical quadrature to (3.9) with these approximated coefficients and exploiting the established symmetry properties yields (3.12), which completes the proof. ■

3.3.2. Properties of Frequency Gramians

We state two important results on the properties of the time-averaged frequency Gramian.

Lemma 4 Suppose none of the Floquet exponents of the system lie on the imaginary axis. Then, for any $\mathcal{J} \subseteq \mathbb{N}_m$, the time-averaged frequency Gramian $\Lambda_{\mathcal{J}}$, introduced in (3.8), is well-defined and positive semidefinite, i.e., $\Lambda_{\mathcal{J}} \geq 0$. Moreover, it is additive, i.e., for any $\mathcal{J}_1, \mathcal{J}_2 \subseteq \mathbb{N}_m$, we have

$$\Lambda_{\mathcal{J}_1 \cup \mathcal{J}_2} = \Lambda_{\mathcal{J}_1} + \Lambda_{\mathcal{J}_2 \setminus \mathcal{J}_1}. \quad (3.17)$$

Proof: We prove each property separately.

Well-definedness: For both stable and unstable time-periodic systems, the frequency Gramian $G_f(t)$ is well-defined since none of the Floquet exponents of the system lie on the imaginary axis. This follows from the fact that each Floquet-transformed Fourier coefficient of the Gramian satisfies a Sylvester equation, as exemplified in [55, Prop. 1]. It follows that $\Lambda_{\mathcal{J}}$ is also well-defined.

Positive Semidefiniteness: Let $x \in \mathbb{R}^n$. Then

$$\begin{aligned} x^\top \Lambda_{\mathcal{J}} x &= \frac{1}{2\pi} \int_{-\infty}^{\infty} \sum_{k \in \mathbb{Z}} x^\top \Psi_k(\gamma, \mathcal{J}) \Psi_k(\gamma, \mathcal{J})^H x \, d\gamma \\ &= \frac{1}{2\pi} \int_{-\infty}^{\infty} \sum_{k \in \mathbb{Z}} \|\Psi_k(\gamma, \mathcal{J})^H x\|^2 \, d\gamma \geq 0. \end{aligned} \quad (3.18)$$

Since the above expression is nonnegative for any $x \in \mathbb{R}^n$, we have $\Lambda_{\mathcal{J}} \geq 0$ for all $\mathcal{J} \subseteq \mathbb{N}_m$.

Additivity: Consider two sets $\mathcal{J}_1, \mathcal{J}_2 \subseteq \mathbb{N}_m$. By the structure of $C_{\mathcal{J}}$, we can write

$$\begin{aligned} \Psi_k(\gamma, \mathcal{J}_1 \cup \mathcal{J}_2) &= [\mathcal{H}(\gamma)^H]_{k,0} C_{\mathcal{J}_1 \cup \mathcal{J}_2}^H \\ &= [\mathcal{H}(\gamma)^H]_{k,0} \begin{bmatrix} C_{\mathcal{J}_1}^H & C_{\mathcal{J}_2 \setminus \mathcal{J}_1}^H \end{bmatrix} \\ &= [\Psi_k(\gamma, \mathcal{J}_1) \quad \Psi_k(\gamma, \mathcal{J}_2 \setminus \mathcal{J}_1)]. \end{aligned} \quad (3.19)$$

Substituting this expression into the definition of the frequency Gramian and using the linearity of integration, we obtain $\Lambda_{\mathcal{J}_1 \cup \mathcal{J}_2} = \Lambda_{\mathcal{J}_1} + \Lambda_{\mathcal{J}_2 \setminus \mathcal{J}_1}$, which completes the proof. ■

Lemma 5 For LTI systems, the time-averaged frequency observability Gramian from (3.9) reduces to the standard observability Gramian.

Proof: Consider a stable LTI system with $A(t) = A$ for all $t \in \mathbb{R}$. For brevity, we write C instead of $C_{\mathcal{J}}$. The standard observability Gramian W_o satisfies $A^\top W_o + W_o A = -C^\top C$. For LTI systems, (3.7) yields a single nonzero Fourier coefficient $\Psi_0(\gamma, \mathcal{J}) = (i\gamma I_m - A^\top)^{-1} C^\top$ since off-diagonal blocks of \mathbf{T} vanish. Substituting into (3.9), we obtain the time-averaged frequency observability Gramian for LTI systems:

$$\Lambda_{\mathcal{J}}^{\text{LTI}} = \frac{1}{2\pi} \int_{-\infty}^{\infty} (i\gamma I_m - A^\top)^{-1} C^\top C (-i\gamma I_m - A)^{-1} \, d\gamma. \quad (3.20)$$

This coincides with the standard observability Gramian for LTI systems, establishing that the time-averaged frequency formulation generalizes the classical result. ■

3.4. Greedy Sensor Selection

THIS section introduces key concepts from submodular optimization theory and analyzes submodularity properties of common energy metrics. We present a greedy algorithm for sensor selection with guaranteed optimality bounds.

3.4.1. Supermodular and Non-Supermodular Set Functions

Definition 4 (Supermodularity) Let \mathbb{N}_m be a finite ground set and $\rho : 2^{\mathbb{N}_m} \rightarrow \mathbb{R}$ be a set function. Function ρ is supermodular if for all $\mathcal{J}_1 \subseteq \mathcal{J}_2 \subseteq \mathbb{N}_m$ and $s \in \mathbb{N}_m \setminus \mathcal{J}_2$, we have

$$\rho(\mathcal{J}_1) - \rho(\mathcal{J}_1 \cup \{s\}) \geq \rho(\mathcal{J}_2) - \rho(\mathcal{J}_2 \cup \{s\}). \quad (3.21)$$

Definition 5 (Non-increasing Set Function) A set function $\rho : 2^{\mathbb{N}_m} \rightarrow \mathbb{R}$ is said to be non-increasing if for all $\mathcal{J}_1, \mathcal{J}_2 \subseteq \mathbb{N}_m$ and $\mathcal{J}_1 \subseteq \mathcal{J}_2$, one has $\rho(\mathcal{J}_1) \geq \rho(\mathcal{J}_2)$.

We provide two key theorems on the supermodularity of two established metrics.

Theorem 6 Let $\Lambda_{\mathcal{J}}$ be the time-averaged frequential Gramian as defined in (3.8), with $\mathcal{J} \subseteq \mathbb{N}_m$. The set function

$$\rho(\mathcal{J}) = \log \det(\Lambda_{\mathcal{J}} + \varepsilon I_m)^{-1} \quad (3.22)$$

is supermodular and non-increasing.

Proof: From Lemma 4, $\Lambda_{\mathcal{J}}$ is well-defined, positive semidefinite (with $\Lambda_{\mathcal{J}} + \varepsilon I_m$ positive definite) and satisfies the additivity property. Following similar steps to those in [57, Prop. 2], we have that (3.22) is a supermodular and non-increasing set function. ■

Remark 6 Let $\tilde{\Lambda}_{\mathcal{J}}$ denote the first term in the right-hand side of (3.12), i.e., the approximate time-averaged frequential Gramian. One can verify that $\tilde{\Lambda}_{\mathcal{J}}$ is well-defined, positive semidefinite, and additive using similar arguments as in the proof of Lemma 4. In (3.22), we can alternatively use $\tilde{\Lambda}_{\mathcal{J}}$. The supermodularity of the resulting set function can be shown following the same steps as in the proof of Theorem 6. ◇

Proposition 7 Let $\Lambda_{\mathcal{J}}$ be the time-averaged frequential Gramian as defined in (3.8), with $\mathcal{J} \subseteq \mathbb{N}_m$. The set function

$$\rho(\mathcal{J}) = \text{tr} \left((\Lambda_{\mathcal{J}} + \varepsilon I_m)^{-1} \right), \quad (3.23)$$

is non-supermodular.

Proof: Lemma 5 provides a case where the time-averaged frequential Gramian reduces to the standard observability Gramian. Following steps similar to those in [58], one can show by counterexample that (3.23) is not supermodular. ■

The regularization parameter ε in (3.22) and (3.23) guarantees the existence of well-defined solutions to $\rho(\mathcal{J})$ for arbitrary $\mathcal{J} \subseteq \mathbb{N}_m$, thereby eliminating the necessity of imposing observability a priori. Moreover, as $\varepsilon \rightarrow 0$, the regularized function approaches the original, unperturbed function [56, 57]. While the precise implications of the regularization parameter ε on observability properties in the context of frequential Gramians are less obvious, it nevertheless serves as an effective mechanism for ensuring invertibility of $\Lambda_{\mathcal{J}} + \varepsilon I_m$.

Algorithm 1 Frequency-Domain Greedy Sensor Placement

Input: constant κ , parameter ε , matrix \mathbf{T} , and function $\rho \in \{\text{tr}(\widehat{\Lambda}^{-1}), \log \det(\widehat{\Lambda}^{-1})\}$, $\widehat{\Lambda} = \widehat{\Lambda}_{\mathcal{J}} + \varepsilon I_m$.

Output: set of sensors $\mathcal{J} \subseteq \mathbb{N}_m$ with $\rho(\mathcal{J}) \leq \kappa$.

```

1:  $\mathcal{J} \leftarrow \emptyset$  and  $\widehat{\Lambda}_{\mathcal{J}} \leftarrow \mathbf{0}_{m \times m}$ 
2:  $\mathbf{T} \leftarrow$  compute  $\mathbf{T} \in \mathbb{R}^{(2r+1)m \times (2r+1)m}$  according to (3.6)
3: while  $\rho(\mathcal{J}) > \kappa$  do
4:   for  $s \in \mathbb{N}_m \setminus \mathcal{J}$  do
5:     for  $\alpha_l \in [0, \omega/2]$ , with  $l \in \{1, \dots, L\}$  do
6:       for  $\ell \in \{-r, \dots, 0, \dots, r\}$  do
7:          $\Psi_{\mathcal{J}, \alpha_l}^{(\ell)} \leftarrow$  solve  $(i\alpha_l I_m - \mathbf{T})^H \Psi_{\mathcal{J}, \alpha_l}^{(\ell)} = \mathbf{C}_s^{(\ell)}$  (3.15)
8:         if  $\gamma_l = \alpha_l + c\omega \geq 0$  then
9:            $|\Psi_{\gamma_l}| \leftarrow$  store  $\sum_{k=-r}^r \widehat{\Psi}_k(\gamma_l, \mathcal{J}) \widehat{\Psi}_k(\gamma_l, \mathcal{J})^H$ 
10:        end if
11:       end for
12:     end for
13:      $\widetilde{\Lambda}_{\mathcal{J} \cup \{s\}} \leftarrow \widetilde{\Lambda}_{\mathcal{J}} + 1/\pi \sum_{\gamma_l} \xi_l \beta_l |\Psi_{\gamma_l}|$  (3.12)
14:   end for
15:    $s^* \leftarrow \text{argmax}_{s \in \mathbb{N}_m \setminus \mathcal{J}} \{\rho(\mathcal{J}) - \rho(\mathcal{J} \cup \{s\})\}$ 
16:    $\mathcal{J} \leftarrow \mathcal{J} \cup \{s^*\}$  and  $\widetilde{\Lambda}_{\mathcal{J}} \leftarrow \widetilde{\Lambda}_{\mathcal{J} \cup \{s^*\}}$ 
17: end while

```

3.4.2. The Greedy Algorithm

For the minimal sensor placement problem with bounded energy (3.3), the greedy algorithm presented in Algorithm 1 iteratively selects sensors that provide the maximum marginal improvement in the selected energy metric. Specifically, at each iteration, we select a sensor that maximizes the decrease in $\rho(\mathcal{J}) - \rho(\mathcal{J} \cup \{s\})$ for each sensor $s \in \mathbb{N}_m \setminus \mathcal{J}$. The algorithm iterates this process until termination, which is when the performance metric $\rho(\mathcal{J})$ meets the desired tolerance κ .

Proposition 8 The computational complexity of Algorithm 1 is $O(Lm^5(2r+1)^4)$, where m is the number of possible sensors, L is the number of samples over $[0, \omega/2]$, and $2r+1$ denotes the range of considered frequency modes.

Proof: The four nested loops in lines 3 to 6 contribute $O(m^2L(2r+1))$ iterations. The most intensive operation occurs at line 7, solving a linear system with a $(2r+1)m \times (2r+1)m$ matrix, requiring $O(m^3(2r+1)^3)$ operations per solution. Combining these factors yields a dominant complexity component of $O(Lm^5(2r+1)^4)$. ■

Remark 7 Supermodularity provides greedy algorithms such as Algorithm 1 with provable approximation guarantees similar to those established for the standard observability Gramian. For the function $\rho(\mathcal{J}) = \log \det((\Lambda_{\mathcal{J}} + \varepsilon I)^{-1})$, the solution to




Figure 3.1: District heating network topology in numerical experiments.

Algorithm 1 terminating at iteration k , i.e., $\mathcal{J}^{\text{greedy}} = \mathcal{J}_k$, satisfies

$$\begin{aligned}
 \frac{|\mathcal{J}_k|}{|\mathcal{J}^*|} &\leq 1 + \log \frac{\rho(\emptyset) - \rho(\mathbb{N}_m)}{\rho(\mathcal{J}_{k-1}) - \rho(\mathbb{N}_m)} \\
 &\leq 1 + \log \frac{\rho(\emptyset) - \rho(\mathbb{N}_m)}{\kappa - \rho(\mathbb{N}_m)} \\
 &= 1 + \log \frac{-m \log(\varepsilon) + \log \det(\Lambda_{\mathbb{N}_m} + \varepsilon I_m)}{\kappa + \log \det(\Lambda_{\mathbb{N}_m} + \varepsilon I_m)},
 \end{aligned} \tag{3.24}$$

with respect to \mathcal{J}^* , the the optimal solution to (3.3). The second inequality follows from the fact that $\rho(\mathcal{J}_{k-1}) > \kappa$. \diamond

3.5. Numerical Experiments and Results

WE compare our frequential Gramian sensor placement against a standard approach using the observability Gramian with time-averaged dynamics. For fair comparison, we first execute Algorithm 1 until threshold K is reached, obtaining sensor configuration $\mathcal{J}_{\text{freq}}$. Subsequently, we run Algorithm 1 with $\bar{A} = \frac{1}{T_p} \int_0^{T_p} A(\tau) d\tau$ until a configuration \mathcal{J}_{sd} with cardinality equal to $|\mathcal{J}_{\text{freq}}|$ is found. For both methods, we employ the trace metric $\rho(\mathcal{J}) = \text{tr}((\Lambda_{\mathcal{J}} + \varepsilon I_m)^{-1})$ with $\varepsilon = 10^{-5}$.

Experiment Setup

We consider the DHN depicted in Figure 3.1 with two topologies: (a) an open-loop configuration i.e., a non-cyclic network of 10 nodes from generation point to consumers, and (b) a closed configuration which additionally features return connections from consumers to generation point. System dynamics are modeled with $\omega = 0.001\pi$ rad/s, $\zeta_e = 0.01$, for any pipe e , and truncation parameters $r = 8$. Flow velocities, i.e., $v(t) = \frac{4}{\pi d^2} q(t)$, follow a periodic pattern

$$v(t) = v_{\text{sc}} \odot (\bar{v} + \alpha s(t)), \tag{3.25}$$

$$s(t) = \sin(2\omega(t-1)) + 0.7 \sin\left(\frac{1}{2}\omega t\right), \tag{3.26}$$

Table 3.1: Frequency-based vs. Standard Sensor Placement

Parameters ($\bar{v}, \alpha, \bar{\lambda}, \beta$)	Sensors	Sensors	GIR	RMSE $\pm\sigma$		PIP (%)
	FB	SD		FB	SD	
			OL			
$(\frac{2}{5}, 0, \frac{1}{2}, \frac{1}{5})$	{9,4,2,6}	{9,4,5,1}	6.55	.133 \pm .016	.144 \pm .022	7.6
$(\frac{3}{5}, \frac{2}{5}, \frac{1}{4}, \frac{1}{10})$	{4,9,6,1}	{9,4,5,1}	2.13	.098 \pm .026	.105 \pm .028	6.7
$(\frac{3}{5}, \frac{2}{5}, \frac{1}{2}, \frac{7}{20})$	{9,4,6}	{9,4,5}	1.92	.127 \pm .024	.130 \pm .021	2.3
$(1, \frac{3}{5}, \frac{3}{4}, \frac{1}{10})$	{9,6,4,1}	{9,4,2,7}	20.95	.103 \pm .017	.102 \pm .014	-1.0
$(\frac{6}{5}, \frac{2}{5}, \frac{1}{2}, \frac{1}{10})$	{9,4,1,7}	{9,4,5,2}	2.70	.076 \pm .006	.082 \pm .007	7.3
$(\frac{6}{5}, 1, \frac{1}{2}, \frac{1}{5})$	{9,4,1,6}	{9,4,5,2}	2.77	.077 \pm .015	.083 \pm .017	7.2
			CL			
$(\frac{2}{5}, 0, \frac{1}{2}, \frac{1}{5})$	{4,7,0}	{4,9,5}	1.48	.160 \pm .023	.164 \pm .028	2.4
$(\frac{3}{5}, \frac{2}{5}, \frac{1}{4}, \frac{1}{10})$	{4,6,0}	{3,9,6}	0.98	.112 \pm .032	.111 \pm .034	-0.9
$(\frac{3}{5}, \frac{2}{5}, \frac{1}{2}, \frac{7}{20})$	{1,7,4}	{4,9,2}	7.95	.128 \pm .020	.139 \pm .025	7.9
$(1, \frac{3}{5}, \frac{3}{4}, \frac{1}{10})$	{9,6,4}	{4,9,2}	1.85	.108 \pm .013	.132 \pm .022	18.2
$(\frac{6}{5}, \frac{2}{5}, \frac{1}{2}, \frac{1}{10})$	{4,6,1,9}	{4,9,5,1}	7.53	.076 \pm .007	.082 \pm .008	7.3
$(\frac{6}{5}, 1, \frac{1}{2}, \frac{1}{5})$	{4,6,1}	{4,9,2}	2.51	.093 \pm .023	.096 \pm .021	3.1

Note: σ =Standard deviation; OL=Open-Loop, CL=Closed-Loop; FB=Frequency-based; SD=Standard.

where the scaling factors are $v_{sc,i} = 1$ for $i \in \{0, 1, 2\}$, $v_{sc,i} = \lambda(t)$ for $i \in \{3, 4\}$, and $v_{sc,i} = 1 - \lambda(t)$ for $i \geq 5$, with scaling factors controlled by distribution parameter $\lambda(t) = \bar{\lambda} + \beta \sin(\frac{\pi t}{500})$. We conduct Monte Carlo simulations with 100 trials across several parameter sets denoted by $(\bar{v}, \alpha, \bar{\lambda}, \beta)$.

Performance Evaluation

In each trial, we implement a standard Kalman filter for each sensor configuration starting at a randomized $t_0 \sim \mathcal{U}(0, T_p)$ and $x_{\text{true}}(t_0) \sim \mathcal{N}(0, I_m)$ to capture as much as possible the variation in the dynamics. We compute the root mean square error (RMSE) using

$$\text{RMSE}(t_k) = \left[\frac{1}{n} \sum_{i=1}^n (x_{\text{true},i}(t_k) - \hat{x}_{\text{KF},i}(t_k))^2 \right]^{\frac{1}{2}}. \quad (3.27)$$

Performance improvement percentage (PIP) is quantified as

$$\text{PIP} (\%) = 100 \cdot \frac{\text{Mean RMSE}_{\text{sd}} - \text{Mean RMSE}_{\text{freq}}}{\text{Mean RMSE}_{\text{sd}}}. \quad (3.28)$$

Furthermore, we compute the Gramian improvement ratio (GIR), i.e., the ratio between the frequential Gramians obtained from the frequency-based set $\mathcal{J}_{\text{freq}}$ and the set obtained using the standard metric \mathcal{J}_{sd} , which is defined as $\text{GIR} = \Lambda_{\mathcal{J}_{\text{sd}}} / \Lambda_{\mathcal{J}_{\text{freq}}}$. Results are shown in Table 3.1.

Discussion

Our analysis comparing frequency-based and standard sensor placement methods shows the frequency-based approach demonstrated positive performance in five of six cases for both open-loop and closed-loop configurations, with PIPs ranging from -1.0% to 7.6% in open-loop and -0.9% to 18.2% in closed-loop. Higher mean velocities generally resulted in improvements exceeding 7%. Interestingly, Gramian ratio values did not reliably predict performance gains; the configuration with the highest ratio of 20.95 showed a performance decrease, while others with lower GIR values sometimes achieved significant improvements. These findings indicate that frequency-based sensor placement offers meaningful advantages over standard approaches, though benefits vary significantly with specific network configurations.

3.6. Conclusion

WE presented a novel frequency-domain framework for sensor placement in district heating networks using frequential Gramians to capture both periodic behaviors and slow thermal transients. We established key theoretical properties, including frequency mode decomposition that enables efficient computation. Our numerical experiments demonstrate that frequency-based sensor placement outperforms standard approaches with improvements of up to 18.2% in state estimation accuracy.

For future work, one may extend the developed methodologies to address uncertainties arising from demand variations. Furthermore, the arguments and discussion in this paper can be generalized to any systems with similar structure.

4

Economic Nonlinear Model Predictive Control for District Heating Networks

In this chapter, we propose an economic nonlinear model predictive control (MPC) algorithm for district heating networks (DHNs). The proposed method features prosumers, multiple producers, and storage systems, which are essential components of 4th generation DHNs. These networks are characterized by their ability to optimize their operations, aiming to reduce supply temperatures, accommodate distributed heat sources, and leverage the flexibility provided by thermal inertia and storage—each crucial for achieving a fossil-fuel-free energy supply. Developing a smart energy management system to accomplish these goals requires detailed models of highly complex nonlinear systems and computational algorithms able to handle large-scale optimization problems. To address this, we introduce a graph-based optimization-oriented model that efficiently integrates distributed producers, prosumers, storage buffers, and bidirectional pipe flows, such that it can be implemented in a real-time MPC setting. Furthermore, we conduct several numerical experiments to evaluate the performance of the proposed algorithms in closed-loop. Our findings demonstrate that the MPC methods achieved up to 9% cost improvement over traditional rule-based controllers while better maintaining system constraints.

This chapter is based on [34].

4.1. Introduction

Classical control methods for DHNs relied on controlling the temperature and the differential pressure at central supply units. For details on these methods, we refer to [6, 7], and the references therein. Recently, graph-based modeling techniques for DHNs have been a useful tool for developing various stabilizing controllers for both hydraulic and thermal management of DHNs. In this context, the pressure and flow regulation problem was addressed in [43, 59], temperature regulation was studied in [60, 61], and the stabilizing control of both storage and temperature in a DHN was studied in [29, 62, 63]. Moreover, the works [29, 63] introduced a multi-producer graph model integrating the dynamic evolution of storage volumes.

Operational optimization of DHNs has been a topic of interest for several decades, with earlier works dating back to the 1990s [20, 64]. These studies highlighted the complexity of managing load distribution while simultaneously minimizing supply temperatures, a challenge stemming from the nonlinear nature of thermal transients and their dependence on flow rates. Recent studies, such as [18], have addressed thermal transients and variable flow rates through the development of open-loop optimization-based controllers using graph-theoretic models based on partial differential equations (PDEs) governing the one-dimensional pipe dynamics. This work also employed a complementary constrained formulation, originating from literature on gas transportation networks [65], to manage switching flow directions. However, closed-loop implementation and important 4th generation DHN features, such as multiple producers or storage, were not studied in [18].

Regarding MPC, the nonlinear nature of DHNs often complicates the design of a real-time implementable controller. Hence, several linearized formulations have been proposed, see, e.g., [13, 15, 66, 67]. Recently, a nonlinear MPC and a mixed-integer nonlinear MPC were introduced for a small-scale network [12, 25], allowing the neglect of thermal transients to maintain a tractable formulation. Similarly, [68] developed an MPC algorithm for prosumer DHNs with storage, leaving out thermal transients and heat losses. While computationally attractive, disregarding thermal transients is not suitable for large-scale DHNs due to the significant time delays [20]. A nonlinear MPC scheme that considers thermal transients, multiple producers, and storage was considered in [69]. Nonetheless, a stabilizing scheme was employed instead of an economic scheme, the management of pressure and Kirchhoff loop constraints was not rigorously addressed, and prosumers were not included. Lastly, [41] developed a nonlinear MPC for the AROMA network, incorporating fixed-volume layered storage, as one subsystem within a multi-carrier energy system. However, this work did not account for multiple producers, prosumers, or bidirectional flows.

While significant advancements have been made in the operational optimization and control of DHNs, current methods often lack consideration of essential features of 4th generation DHNs, such as multiple producers, prosumers, bidirectional flows, and thermal transients. Additionally, many approaches do not consider an economic objective or do not consider the complexity and scalability necessary to manage large-scale networks in real time. These gaps highlight the need for further research to develop comprehensive solutions that integrate these elements and address the computational challenges inherent in MPC for DHNs.

Contributions. In this chapter, we focus on economic model predictive control of district heating networks. Hence, we consider the problem of scheduling and management for economic operation. We state our contributions as follows:

1. We provide a novel efficient extension to the DHN model to include prosumers, formerly only considered without thermal transients in [68], and a multi-producer generalization of the Kirchhoff loop convexification approach, previously limited to single producer DHNs [70].
2. We introduce a novel economic nonlinear MPC algorithm for DHNs and we provide numerical analysis on the convergence properties of the proposed controller.
3. We conduct a comprehensive study into the numerical performance of the proposed control methods in a closed-loop setting, specifically examining the added value of incorporating storage and multiple producers. Additionally, we provide an in-depth analysis of the computational efficiency of our algorithms.

4.2. Treatment of the Momentum Equation

IN practice, due to the significant time scale separation between the hydraulic and thermal dynamics, the control of corresponding components within the DHN is executed over different time intervals. Thus, we consider the hydraulics to be in steady-state. Nonetheless, one needs to certify feasibility for the hydraulic operation. More precisely, we want to identify a set \mathcal{Q} , such that for all $q \in \mathcal{Q}$, pressure remains within the limits at all nodes of the network, and Kirchhoff's second law, stating that the sum of pressure differences along each loop in the network equates to zero, is satisfied.

One can approximate the conservation of momentum equation (2.7) by substituting v_e with $v_e = \frac{4q_e}{\pi d_e^2}$ and discretizing $\partial_x p = \frac{\Delta p_e}{L_e}$, where L_e denotes the pipe length, for each $e \in \mathcal{E}$. Accordingly, one obtains the equation describing pressure drop over pipe segment e caused by friction as

$$\Delta p_e = 8\rho L_e \frac{K_e}{\pi^2 d_e^5} |q_e| q_e = R_{\mu,e} q_e^2, \quad \forall e \in \mathcal{E}, \quad (4.1)$$

where $R_{\mu,e}$ is a combined constant term representing the frictional resistance in the pipe. Additionally, the sign dependency for q_e is removed according to (2.12), i.e., the flow direction on each edge is fixed.

We define the sets $\mathcal{P} \subset \mathcal{E}$ and $\mathcal{V} \subset \mathcal{E}$ that contain the edges with a pump and a valve, respectively. Then, for each edge $e \in \mathcal{P}$, the pressure change over e is described by

$$\Delta p_e = R_{\mu,e} q_e^2 - h_e(r_e), \quad (4.2)$$

with $h_e(r_e) = c_e r_e$ being the pressure difference induced by the pump, c_e being the maximum pump head, and $r_e \in [0, 1]$ being the normalized pump head. Additionally, for each edge $e \in \mathcal{V}$, the pressure drop over e is described by

$$\Delta p_e = R_{\mu,e} q_e^2 + R_{\nu,e} (v_e) q_e^2, \quad (4.3)$$

where $0 \leq R_{v,e}(v_e) < \infty$ is a time-varying control variable and $v_e \in [0, \infty)$ is the position of the valve such that $R_{v,e}(0) = 0$ corresponds to the full opening of the valve and $\lim_{v_e \rightarrow \infty} R_{v,e}(v_e) \rightarrow \infty$ corresponds to the full closing of the valve. Here, we are not considering specific valve types and characteristics. It is sufficient to only assume that any increase in v_e will increase the resistance over the valve such that $R_{v,e}$ is strictly monotone. Nonetheless, for the ease of discussion, one can consider $R_{v,e}$ to be linear in $v_e \in [0, \infty)$, i.e., we have $R_{v,e}(v_e) = R_{v,e}v_e$, where $R_{v,e}$ is a positive scalar indicating the resistance coefficient of the valve.

4.3. Convex Reformulation of the Kirchhoff Loop Constraints

THE constraints imposed by the equations (4.1)-(4.3) are nonconvex with respect to the flow rate due to the quadratic friction terms, and therefore, they need to be reformulated. More precisely, defining $\mathcal{S}^\dagger(\mathcal{G})$ as the set of *all* cycles in \mathcal{G} independent of the direction of the edges, Kirchhoff's second law states that

$$\sum_{e \in \mathcal{C}_i} \Delta p_j = 0, \quad \forall \mathcal{C}_i \in \mathcal{S}^\dagger(\mathcal{G}), \quad (4.4)$$

i.e., for each cycle, we need to satisfy a quadratic equality constraint, which is essentially nonconvex.

The pressure is defined pointwise in space, and therefore, we consider a pressure variable p_n , for each $n \in \mathcal{N}$. Let $\mathbf{p}_n \in \mathbb{R}^{|\mathcal{N}|}$ be the vector of nodal pressures and $\Delta \mathbf{p}_e \in \mathbb{R}^{|\mathcal{E}|}$ be the vector of pressure differences over all edges. Combining (4.1)-(4.3), we have

$$-\mathbf{E}^\top \mathbf{p}_n = \Delta \mathbf{p}_e = R(v)(q \odot q) - H(r), \quad (4.5)$$

where \odot denotes the element-wise product of two vectors, $R(v) = R_\mu + R_v(v)$ is a diagonal matrix that has on its diagonal the sum of resistances due to friction and valve effects on each edge, and $H(r)$ is the vector of induced pump pressure difference for each edge. As in [29], we multiply (4.5) from the left by F to obtain the sum of pressure difference over all fundamental cycles as

$$-\mathbf{F}\mathbf{E}^\top \mathbf{p}_n = \sum_{e \in \mathcal{C}_i} \Delta p_e = 0, \quad \forall \mathcal{C}_i \in \mathcal{S}_f(\mathcal{G}). \quad (4.6)$$

Subsequently, combining (4.5) and (4.6) leads to

$$\mathbf{F}R(v)(q \odot q) = \mathbf{F}H(r), \quad (4.7)$$

implying that the sum of all pressure drops due to valve effects and friction in any directed cycle of the graph should be equal to the sum of all induced pump pressure differences within that same cycle. Similarly, we introduce the inequality

$$\mathbf{F}_r R_\mu(q \odot q) \leq \mathbf{F}_r H(\mathbf{1}), \quad (4.8)$$

representing an upper bound on the induced pressure difference along each loop and corresponds to the scenario where all valves are maximally open and all pumps are operated at maximum capacity, i.e., $v_e = 0, \forall e \in \mathcal{V}$, and $r_e = 1, \forall e \in \mathcal{P}$.

Before proceeding to the main result of this section, we need to introduce the following assumption.

Assumption 1 There are $|\mathcal{V}| \geq m_f$ edges with a valve placed throughout the network, and the position of these edges is such that $\text{rank}(\mathbf{F}\Pi) = m_f$, where $\Pi \in \{0, 1\}^{|\mathcal{E}| \times |\mathcal{V}|}$ is a selection matrix with entries

$$\Pi_{ek} = \begin{cases} 1, & \text{if valve } k \text{ is on edge } e, \\ 0, & \text{otherwise,} \end{cases}$$

for any $e \in \{1, 2, \dots, |\mathcal{E}|\}$ and any $k \in \{1, 2, \dots, |\mathcal{V}|\}$. Since $\mathbf{F}\Pi$ has the same rank as \mathbf{F} , it can be mapped back to \mathbf{F} through a linear transformation, i.e., we have

$$\mathbf{F} = \mathbf{F}\Pi\Upsilon, \quad (4.9)$$

where $\Upsilon = (\mathbf{F}\Pi)^\dagger \mathbf{F}$. Let m_Θ denote the dimension of the null space of matrix $\mathbf{F}\Pi$. We assume that there exist matrices $Z_1, Z_2 \in \mathbb{R}^{m_\Theta \times |\mathcal{E}|}$ such that

$$(\Upsilon + \Theta_{\mathbf{F}\Pi} Z_1) H(\mathbf{1}) \geq (\Upsilon + \Theta_{\mathbf{F}\Pi} Z_2) R_\mu(q \odot q) \quad (4.10)$$

where $\Theta_{\mathbf{F}\Pi} := I - (\mathbf{F}\Pi)^\dagger \mathbf{F}\Pi$ denotes the kernel of $\mathbf{F}\Pi$.

Before proceeding further, we need to highlight several remarks about the introduced assumption.

Remark 8 Multiplying (4.10) from the left by $\mathbf{F}_r \Pi$ yields (4.8). In the majority of cases, the validity of (4.8) implies that (4.10) is also satisfied. Nonetheless, (4.10) can be explicitly guaranteed by imposing it as a constraint within the MPC formulation introduced in Section 4.5. Note that, to ensure convexity, it may be required to determine Z_2 a priori such that $\Upsilon + \Theta_{\mathbf{F}\Pi} Z_2$ is a matrix with nonnegative entries. \diamond

Remark 9 One can easily ensure the existence of matrix Z_2 such that the right-hand side of (4.10) is element-wise nonnegative. To this end, we need to consider only the supply section of the DHN. Then, the corresponding graph consists of source nodes coming from producers, sink nodes reaching consumers, and intermediate nodes that represent junctions. Additionally, we consider networks where the degree of each intermediate node is at most three, meaning that each junction is either a *splitting node* or a *merging node*. Valves are positioned in reverse cascading order from consumers to intermediate nodes. At splitting nodes, valves are required on outgoing edges, except when an edge connects directly to another splitting node. For merging nodes, valve placement is only necessary on edges originating from producer nodes; otherwise, the analysis proceeds to the subsequent node. This process is repeated until all edges are traversed. We illustrate this procedure in Figure 4.1. The rationale for this valve checking strategy is that it ensures the flow on any edge lacking a valve becomes a nonnegative linear combination of the flows on edges equipped with valves. Formally, this implies the existence of a matrix $\Upsilon + \Theta_{\mathbf{F}\Pi} Z_2$ with exclusively nonnegative entries. \diamond

Remark to practitioners. While the valve placement strategy may appear extensive, it reflects the fundamental requirement that a high degree of freedom in flow control requires a sufficiently high number of strategically placed pumps or valves. In practical situations, where valve placement is restricted by physical or economic constraints,

analyzing (4.10) helps determine where fewer valves can still preserve convexity. In cases where achieving convexity is not feasible, it serves as a practical tool for practitioners to uncover inherent system limitations. These insights can support the design phase by identifying the most impactful valve placements or guide the development of operational constraints that satisfy (4.10) within practical limitations. \diamond




Figure 4.1: The valve checking procedure, described in Remark 9, illustrated by showing the flow from producers (p) to consumers (c) going through the supply network consisting of merging nodes (m) and splitting nodes (s).

Remark 10 The valve placement approach described in Remark 9 extends naturally to networks with bidirectional edges, where nodes connect to edges that can experience flow in both directions are treated as both merging and splitting nodes to account for all operational modes. Nevertheless, for bidirectional edges, a single physical valve suffices since only one flow direction is active at any given time. Furthermore, one should note that the introduced assumption on the total number of required valves aligns with the literature on DHNs [29, 43], where each chord of the graph's spanning tree is equipped with a valve. \diamond

The following proposition guarantees that (4.8) is a sufficient condition for (4.7).

Proposition 9 Let Assumption 1 hold. Then, for any q that satisfies (4.8), there exist $v \in \mathbb{R}_+^{|\mathcal{V}|}$ and $r \in [0, 1]^{|\mathcal{P}|}$ such that (4.7) is satisfied.

Proof: Recall that, for any valve openings v , we have $R(v) = R_\mu + R_v(v)$, where $R(v)$ is a diagonal matrix that has on its diagonal the sum of resistances due to friction and valve effects on each edge. Hence, we have

$$\mathbf{FR}(v)(q \odot q) = \mathbf{FR}_\mu(q \odot q) + \mathbf{FR}_v(v)(q \odot q). \quad (4.11)$$

Furthermore, we know that, for any v , q , and Π , there exists a positive semidefinite diagonal scaling matrix $D_{R,q}$ and a vector $y = D_{R,q}v$, where y represents a scaled substitute of the valve openings, such that we can write $R_v(v)(q \odot q) = \Pi y$. Therefore, we have

$$\mathbf{FR}(v)(q \odot q) = \mathbf{FR}_\mu(q \odot q) + \mathbf{F}\Pi y. \quad (4.12)$$

Accordingly, to show that (4.7) holds for some y and r , we need to verify the same argument for the following equation

$$\mathbf{F}H(r) = \mathbf{F}R_\mu(q \odot q) + \mathbf{F}\Pi y. \quad (4.13)$$

If we set $r = \mathbf{1}$, then it is enough to show that there exists $y \in [0, \infty)^{|\mathcal{V}|}$ such that

$$\mathbf{F}\Pi y = \mathbf{F}(H(\mathbf{1}) - R_\mu(q \odot q)). \quad (4.14)$$

From (4.8), we have that the right-hand side of (4.14) is element-wise nonnegative, which is necessary for the existence of nonnegative solutions to y . Following Assumption 1, we know that (4.14) can be written as

$$\mathbf{F}\Pi(y - \Upsilon H(\mathbf{1}) + \Upsilon R_\mu(q \odot q)) = 0. \quad (4.15)$$

Hence, the solutions to (4.15) lie inside the null space of $\mathbf{F}\Pi$. Accordingly, we know that, for each y satisfying (4.15), there exists $z \in \mathbb{R}^{m_\Theta}$ such that

$$\begin{aligned} y &= \Upsilon H(\mathbf{1}) - \Upsilon R_\mu(q \odot q) + \Theta_{\mathbf{F}\Pi} z, \\ &= (\Upsilon + \Theta_{\mathbf{F}\Pi} Z_1) H(\mathbf{1}) - (\Upsilon + \Theta_{\mathbf{F}\Pi} Z_2) R_\mu(q \odot q), \end{aligned} \quad (4.16)$$

where $Z_1, Z_2 \in \mathbb{R}^{m_\Theta \times |\mathcal{E}|}$ are the matrices introduced in Assumption 1. Thus, due to (4.10), we have that y is a vector with nonnegative entries, which implies that $v \geq 0$. More precisely, there exist $v \in \mathbb{R}_+^{|\mathcal{V}|}$ and $r \in [0, 1]^{|\mathcal{P}|}$ such that (4.7) is satisfied. This concludes the proof. \blacksquare

The previous proposition guarantees, for any q satisfying (4.8), the existence of v and r satisfying (4.7), and, therefore, satisfying (4.6). In the following proposition, we show that (4.6) implies (4.4), meaning that Kirchhoff's second law is fulfilled.

Proposition 10 Under Assumption 1, we have that (4.6) is a necessary and sufficient condition for (4.4), i.e., if there exists a (q, v, r) such that (4.6) holds, then (4.4) is also satisfied.

Proof: Considering $\mathcal{S}_f \subseteq \mathcal{S}^\dagger$, from the definition of (4.4) and (4.6), we know that (4.4) implies (4.6). Therefore, we only need to prove the sufficiency part of the claim.

Let \mathbf{F}_i be the i^{th} row of \mathbf{F} , i.e., \mathbf{F}_i is a vector with elements equal to $\mathbf{F}_{ij} = 1$ if $j \in \mathcal{C}_i$ and 0 otherwise. All cycles $\mathcal{C}_1, \mathcal{C}_2, \dots, \mathcal{C}_{m_f}$ form a cycle basis for \mathcal{G} , which means that there are m_f linearly independent basis vectors $\mathbf{F}_1, \dots, \mathbf{F}_{m_f}$ spanning the basis of all other cycles in the graph. As a result, any vector $\mathbf{F}^{(\mathcal{C})}$, which corresponds to a cycle $\mathcal{C} \in \mathcal{S}^\dagger \setminus \mathcal{S}_f$, can be constructed by $\mathbf{F}_1, \dots, \mathbf{F}_{m_f}$ as the following integer linear combination

$$\mathbf{F}^{(\mathcal{C})} = \sum_{i=1}^{m_f} a_i \mathbf{F}_i, \quad (4.17)$$

where $a_1, \dots, a_{m_f} \in \mathbb{Z}$. Therefore, for the sum of pressure differences over \mathcal{C} , we have

$$\begin{aligned} \sum_{e \in \mathcal{C}} \Delta p_e &= \mathbf{F}^{(\mathcal{C})} \Delta \mathbf{p}_e \\ &= \sum_{i=1}^{m_f} a_i \mathbf{F}_i \Delta \mathbf{p}_e = \sum_{i=1}^{m_f} a_i \sum_{e \in \mathcal{C}_i} \Delta p_e. \end{aligned} \quad (4.18)$$

Subsequently, from (4.6), it is implied that

$$\sum_{e \in \mathcal{C}} \Delta p_e = 0, \quad (4.19)$$

which concludes the proof. \blacksquare

In the following proposition, we present the key result regarding the convexity of (4.8).

Proposition 11 The equation (4.8) is convex with respect to flow vectors q and q_r .

Proof: For $i = 1, \dots, m_f$, let matrix \mathbf{R}_μ^i be defined as

$$\mathbf{R}_\mu^i = \text{diag}(\mathbf{F}_{r,i}) R_\mu, \quad (4.20)$$

which is a diagonal matrix with nonnegative entries, and thus, positive semidefinite. According to the definition of \mathbf{F}_r , R_μ , h_e , and equation (4.8), we have

$$q^\top \mathbf{R}_\mu^i q = \sum_{e \in \mathcal{C}_i} R_{\mu,e} q_e^2 \leq \sum_{e \in \mathcal{C}_i} h_e(1) = \sum_{e \in \mathcal{C}_i} c_e, \quad (4.21)$$

for any $i = 1, \dots, m_r$, which is equivalent to (4.8) and implies that it is convex with respect to q . For any $i = 1, \dots, m_r$, define matrix \mathbf{Z}^i as

$$\mathbf{Z}^i = \mathbf{F}_r \mathbf{R}_\mu^i \mathbf{F}_r^\top. \quad (4.22)$$

Note that the positive semidefiniteness of \mathbf{R}_μ^i implies the same property for \mathbf{Z}^i , for each $i = 1, \dots, m_r$. From $q = \mathbf{F}_r^\top q_r$, one has

$$q_r^\top \mathbf{Z}^i q_r = (\mathbf{F}_r^\top q_r)^\top \mathbf{R}_\mu^i (\mathbf{F}_r^\top q_r) = q^\top \mathbf{R}_\mu^i q. \quad (4.23)$$

Accordingly, we can write (4.21), or equivalently (4.8), as

$$q_r^\top \mathbf{Z}^i q_r \leq \sum_{e \in \mathcal{C}_i} c_e, \quad \forall i = 1, \dots, m_r, \quad (4.24)$$

which implies the convexity of (4.8) with respect to q_r . This concludes the proof. \blacksquare

The proposed propositions indicate the presence of a convex reformulation of Kirchhoff's second law through (4.8). Proposition 9 shows that this reformulation ensures the existence of a feasible set of valve openings necessary to achieve the specified flow vector q . Moreover, Proposition 10 describes the sufficiency of (4.6) for satisfying (4.4). Lastly, we demonstrate the convexity of (4.8) through a straightforward transformation as detailed in Proposition 11.

Remark 11 More generally, we expect the results to hold for other convex pressure-flow relationships, e.g., $\Delta p = \mathcal{R}_\mu(q)$, not just the quadratic one typical for turbulent flow in DHNs. When replacing (4.1) with such a function, the optimization problem preserves convexity, since the left-hand side of (4.8) remains a sum of convex functions. Since Propositions 9 and 10 are independent of the specific form of $\mathcal{R}_\mu(q)$, the arguments should still apply. \diamond

4.4. Time Discretization of the Thermal Dynamics

A discrete-time model is required for the MPC problem. Due to the slow thermal dynamics of DHNs, online computation constraints, and demand measurement intervals, practical implementations typically use time steps Δt ranging from 15 minutes to 1 hour [18]. For explicit discretization schemes, satisfying the Courant-Friedrichs-Lewy (CFL) condition, i.e., $q_i(k)\Delta t \leq V_i \forall i, k \in \mathbb{N}$, is difficult when modeling with large time steps. Hence, we employ the numerically stable implicit Euler method:

$$x(k+1) = x(k) + \Delta t f_{\text{ct}}(x(k+1), u(k+1)), \quad (4.25)$$

where f_{ct} is the continuous-time dynamics. The application of the implicit Euler method to (2.23) results in the following discrete-time system

$$V \tilde{T}(k+1) = V \tilde{T}(k) + \Delta t [A(q(k)) \tilde{T}(k+1) + B w(k)], \quad (4.26)$$

or, equivalently, we can split the algebraic part from the equation by introducing a subset of variables $x^c(k) \in \mathbb{R}^{m_c}$ to denote the state of thermal node cells with positive volume and $\tilde{T}^j(k) \in \mathbb{R}^{|\mathcal{N}_0|}$ to denote the state of junctions with zero volume. As a result, we can describe the system through a discrete-time differential algebraic equation (DAE) as

$$\begin{aligned} \tilde{T}^c(k+1) &= f(\tilde{T}^c(k), \tilde{T}^j(k), u(k), d(k)), \\ 0 &= g(\tilde{T}^c(k), \tilde{T}^j(k), u(k), d(k)). \end{aligned} \quad (4.27)$$

The system's dynamics are governed by the functions f and g , which relate the state vector $\tilde{T}(k) = [\tilde{T}^c(k)^\top \tilde{T}^j(k)^\top]^\top$, input vector $u(k) = [q_r(k)^\top P(k)^\top]^\top$, and disturbance $d(k)$. A description of these variables is detailed in Table 4.1. Note that, in (4.27), the previously mentioned external input $w(k)$ has been split into a controllable power injection part $P(k)$ considered in $u(k)$ and a disturbance part considered in $d(k)$, which represents consumer demand.

Table 4.1: Description of system variables.

Symbol	Description	Dimension
$\tilde{T}^c(k)$	State: Thermal node temperature	m_c
$\tilde{T}^j(k)$	State: Junction temperature	$ \mathcal{N}_0 $
$q_r(k)$	Control input: Volume flow rate	m_r
$P(k)$	Control input: Power injection	$ \mathcal{W}_P $
$d(k)$	Disturbance: Consumer demand	$ \mathcal{W}_C $
$\tilde{T}(k)$	Combined state vector	$ \mathcal{N} $
$u(k)$	Combined input vector	$m_r + \mathcal{W}_P $

4.5. Economic MPC Formulation for District Heating Networks

IN this section, we introduce our economic MPC scheme for optimizing the performance of DHNs. To this end, we formulate the MPC optimization problem, discuss relevant design choices for the objective function, and present various computational techniques to improve the numerical performance of the employed solvers.

4.5.1. Problem Formulation

Consider the discrete-time dynamics equation (4.27). The receding horizon optimal control problem at time k is defined as

$$\begin{aligned}
 \min_{(x(i))_{i=1}^N, (u(i))_{i=0}^{N-1}} \quad & J_N((x(i))_{i=1}^N, (u(i))_{i=0}^{N-1}) \\
 \text{s.t.} \quad & x^c(i+1) = f(x^c(i), x^j(i), u(i), d(i+k)), \\
 & 0 = g(x^c(i), x^j(i), u(i), d(i+k)), \\
 & x(i) \in \mathbb{X}(i+k), \\
 & u(i) \in \mathbb{U}(i+k), \\
 & \forall i \in \{0, \dots, N-1\}, \\
 & x(0) = \tilde{T}(k),
 \end{aligned} \tag{4.28}$$

where $\tilde{T}(k)$ is the measured state at time k and $x(i)$ are the state decision variables in the MPC optimization, the sets $\mathbb{X}(i) := \{x : \mathcal{F}(i, x) \leq 0\}$ are time-varying constraint sets representing operational bounds on the state such as

$$\begin{aligned}
 \underline{T}_{\text{sup}} - T_a &\leq \tilde{T}_i \leq \overline{T}_{\text{sup}} - T_a, \quad i \in \mathcal{N}_{\text{sup}} \text{ (supply nodes)}, \\
 \underline{T}_{\text{ret}} - T_a &\leq \tilde{T}_i \leq \overline{T}_{\text{ret}} - T_a, \quad i \in \mathcal{N}_{\text{ret}} \text{ (return nodes)},
 \end{aligned}$$

and the sets $\mathbb{U}(i) := \{u : \mathcal{H}(i, u) \leq 0\}$ representing time-varying operational constraints on the inputs, including the hydraulic constraints as defined in equations equation (2.12), equation (2.13), and equation (4.24). Let the solution of the optimal control problem for horizon N be denoted by $(x_{N,k}^*, u_{N,k}^*)$, where the subscripts are included to stress the dependence of optimal solutions on N and k . In MPC, the optimal control problem in equation (4.28) is solved iteratively. In each step, the feedback control law is

$$\mu_N(x(k)) = u^*(0),$$

where $u^*(0)$ is the first element of the optimal sequence $u_{N,k}^*$.

Remark 12 The optimal control problem yields an optimal volume flow rate sequence, $q(k)$. However, direct actuation of valves by the high-level MPC is impractical. Therefore, $q(k)$ serves as a setpoint for a lower-level valve controller. \diamond

4.5.2. Objective Function

The objective function is a critical component in the design of economic MPC schemes, allowing for the selection of a cost function that accurately represents our practical goals or specifications. Given the current objective of optimizing the performance of the DHN, we define the objective function as

$$J_N = J_N^{\text{price}} + J_N^{\text{temp}} + J_N^{\text{diff}} + J_N^{\text{sto}} + J_N^{\text{slack}}, \quad (4.29)$$

where each of these terms reflects a desirable operational feature or aspect, discussed below.

- *Price term.* Operational management requires minimization of operational costs. Additionally, we assume part of the generating mechanism is linked to a market, such as a heat pump purchasing from the electricity grid. Hence, we use a linear cost function

$$J_N^{\text{price}} = \sum_{i=0}^{N-1} R^{\text{price}}(i)P(i), \quad (4.30)$$

where $R^{\text{price}}(i)$ represents the time-varying price, or relative price, of generating $P(i)$.

- *Temperature term.* It is desired to operate DHNs at low temperatures to improve their efficiency. Therefore, we include a state term in the objective as

$$J_N^{\text{temp}} = \sum_{i=1}^N R^{\text{temp}}(i)x(i)^p \quad (4.31)$$

with the penalty coefficient $R^{\text{temp}}(i)$, and power index $p \in \{1, 2\}$, suggesting that the function can be linear or quadratic in x .

- *Input variation term.* It is undesirable to have fast switching in supply temperatures in DHNs, primarily due to the pipeline deterioration from the resulting thermal stress [71]. Therefore, we consider a cost term as

$$J_N^{\text{diff}} = \sum_{i=0}^{N-2} R^{\text{diff}}(i)(P(i+1) - P(i))^2, \quad (4.32)$$

which penalizes input deviations between any successive timesteps.

- *Storage term.* Minimization of operating costs and temperature usually does not favor charging of a storage buffer. Therefore, without any storage term, the MPC typically stays on discharging mode for all of the storage units. To address this issue, we introduce a terminal tracking cost on the temperature of storage nodes as

$$J_N^{\text{sto}} = \|x^{\text{sto}}(N) - \bar{x}^{\text{sto}}\|_{R^{\text{sto}}}^2, \quad (4.33)$$

which is similar to the one employed in [41]. This term will encourage the MPC to charge situationally depending on the size of R^{sto} .

- *Slack term.* The optimization is a large-scale nonlinear program, where the dynamics and constraints change in each iteration due to time-varying elements. Accordingly, guaranteeing feasibility in every iteration is not always possible. On the other hand, certain constraints, e.g., temperature bounds and demand satisfaction, may not be *hard* constraints, meaning that a certain degree of violation is allowed. For this, we introduce variables σ and a cost term as

$$J_N^{\text{slack}} = \|\sigma\|_{R^{\text{slack}}}^2, \quad (4.34)$$

with R^{slack} chosen to be large enough penalizing undesired constraint violations.

4.5.3. Improving Numerical Performance

The design of MPC algorithms, which inherently involve the iterative solution to an optimal control problem, depend on the convergence speed of the numerical solvers used to obtain the solution. To this end, we implement various methods to improve the computational performance, as discussed below.

- *Warm starting.* This technique can improve computational speed by providing the numerical solver with a near-optimal initial guess. After a single iteration, we possess the N -step prediction of the state x . Thus, in the next iterate, we initialize the solver with $x^{k+1}(0), x^{k+1}(1), \dots, x^{k+1}(N-1) \leftarrow x^{k*}(0), x^{k*}(1), \dots, x^{k*}(N-1)$, where k denotes the index for current iteration. This process is repeated for all subsequent iterations.
- *Objective terms.* Certain objective terms can assist the numerical solver in finding a solution. In particular, we noticed that the input variation term equation (4.32) significantly improves the speed of convergence.

Other approaches to reduce complexity include implementing a *control horizon* or *move blocking*, which fix control inputs over certain time periods, but may lead to restrictive solutions; see, e.g., [14, 72].

4.6. Numerical Experiments and Results

IN this section, to assess and verify the economic and computational performance of our proposed methods, we perform suitably designed numerical experiments and simulation studies. To this end, we compare the proposed method to existing control strategies in the literature, including single-producer MPCs (SP-MPC) algorithms, which are based on, or similar to, optimization-based controllers used in [18] and [41].

Additionally, we compare with a rule-based control (RBC) scheme implementation, which is close to the widely adopted approach in practice for the control of DHNs. Furthermore, we perform a numerical study to evaluate the computational tractability of the proposed methods. In particular, we assess the impacts of spatial oversampling and changing the prediction horizon on the computational load and performance of the algorithms. The district heating network considered in our numerical experiments is the AROMA network, introduced in Section 2.1.4 and illustrated in Figure 4.2.

For our numerical experiments, we employ a standard laptop with an *Intel i7-1185G7* processor to run the simulations and Julia to create the models. We use the mathematical programming package *JuMP.jl* [73] to build the optimization problems, *Ipopt* [74] to solve the problems, and *DifferentialEquations.jl* [75] to simulate the DHN between iterations.




Figure 4.2: The AROMA network labeled with five consumers (labeled C), two producers (labeled P), and a storage (labeled S). Dashed lines indicate pipes on which bidirectional flows are allowed.

Before proceeding with the results, we introduce the ground truth reference model to which we apply the generated inputs, and, secondly, we discuss the physical parameters, load profiles, electricity price profiles that were used in all simulations.

4.6.1. High-fidelity Model

To assess the performance of the controller, we will apply the inputs $\mu_N(x(k))$ to a high-resolution simulator that accurately describes the system. This simulator model, called here the *high-fidelity model*, acts as a representation of the real system. In [22], a detailed study on the accuracy of these graph-theoretic simulation models for district heating systems is provided along with a comparison to the other high-fidelity simulators that have been verified using real measurements. It is shown in [22] that, even for reduced-order models, their method exhibits high accuracy.

We refer to the high-fidelity model by P_{CT}^{HF} , which is a continuous-time system obtained using the same method as in equation (2.23). Nonetheless, the dimension of x^{HF} , the state vector in the high-fidelity model, is equal to $|\mathcal{N}_0| + \beta m_c$, where $\beta \in \mathbb{N}$, and spatial discretization step $\Delta x_i^{HF} = \Delta x_i / \beta$ to compensate for pipe length. This change suggests that for large values of β , we achieve a much higher spatial resolution, which leads to a better approximation of the original system. Finally, the resulting system P_{CT}^{HF} is a system of differential algebraic equations that we solve using dedicated solvers in Julia [75]. In each iteration, the solver computes the evolution of the states for Δt seconds. Every iteration is initialized using the final step of the previous simulation $\tilde{T}^{HF}(k\Delta t - \Delta t)$ and the optimal inputs $\mu_N(\tilde{T}(k))$ remain constant for the duration of the simulation, i.e., on the time interval $[k\Delta t - \Delta t, k\Delta t]$. After completion, the current state $\tilde{T}(k) = \tilde{T}^{HF}(k\Delta t)$ is fed into the MPC controller; see Figure 1.3 for a schematic illustration of this feedback loop.

4.6.2. Parameters and Data

In Table 4.2, we list the physical parameters used in our simulations. System dimensions such as pipe diameters and lengths are the same as in [18]. We obtain the heat

transmission coefficient U_{pipe} from [6, p. 77], and the friction coefficient K_{pipe} from [6, p. 444].

Table 4.2: List of physical system parameters.

Parameter	Value	Units
T_a	10	$^{\circ}\text{C}$
ρ	981	$\text{kg}\cdot\text{m}^{-3}$
c_p	4182	$\text{J}\cdot\text{kg}^{-1}\cdot\text{K}$
U_{pipe}	0.4	$\text{W}\cdot\text{m}^{-2}\cdot\text{K}^{-1}$
K_{pipe}	0.02	-
d_{pipe}	70 – 107	mm
d_{sto}	2000	mm
L_{pipe}	300 – 600	m
L_{sto}	8	m
$(\underline{T}_{\text{sup}}, \overline{T}_{\text{sup}})$	(70, 100)	$^{\circ}\text{C}$
$(\underline{T}_{\text{ret}}, \overline{T}_{\text{ret}})$	(30, 50)	$^{\circ}\text{C}$

The storage unit has a capacity of 25.13 m^3 , corresponding to approximately 1.2 MWh of usable thermal energy (at $\Delta T = 40 \text{ K}$), which represents roughly 10% of the total daily heat demand. The storage is thus sized as a short-term flexibility asset, enabling peak shaving and limited load shifting to exploit time-varying electricity prices.

The demand profile employed here is an approximation of the one used in [18]. The electricity prices are acquired from Ember [76], where hourly electricity spot prices are provided for the Netherlands. We use price data from March 14, 2024 to determine the relative price term in equation (4.30), i.e., $R^{\text{price}}(i)$, for all i . Figure 4.3 illustrates the demand and electricity prices for 24 hours, where only the net demand of consumers is shown. The demand of each individual consumer is computed as a fraction of the total demand as shown in Table 4.3, assuming the same load distribution as in [18].

Table 4.3: Fraction of total demand.

C1/P2	C2	C3	C4	C5
0.08	0.34	0.11	0.08	0.38

4.6.3. Economic Performance and Comparisons

We conduct a comparative analysis of our proposed method against several established control strategies, including rule-based control and MPC for single producer DHNs without storage capabilities. The rule-based control strategy employs a heating curve to determine supply temperature setpoints based on outdoor conditions, which are tracked

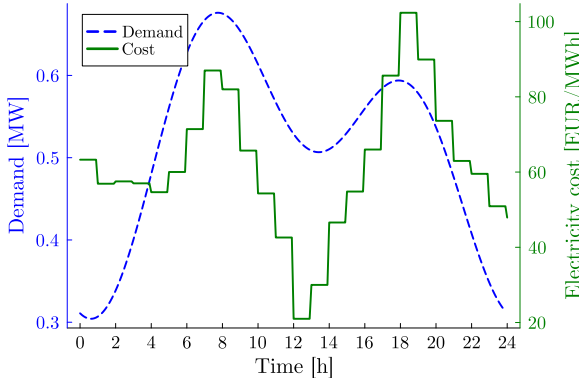


Figure 4.3: Total demand profile and electricity price profile for 24 hours. Electricity prices are from March 14, 2024, in the Netherlands.

locally using a PI controller. The existing rule-based control strategy is referred to as RBC, while the single producer MPC strategy is denoted as SP-MPC. Our proposed methods are referred to as follows: the single-producer model with storage is denoted as SPS-MPC, the multi-producer model without storage is referred to as MP-MPC, and the multi-producer model with storage is indicated as MPS-MPC.

To quantify the economic value of the proposed methods, we analyze the financial implications for network operators under representative operating conditions. Here, prosumer C1/P2 generates a heat surplus of 100 kW during the interval 12:00-17:00, resulting in negative net demand. To maintain a similar total demand for comparative analysis, we introduce a compensatory 80 kW load increase at consumer C4.

Cost comparison

In our numerical experiments, we set the control interval to $\Delta t = 15$ minutes. The cost function J_N is formulated to prioritize economic performance, with the relative price term $R^{\text{price}}(i)$ weighted substantially higher than other objective function terms, excluding slack variable penalties. A general overview of the economic performance results are presented in Table 4.4.

Table 4.4: Performance Comparison of Control Strategies. All results are based on 24 hours of simulation. Results from left to right are for: rule-based control (RBC), single-producer MPC (SP-MPC), single-producer MPC with storage (SPS-MPC), multi-producer MPC (MP-MPC), and multi-producer MPC with storage (MPS-MPC). The average runtime for each iteration of the MPC algorithms was under 10 seconds.

Cost (€)	RBC	SP-MPC	SPS-MPC	MP-MPC	MPS-MPC
$N = 16$ (4h)		902	891	868	865
$N = 32$ (8h)	922	907	882	867	840

Quantifying the effect of slack variables

To assess the true performance of the proposed methods, we quantify the extent to which the constraints have been violated as a consequence of the effects of the slack variables.

In our analysis, we observe distinct patterns in constraint violations. Pumping capacity bottlenecks manifest primarily as demand violations, where the network cannot deliver sufficient flow rates to meet consumer requirements. Conversely, temperature violations, defined as failures to maintain minimum required supply temperatures at consumer substations, occur predominantly in scenarios where pumping capacity is sufficient to meet demand. While theoretical coupling between these constraint violations is possible, our simulations suggest they tend to be mutually exclusive, with temperature violations emerging only when demand is fully satisfied and demand violations occurring only under pumping capacity constraints.

Temperature requirements, typically mandated by regulatory frameworks, remain critical operational constraints that DHN operators must prioritize. For a simulation of t_f time steps, we quantify these average temperature violations (ATV) in degrees Celsius as

$$\text{ATV} = \frac{1}{t_f |\mathcal{W}_C|} \sum_{k=1}^{t_f} \sum_{c \in \mathcal{W}_C} \max(0, \tilde{T}_{\text{sup,min}} - \tilde{T}_c(k)). \quad (4.35)$$

On the other hand, network operators must ensure sufficient heat delivery to meet consumer demand requirements. We quantify demand violation (DV) over t_f steps as

$$\text{DV} = 100\% \times \frac{\sum_{k=1}^{t_f} \sum_{j \in \mathcal{W}_C} d_j(k) - d_j^{\text{true}}(k)}{\sum_{k=1}^{t_f} \sum_{j \in \mathcal{W}_C} d_j(k)}, \quad (4.36)$$

where the true heat exchange is defined as

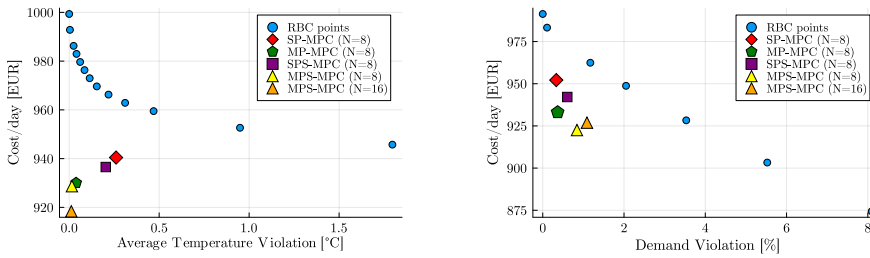
$$d_j^{\text{true}}(k) = q_j(k) (\tilde{T}_{j-1}(k) - \tilde{T}_j(k)), \quad j \in \mathcal{W}_C.$$

In Figure 4.4, we compare these violation metrics against the operational costs obtained in each method.

Remark 13 Demand and temperature violations under the RBC strategy arise from insufficient flow capacity or the inability to maintain adequate temperatures at network endpoints, as the controller lacks anticipatory capability and responds only to local measurements. \diamond

Discussion: Several key observations emerge from Table 4.4 and Figure 4.4. First, there is a noticeable improvement trend in cost reduction, with each feature added to the algorithm providing incremental benefits. The MPS-MPC algorithm with $N = 32$ achieves the highest performance, demonstrating a 9% cost reduction compared to the rule-based controller.

Our analysis of constraint violations reveals that the MPC-based approaches actively optimize network operations based on predicted demand and price signals, resulting in dynamic temperature management rather than simple static setpoint tracking. This continuous optimization allows the controller to systematically reduce average network



(a) Daily operating costs versus temperature violations for a DHN with sufficient pumping capacity.

(b) Daily operating costs versus demand violations for a DHN operating with pumping capacity constraints.

Figure 4.4: Relationship between operational costs and constraint violations under different pump configurations measured over a four day period.

temperatures when beneficial while maintaining required service levels, a capability that rule-based controllers inherently lack. This dynamic temperature management is key to achieving both cost reductions and improved constraint satisfaction compared to rule-based approaches.

Additionally, when analyzing demand violations under pumping capacity constraints, the data suggest that prosumer-based control strategies, particularly MP-MPC and MPS-MPC, can better manage network limitations. The ability to actively coordinate multiple producers and redistribute flow patterns allows these approaches to excel at balancing competing objectives, demonstrating superior performance in maintaining service quality while optimizing operational costs.

4.6.4. Added Value of Storage

The aim of this section is to demonstrate the benefits of storage in the DHN. As can be seen in Table 4.4, the integration of storage yields significant economic benefits. In the top part of Figure 4.5, the production schedules of the SP-MPC and SPS-MPC are plotted against each other. Secondly, the lower portion of the figure illustrates the corresponding storage charging and discharging periods. It is worth noting that, in this case, we have additionally implemented a constraint that ensures that the total charging volume approximately matches the total discharging volume over the course of the day. Further specific details regarding this constraint are provided in the subsequent discussion.

Discussion: The results demonstrate that the controller effectively utilizes storage capacity to enable operational flexibility, as illustrated in Figure 4.5. The comparison between SPS-MPC and SP-MPC reveals that the former shifts producer load by charging storage during early hours and discharging after approximately 7.5 hours, thereby reducing producer load and achieving cost savings during high-price periods.

Several limitations should also be acknowledged. While the storage term in the cost function, introduced in Section 4.5.2, aims to maintain adequate hot water levels in the top storage layer, its effectiveness is highly dependent on layer volumes and




Figure 4.5: The upper plot compares the power injection relative to demand extraction for both SP-MPC and SPS-MPC methods, with the normalized heat generation cost shown in green for reference. The lower plot illustrates the storage buffer's charging and discharging modes in SPS-MPC, maintaining zero net mass flow throughout the simulation period.

objective weight selection. Small weights lead to continuous discharge, while only substantially large weights induce charging behavior. Hence, in this study, we have opted to manually constrain the charging rates to ensure the total charging volume is equal to the discharging volume and emulate a more equal scenario.

4.6.5. Added Value of Multiple Producers

In addition to distributing heat production, which allows facilities like waste incineration plants or data centers to contribute to DHNs, multiple producers can reduce pressure load on the central plant by dividing supply streams. Since pressure drop increases significantly with flow velocity, each pipeline has a limited flow capacity. Distributing heat supply flows among different producers enables higher mass flows at consumer stations, providing better operational margins. For this example, the setup has been modified in three ways. First, prosumer C1/P2 now acts as a full producer, consistently generating 100 kW of heat for the DHN, similar to a waste incineration plant or data center. Second, the available pump head from each pump has been scaled down to emphasize hydraulic constraints. Third, temperature bounds are tightened to $70 \leq T \leq 90$ for all system states, with any values outside this range indicating the optimization could not find a feasible solution within these bounds.

In Figure 4.6, the consumer inlet temperatures and mass flows for the single-producer case are shown. On the other hand, Figure 4.7 shows the temperatures and mass flows at consumers when C1/P2 contributes 100 kW constantly and P1 is freely controllable, corresponding to the multi-producer case.

Discussion: The results show that in the multi-producer scenario the network is able to deliver more heat to large consumers C2 and C5. In particular, the peak flow rate of C2 in Figure 4.7 lies about 40% higher than in Figure 4.6. The increase can be attributed to the fact that C2 receives heat from both P1 and C1/P2 sources. Consequently, P1's contribution to C2 is reduced, enabling it to allocate a greater share of its flow to C5, the largest consumer. Overall, the aggregate energy requirement in the multi-producer scenario, comprising the combined outputs of P1 and C1/P2, is roughly equivalent to the total production of P1 in the single-producer scenario. Therefore, the MP-MPC shows improved performance compared to the SP-MPC in staying within operational limits and fairly distributing heat.




Figure 4.6: Consumer inlet temperatures and mass flows for a single-producer (SP-MPC) case.




Figure 4.7: Consumer inlet temperatures and mass flows for a multi-producer (MP-MPC) case.

4.6.6. Computational Study

Finally, we examine the computational performance and scalability of our algorithms for various model resolutions and prediction horizons. In Figure 4.8, the median computational times of the solver iterations for different model resolutions (left) and prediction horizons (right) are presented.

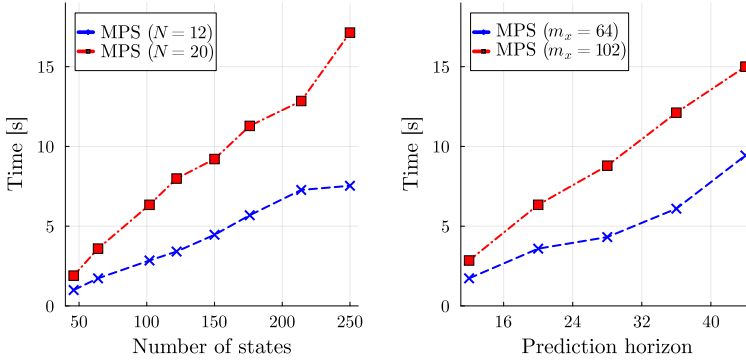


Figure 4.8: The left figure shows *median solver* times for the MPS-MPC algorithm using a prediction horizon $N = 12$ and $N = 20$ for different number of state variables. The right figure shows for two models, one with 64 states and one with 102 states, the *median solver* time based on the chosen prediction horizon.

Discussion: In general, the computational cost associated with increasing model complexity remains manageable from an operational perspective, with even the most complex models requiring, on average, less than a minute per iteration. However, our observations indicate that the consistency of the solver in finding solutions within acceptable times diminishes under certain conditions, particularly for more complex models. These include models with over 200 states or with a prediction horizon greater than or equal to 40, where solution times can vary significantly.

Assuming the existence of a feasible solution, a pragmatic approach to deal with varying solver times involves constraining the number of solver iterations or the solver time to ensure that controls are computed within the required timeframe, even if this may result in a decrease in solution quality. Nonetheless, increasing model complexity does not necessarily enhance performance. In certain scenarios, reducing the model resolution can actually decrease model mismatch. This is due to the fact that truncation error is influenced by the CFL condition, and this error is minimized when the CFL value is close to one [22]. Consequently, the determination of model resolution is not straightforward and should be carefully considered on a case-by-case basis.

4.6.7. Limitations and Future Work

While our results demonstrate the effectiveness of MPC approaches for DHNs, several limitations should be acknowledged. The performance of the proposed methods relies on accurate demand predictions, with prediction errors affecting controller performance. Our implementation assumes both perfect knowledge of system parameters and access to full state measurements throughout the network. In practice, these assumptions may not hold, as many temperature and flow measurements might be unavailable or inaccurate. From a computational perspective, solution times increase with model complexity and prediction horizon length, which may impact scalability for larger networks. Additionally, the nonconvex nature of the optimization problem, particularly from bidirectional flow constraints, can cause the solver to struggle in finding optimal trajectories. This sometimes requires warm-starting strategies to guide the optimization toward desired flow configurations, an approach that introduces additional implementation complexity. Future work could address these aspects through robust MPC formulations to handle demand uncertainty, state estimation techniques for networks with limited measurements, and development of computationally efficient optimization strategies.

4.7. Conclusion

We developed an economic model predictive control algorithm designed for the operational management of district heating networks incorporating essential elements of 4th generation district heating networks such as multiple distributed heat sources, prosumers, and storage. A key aspect of our algorithm is its innovative treatment of hydraulic constraints through a convexification approach. We have conducted comprehensive numerical experiments to evaluate the proposed features, demonstrating that MPC approaches significantly outperformed conventional rule-based controllers, yielding up to 9% cost reduction alongside reduced constraint violations, with computation times remaining within practical limits. Furthermore, the integration of storage capabilities and multiple-producer configurations enhanced these performance metrics.

5

Stability Analysis of Economic MPC for District Heating Networks Using Dissipativity Theory

The inherently nonlinear, large-scale, and time-varying nature of district heating systems pose significant challenges from a control perspective. In this chapter, we address these challenges by applying an economic MPC. Economic MPC is a dynamic real-time optimization method, enabling both optimal planning and stability of the closed-loop system. Our strategy constitutes several steps. First, we introduce a discrete-time modular framework for the district heating system, establishing its strict dissipativity with respect to a desired, potentially time-varying, equilibrium. We identify a set of meaningful objective functions for the district heating systems, preserving this property. Second, we show how strict dissipativity implies the turnpike property, which, in turn, guarantees approximate optimality, practical stability, and recursive feasibility for the EMPC closed-loop. Finally, we provide numerical simulations to demonstrate the effectiveness of our work.

This chapter is based on [36].

5.1. Introduction

ECONOMIC model predictive control (EMPC) is a generalization of stabilizing MPC, where the cost function is not necessarily a penalization of the distance to a specified equilibrium, but can encompass a broader range of objectives, such as energy minimization [32]. Nonetheless, closed-loop performance of EMPC does not inherently share the same guarantees as stabilizing MPC. Furthermore, for nonlinear or time-varying systems, it is often difficult to find suitable terminal conditions that ensure stability and recursive feasibility of the closed-loop [77]. Therefore, we require that the EMPC finds the optimal trajectory by itself. A key property for optimal control problems that ensures that the EMPC closed-loop solution remains close to the optimal trajectory is called the *turnpike property*. Evidence for the existence of this property can often be obtained empirically through simulation, however, there has been a substantial development in the literature indicating that the turnpike property can be implied from another property, called strict dissipativity; see [32, 33, 77–80].

For physical systems, strict dissipativity may be interpreted as the dissipation of energy to the environment. However, in the context of optimal control, even for the thermodynamic systems considered in this work, it does not have to serve this interpretation [78]. Particularly, this is because strict dissipativity is respective to a specific cost function and equilibrium, and this equilibrium does not need to be the one where the total energy stored is equal to zero. Generally, verification of strict dissipativity is not an easy task. It demands the existence of a storage function such that the dissipation inequality is satisfied at all times, which is difficult to check analytically in general. Therefore, a numerical approach was developed to verify strict dissipativity for continuous-time optimal control problems in [80] and for discrete-time problems in [81] and [82].

In this chapter, we present a modeling framework for a DHN, demonstrating, both analytically and numerically, its strict dissipativity for a range of cost functions. Subsequently, we exploit this dissipativity to establish the turnpike property, which, in turn, implies approximate optimality, practical stability, and recursive feasibility for the closed-loop. To the best of our knowledge, this type of dissipativity analysis for EMPC has not been studied in the context of DHN. In the simulations section, we show empirical evidence for the turnpike property and demonstrate the performance of the EMPC.

The chapter is structured as follows. In Section 5.2, we provide the results on strict dissipativity of the thermal networks. In Section 5.3, we summarize implications this property has for EMPC, and finally, in Section 5.4, we demonstrate the performance of the EMPC algorithm.

5.2. Strict Dissipativity for Thermal Networks

IN this section, we establish strict dissipativity, both analytically and numerically, for open and closed thermal networks introduced in the previous section. We require this property to ensure certain guarantees for the EMPC algorithm for DHN given in the next section.

Definition 6 (Strict dissipativity) A discrete-time dynamical system

$$x(k+1) = f(x(k), u(k)), \quad x(0) = x_0, \quad (5.1)$$

with $x(k) \in \mathbb{R}^n$, $u(k) \in \mathbb{R}^m$, $f: \mathbb{R}^n \times \mathbb{R}^m \rightarrow \mathbb{R}^n$, and $(x(k), u(k)) \in \mathbb{X} \times \mathbb{U}$ is strictly dissipative w.r.t. the supply rate $s(x, u) = \ell(x, u) - \ell(x^*, u^*)$ and the optimal steady state $(x^*, u^*) \in \mathbb{X} \times \mathbb{U}$, if there exists a storage function $\lambda: \mathbb{X} \rightarrow \mathbb{R}$ bounded from below on \mathbb{X} and a function $\zeta \in \mathcal{K}_\infty$ such that for all $(x, u) \in \mathbb{X} \times \mathbb{U}$,

$$\lambda(f(x, u)) - \lambda(x) \leq s(x, u) - \zeta(|x - x^*|). \quad (5.2)$$

5.2.1. Dissipativity of Closed Thermal Networks

We apply an explicit Euler discretization with time step equal to one to system (2.23) to obtain:

$$Vx(k+1) = Vx(k) + A(q(k))x(k) + B_p w_p(k) - B_d w_d, \quad (5.3)$$

where we decomposed $Bw(k)$ into a controllable part $B_p w_p(k)$ and a known (constant) disturbance $-B_d w_d$, where $w_p \in \mathbb{R}^{n_p}$ and $w_d \in \mathbb{R}^{n_d}$ represent the producer nodes and demand nodes, respectively, and $B_p \in \mathbb{R}^{n \times n_p}$ and $B_d \in \mathbb{R}^{n \times n_d}$ consist of augmented column vectors that have exactly one element equal to 1 at the node that is active.

Theorem 12 The discrete-time *closed* thermal network (5.3) is strictly dissipative with respect to the supply rate $\hat{\ell}(x, w_p, q) = \ell(x, w_p, q) - \ell(x^*, w_p^*, q^*)$ and the optimal steady state $(x^*, w_p^*, q^*) \in \mathcal{Z} = \mathbb{X} \times \mathbb{U} \times \mathbb{Q}$ with a linear storage function $\lambda(x) = \gamma \mathbf{1}^\top Vx$ and a linear cost function $\ell(x, w_p, q) = r^\top w_p$, where $r = \gamma \mathbf{1}^\top B_p$ as long as $\gamma \neq 0$.

Proof: We apply Definition 6 to system (5.3) to obtain

$$\begin{aligned} \lambda(V^{-1}(A(q)x + B_p w_p - B_d w_d)) \\ \leq r^\top w_p - r^\top w_p^* - \zeta(|x - x^*|), \end{aligned}$$

which then can be separated into two inequalities, where we apply $\lambda(x) = \gamma \mathbf{1}^\top Vx$:

$$\begin{aligned} \gamma \mathbf{1}^\top A(q)(x - x^*) &\leq -\zeta(|x - x^*|), \\ \gamma \mathbf{1}^\top A(q)x^* + \gamma \mathbf{1}^\top (B_p w_p - B_d w_d) &\leq r^\top (w_p - w_p^*). \end{aligned}$$

Due to mass conservation, we have that multiplying from the left with $\mathbf{1}^\top$ cancels out all dependency on q such that $\mathbf{1}^\top A(q) = -\mathbf{1}^\top D_\alpha$. Furthermore, it is easy to check that any equilibrium $(x_e, w_{p,e}, q_e) \in \mathcal{Z}$, not necessarily optimal, has to satisfy $\sum_{i=1}^{n_p} w_{p,e,i} = \sum_{i=1}^{n_d} w_{d,i} + \sum_{i=1}^n \alpha_i x_{e,i}$ (energy balance). With this, we are able to state the following two conditions required for strict dissipativity:

$$-\gamma \sum_{i=1}^n \alpha_i (x_i - x_i^*) \leq -\zeta(|x - x^*|), \quad (5.4)$$

and

$$\underbrace{-\gamma \sum_{i=1}^n \alpha_i x^* + \gamma \left(\sum_{i=1}^{n_p} w_{p,i}^* - \sum_{i=1}^{n_d} w_{d,i} \right)}_{=0} + \gamma \sum_{i=1}^{n_p} (w_{p,i} - w_{p,i}^*) \leq r^\top (w_p - w_p^*), \quad (5.5)$$

which is satisfied for $r^\top = \gamma \mathbf{1}^\top B_p$. Then, the optimal equilibrium temperature must be minimal if $\gamma > 0$, and maximal if $\gamma < 0$ for (5.4) to hold for any $x \in \mathbb{X}$.¹ To illustrate, in Fig 5.1, we compare the two planes for any (arbitrary) equilibrium point $x_{\text{eq}} \in \mathbb{X}$. When $x_{\text{eq}} > x_{\text{min}}$, we find that the dissipation inequality is not satisfied for part of the feasible set, and we do not have strict dissipativity. Hence, we require $x_{\text{eq}} = x_{\text{min}} = x^*$ for strict dissipativity. ■




Figure 5.1: Illustrating hyperplanes of the (difference in) storage function and supply rate. Strict dissipativity holds when the blue plane is above the green plane within the feasible set. Note that the slope of the red line, which encodes the effect of heat losses, has been greatly exaggerated for illustrative purposes. The distance measure $\Delta x_{\text{eq}} > 0$ suggests a subset of the feasible set does not satisfy the dissipation inequality.

5.2.2. Dissipativity of Open Thermal Networks

Here we apply again an explicit Euler discretization to equation (2.23); however, we now explicitly consider the open thermal network variant to obtain

$$Vx(k+1) = Vx(k) + A(q(k))x(k) + B_o u_o(k) - B_d w_d, \quad (5.6)$$

¹For $\gamma = 0$ we have only dissipativity, not strict dissipativity.

and we assume $w_p = 0$, since we can control heat insertion through $u_o(k) \in \mathbb{R}^{n_o}$.

Theorem 13 The discrete-time *open* thermal network (5.6) is strictly dissipative with respect to the supply rate $\hat{\ell}(x, u_o, q) = \ell(x, u_o, q) - \ell(x^*, u_o^*, q^*)$ and the optimal steady state $(x^*, u_o^*, q^*) \in \mathcal{Z} = \mathbb{X} \times \mathbb{U} \times \mathbb{Q}$ with a linear storage function $\lambda(x) = \gamma \mathbf{1}^\top V x$ and a linear cost function $\ell(x, u_o, q) = \gamma \mathbf{1}^\top (B_o u_o - B_d B_d^\top Q x)$.

Proof: The result is similar to the proof for closed thermal networks (Theorem 12). The main difference is that in this case, we get an additional term from $\mathbf{1}^\top A(q)x = -\sum_{i=1}^n \alpha_i x_i - \mathbf{1}^\top B_d B_d^\top Q x$. Nonetheless, just like in the closed thermal network case, we can write the first inequality centered around x^* to obtain

$$-\gamma \sum_{i=1}^n \alpha_i (x_i - x_i^*) \leq -\zeta(\|x - x^*\|).$$

Then, the remaining parts are used to form the second inequality

$$\begin{aligned} & -\gamma \left(\underbrace{\sum_{i=1}^n \alpha_i x_i^* + \mathbf{1}^\top B_d B_d^\top Q^* x^* + \mathbf{1}^\top B_d w_d - \mathbf{1}^\top B_o u_o^*}_{=0} \right) \\ & + \gamma \mathbf{1}^\top (B_o (u_o - u_o^*) - B_d B_d^\top (Q x - Q^* x^*)) \leq \hat{\ell}(x, u_o, q), \end{aligned}$$

which we simplify by defining $\Delta u := \mathbf{1}^\top (B_o u_o - B_d B_d^\top Q x)$. This change results in the following inequality:

$$\gamma(\Delta u - \Delta u^*) \leq \hat{\ell}(x, u_o, q),$$

and, consequently, we have to satisfy this with equality if $\Delta u_{\min} < \Delta u^* < \Delta u_{\max}$. Therefore, in analogy with the closed example, we have strict dissipativity with respect to the state whenever x^* is minimal (for $\gamma > 0$) and maximal (for $\gamma < 0$). ■

Remark 14 This particular cost function represents the difference between the sum of energy at each pipe inlet and the sum of energy of the flow exiting the pipe after each consumer.

Remark 15 Both results for closed and open thermal networks rely on using a specific cost function that is linear in the input energy. Moreover, we see that this assumption yields strict dissipativity for a specific equilibrium on the boundary of our constraint set, and with respect to x only. This begs the question whether other cost functions, or strictness with respect to the input, are a possibility. Analytically, it is difficult to construct higher order storage functions to allow for various meaningful cost functions, but numerically, we can increase the complexity of the storage functions. Motivated by this, we present a numerical approach in the following section.

5.2.3. Numerical Verification of Strict Dissipativity

In this section, we provide a numerical approach to verifying strict dissipativity. We extend the definition to strictness with respect to both state and input, therefore, we introduce the notation $|(a, b)|_{a^*, b^*} = \|a - a^*\| + \|b - b^*\|$.

Method. Contrary to an analytical approach, in a numerical approach, we have the ability to expand the range of potential storage functions, which enhances our prospects of identifying a storage function that satisfies strict dissipativity for any specified cost function. On the other hand, we are restricted to evaluate systems of low dimension due to high computational demands.

In short, the objective is to find a function $\lambda : \mathbb{R}^n \rightarrow \mathbb{R}$ from a given set $\Lambda = \{\lambda\}$ and a scalar $\sigma \in \mathbb{R}$ that satisfy the following inequality:

$$\ell(x, u) - \sigma + \lambda(x) - \lambda(f(x, u)) - \zeta(|(x, u)|_{x^*, u^*}) \geq 0, \quad (5.7)$$

for a discrete-time system (5.1) by writing an optimization problem as follows:

$$\begin{aligned} & \max_{\sigma \in \mathbb{R}, \lambda \in \Lambda} \sigma, \\ \text{s.t.} \quad & (5.7), \\ & g_i(x, u) \geq 0, \quad \forall i = \{1, \dots, l\}, \end{aligned} \quad (5.8)$$

where $g : \mathbb{R}^{n+m} \rightarrow \mathbb{R}^l$ comprises all state and input constraints, i.e., $(x, u) \in \mathbb{X} \times \mathbb{U}$. Problem (5.8) is a semi-infinite optimization, which is generally hard to solve. Instead, we implement a relaxation based on the S-procedure, as was done in [81]. This procedure makes use of sum of squares (SOS) programming by defining Λ as a set that contains parameterized polynomial basis functions of x up to a certain degree d . Subsequently, rather than proving the inequality in (5.8) directly, the goal is to prove that there exists a $\lambda \in \Lambda$ such that the left hand side of (5.7) can be written in SOS formulation [82]. Furthermore, by introducing SOS multiplier variables $\mu_i : \mathbb{R}^{n+m} \rightarrow \mathbb{R}$ for each constraint $g_i(x, u)$, $i \in \{1, \dots, l\}$, which are polynomial basis functions themselves, we can rewrite all inequalities $g_i(x, u) \geq 0$ within the dissipativity inequality.

As a result, we get the following polynomial optimization problem:

$$\begin{aligned} & \max_{\sigma \in \mathbb{R}, \lambda \in \Lambda} \sigma, \\ \text{s.t.} \quad & \ell(x, u) - \sigma + \lambda(x) - \lambda(f(x, u)) \\ & - \zeta(|(x, u)|_{x^*, u^*}) - g(x, u)^\top \mu \geq 0, \end{aligned} \quad (5.9)$$

which can be solved efficiently using sum of squares. The solution σ^* represents a lower bound on the optimal stage cost $\ell(x^*, u^*)$; see [82]. Simultaneously, we require $\sigma^* \geq \ell(x^*, u^*)$ to satisfy the inequality in (5.7). Thus, concluding that we need $\sigma^* = \ell(x^*, u^*)$.

Results. We tested a set of cost functions for the following open thermal network with two states:

$$\begin{bmatrix} x_1 \\ x_2 \end{bmatrix}^+ = \begin{bmatrix} 1 - \frac{q+\alpha}{V} & 0 \\ \frac{q}{V} & 1 - \frac{q+\alpha}{V} \end{bmatrix} \begin{bmatrix} x_1 \\ x_2 \end{bmatrix} + \begin{bmatrix} \frac{q}{V} \\ 0 \end{bmatrix} x_o + \begin{bmatrix} 0 \\ \frac{1}{V} \end{bmatrix} w_d, \quad (5.10)$$

where we substituted $u_o(k) = q(k)x_o(k)$ in (5.6), where $q(k)$ is the mass flow and $x_o(k)$ is the pipe inlet temperature. Also, w_d is the extracted heat. The system is subject to box constraints $g(x, x_o, q) \leq 0$.

We choose a set of storage functions $\Lambda = \{\theta \in \mathbb{R}^6 | \lambda(x) = \sum_{d=1}^3 \theta_d x_1^d + \theta_{d+3} x_2^d\}$ and, similarly, define a set of multiplier functions of degree 1. Then, we solve problem

(5.9) to find σ^* . To confirm if σ^* guarantees strict dissipativity, we solve a separate optimization problem to find the value of the cost at the optimal equilibrium, i.e., $\ell(x^*, x_o^*, q^*) = \min \{ \ell(x, x_o, q) \mid x = f(x, x_o, q), \text{ and } (x, x_o, q) \in \mathcal{Z} \}$.

Thus, we are able to confirm strict dissipativity with respect to state and input, whenever $\sigma^* \geq \ell(x^*, x_o^*, q^*)$. Table 5.1 contains four possible cost functions with different equilibrium points for which this is the case.

Table 5.1: Examples of cost functions that guarantee strict dissipativity.

Stage cost	Equilibrium
$\ell = Rx_o^2 + q^2, R \geq 1/\rho^2$	Minimum inputs
$\ell = (x_o - x_{o,R})^2 + (q - q_R)^2$	Reference setpoints $(x_{o,R}, q_R)$
$\ell = -qx_1$	Maximum supply temp. and flow
$\ell = (qx_1 - E_R)^2$	Satisfies reference $q^* x_1^* = E_R$

Remark 16 We used $\zeta = \epsilon |x, x_o, q|_{x^*, x_o^*, q^*}$ as a \mathcal{K}_∞ function with $\epsilon = 10^{-7}$ to enforce strictness of the inequality with respect to states and inputs.

5.3. Economic Nonlinear Model Predictive Control

IN this section, we summarize a set of results from the literature on EMPC. In particular, we refer the reader to works such as [32, 33, 77, 78], for details on the methods and proofs. Here, we only state the important concepts relevant for this work, which are mostly related to the time-varying case, as presented in [77].

Consider the *time-varying* discrete-time dynamics

$$x(k+1) = f_k(x(k), u(k)), \quad x(0) = x_0, \quad (5.11)$$

with $x(k) \in \mathbb{R}^n$, $u(k) \in \mathbb{R}^m$, $k \in \mathbb{N}$, and $f_k: \mathbb{R}^n \times \mathbb{R}^m \rightarrow \mathbb{R}^n$.

The receding horizon optimal control problem without terminal conditions is defined as

$$\begin{aligned}
 V_{N,k}(x(k)) &= \min_{\substack{\{\mathbf{x}(i)\}_{i=0}^N \\ \{\mathbf{u}(i)\}_{i=0}^{N-1}}} \sum_{i=0}^{N-1} \ell_{k+i}(\mathbf{x}(i), \mathbf{u}(i)), \\
 \text{s.t. } \mathbf{x}(i+1) &= f_{k+i}(\mathbf{x}(i), \mathbf{u}(i)), \\
 \mathbf{x}(i) &\in \mathbb{X}(k+i), \quad \mathbf{u}(i) \in \mathbb{U}(k+i), \\
 i &= \{0, \dots, N-1\}, \\
 \mathbf{x}(0) &= x(k).
 \end{aligned} \quad (5.12)$$

The solution of the optimal control problem is an optimal trajectory $(\mathbf{x}^*, \mathbf{u}^*)$ for a specified horizon $N \in \mathcal{T}$. In MPC, the optimal control problem in (5.12) is solved iteratively. In each step, the feedback control law

$$\mu_{N,k}(x(k)) = \mathbf{u}^*(0)$$

is applied to the system. Generally, to ensure stability of the closed-loop system, a terminal cost and constraint are required. These conditions will ensure that the closed-loop converges to the infinite-horizon optimal trajectory. However, for nonlinear time-varying systems, it is not straightforward to compute these terminal ingredients, and in EMPC we typically do not have any knowledge of optimal trajectories. Instead, we are interested in showing whether the solution of the closed-loop MPC without terminal conditions approximates the solution of the infinite horizon problem. To this end, we introduce the following definitions.

Definition 7 (Time-varying strict dissipativity [77]) A discrete-time system (5.11) is strictly dissipative with respect to the supply rate $s_k(x, u) = \ell_k(x, u) - \ell_k(x^*(k), u^*(k))$ and the optimal trajectory $(x^*(k), u^*(k)) \in \mathbb{X}(k) \times \mathbb{U}(k)$ if there exists a storage function $\lambda_k : \mathbb{N} \times \mathbb{X}(k) \rightarrow \mathbb{R}$ bounded from below on $\mathbb{X}(k)$ with $\zeta \in \mathcal{K}_\infty$ such that for all $(x, u) \in \mathbb{X}(k) \times \mathbb{U}(k)$ and for all $k \in \mathbb{N}$, we have

$$\begin{aligned} & \ell_k(x, u) - \ell_k(x^*(k), u^*(k)) \\ & + \lambda_k(x) - \lambda_k(f_k(x, u)) \geq \zeta(|(x, u)|_{x^*(k), u^*(k)}). \end{aligned} \quad (5.13)$$

Assumption 2 (Continuity of $V_{N,k}$ at x^* [77]) There exists a function β_V such that for each $x \in \mathbb{X}$, $k \in \mathbb{N}$, $N \in \mathcal{T}$, the following holds:

$$|V_{N,k}(x) - V_{N,k}(x^*(k))| \leq \beta_V(|x|_{x^*(k)}). \quad (5.14)$$

This assumption also means that the optimal value function is bounded.

Strict dissipativity and boundedness of the optimal value functions are sufficient conditions for the existence of the turnpike property in time-varying systems [79]. Intuitively, the turnpike property suggests that the closed-loop EMPC solution remains close to the optimal trajectory for the majority of the time. As it turns out, the turnpike property will guarantee us three key properties regarding the performance of the EMPC [83]. These are as follows: approximate optimality, convergence to the optimal trajectory, and recursive feasibility.

5.4. Numerical Experiments and Results

IN this section, we consider a three consumer, one producer district heating network. The network is modeled as a tree structure, where all the flow from the producer is distributed to the consumers through a supply pipe. Each consumer has its own time-varying demand $w_{d,i} > 0$. The layout is illustrated in Fig 5.2.

We assume the return pipe has a constant temperature, and therefore, we do not include it in the model. Thus, the model of the network is derived from the open thermal network (2.23). A spatial discretization step of 500 meter is used, and, as a result, we have $n = 40$ states in the model.

Discretization. In order to solve the optimal control problem we require a discrete-time model of the system. Thermal systems typically have large inertia, and their dynamics change relatively slowly. Moreover, it is desirable to have sufficient time for solving the




Figure 5.2: DHN layout for simulations. The numbers on the edges denote the pipe length in kilometers.

optimal control problems to ensure tractability. Therefore, it is preferable to work with large time steps (i.e., 15 minutes to 1 hour); see [18] for more details. Accordingly, an implicit discretization of (2.23) yields the discrete-time dynamics (4.26).

Numerical simulation setup. The cost function is chosen as a weighted norm of the temperature at the inlet and the mass flow at the inlet, i.e., $\ell(x_o, q_o) = 50\alpha(k)x_o^2 + \rho^2 q_o^2$, with time-varying parameter $\alpha(k) = 0.5$ if $k \in [10, 15]$ and $\alpha(k) = 1$ otherwise. This cost function encourages minimization of temperature in the network, while also penalizing the flow rate, but also allows flexibility to build up a buffer during off peak hours. To give the controller sufficient time to react to demand and/or price variations, we choose a prediction horizon of $N = 8$ and time step $\Delta t = 3600$ s. Furthermore, constraints have been added to restrict large changes in inputs between time steps, which avoids fast switching. We solve the nonlinear EMPC problem (5.12) using *Ipopt* and, subsequently, apply the first element of the optimal sequence of inputs computed in the optimization to the plant, which is simulated using the approach described in Section 4.6.1, using a much finer spatial discretization with $\beta = 10$, simulated with a time step of one second using an ODE solver.

The results of Fig 5.3 show that the controller is able to maintain sufficient supply temperature for each of the consumers. At the same time, there is a period during the middle of the day where the heat production is ramped up to account for the evening peak. In addition, we find that in our simulations the open-loop trajectories generated by the EMPC stay near the closed-loop optimal trajectory, but diverge towards the end, as can be seen in Fig 5.4. This observation further supports our conjecture that the system exhibits the turnpike property.

5.5. Discussion and Conclusion

WE presented two different models that can be used in the context of MPC for DHNs, derived from partial differential equations governing the heat transport dynamics. We then provided both theoretical guarantees and numerical evidence for the existence of strict dissipativity in these networks, accompanied by a set of practically motivated objective functions that enable approximately optimal economic operation of the controller.

The results from this work have meaningful implications for DHN. In practice,




Figure 5.3: The pipe inlet temperature and the observed outlet temperatures at the consumers.




Figure 5.4: The predicted outlet temperature at consumer 3: the closed loop solution versus the open loop solutions at each iterate.

implementing MPC for DHN should remain simple, while also providing users with the ability to customize controller parameters to suit their preferences. Therefore, it is relevant to look at the case of EMPC without terminal conditions. We have provided evidence, and in some cases demonstrated theoretical validity, that the closed-loop solutions of the EMPC are approximately optimal and practically converge towards the optimal trajectory. Hence, these findings can assist practitioners when determining how to design or tune, for example, their objective function through a systematic verification approach. Nonetheless, we should acknowledge that the bridge to actual implementation has not been crossed yet. The verification results have potentially broader applicability to more general models and architectures. Furthermore, our modeling framework and verification approach may prove to be inadequate for certain DHN configurations, e.g., when there are discrete decisions involved. Despite these challenges, the results indicate an important first step towards the integration of EMPC theory for DHN applications.

For future work, the results can be extended to incorporate even more generalized sets of objective functions and model architectures. To this end, one could address the challenge of scalability and extend the application of our theory to larger, more complex real-world systems. Additionally, considering the emergence of smart energy systems, we expect that the analysis of a multi-producer network with enhanced emphasis on flexibility, such as through more detailed models for buffers or heat exchangers, would be of substantial value.

6

Scalable Predictive Control for District Heating Networks: Value Function Approximation using Dual Dynamic Programming

In this chapter, we use dual dynamic programming to address the myopic nature of MPC for scheduling of district heating networks by designing value functions that can approximate the effects of time-varying elements on the objective function beyond the initial prediction horizon. To this end, we formulate the control problem as a two-level MPC. More precisely, in the first-level, we consider a short-horizon nonlinear MPC equipped with a terminal cost approximating the value function. Subsequently, a long-horizon linear MPC is solved in the second-level to establish a lower bound on the terminal cost function from the first-level, thereby improving the value function approximation. Specifically, we consider scheduling of thermal and hydraulic components within district heating networks. Our numerical example demonstrates that our method can anticipate demand variations beyond the prediction horizon while maintaining computational efficiency.

This chapter is based on [37].

6.1. Introduction

SCHEDULING problems are common in various applications, including chemical processes [84], supply chain decision-making [85], manufacturing [86], and energy systems [87]. These problems involve a hierarchical decision-making approach, ranging from long-term planning to short-term control [84]. Typically, these systems are driven by a supply and demand dynamic that requires quick and timely decisions to meet operational demands. The complexity of these problems poses a significant challenge as they often involve multiple production and consumption units connected through transportation lines, such as pipe or transmission lines, that enable the flow of goods or quantities throughout the network. Furthermore, particularly in energy systems, there are additional challenges due to high fluctuations in demand and access to renewable energy generation. Effective decision-making is crucial in such systems at multiple levels.

Specifically, we concentrate on the scheduling problem for district heating networks (DHN). However, it is worth noting that our techniques can, in principle, be applied to other one-dimensional fluid-based scheduling processes. DHN can be categorized by two components: the temperature of the water that is being transported, which we call the state, and the speed at which it travels. We use the advection equation, a partial differential equation (PDE), to describe the one-dimensional dynamic evolution of the state throughout the network. Additionally, we assume that flow merging and splitting adhere to conservation laws (i.e., no leakage, incompressible flows). These components lend themselves to graph-based modeling approaches where edges represent transportation lines and nodes enforce conservation laws; see, e.g., [18] or [29].

Historically, the scheduling and control of DHN happen separately, using mainly decentralized rule-based controllers based on heuristics and common practices [12] with limited communication between central heat stations, substations, and consumers. We introduce a nonlinear model predictive control (MPC) algorithm that aims to capitalize on the mismatch between demand and supply in the district heating network by cleverly exploiting the thermal inertia of the system. Thus improving its efficiency and economic gain.

To ensure that we accomplish this target, it is crucial to have access to highly accurate DHN models. Unfortunately, high-resolution models are often not scalable, and thus could jeopardize the tractability of the optimal control problem. Some works have employed discretization catalogs [17, 23], i.e., sets of rules that adaptively determine the model resolution and complexity. Another study [42] addresses the issue of tractability through structure-preserving model order reduction techniques. However, neither study has tackled the problem in an online setting.

In an online setting, such as MPC, we are additionally constrained by the timeframe between control intervals. As a consequence, tractability is a more urgent issue. Often, especially for (mixed-integer) nonlinear MPC, we are limited by the number of steps we can predict ahead due to computational complexity and real-time constraints. Simultaneously, the advection of heat throughout the network is limited by the maximum allowable velocity of the fluid, suggesting an issue with controllability: if the time it takes for the fluid to reach a consumer is greater than the prediction horizon of the MPC, we cannot exert any control over that consumer. Hence, we cannot guarantee recursive feasibility of the closed-loop MPC. In some instances, this can be solved by adding

terminal ingredients to the problem; however, for time-varying nonlinear systems, it is not apparent how to find these ingredients [32]. For scheduling problems, this issue is often called *short-sighted* or *myopic* MPC, as it disregards any time-varying information beyond the prediction horizon, often leading to greedy short-term behavior [86, 87].

In this chapter, we introduce a dual dynamic programming formulation to approximately solve a long-horizon nonlinear MPC problem for district heating networks. The dual dynamic programming approach is based on generalized Benders' decomposition [88] and has been studied specifically for MPC by Flamm et al. [83]. Our contributions are as follows: *a)* we apply similar methods to those from [83] to district heating networks and show how they can effectively account for demand variations beyond the initial prediction horizon, *b)* our algorithm handles infeasible solutions to the second-level by generating a feasible cut for the first-level problem, and *c)* we demonstrate the computational strengths of the algorithm compared to state of the art solvers for the full nonlinear problem.

In the following, Section 6.2 describes the problem and introduces the model and MPC formulation. In Section 6.3, the two-level dual dynamic programming method and algorithm for long horizon MPC problems is explained. Subsequently, numerical examples demonstrating the efficacy of our methods are shown in Section 6.4. Finally, Section 6.5 concludes the section.

6.2. Short-Sightedness of MPC for DHN

OPERATIONAL optimization of district heating networks requires simultaneous optimization of load distribution and minimization of supply temperatures. However, these two controls are out of sync due to delays in heat transport. In fact, in larger systems, time delays can exceed ten hours [20]. In the Rotterdam heat network, for example, existing pipelines cover over 30 km, and the network may expand to over 50 km. As a result, it can take five or more hours for supply temperature changes to reach distant customers.

In essence, this means that any dynamic optimization-based controller should be capable of looking sufficiently far ahead in time to account for both load satisfaction and supply temperature minimization. However, the computational complexity of the nonlinear optimization problem is a function of the horizon length, and thus scales poorly for high-resolution models with many decision variables.

In this section, we will elaborate further on the problem setting. To this end, we need to first discuss the system model and introduce the MPC problem formulation.

MPC Problem. We introduce the nonlinear model predictive control problem which solves an optimal control problem with varying initial conditions in a receding horizon fashion. For a district heating network modeled by a graph $\mathcal{G} = (\mathcal{N}, \mathcal{E})$ as described in Chapter 2, we consider a specific variant of the open thermal network model from (2.23) where the flow velocity $v(k)$ (a control input) is isolated from the temperature inlet variable $u(k)$ (also a control input) leading to a bilinear interaction between the two control inputs. The full dynamical system, for any $\tilde{T}(k) := [\tilde{T}_i(k)]_{i \in \mathcal{N}}$, is described by the

nonlinear discrete-time dynamics

$$\begin{aligned}\tilde{T}(k+1) &= A(v(k))\tilde{T}(k) + B(v(k))u(k), \\ y(k) &= \tilde{T}(k).\end{aligned}\tag{6.1}$$

The heat demand satisfaction constraint, which is otherwise implicit to the formulation in (2.23), is now considered explicitly through a constraint

$$\rho c_p \Phi_i v_{e \rightarrow i}(k) (\tilde{T}_i(k) - \tilde{T}_R(k)) = w_{d,i}(k)\tag{6.2}$$

for any $i \in \mathcal{W}_C$, where $\tilde{T}_R(k)$ denotes a known time-varying return temperature and the subscript in $v_{e \rightarrow i}$ indicates the velocity of the water from the incoming edge to node i (where node i has at most one incoming edge since it is a consumer node).

We define a demand and return temperature vector $p(k) = [w_d(k)^\top, \tilde{T}_R(k)^\top]^\top$, which we consider to be a perfect forecast, i.e., uncertainty is not accounted for. Accordingly, we can formulate the optimization problem, henceforth using subscript time notation for internal optimization variables, as

$$\min_{x,u,v} \sum_{i=0}^{N-1} \ell(i, x(i), u(i), v(i)) + \ell(N, x(N))\tag{6.3a}$$

$$\text{s.t. } x(i+1) = f(x(i), u(i), v(i)),\tag{6.3b}$$

$$g(x(i), u(i), v(i), p(i)) \leq 0,\tag{6.3c}$$

$$(x(i), u(i), v(i)) \in \mathbb{X}(i) \times \mathbb{U}(i) \times \mathbb{V}(i),\tag{6.3d}$$

$$i = 0, \dots, N-1,\tag{6.3e}$$

$$x(0) = \tilde{T}(k),\tag{6.3f}$$

where (6.3b) corresponds to (6.1) and (6.3c) to (6.2). The optimal control problem is a nonlinear optimization problem and is nonconvex due to (6.1). Therefore, a major bottleneck consisting of three parts appears. Firstly, state x has dimension $|\mathcal{N}|$, meaning that the dimension scales with the number of pipes, and the district heating network typically has a large number of pipes. Secondly, the error between the discretized model and the original advection PDE is dictated by the truncated higher order terms appearing in the Taylor series expansion, which depend polynomially on the discretization step size. Hence, one should take care to choose stepsizes that are not too large to avoid significant model errors. Finally, due to the time delays from thermal transport it can take up to several hours before water inserted at the production side arrives at distant consumers. This implies that we need a sufficiently large prediction horizon N to exert any control over distant consumers. The size and nonlinear nature of the problem means that (6.3) will quickly become intractable for large N .

6.3. Dual Dynamic Programming

THIS section presents a dual dynamic programming approach to solving complex MPC problems for long horizons, originally introduced by Flamm *et al.* [83]. The idea

is as follows: the original nonlinear optimization problem is split into two levels where the first level solves the original nonlinear problem over a short horizon N_1 , i.e., we have

$$\begin{aligned}
\min_{x,u,v} \quad & \sum_{i=0}^{N_1-1} \ell(i, x(i), u(i), v(i)) + V(x(N_1)) \\
\text{s.t.} \quad & x(i+1) = f(x(i), u(i), v(i)), \\
& g(x(i), u(i), v(i), p(i)) \leq 0, \\
& (x(i), u(i), v(i)) \in \mathbb{X}(i) \times \mathbb{U}(i) \times \mathbb{V}(i), \\
& i = 0, \dots, N_1 - 1, \\
& x(0) = \tilde{T}(k).
\end{aligned} \tag{M}_{\text{NLP}}$$

We call this problem the *master problem*. The second-level problem is a linear approximation of the first problem around a possibly time-varying operating velocity $\bar{v}(k)$ at any time instance k . The first order Taylor series expansion of (6.1) has the form

$$\begin{aligned}
f(x(k)) &= f(\bar{x}) + \frac{\partial f}{\partial x}(\bar{x})(x(k) - \bar{x}) + \mathcal{O}((x(k) - \bar{x})^2) \\
&\approx \bar{x} + A(\bar{v}(k))(x(k) - \bar{x}) + B(\bar{v}(k))(u(k) - \bar{u}) \\
&\quad + \frac{\partial A(v(k))}{\partial v} \bar{x}(v(k) - \bar{v}(k)) + \frac{\partial B(v(k))}{\partial v} \bar{u}(v(k) - \bar{v}(k)),
\end{aligned} \tag{6.4}$$

for $x(k)$ near \bar{x} and $v(k)$ near $\bar{v}(k)$. Any equilibrium satisfies $\bar{u} = \bar{x}^1 = \bar{x}^2 = \dots = \bar{x}^m$, and due to the specific structure of $A(v(k))$ and $B(v(k))$, deriving from energy conservation laws, we find that

$$\frac{\partial A(v(k))}{\partial v} \bar{x} + \frac{\partial B(v(k))}{\partial v} \bar{u} = 0, \tag{6.5}$$

for any equilibrium point. Therefore, the linearization of (6.1) is equilibrium-independent with respect to \bar{x} and \bar{u} , because we can shift the state and input by arbitrary constants and keep the same linearization. As a result, we retain the same interpretation of the state and input, in this case, temperature, for the linear MPC problem. The only difference is that we substitute $v(k) = \bar{v}(k)$ which gives linear system matrices $\bar{A}(k)$ and $\bar{B}(k)$. The second-level problem is defined as

$$\begin{aligned}
V(x(N_1)) &= \min_{x,u} \sum_{i=N_1}^{N-1} (c(i)^\top x(i) + d(i)^\top u(i)) + c(N)x(N) \\
\text{s.t.} \quad & x(i+1) = \bar{A}(i)x(i) + \bar{B}(i)u(i), \\
& E_i x(i) + F_i u(i) \leq h_i(\bar{v}(i), p(i)), \\
& (x(i), u(i)) \in \mathbb{X}(i) \times \mathbb{U}(i), \\
& i = N_1, \dots, N - 1,
\end{aligned} \tag{S}_{\text{LP}}$$

where $h_i(\bar{v}(i), p(i))$ does not depend on any decision variables. Furthermore, cost function parameters $c(i)$ and $d(i)$ are time-varying to represent changes in energy prices.

6.3.1. Prediction Form MPC

The second-level problem has linear time-varying dynamics which means that we can reformulate the optimization problem to eliminate the state variables $\tilde{x} = (x(N_1 + 1)^\top, \dots, x(N)^\top)^\top$ by substituting

$$\tilde{x} = \mathbf{A}x(N_1) + \mathbf{B}\tilde{u} \quad (6.6)$$

with $\tilde{u} = (u(N_1)^\top, \dots, u_{N-1}^\top)^\top$,

$$\mathbf{A} = \begin{bmatrix} I \\ \bar{A}(N_1) \\ \bar{A}(N_1 + 1)\bar{A}(N_1) \\ \vdots \\ \prod_{i=N_1}^{N-1} \bar{A}(i) \end{bmatrix}, \quad \text{and}$$

$$\mathbf{B} = \begin{bmatrix} 0 & 0 & \cdots & 0 \\ \bar{B}(N_1) & 0 & \cdots & 0 \\ \bar{A}(N_1 + 1)\bar{B}(N_1) & \bar{B}(N_1 + 1) & \ddots & \vdots \\ \vdots & \ddots & \ddots & 0 \\ \prod_{i=N_1+1}^{N-1} \bar{A}(i)\bar{B}(N_1) & \cdots & \bar{A}(N-1)\bar{B}(N-2) & \bar{B}(N-1) \end{bmatrix},$$

such that (S_{LP}) can be rewritten in terms of complicating variable $x(N_1)$ and input vector \tilde{u} :

$$\begin{aligned} V(x(N_1)) &= \min_{\tilde{u} \in \mathcal{U}} c^\top (\mathbf{A}x(N_1) + \mathbf{B}\tilde{u}) + d^\top \tilde{u} \\ \text{s.t.} \quad & \mathbf{E}(\mathbf{A}x(N_1) + \mathbf{B}\tilde{u}) + \mathbf{F}\tilde{u} \leq \mathbf{h}(v), \end{aligned} \quad (6.7)$$

with $\mathbf{E} = \text{diag}(E_{N_1}, \dots, E_{N-1})$, $\mathbf{F} = \text{diag}(F_{N_1}, \dots, F_{N-1})$, $\mathbf{h} = \text{col}(h_{N_1}, \dots, h_{N-1})$, and $v = (\bar{v}(i), p(i))$.

6.3.2. Approximate Lower Bound on Value Function

Dual dynamic programming uses the dual solution of the second-level problem to find a lower bound on the value function. By iteratively generating lower bounds on the value function for different values of $x(N_1)$, convergence to the true value function can be achieved.

Here, we show the steps to find the value function. First, we denote an iteration of the algorithm by l , where we initiate $l = 1$ and update $l = l + 1$ after each iteration. This is repeated until a desired convergence tolerance is met. Second, let $(\cdot)^{(l)}$ denote a variable or function at iteration l , and $(\cdot)^{* (l)}$ denote its optimal value at iteration l . When adding iteration indices or optimal value indicators, we move function arguments to subscripts for compactness: e.g., the value of $x(N_1)$ at iteration l is written $x_{N_1}^{(l)}$, and its optimal value is $x_{N_1}^{* (l)}$. We write then the second-level problem by rearranging the terms in (6.7) to

obtain the linear program

$$\begin{aligned} V^{(l)}(x(N_1)) = \min_{\tilde{u} \in \mathcal{U}} & \quad c^\top \mathbf{A}x(N_1) + (c^\top \mathbf{B} + d^\top) \tilde{u} \\ \text{s.t.} & \quad \mathbf{E}\mathbf{A}x(N_1) + (\mathbf{E}\mathbf{B} + \mathbf{F}) \tilde{u} \leq \mathbf{h}(v). \end{aligned} \quad (\mathbf{S}_{\text{LP}}^{(l)})$$

If there exists an optimal solution $\tilde{u}^{*(l)}$ to problem $(\mathbf{S}_{\text{LP}}^{(l)})$ for some $x(N_1)$, we can use the dual solution $\lambda^{*(l)}$ to find a linear lower bound on the primal problem. The steps to arrive at this lower bound are presented in Appendix 6.A. We note that the derivation is very similar to [83], but not exactly the same (see Remark 17).

Finally, we can derive the following function as

$$\begin{aligned} L^{(l)}(x(N_1)) = & (c^\top \mathbf{A} + (\lambda^{*(l)})^\top \mathbf{E}\mathbf{A})x(N_1) \\ & - (\lambda^{*(l)})^\top \mathbf{h}(v) \\ & + \left(c^\top \mathbf{B} + d^\top + (\lambda^{*(l)})^\top (\mathbf{E}\mathbf{B} + \mathbf{F}) \right) \tilde{u}^{*(l)}, \end{aligned} \quad (6.8)$$

for which we know that the maximum over all iterations is a lower bound on the true value function, i.e.,

$$\hat{V}^{(l)}(x(N_1)) = \max_{1 \leq j \leq l} L^{(j)}(x(N_1)) \leq V(x(N_1)). \quad (6.9)$$

Consequently, this lower bound is added to $(\mathbf{M}_{\text{NLP}})$ by introducing a new variable σ to the objective function and constraining it to be greater than or equal to the approximate value function; see $(\mathbf{M}_{\text{NLP}}^{(l)})$.

Remark 17 The authors of [83] assume the lower bound is obtained whenever the *partial* Lagrangian satisfies stationarity conditions, which neglects any non-complicating constraints that may be active at the optimal solution. Therefore, we have an additional term show up in the lower bound (bottom line in (6.8)). \diamond

6.3.3. Infeasible Subproblems

If $(\mathbf{S}_{\text{LP}}^{(l)})$ is infeasible for $x_{N_1}^{*(l)}$ found by the master problem, we generate a Benders' cut that acts as a constraint on the master problem $(\mathbf{M}_{\text{NLP}})$ for the next iterate. More precisely, we want to determine a solution $\tilde{u}^{*(l)}$ and vector $\eta^{*(l)} \geq 0$ such that there exists a convex combination of constraints greater than zero, i.e.,

$$(\eta^*)^\top (\mathbf{E}\mathbf{A}x_{N_1}^* + (\mathbf{E}\mathbf{B} + \mathbf{F}) \tilde{u}^* - \mathbf{h}(v)) > 0, \quad (6.10)$$

where $\sum_{i=1}^{n_c} \eta(i) = 1$. Subsequently, we cut this part out of the feasible set of the master problem by introducing the constraint

$$(\eta^*)^\top (\mathbf{E}\mathbf{A}x(N_1) + (\mathbf{E}\mathbf{B} + \mathbf{F}) \tilde{u}^* - \mathbf{h}(v)) \leq 0. \quad (6.11)$$

The new solution to the master problem will ensure that there exists at least one feasible solution to the second-level problem $V^{(l+1)}(x_{N_1}^{*(l+1)})$ for $\tilde{u} = \tilde{u}^{*(l)}$. However, it could be




Figure 6.1: Illustration of the iterative feasibility recovery mechanism in the hierarchical scheme. Starting from x_0 , the first-level problem ($M_{\text{NLP}}^{(1)}$) is solved, yielding a terminal state x_{N_1} at the handover point. Subsequently, the second-level problem ($S_{\text{LP}}^{(1)}$) is solved, but may return an infeasible solution. To recover feasibility, a terminal constraint is imposed on the first-level problem, requiring an elevated state $\tilde{x}_{N_1,A}$ or $\tilde{x}_{N_1,B}$ at N_1 . The required elevation, however, depends on the minimal feasible trajectory achievable by the second-level problem.

undesirable to choose any arbitrary \tilde{u}^* . For instance, as illustrated in Figure 6.1, if we choose \tilde{u}^* to maximize the gap (to zero) in (6.10), which can be achieved by choosing \tilde{u}^* as small as possible, then the constraint on $x(N_1)$ added to master problem (M_{NLP}) will be very tight, resulting in conservative or even infeasible solutions for the master problem.

Therefore, we design a linear program to compute \tilde{u} such that the constraint gap (Δ_A , Δ_B in Figure 6.1) is as small as possible. Specifically, we minimize the sum of all constraints by solving the following linear program

$$\begin{aligned} \min_{\tilde{u} \in \mathcal{U}, \gamma \in \mathbb{R}} \quad & \gamma \\ \text{s.t.} \quad & \mathbf{1}^\top (\mathbf{E}\mathbf{A}x_{N_1}^* + (\mathbf{E}\mathbf{B} + \mathbf{F})\tilde{u} - \mathbf{h}(v)) \leq \gamma. \end{aligned} \quad (\text{S}_{\text{inf}})$$

Problem (S_{inf}) is guaranteed to find an optimal solution which we will denote $\tilde{u}_{\text{inf}}^{(l)}$. Subsequently, we add the following constraint to the master problem:

$$\mathbf{g}_{\text{inf}}^{(l)}(x(N_1)) = \mathbf{E}\mathbf{A}x(N_1) + (\mathbf{E}\mathbf{B} + \mathbf{F})\tilde{u}_{\text{inf}}^{(l)} - \mathbf{h}(v) \leq 0. \quad (6.12)$$

Algorithm. In short, the second-level problem always generates a constraint for the master problem, even when it is infeasible. More precisely, we can describe the master problem at iteration l by

$$\begin{aligned}
M^{(l)}(\tilde{T}(k)) &= \min_{x,u,v,\sigma} \sum_{i=0}^{N_1-1} \ell(i, x(i), u(i), v(i)) + \sigma \\
\text{s.t. } & \mathbf{C}(\mathbf{M}_{\text{NLP}}), \\
& \sigma \geq L^{(j)}(x(N_1)), \quad j = 0, \dots, l-1, \\
& g_{\text{inf}}^{(j)}(x(N_1)) \leq 0, \quad j = 1, \dots, l-1,
\end{aligned} \tag{M_{\text{NLP}}^{(l)}}$$

where $\mathbf{C}(\mathbf{M}_{\text{NLP}})$ is the set of constraints of the original master problem. The previously described steps of the algorithm are summarized in Algorithm 2.

Algorithm 2 Value function approximation for nonlinear MPC

Input: Initial state $\tilde{T}(k)$, dataset $p_{k:k+N}$, prediction $\bar{v}_{k+N_1:k+N}$, cost functions $\ell_{k:k+N_1}$, $c_{k+N_1:k+N}$, and $d_{k+N_1:k+N-1}$.

Output: Optimal control (u_0^*, v_0^*)

for $k \in [t_0, t_{\text{end}}]$ **do**

Initialize: $l \leftarrow 1$, $\text{LB} \leftarrow -\infty$, $\text{UB} \leftarrow \infty$, $\epsilon \leftarrow 10^{-1}$, $L^{(0)} \leftarrow -10^5$

while $\text{UB} \geq \text{LB} + \epsilon$ **do**

Solve master problem $(\mathbf{M}_{\text{NLP}}^{(l)})$ to find $x_{N_1}^{*(l)}$, $\sigma^{*(l)}$, and $M^{*(l)}(\tilde{T}(k))$

Set $\text{LB} = M^{*(l)}(\tilde{T}(k))$

Solve subproblem (\mathbf{S}_{LP}) for $x(N_1) = x_{N_1}^{*(l)}$

if (\mathbf{S}_{LP}) is optimal **then**

Set $\text{UB} = M^{*(l)}(\tilde{T}(k)) - \sigma^{*(l)} + V^{*(l)}(x(N_1))$

Add $\sigma \geq L^{(l)}(x(N_1))$ to $(\mathbf{M}_{\text{NLP}}^{(l)})$

else if (\mathbf{S}_{LP}) is infeasible **then**

Solve $(\mathbf{S}_{\text{inf}})$ to find $\tilde{u}_{\text{inf}}^{(l)}$

Add $g_{\text{inf}}^{(l)}(x(N_1)) \leq 0$ to $(\mathbf{M}_{\text{NLP}}^{(l)})$

end if

$l \leftarrow l + 1$

end while

Return (u_0^*, v_0^*) from $(\mathbf{M}_{\text{NLP}}^{(l)})$

end for

6.4. Numerical Examples

IN this section, we provide numerical examples to demonstrate the performance of Algorithm 2. All simulations were performed using *The Julia Programming Language* with *Ipopt* used to solve nonlinear problems and *HIGHS* used for the linear programs.

We consider a system consisting of one producer and two consumers, where one consumer is located much farther from the producer than the other, see Figure 6.2. Furthermore, we have $x(k) \in \mathbb{R}^{10}$, $\Delta t = 300 \text{ s} = 5 \text{ min}$, and $\Delta x = 500 \text{ m}$. Consequently, the network reaches 3.5 km in length. For practical reasons, the backward Euler method was used to discretize the system and a slack variable was added to the demand constraint to avoid infeasible solutions from the master problem. Nevertheless, this slack variable is

used to quantify the level of demand violation that occurs in the myopic MPC. Practically, however, this slack variable can be interpreted to be the heat provided by a gas boiler that covers any demand gaps locally, but incurs very high costs.




Figure 6.2: The layout of the small-scale DHN with two consumers. Consumer c_1 is located significantly further from the inlet than consumer c_2 .

We consider the following two scenarios:

- *Scenario 1:* A scenario where the MPC has to deal with a sudden, but predicted, increase in demand at consumer c_1 . We show that by adding the second-level value function approximation we can prevent infeasibility with minimal added computational cost.
- *Scenario 2:* A generic scenario with randomized fluctuating demand profile to estimate the relative performance and computational efficiency of Algorithm 2 compared to fully nonlinear MPC for different prediction horizons.

Scenario 1. In this example, we compare the performance of a short-horizon nonlinear MPC to the performance of Algorithm 2 with the same initial horizon, but now augmented with the second-level problem. We have an objective function $\ell^1(u, v, s) = \mathbf{1}^\top u + \|v\|^2 + 10^4 \|s\|^2$ for the master problem, where s represents a slack variable. For the subproblem, we have chosen a cost function $\ell^2(\tilde{u}) = \mathbf{1}^\top \tilde{u}$. For the velocity of the second-level we assume $\bar{v}(k) = 0.5$ for all k . In Figure 6.3, the performance of full nonlinear MPC with $N_1 = N = 4$ is shown when there is a sudden demand increase at the distant consumer at $k = 13$. Similarly, Figure 6.4 shows the performance of the algorithm for the same scenario when $N_1 = 4$ and $N = 20$. We note that the reason for the difference in lower bound $g_k(\bar{T}_{c_1}(k)) \leq 0$ between the two cases appears due to the implicit dependence on velocity decision variable v_k .

Discussion: Algorithm 2 successfully prevented short-sighted behavior by the MPC without requiring the solution of the entire nonlinear problem. In Figure 6.3, we see precisely the problem depicted in Figure 6.1. In Figure 6.4, the controller is able to mitigate this problem and satisfy the consumer's demand almost perfectly (the constraint violation also persists for the full nonlinear model). While this example demonstrates a good approximation of the full nonlinear problem, it is important to acknowledge the limitations of this method. In particular, it remains unclear how to select the




Figure 6.3: The input and consumer state trajectory simulated for 25 time steps with $N_1 = N = 4$. The black dashed line represents the minimum required temperature at consumer c_1 , when it exceeds $T_{c_1}(k)$ (denoted in the legend as $x_k^{c_1}$) the consumer's demand is not satisfied.




Figure 6.4: The input and consumer state trajectory simulated for 25 time steps with $N_1 = 4$ and $N = 20$. The black dashed line represents the minimum required temperature at consumer c_1 .

operating velocity for linearization in the second-level problem. This velocity choice can significantly impact the feasibility and solution of the second-level. Additionally, if there is no one-to-one correspondence of the cost functions between the first and second level, closed-loop solutions can differ greatly compared to the full nonlinear problem.

Scenario 2. In this scenario, we simulate the problem for a range of different prediction horizons. In each simulation, we run a total of 60 time steps, which is equal to 5 hours. Both the demand $w_d(k)$ and relative cost parameter $d(k)$ are generated as periodic time-varying signals with random disturbances. For the first-level cost function $\ell_k^1(u, v, s) = [d_j]_{j=k:k+N_1-1}^\top u + 100\|v\|^2 + 100\|s\|^2$ we reduce the penalty on the slack variable so that it does not dominate the final cost, and add relative pricing on the input weight. For the subproblem we drop the last two terms to obtain $\ell_k^2(\tilde{u}) = [d_j]_{j=k+N_1:k+N-1}^\top \tilde{u}$. Finally, we simulate the full nonlinear MPC from $N = 5$ to $N = 50$ and compare the average time taken for each iteration k and the closed-loop cost (the cost evaluated at the closed-loop solutions) to the solutions of Algorithm 2 when we fix $N_1 = 5$ and $N_1 = 10$. The results are shown in Figure 6.5.

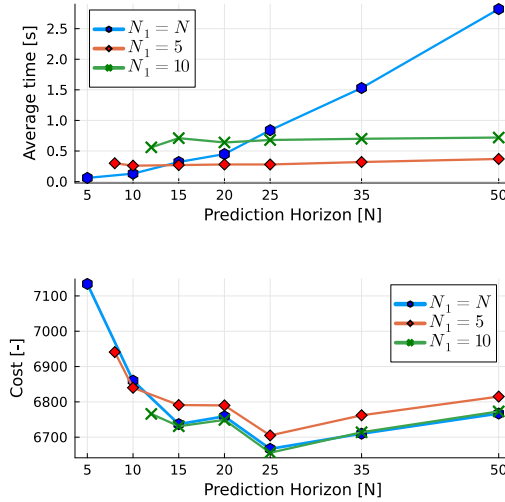


Figure 6.5: The top figure shows the average time per iteration of the MPC and the bottom figure shows the closed-loop cost for the full nonlinear case and for the solution of Algorithm 2 with $N_1 = 5$ and $N_1 = 10$.

Discussion: The findings presented in Figure 6.5 demonstrate that Algorithm 2 achieves a noticeable improvement in computation time for larger prediction horizons compared to the fully nonlinear problem. Additionally, we observed only slightly reduced costs for the fully nonlinear case. Nevertheless, it is essential to note that the scale and complexity of full-scale DHNs are much larger than the ten-state example we analyzed in this study. Furthermore, many control decisions in DHNs are subject to combinatorial constraints,

such as dwell-time constraints, which result in mixed-integer formulations; see, e.g., [25]. These problems are notably more difficult to solve than the nonlinear solver used for this example, *Ippopt*, making them less computationally efficient. Therefore, we anticipate that the need for tractable formulations will be even more pronounced for full-scale model representations of DHNs. These considerations further underline the potential of our approach and present a promising direction for future research. As computational demands continue to grow with increasing network complexity and the integration of renewable energy sources, methods that systematically reduce problem complexity while preserving solution quality will become increasingly valuable for enabling real-time control of next-generation district heating systems.

6.5. Conclusion

WWE have developed a two-level algorithm for controlling fluid-based scheduling problems using model predictive control, focusing on district heating networks. Our approach involves iteratively generating a lower bound on the terminal cost that captures time-varying elements beyond the prediction horizon. By doing this, the controller can cheaply compute optimal inputs to the system that implicitly account for long-term effects. Our method shows great promise in search of tractable solutions to complex scheduling problems over networks.

6.A. Appendix: Lower Bound on Value Function

We show the steps to arrive at a lower bound for $V(x(N_1))$. Firstly, problem $(S_{LP}^{(l)})$ has a linear programming form

$$\begin{aligned} V(y) = \min_{x \in \mathbb{R}^n} \quad & a^\top x + b^\top y \\ \text{s.t.} \quad & Gx + Hy \leq J. \end{aligned} \tag{6.13}$$

The Lagrangian is

$$\begin{aligned} \mathcal{L}(x, y, \lambda) &= a^\top x + b^\top y + \lambda^\top (Gx + Hy - J) \\ &= -\lambda^\top J + (a^\top + \lambda^\top G)x + (b^\top + \lambda^\top H)y, \end{aligned} \tag{6.14}$$

with $\lambda \geq 0$. The dual problem is defined as

$$\begin{aligned} D(y) &= \max_{\lambda} \min_{x \in \mathbb{X}} \mathcal{L}(x, y, \lambda) \\ &= \max_{\lambda} (b^\top + \lambda^\top H)y - \lambda^\top J + (a^\top + \lambda^\top G)x_s, \end{aligned} \tag{6.15}$$

and for some y_s that solves $D(y_s)$ with optimal primal values x_s and optimal dual variables λ_s , we know that λ_s is feasible for any y . However, λ_s is not necessarily optimal for other choices of y . Thus, we have

$$D(y) \geq (b^\top + \lambda_s^\top H)y - \lambda_s^\top J + (a^\top + \lambda_s^\top G)x_s. \tag{6.16}$$

Additionally, from strong duality we have

$$V(y) = D(y). \quad (6.17)$$

As a result, the linear function

$$(b^\top + \lambda_s^\top H)y - \lambda_s^\top J + (a^\top + \lambda_s^\top G)x_s \quad (6.18)$$

is a valid lower bound for $V(y)$.

7

Scalable Predictive Control for District Heating Networks: A Physics-Guided Koopman Operator Approach

We develop a physics-guided systematic methodology to construct meaningful Koopman observables by integrating the graph topology and fundamental thermodynamic conservation laws governing the district heating network. Our approach ensures that critical nonlinear interactions and energy transport phenomena are accurately captured and provides physically interpretable lifted representations. Using the obtained linear model, we formulate the control problem in a scalable Koopman-based receding horizon fashion, where at each step a convex quadratic program is solved. We perform extensive numerical validation on benchmark DHNs, demonstrating that the proposed Koopman predictive control strategy achieves a computational speedup exceeding an order of magnitude over state-of-the-art nonlinear MPC schemes, while maintaining superior or comparable control performance, thereby establishing its practical applicability for real-time control of large-scale DHNs. Furthermore, we extend the methodology to DHNs with bidirectional-flow pipes, providing a tractable optimization framework with superior control performance compared to nonlinear MPC.

This chapter is based on [38].

7.1. Introduction

MODEL predictive control represents a promising approach for DHNs through its ability to optimize control actions over future horizons while respecting operational constraints, which is valuable for leveraging thermal inertia and for satisfying heat demand in district heating networks. Nonetheless, applying MPC to DHNs is challenging due to the computational intractability of solving nonlinear programs (NLPs) for large-scale DHN models, which are derived from discretized partial differential equations (PDEs) [18, 89].

Early approaches utilized simplified linear models or decoupled hydraulic and thermal dynamics [13, 15, 66], leading to oversimplification and inferior control performance. Recent approaches to DHN optimization rely on complex, nonconvex, and large-scale NLPs [12, 18, 25, 34, 41, 68, 69, 90], which are substantially computationally demanding and often result in spurious local minima and suboptimal solutions, imposing a significant hurdle for real-time implementation and causing considerably lower operational performance.

To mitigate the computational complexity of the problem various approaches have been developed, including model reduction techniques that simplify the system dynamics [42] and problem decomposition strategies that break down the large optimization problem into smaller, more manageable subproblems [37, 91, 92]. Besides these existing methods, Koopman operator theory [93] presents promising approaches as it has potential to offer particularly effective data-driven schemes for modeling complex nonlinear systems, transforming them into linear ones, and enabling more computationally efficient MPC.

Koopman operator theory provides a rigorous framework for representing nonlinear dynamics as linear systems by lifting them to a higher-dimensional space through transformations called *observables* [94–96]. Leveraging the resulting system characterization, Koopman-based predictive control has been developed and proven as an effective data-driven method for various nonlinear control problems, including automated driving [97], synchronization of coupled pendulums [98], waste-heat recovery in organic Rankine cycles [99], thermal management in electric vehicles [100], and laboratory-scale thermal systems [101]. For DHN control, other data-driven methods have been developed, including neural network-based modeling for district cooling networks [102], physics-informed recurrent neural networks that embed network topology [103], and graph neural networks for capturing geometric and topological structure [104]. Nevertheless, the application of Koopman-based predictive control to large-scale DHNs remains generally unexplored, despite its considerable potential for addressing the aforementioned computational challenges while providing data-driven and interpretable control frameworks.

Contributions. In this chapter, we present a physics-guided Koopman predictive control (KPC) strategy specifically designed for district heating networks. The main contributions are stated as follows:

- **Physics-informed Koopman observable design**

We develop a systematic methodology for constructing physically meaningful Koopman observables for large-scale networked systems by leveraging graph topology and thermodynamic conservation laws. Our approach exploits the

network structure to identify only those bilinear interactions corresponding to physical energy transport at a reduced set of nodes derived from the graph, yielding a sparse observable set that scales linearly with network size and ensures that critical nonlinear interactions are accurately captured while maintaining interpretable lifted representations suitable for real-time control. This construction avoids including all possible bilinear combinations, which would scale poorly with discretization refinement, and instead yields a dictionary that scales linearly with network size.

- **Scalable predictive control for DHNs**

We integrate the Koopman predictor into a receding horizon control framework that achieves the computational efficiency necessary for real-time control, establishing what is, to the best of the authors' knowledge, the first viable path toward scalable predictive control for large-scale district heating networks.

- **Novel bidirectional flow formulation for DHNs**

We extend the convex Koopman framework to networks with bidirectional flow pipes. Modeling flow direction as a discrete decision in the original bilinear dynamics leads to an intractable formulation; we show that the Koopman formulation yields a more tractable mixed-integer quadratic program that outperforms the complementarity reformulation in closed-loop performance.

Outline. The remainder of this chapter is structured as follows. Section 7.2 establishes the coupled thermal-hydraulic model. Section 7.3 presents the Koopman operator-based lifting framework for DHN modeling, encompassing theoretical foundations and data-driven approximation techniques. Section 7.4 details the physics-guided observable selection methodology for DHNs. Accordingly, Section 7.5 formulates the resulting Koopman predictive control problem for DHNs. Section 7.6 presents numerical results and case studies. Finally, Section 7.7 concludes with a summary of contributions and future research directions.

7.2. Thermo-Hydraulic Dynamics

WHILE this chapter employs the same foundational DHN model from Chapter 2, we do not impose the steady-state hydraulics assumption. This introduces an additional time-derivative term in equation (2.7) leading to the following PDE for any $e \in \mathcal{E}$:

$$\partial_t q_e + \frac{\pi d_e^2}{4\rho L_e} \Delta p_e = -\frac{2K_e}{\pi d_e^3} |q_e| q_e. \quad (7.1)$$

We augment this model with pressure changes from pumps and valves. For pump edges $e \in \mathcal{P} \subset \mathcal{E}$, we have $\Delta p_e = R_{\mu,e} q_e^2 - h_e$. Similarly, for each valve edges $e \in \mathcal{V} \subset \mathcal{E}$, one has $\Delta p_e = R_{\mu,e} q_e^2 + R_{\nu,e}(v_e) q_e^2$, where h_e and v_e are pump and valve control variables, respectively.

Combining the pipe, pump, and valve equations in matrix form and multiplying from the left by the incidence matrix \mathbf{F} gives

$$\mathbf{F}\mathbf{F}^\top \dot{q}_f = -\mathbf{F}\mathbf{R}(v) \text{diag}(\mathbf{F}^\top q_f q_f^\top \mathbf{F}) + \mathbf{F}\mathbf{B}_{\mathcal{P}} h, \quad (7.2)$$

where $R(v) = R_\mu + R_\nu(v)$ contains resistance factors and $B_{\mathcal{P}}$ relates pumps to edges. Since \mathbf{F} has full row rank, $\mathbf{F}\mathbf{F}^\top$ is invertible, giving the fundamental flow dynamics:

$$\dot{q}_f = -\zeta(q_f, v) + Ph, \quad (7.3)$$

where $\zeta(q_f, v) = (\mathbf{F}^\top)^\dagger R(v) \text{diag}(\mathbf{F}^\top q_f q_f^\top \mathbf{F})$ and $P = (\mathbf{F}^\top)^\dagger B_{\mathcal{P}}$ with $(\mathbf{F}^\top)^\dagger = (\mathbf{F}\mathbf{F}^\top)^{-1}\mathbf{F}$.

Given the complexity of DHN operations and the need for reliable temperature regulation, practical implementations typically employ hierarchical control architectures [105]. At the operational level, local controllers manage individual components (pumps and valves) to track setpoints, while higher-level supervisory control coordinates these setpoints for system-wide optimization. Specifically, for any desired setpoint $\bar{q}_f \in \mathcal{Q} \subset \mathbb{R}^{|\mathcal{S}|}$, where \mathcal{Q} is a polytopic set of nonnegative flows, we assume the existence of local control systems capable of achieving $\zeta(\bar{q}_f, v(\bar{q}_f)) = Ph(\bar{q}_f)$ at steady state, with \bar{q}_f serving as a locally asymptotically stable equilibrium. To formalize this control structure, we consider the tracking error dynamics between actual flows q_f and desired setpoints \bar{q}_f . Define the closed-loop error $e = q_f - \bar{q}_f$, where the local controllers are designed to ensure the error dynamics satisfy

$$\dot{e} = -Ke \quad (7.4)$$

where $-K$ is Hurwitz. This structure arises naturally from decentralized proportional-type controllers acting on individual flows, where K is diagonal with positive entries corresponding to local controller gains.

Coupled Thermal-Hydraulic State-Space Description

Let us define a shifted temperature state, i.e., $\tilde{T}_i = T_i - T_a$, for $i \in \mathcal{N}$, and the corresponding state vector $\tilde{T} = [\tilde{T}_i]_{i \in \mathcal{N}}$, representing the exergy (available energy) of a node with respect to the ambient temperature. We construct a combined dynamics for the coupled thermal-hydraulic system where the thermal dynamics are as in (2.23). We have

$$\begin{bmatrix} \dot{q}_f \\ V\dot{\tilde{T}} \end{bmatrix} = \begin{bmatrix} -\zeta(q_f, v) \\ A(q_f)\tilde{T} \end{bmatrix} + \begin{bmatrix} P & 0 \\ 0 & B \end{bmatrix} \begin{bmatrix} h \\ w \end{bmatrix}, \quad (7.5)$$

which can be written considering the hydraulic part in closed-loop as follows:

$$\begin{bmatrix} \dot{q}_f \\ V\dot{\tilde{T}} \end{bmatrix} = \begin{bmatrix} -K & 0 \\ 0 & A(q_f) \end{bmatrix} \begin{bmatrix} q_f \\ \tilde{T} \end{bmatrix} + \begin{bmatrix} K & 0 \\ 0 & B \end{bmatrix} \begin{bmatrix} \bar{q}_f \\ w \end{bmatrix}. \quad (7.6)$$

Henceforth, we denote the combined input variable as $u = [\bar{q}_f^\top w^\top]^\top$, where distinguish between flow setpoints, producer inputs, and consumer disturbances using u^q , u^p , and u^c , respectively.

7.3. Koopman Operator Theory

OUR goal is to obtain a predictive model suitable for linear MPC despite the inherently nonlinear dynamics of DHNs. To this end, we employ Koopman operator theory, which provides a mathematically rigorous framework for lifting nonlinear systems into a

higher-dimensional space where their evolution can be modeled by linear dynamics. Following [94], we first provide the theoretical foundation of Koopman operators in Section 7.3.1, then develop a data-driven identification algorithm using extended dynamic mode decomposition (EDMD) in Section 7.3.2 and Section 7.3.3.

7.3.1. Koopman Operator Framework

The Koopman operator is a linear infinite-dimensional operator that describes the evolution of so-called observable functions along trajectories of a dynamical system [93]. Consider an autonomous nonlinear dynamical system

$$x^+ = f(x), \quad (7.7)$$

where $x \in \mathbb{R}^n$ is the state vector, $f: \mathbb{R}^n \rightarrow \mathbb{R}^n$ denotes the vector field governing the dynamics, and x^+ denotes the next state obtained by applying f to x . Furthermore, let \mathcal{H} be a set of scalar-valued functions defined on the state space of (7.7). Then, the Koopman operator $\mathcal{K}: \mathcal{H} \rightarrow \mathcal{H}$ is defined as a linear map that assigns to each $\psi \in \mathcal{H}$ the composition of ψ with f , i.e., $\mathcal{K}\psi = \psi \circ f$. More precisely, for any $\psi \in \mathcal{H}$ and any $x \in \mathbb{R}^n$, we have

$$(\mathcal{K}\psi)(x) = \psi(f(x)). \quad (7.8)$$

One can easily see that, despite the nonlinearity of f , the Koopman operator \mathcal{K} is linear in the space of observables. This fundamental property enables us to reformulate nonlinear dynamics in an infinite linear framework.

To apply Koopman theory to the DHN system, the framework needs to be extended to accommodate dynamics with control inputs [94], i.e., to nonlinear controlled systems described by

$$x^+ = f(x, u), \quad (7.9)$$

where x , x^+ , and f are defined as above, and $u \in \mathcal{U} \subset \mathbb{R}^m$ denotes the control input, with \mathcal{U} representing the set of admissible control actions. Note that, once we obtain a time-discretized dynamical model from (7.6), in the resulting formulation, x represents the vector of flow q and shifted temperature \tilde{T} , and u is the vector of flow setpoints u^q , producer inputs u^p , and consumer disturbances u^c . Following [94], the *extended state* is then defined as $\chi = [x^\top, \mathbf{u}^\top]^\top$, where $\mathbf{u} = (u_i)_{i=0}^\infty$ represents a control sequence within the space of all sequences \mathcal{U}^∞ . Subsequently, the *extended system* describes the dynamics of χ as

$$\chi^+ = \begin{bmatrix} f(x, \mathbf{u}(0)) \\ \mathcal{S}\mathbf{u} \end{bmatrix}, \quad (7.10)$$

where $\mathbf{u}(0)$ denotes the first element of \mathbf{u} and \mathcal{S} represents the shift operator that advances the control sequence one step. Note that $\mathbf{u}(0)$ is indeed control input u at the current time. Given the introduced settings, the Koopman operator $\mathcal{K}: \mathcal{H} \rightarrow \mathcal{H}$ is defined similarly as before but for the extended system (7.10), i.e., we have

$$(\mathcal{K}\varphi)(\chi) = \varphi(\chi^+), \quad (7.11)$$

for any observable function $\varphi: \mathbb{R}^n \times \mathcal{U}^\infty \rightarrow \mathbb{R}$ belonging to observable space \mathcal{H} .

7.3.2. Data-Driven Model Approximation for DHNs

In practice, we need a finite-dimensional approximation of the Koopman operator suitable for trajectory prediction of DHNs. To this end, we define the combined state variable at time $k \in \mathbb{Z}$ as

$$x(k) = \begin{bmatrix} q_f(k) \\ \tilde{T}(k) \end{bmatrix}. \quad (7.12)$$

The Koopman operator for the augmented state-input system retains linearity in the space of observables, allowing us to identify linear predictors of the form

$$z(k+1) = Az(k) + Bu(k), \quad (7.13)$$

$$x(k) = Cz(k), \quad (7.14)$$

where $z(k) := \psi(x(k)) \in \mathbb{R}^{n_K}$ is the lifted state vector. In practical settings, n_K is chosen to be sufficiently large so that the lifted dynamics provide an accurate representation of the nonlinear system. It is also noted that $x(k)$ can be recovered exactly whenever it is contained within $\psi(x(k))$.

The key is to select appropriate lifting functions $\psi(x) = [\psi_1(x), \psi_2(x), \dots, \psi_{n_K}(x)]^\top$ that transform the original state into the lifted space. Then, the initial condition for this predictor is

$$z(0) = \psi(x(0)) = \begin{bmatrix} \psi_1(x(0)) \\ \vdots \\ \psi_{n_K}(x(0)) \end{bmatrix}. \quad (7.15)$$

To realize the discussed approach for DHNs, we opt for the extended dynamic mode decomposition EDMD algorithm, implementing a structure where observable functions take the form as

$$\varphi_i(x, \mathbf{u}) = \psi_i(x) + L_i(\mathbf{u}), \quad (7.16)$$

where $\psi_i : \mathbb{R}^n \rightarrow \mathbb{R}$ are general nonlinear functions of the state, and $L_i : \mathcal{U}^\infty \rightarrow \mathbb{R}$ are linear functionals of the control sequence. For tractability, we focus on a finite vector of observable functions, i.e.,

$$\varphi(x, \mathbf{u}) = \begin{bmatrix} \psi(x) \\ \mathbf{u}(0) \end{bmatrix}, \quad (7.17)$$

where $\mathbf{u}(0) \in \mathbb{R}^m$ is the first component of \mathbf{u} .

7.3.3. Numerical Identification Algorithm

To identify the matrices A and B in (7.13) for our DHN model, we use data collected from system simulations or historical data from operational measurements. Given a dataset consisting of state transitions and control inputs as

$$\mathcal{D} = \{(x(k), u(k), x(k+1)) : k = 0, 1, \dots, K-1\}, \quad (7.18)$$

we formulate and solve the least-squares regression problem given by

$$\min_{A, B} \sum_{k=0}^{K-1} \|\psi(x(k+1)) - A\psi(x(k)) - Bu(k)\|_2^2. \quad (7.19)$$

The regression problem can be written in matrix form as

$$[A^*, B^*] = \underset{A, B}{\operatorname{argmin}} \|Y - AX - BU\|_F^2, \quad (7.20)$$

where we have the matrices described as

$$X = [\psi(x(0)), \dots, \psi(x(K-1))], \quad (7.21)$$

$$Y = [\psi(x(1)), \dots, \psi(x(K))], \quad (7.22)$$

$$U = [u(0), \dots, u(K-1)], \quad (7.23)$$

for which the optimal solution satisfies the closed form expression

$$[A^*, B^*] = YZ^\dagger, \quad Z = [X^\top, U^\top]^\top, \quad (7.24)$$

where \dagger denotes the Moore-Penrose pseudoinverse. Alternatively, to promote simplicity and improve conditioning a Tikhonov regularization term can be added to the objective function, leading to the regularized solution

$$[A^*, B^*] = YZ^\top (ZZ^\top + \lambda I)^{-1}. \quad (7.25)$$

The steps to obtain A^* and B^* are outlined in Algorithm 3.




Figure 7.1: Physical selection of Koopman observables for thermal networks. Flow rates (left) and temperatures (right) are combined into bilinear observables representing thermal energy transport. This physics-informed lifting preserves the underlying conservation principles in the transformed coordinate system.

7.4. Observable Selection for District Heating Networks: A Physics-Guided Approach

THE construction of an effective lifted coordinate system for DHNs requires careful selection of observables that capture both the fundamental linear dynamics and the essential nonlinear thermal-hydraulic interactions. This section presents a systematic

Algorithm 3 Learning Koopman Operator Model**Input:** Dataset $\mathcal{D} = (x(k), u(k), x(k+1))_{k=1}^K$ **Output:** Koopman matrices A, B define observable vector $\psi(x)$ **for** $k = 1$ to K **do** \diamond $X_k \leftarrow \psi(x(k))$ \diamond $Y_k \leftarrow \psi(x(k+1))$ \diamond $U_k \leftarrow u(k)$ **end for** $Z \leftarrow [X^\top \ U^\top]^\top$ $\lambda_{\text{opt}} \leftarrow$ find through cross-validation with (Z, Y) $[A^*, B^*] \leftarrow$ solve (7.25) using $(Z, Y, \lambda_{\text{opt}})$

approach to observable selection guided by the underlying physics of heat and mass transport in DHNs.

As illustrated in Figure 7.1, we construct the observable vector $\psi(x)$ to include the following terms:

1. **Linear state components.** We include all original state variables from the physical system as observables, i.e., we have

$$\psi_x(x) = x, \quad (7.26)$$

which is commonly used in Koopman predictive control in order to retrieve the original state through $x = Cz$, where $C = [I_n \ \mathbf{0}]$ and avoid any lifting of the objective function;

2. **Bilinear advection terms.** Rather than including all $|\mathcal{N}| \times |\mathcal{S}|$ possible bilinear combinations $q_\ell \tilde{T}_i$, we exploit the network's topological structure to select only physically meaningful couplings. We define the outgoing edge matrix as

$$\mathbf{E}_F^{\text{out}} = \frac{1}{2} (|\mathbf{E}|_o - \mathbf{E}) \mathbf{F}^\top, \quad (7.27)$$

where $\frac{1}{2} (|\mathbf{E}|_o - \mathbf{E})$ captures outgoing edges. The matrix $\mathbf{E}_F^{\text{out}}$ relates fundamental loops to nodes, with elements defined by

$$\mathbf{E}_{F,i\ell}^{\text{out}} = \begin{cases} 1, & \text{if } i \in \mathcal{N}_{\mathcal{S},\ell}, \\ 0, & \text{otherwise,} \end{cases} \quad (7.28)$$

where set $\mathcal{N}_{\mathcal{S},\ell}$ is defined as

$$\mathcal{N}_{\mathcal{S},\ell} := \{i \in \mathcal{N} \mid j \in \mathcal{N}, (i, j) \in \mathcal{S}_\ell\}. \quad (7.29)$$

We include bilinear coupling terms $q_\ell \tilde{T}_i$ only when $\mathbf{E}_{F,i\ell}^{\text{out}} = 1$, i.e., we have the lifting function $\psi_{\text{th}}(x) \in \mathbb{R}^{\mathcal{N}_{\mathcal{S},\ell}}$ with elements

$$\psi_{\text{th},i\ell}(x) = q_\ell \tilde{T}_i, \quad i \in \mathcal{N}_{\mathcal{S},\ell}, \mathcal{S}_\ell \in \mathcal{S}, \quad (7.30)$$

ensuring that our observables reflect actual physical connections related to locations in the network where flow q_ℓ influences temperature \tilde{T}_i . The introduced topology-based selection significantly reduces the number of bilinear terms while preserving the essential convective coupling effects that govern DHN dynamics.

The developed observable selection strategy offers key advantages by aligning selection with network topology and transport physics, ensuring the lifted system respects underlying physical constraints. One should note that, leveraging *network sparsity* significantly reduces bilinear terms, improving computational tractability for large-scale DHNs while preserving accuracy. The resulting observables possess clear physical meaning, facilitating analysis and control design with intuitive connections to thermal-hydraulic processes.

Remark 18 During the training phase, the demand u^c is treated as an input variable, resulting in a Koopman model that is linear with respect to demand. While this linearity is not strictly necessary since demand acts as a disturbance rather than a control input within the MPC framework and could therefore be incorporated as a nonlinear combination in the lifted state vector, we maintain this linear relationship to preserve the representational integrity of demand as an input. \diamond

7.5. Koopman Predictive Control for District Heating Networks

THIS section presents the Koopman predictive control framework. We begin by developing the standard linear KPC approach, which forms the foundation of our approach. Subsequently, we explore a promising extension to mixed-integer formulations, which become computationally tractable through the proposed Koopman reformulation.

7.5.1. Koopman Predictive Control for Unidirectional DHNs

Given the lifted linear model (7.13), we formulate an MPC problem that optimizes the DHN operation over a finite horizon N , e.g., several hours, discretized into N steps of length Δt . Let $z(i)$ represent the lifted state at time step i , with $z(0)$ obtained by lifting the physical state $x(k)$ obtained through measurement or estimation. We define the linear MPC problem with quadratic cost as

$$\begin{aligned}
 & \min_{\substack{z(i), u(i), \sigma(i) \\ i=0, \dots, N-1}} \sum_{i=0}^{N-1} \ell(Cz(i), u(i), \sigma(i)) + \ell_f(Cz(N)) \\
 & \text{s.t. } z(i+1) = Az(i) + Bu(i), \\
 & \quad E_z z(i) + F_u u(i) \leq G(i) + H_\sigma \sigma(i), \\
 & \quad Cz(i) \in \mathcal{X}, u(i) \in \mathcal{U}, \sigma(i) \geq 0, \\
 & \quad i = 0, \dots, N-1, \\
 & \quad z(0) = \psi(x(k)),
 \end{aligned} \tag{7.31}$$

where $\ell(\cdot), \ell_f(\cdot), \mathcal{X}$, and \mathcal{U} are the same as in the original nonlinear formulation of equation (1.4), $\sigma(0), \dots, \sigma(N-1)$ are slack variables, introduced to softly enforce operational constraints ensuring feasibility of equation (7.31), and H_σ is a suitable selector matrix. Typically, $\ell(\cdot)$ is a stage cost reflecting control objectives. As is standard in MPC applications, we define the stage cost as a quadratic function over the states and inputs, i.e.,

$$\begin{aligned} \ell(x(k), u(k), \sigma(k)) = & \|x(k) - x_{\text{ref}}(k)\|_Q^2 \\ & + \|u(k)\|_R^2 + \|\sigma(k)\|_{R_\sigma}^2, \end{aligned} \quad (7.32)$$

where $x_{\text{ref}}(k)$ denotes a time-varying reference and R_σ is the matrix penalty term considered for the slack variables. While we employ a quadratic stage cost, the framework readily extends to any convex cost function. Moreover, to aid the convergence of the underlying solver we augment the stage cost in the KPC with a small regularization on the lifted state vector, i.e., we add $\alpha \|z(k)\|^2$ to the stage cost, where α is a small positive constant.

The KPC implementation is summarized in Algorithm 4.

Remark 19 The slack variables are not merely an implementation detail, but are practically necessary in DHN operation. In particular, heat demands act as exogenous disturbances, typically available only through forecasts, and may exhibit abrupt variations or prediction errors; together with actuator limits, this can render the hard operational temperature bounds infeasible. Furthermore, introducing slack variables allows practitioners to explicitly quantify and penalize such violations, and to trade a controlled level of constraint violation against improved efficiency or reduced costs when needed.

Remark 20 Since the lifted dynamics and the projection from z to x are linear, affine constraints in (x, u) become affine in (z, u, σ) via $x = Cz$, and convex quadratic stage costs are preserved with $Q \geq 0$, $R \geq 0$, and $R_\sigma \geq 0$. Consequently, at each MPC step the resulting convex QP admits a global optimum. This optimality is with respect to the surrogate model. The model-plant mismatch can be mitigated through the receding horizon implementation, which re-optimizes at each step using updated state measurements.

7.5.2. Extension to Bidirectional DHNs

Modern DHNs frequently incorporate loops, thermal energy storage units and multiple generation points, meaning certain pipes within the network need to operate bidirectionally. The mathematical formulation for bidirectional flow in DHNs presents considerable challenges. First, the thermal dynamics, introduced in Section 2.3, have inherent directionality and are only valid for nonnegative flows. Second, when flows approach zero, the mathematical problem becomes ill-conditioned, leading to numerical issues during optimization.

We associate flow rates with edges in a directed graph, which allows us to model bidirectional flows by introducing opposing directed edges for each physical pipe segment. Mathematically, for any $k \in \mathbb{N}$, we can formulate this as

$$q_{e_+}(k) q_{e_-}(k) = 0, \quad (7.33)$$

Algorithm 4 Koopman Predictive Control Implementation

Input: Koopman model (A, B) , Weight matrices Q, R and R_σ , time step Δt , reference x_{ref} , and initial condition $x(0)$

Set $k = 0$ and $t = 0$

for KPC step k **do**

- ◊ Update the initial state $x(k)$
- ◊ Update demand prediction $d(k+j)$, $j = 0, \dots, N-1$
- ◊ $u^c(j) \leftarrow d(k+j)$, for $j = 0, \dots, N-1$
- ◊ $u^*(k) \leftarrow \text{solve (7.31)}$
- ◊ Set the value of control input to $u^*(k)$ for time period $[k\Delta t, (k+1)\Delta t]$
- ◊ $k \leftarrow k+1$

end for

where $e_+ = (u, v) \in \mathcal{E}$, $e_- = (v, u) \in \mathcal{E}$, and $q_{e_+}, q_{e_-} \in \mathbb{R}_+$. Thus, a logical equality constraint, as in (7.33), renders the NLP a mathematical program with equilibrium constraints (MPEC), for which the feasible region is characterized by sharp corners and discontinuous gradients, increasing the likelihood of a solver ending up in a suboptimal local minimum. For any $k \in \mathbb{N}$, one can reformulate the complementarity constraint as follows:

$$\begin{aligned} q_{e_+}(k) &\leq \delta_e(k) \bar{q}_e(k), \\ q_{e_-}(k) &\leq (1 - \delta_e(k)) \bar{q}_e(k), \\ \delta_e(k) &\in \{0, 1\}, \end{aligned} \tag{7.34}$$

transforming the original problem into a mixed-integer nonlinear program that is computationally intractable for systems of the dimensions considered in this study. In contrast, the Koopman formulation allows, to a certain extent, for the inclusion of combinatorial complexity while maintaining computational tractability. To this end, the inclusion of the constraint (7.34) leads to the following adjustment to the second constraint of the MPC formulation (7.31):

$$D_\delta \delta(i) + E_z z(i) + F_u u(i) \leq G, \tag{7.35}$$

where $\delta(i) \in \mathbb{B}$, for any $i \in \mathbb{N}$, rendering the problem a MIQP. In Section 7.6.5, we demonstrate both the tractability and the potential for improved performance of the mixed-integer Koopman formulation in comparison to the standard NLP formulation with complementarity constraints.

7.6. Application of KPC to a DHN Benchmark: Numerical Experiments and Validations

IN this section, we demonstrate the effectiveness of our Koopman predictive control approach through several benchmark numerical experiments. To this end, we first assess the learning and prediction efficacy of the Koopman model in open-loop, i.e., we generate a representative set of inputs and initial states, learn the Koopman operator, and evaluate the performance of resulting Koopman-based lifted model in

predicting the true trajectory. Subsequently we assess the closed-loop and computational performance of the proposed KPC algorithm through numerical experiments examining two distinct problem formulations. These experiments are designed to demonstrate both the computational efficiency achieved through Koopman lifting for district heating control problems with unidirectional flow variables and, in addition, to validate the effectiveness of KPC extended formulation including bidirectional flows.

We evaluate the control methods on the AROMA DHN model [18, 34, 41], shown in Figure 7.2. This version of the AROMA network consists of one heat producer and five consumers arranged in a branched topology. We use a model of the form described in (2.23) consisting of two discretized temperature states for each edge representing a pipeline, leading to a total of 44 states, 7 flow states and 37 temperature states. The volume of each temperature state is set to 250 m^3 , i.e. $V = 250I_{|\mathcal{N}|}$. The vector of observables for KPC, constructed based on the physics-guided approach introduced in Section 7.4, is of dimension 130.

Remark 21 While pipe diameter and length are not explicitly modeled, the segment volume of 250 m^3 corresponds to realistic values for large-scale DHNs. A typical diameter of 500 mm implies a segment length of approximately 1250 m, resulting in a total pipe length of around 46 km across 37 segments.




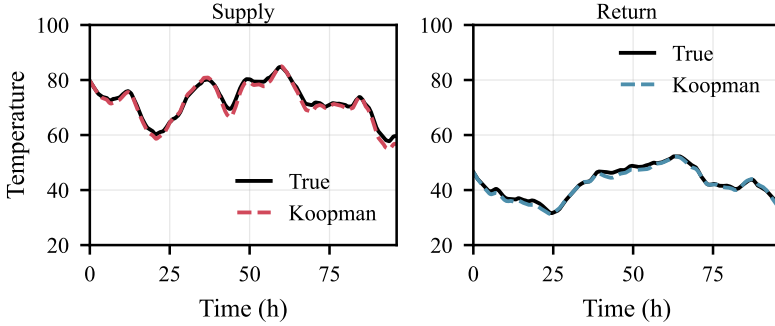
Figure 7.2: The AROMA network considered in simulation. The dashed line indicates the bidirectional edge in the extended example.

7.6.1. Closed-Loop Experiments

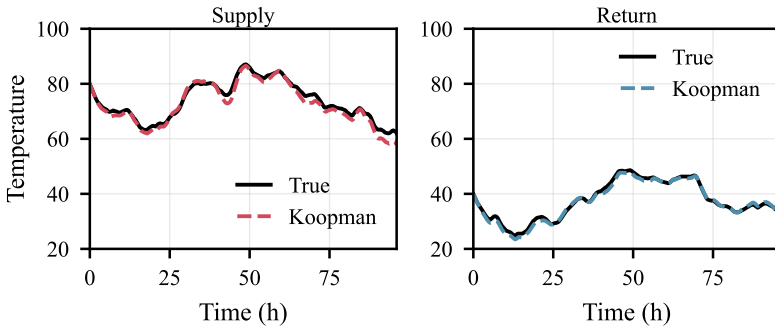
The **unidirectional DHN formulation** considers continuous operational variables within the framework of NLPs. Traditional nonlinear MPC directly handles the nonlinear dynamics, resulting in an NLP, while KPC transforms the problem into a QP through lifting to higher-dimensional linear dynamics. Thus, the resulting reformulation enables direct comparison of control performance while demonstrating the computational advantages of avoiding nonlinear solvers. Moreover, the resulting QP formulation allows one to utilize efficient optimization algorithms and readily available high-performance solvers, such as *Gurobi* [106], *Ipopt* [74], and *OSQP* [107].

The **bidirectional DHN formulation** introduces discrete decision variables representing operational modes related to reversible flows in the network, namely as in (7.33). In the nonlinear MPC, the complementarity structure introduces degenerate

feasible regions making high-quality solutions difficult to obtain. The Koopman lifting naturally accommodates these constraints through binary variables, yielding a mixed-integer quadratic program (MIQP) with established solution methods and optimality properties. For instance, one can use specialized high-performance solvers such as *miOSQP* [108].



(a) Koopman temperature predictions: model-free.



(b) Koopman temperature predictions: model-based.

Figure 7.3: Temperature trajectory comparison for the AROMA model for different operating conditions. Representative node temperatures shown for supply (near the first consumer) and return (near the producer) over a 96-hour simulation period. The solid lines (black) represent the true system response, while the dashed lines (orange) show the Koopman operator model predictions.

7.6.2. Koopman Predictive Control: Data Acquisition, Learning, and Implementation Details

The process of obtaining data, learning a linear model, and implementing KPC is described below.

Training data: Data for learning the Koopman operator can be acquired through two

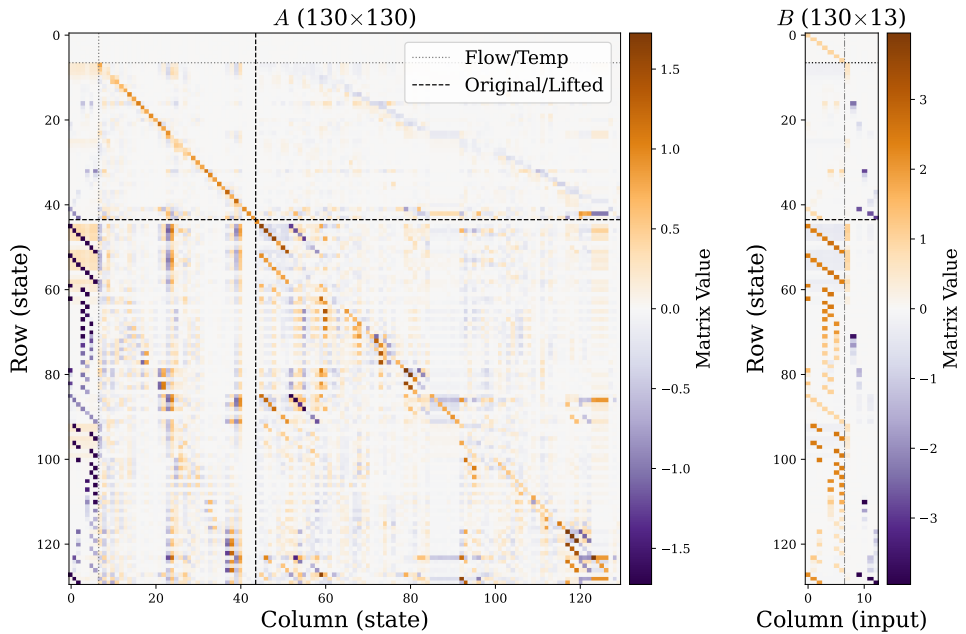


Figure 7.4: Learned Koopman state-space matrices $A (130 \times 130)$ and $B (130 \times 13)$. Dashed lines separate flow and temperature states, dotted lines separate original physical states from lifted Koopman observables. In matrix B , the dash-dotted line distinguishes flows from producer and consumer inputs.

main methods, namely collected from normal operation of the system, or synthetically generated through given high-fidelity models. In the former case, real-world data is collected via carefully designed experiments that balance the need for sufficient system excitation with practical constraints like operational limits and thermal comfort. To this end, we can slightly modify the inputs generated by traditional control during normal operation, using a hybrid approach of rule-based control policies and tracking control. Thus, we have an entirely model-free data acquisition method, requiring no prior system knowledge, and ensuring data is collected without violating critical constraints.

Alternatively, when a high-fidelity model of the DHN, e.g., in the form of (7.6), is available, synthetic data can be generated through extensive numerical simulations, allowing for comprehensive exploration of the operational envelope without the risks of real-world testing and safety-critical experiments. Thus, we can create extensive set of data, exploring various scenarios, even potentially unsafe ones, including seasonal and disturbance variations, providing a rich and informative dataset that captures the full spectrum of system behavior.

For the numerical experiments discussed below, the sampling time is set to 20 minutes and the latter approach is employed for the data acquisition.

Remark 22 The offline identification approach adopted in this work is motivated by the operational characteristics of DHNs, which exhibit fixed topology and slowly-varying parameters. In such settings, a model identified offline remains valid over extended periods, and periodic re-identification can accommodate seasonal variations. We note that online or adaptive extensions are readily accommodated within this framework: the computational cost of EDMD-based Koopman identification is on the order of seconds, enabling frequent model updates if operating conditions require.

Learning the Koopman operator: Given the training dataset, the observables are constructed from the original system states, a bias term, and bilinear interaction terms designed according to the physics-guided approach discussed in Section 7.4. Subsequently, we obtain the Koopman-lifted linear state and input matrices as summarized in Algorithm 3. Figure 7.4 shows the sparsity structure of the resulting A and B matrices.

Koopman Predictive Control: The performance of KPC, described in Algorithm 4, and nonlinear MPC strategies is systematically evaluated through closed-loop simulations, where the same objective functions and constraints are considered for fair comparison. Consumer demand disturbances are generated using multi-harmonic periodic patterns with randomized daily variations to simulate realistic 24-hour load fluctuations. Cost weights are

$$Q = \begin{bmatrix} Q_q & 0 \\ 0 & Q_{\bar{T}} \end{bmatrix}, \text{ and } R = \begin{bmatrix} R_q & 0 \\ 0 & R_p \end{bmatrix},$$

where we set $Q_q = 0I_{|\mathcal{S}|}$, $Q_{\bar{T}} = I_{|\mathcal{N}|}$, $R_q = 10^{-5}I_{|\mathcal{S}|}$, and $R_p = 2 \cdot 10^{-2}I_{|\mathcal{W}_p|}$. Furthermore, we have introduced slack variables, penalized by R_σ with $R_\sigma = 10^4 I_4$, added for each temperature bound; namely supply upper and lower and return upper and lower. Reference values x_{ref} are set to zero to prioritize temperature minimization over tracking.

Moreover, we set flow input bounds to $0.1 \leq u_e^q(k) \leq 0.6$ [m³/s] for all $e \in \mathcal{S}$, and we restrict the scaled producer heat input to $0 \leq u^p(k) \leq 250$. Finally, to assess the

controller's ability to deal with time-varying constraints, we impose a varying supply temperature lower bound satisfying $\tilde{T}^{\text{LB}}(k) \leq \tilde{T}_i(k)$, for all $k \in \mathbb{Z}$ and $i \in \mathcal{N}^{\text{sup}}$, where \mathcal{N}^{sup} is the set of all nodes in the supply network, with

$$\tilde{T}^{\text{LB}}(k) = \begin{cases} 75, & \text{if } 24 \leq (k \bmod 72) \leq 60, \\ 55, & \text{otherwise.} \end{cases} \quad (7.36)$$

7.6.3. Predictive Performance of the Koopman Operator

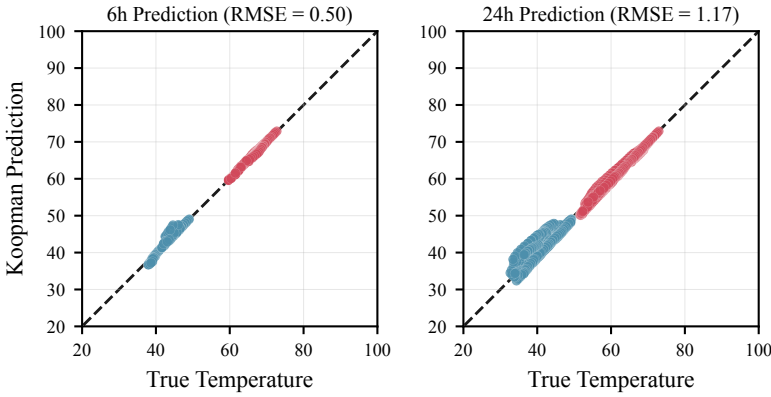
We perform open-loop simulations to assess the predictive accuracy of the Koopman model by comparing predicted state trajectories against ground truth. In Figure 7.3, we show examples of several trajectories and their associated Koopman predictions over a period of four days, i.e., 96 hours. In support of Figure 7.3, parity plots are provided in Figure 7.5 showing predicted temperature state values for all prediction steps from 1 up to the specified maximum, namely 18 and 72 steps ahead, along with the corresponding root mean square error (RMSE). These results demonstrate that the Koopman model retains its predictive capabilities over extended forecast horizons.

Discussion: Both methodologies demonstrate effective Koopman operator learning for DHN trajectory prediction, achieving accurate predictions with low RMSE values. The similarity of results between model-free and model-based methods validates the robustness of Koopman operator theory for DHN systems. The model-free approach learns dynamics directly from data without requiring physical knowledge, while the model-based approach leverages domain expertise offering more excitory data for learning. These findings indicate that practitioners can confidently select either methodology based on their specific constraints such as data availability, computational resources, or interpretability requirements. Both approaches provide reliable foundations for DHN trajectory prediction, with methodology choice depending on application-specific factors rather than fundamental performance limitations.


Comparison with other observable selection methods

Alternative approaches for reducing the lifting dimension of the Koopman dictionary include pruning observables based on subspace metrics [109, 110]. In particular, these approaches require the user to determine a hyperparameter $\varepsilon \in [0, 1]$, which trades off *accuracy* ($\varepsilon = 0$) with *expressiveness* ($\varepsilon = 1$) of a Koopman dictionary. The maximally expressive case corresponds to using the full dictionary, while the most accurate case achieves maximal Koopman-invariance within the given function space.

Specifically, the accuracy measure quantifies the proximity between the range spaces of lifted data vectors, where $\varepsilon = 0$ corresponds to a dictionary for which these range spaces are identical. This condition ensures the existence of an exact finite-dimensional linear representation of the Koopman operator within the chosen function space. However, for many nonlinear systems, such maximally invariant subspaces either do not exist within practical dictionaries or may not include all state variables of interest. In these cases, including the case of the DHN dynamics, ε serves as a systematic method for pruning the observable space, with larger values of ε allowing for larger, less invariant dictionaries that can potentially represent more aspects of the dynamics.



(a) Parity plot for Koopman predictions: model-free.



(b) Parity plot for Koopman predictions: model-based.

Figure 7.5: Multi-step prediction validation for the AROMA DHN model. Parity plots aggregate prediction accuracy across increasing forecast horizons, where each subplot includes all prediction steps from 1 up to the specified maximum: 18 (6h) and 72 (24h) steps ahead. Each red (blue) point represents a single supply (return) temperature state prediction at any step within the horizon across all 37 temperature states, compared against true values (dashed line).

Table 7.1: Performance Comparison of Dictionary Learning Approaches

Metric	Full Dictionary	Pruned $\varepsilon = 0.99$	Pruned $\varepsilon = 0.9$	Physics-guided Dictionary
Dictionary Size	184	155	113	97
MSE	0.3537	0.5981	0.7715	0.3556

Table 7.1 compares the predictive performance of Koopman models constructed using pruning-based and physics-guided observable selection. For district heating networks with known governing equations, the physics-guided approach yields superior one-step-ahead prediction accuracy (measured by MSE) while requiring substantially

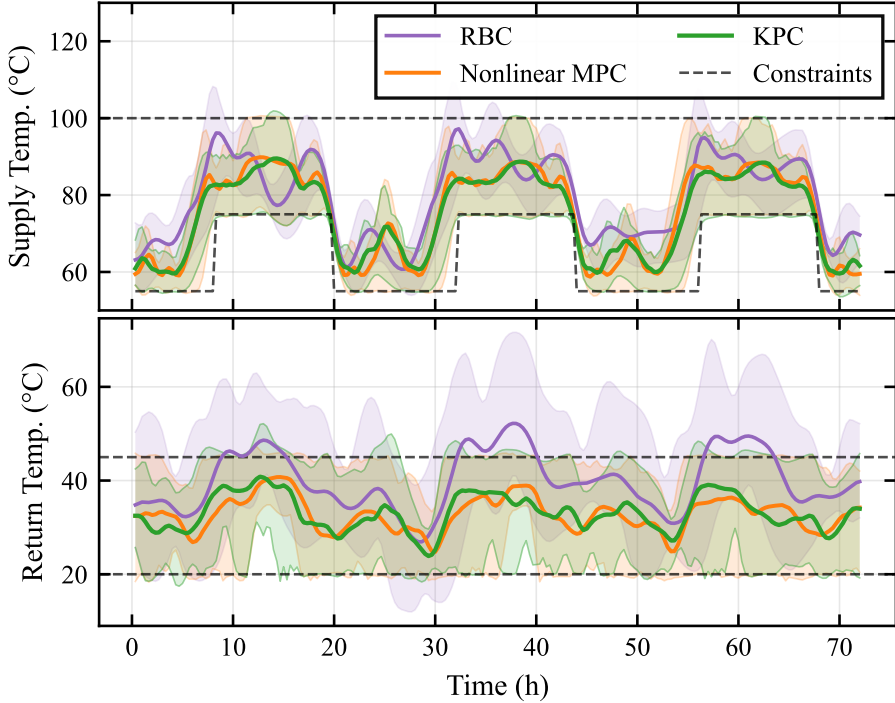


Figure 7.6: Closed-loop temperature trajectories for DHN control. (Top) Supply temperature envelope showing mean trajectory (solid line) with minimum and maximum bounds across all network nodes (shaded region) for each MPC. (Bottom) Return temperature envelope with the same visualization approach. Black dashed lines indicate operational constraints: time-varying lower bound for supply temperature (alternating between 55°C and 75°C in a 72-step period) with fixed upper bound at 100°C; fixed bounds at 20°C and 45°C for return temperature.

fewer observables. The reduction in dictionary dimension directly translates to lower computational burden in the resulting MPC formulation, confirming that exploiting physical structure in basis construction offers concrete advantages over purely data-driven dictionary pruning methods.

Remark 23 While the comparison of the use of these models in closed-loop remains an important area for future research, it falls beyond the scope of the present work. \diamond

7.6.4. Predictive Control for Unidirectional DHNs

In Table 7.2, we provide the statistics of the numerical experiments performed based on different control methods, i.e., rule-based control (RBC), nonlinear MPC, and KPC, where the latter two are evaluated for different prediction horizons ranging from 3 to 24 hours. The table contains information on the computational properties of the solvers as well as performance-based statistics such as average constraint violation and average cost

Table 7.2: Performance of RBC, Nonlinear MPC, and KPC Methods Across Different Prediction Horizons

Metric	RBC	NMPC				KPC			
		9 (3h)	18 (6h)	36 (12h)	72 (24h)	9 (3h)	18 (6h)	36 (12h)	72 (24h)
Avg. solve time (s)	–	1.79	5.24	15.19	37.73	0.38	0.61	1.32	2.89
Median solve time (s)	–	1.71	5.10	14.87	37.60	0.37	0.60	1.25	2.63
Min solve time (s)	–	0.96	2.80	7.19	19.83	0.24	0.48	0.99	1.91
Max solve time (s)	–	3.32	10.04	26.30	52.00	0.82	1.03	2.52	5.20
Avg. violation (°C)	0.42	0.85	0.16	0.12	0.11	1.13	0.28	0.23	0.22
Avg. closed-loop cost ($\times 10^3$)	2.14	1.59	1.88	1.96	1.96	1.38	1.70	1.77	1.77

Table 7.3: Performance of Nonlinear MPC and KPC Methods Across Different System Dimensions for $N = 72$

	NMPC				KPC			
	1x	2x	3x	4x	1x	2x	3x	4x
Avg. solve time (s)	14.74	96.99	331.45	922.65	2.55	5.65	10.15	16.27
Median solve time (s)	14.42	90.72	313.35	804.22	2.46	5.56	10.09	16.25
Min solve time (s)	10.26	58.56	209.77	503.65	2.18	5.17	9.33	15.19
Max solve time (s)	23.43	219.02	869.13	3052.74	3.86	6.54	11.60	17.79
State dimension	44	89	134	189	130	175	220	265

information. In Figure 7.8, we depict, for each of the control methods and a prediction horizon of $N = 18$, the associated Pareto front showing the trade-off between closed-loop cost and average constraint violation.

To further assess the scalability of the Koopman approach, we conduct simulations with increasing network complexity. Specifically, we augment the original AROMA network by inserting $n = 0, 1, 2, 3$ additional nodes per edge, creating four model variants denoted by ‘ $(n + 1)x$ ’. For each variant, we collect new training data following the same procedure as before and train the corresponding Koopman operators. Notably, while the added states contribute to the linear component of the observable vector, we do not introduce additional bilinear terms to the dictionary. This is for two reasons: larger dictionaries yield no improvement or even degrade predictive performance, often exhibiting unstable behavior for very large dictionaries; and bilinear terms substantially increase dictionary size (up to 7 additional terms per node) and thus problem complexity. The results are presented in Table 7.3.

Discussion: Table 7.2 reveals substantial computational advantages achieved by the proposed approach, which is further illustrated in Figure 7.7. Most notably, the KPC method delivers an impressive average speedup ratio exceeding one order of magnitude, i.e., 10 times, compared to nonlinear MPC. This considerable computational improvement is particularly relevant, as such performance gains directly support greater scalability and real-time implementability. The cost and violation statistics reveal that while both MPC approaches clearly outperform the RBC baseline, KPC tends to achieve lower state costs, i.e., reduced temperature deviations, at the expense of increased constraint violations. More precisely, Figure 7.8 illustrates this trade-off behavior, showing

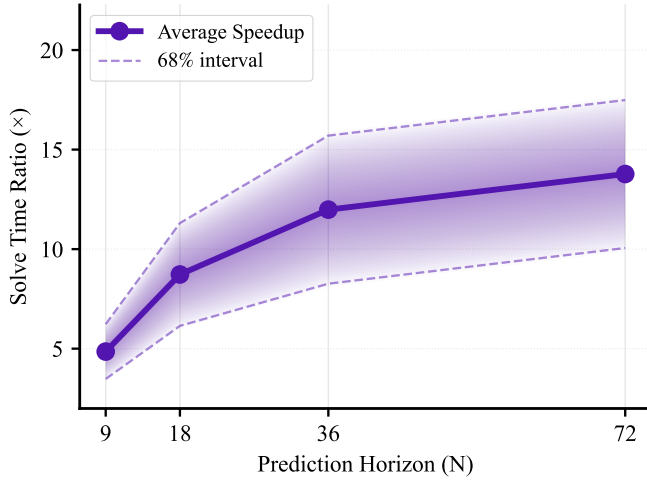


Figure 7.7: Relative solve time improvement of KPC versus nonlinear MPC as a function of prediction horizon length. The plot shows the percentage reduction in computation time achieved by KPC compared to nonlinear MPC (solid line: average speedup, shaded region: 68% confidence interval) across different prediction horizons ($N = 9$ to 72). Results are based on 216 iterations from the same datasets used for Table 7.2.

that the KPC solutions are positioned toward slightly higher temperature violation values while remaining below the Pareto front established by the nonlinear MPC.

7.6.5. Predictive Control for Bidirectional DHNs

In this numerical example, we include the bidirectional flow constraints from (7.34) and compare the resulting optimal trajectories of KPC and nonlinear MPC. Both models are assessed with a prediction horizon of $N = 36$. To ensure the comparison reflects the underlying model performance rather than a mix of performance and violation objectives, we primarily focus on the temperature tracking scenario by increasing the temperature tracking weight to $Q_{\bar{T}} = 10I_{|\mathcal{N}|}$, relaxing the temperature constraints by reducing the slack weight to $R_{\sigma} = 10I_4$, and using nonzero reference values; 75°C for supply nodes and 35°C for return nodes.

One limitation of the mixed-integer reformulation is the inherent exponential complexity growth with the prediction horizon. To mitigate this computational burden, we implement two complementary strategies: (i) a move-blocking scheme that restricts flow direction changes to hourly intervals (every 3 time steps), and (ii) relaxation of binary variables to continuous values in the latter portion of the horizon. While this approach ensures binary logic is always properly enforced in the next iteration, it introduces a model-plant mismatch that may degrade MPC performance. We denote the number of active binary variables by $N_b \in \{1, \dots, 12\}$, where the upper bound of 12 results from partitioning the 36-step prediction horizon into move-blocking intervals of three consecutive time steps. Figure 7.9a demonstrates the computational performance and

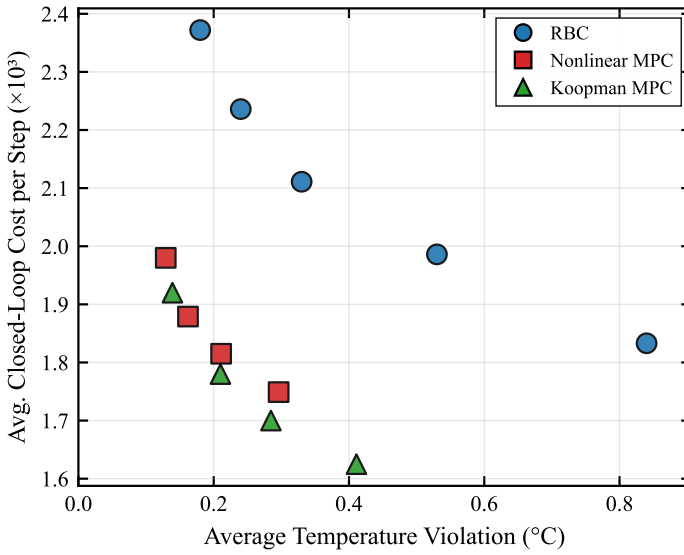
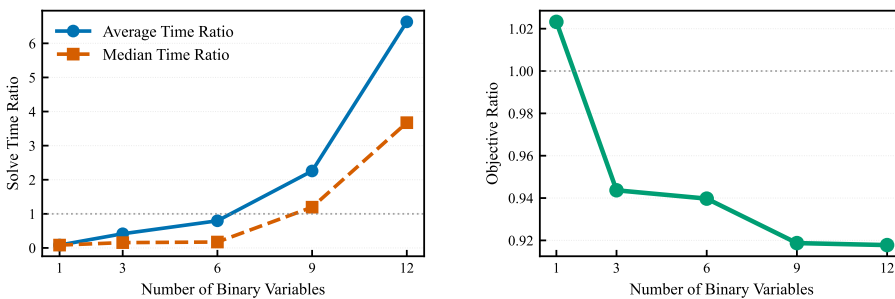


Figure 7.8: Pareto analysis comparing control performance trade-offs between operational cost and constraint violations. The plot shows average closed-loop cost per step (y-axis, excluding slack penalties) versus average temperature violation (x-axis) for KPC, nonlinear MPC, and rule-based control (RBC). MPC algorithms used prediction horizon $N = 18$ with varying slack weights (100, 200, 500, 5000) to generate different operating points. RBC points were created by varying supply and return temperature reference values. Points closer to the origin indicate better performance with lower costs and fewer violations.

Figure 7.9b shows the closed-loop cost of the KPC relative to the nonlinear MPC.



(a) Computational performance metrics for various relaxed problem formulations compared against the nonlinear MPC benchmark.

(b) Relative closed-loop cost, quantifying the performance degradation incurred by computational simplifications.

Figure 7.9: Computational-performance trade-off for relaxed bidirectional KPC formulations with binary variable dimension up to N_b .

Discussion: The proposed KPC exhibits a substantial performance improvement of up to 8% compared to the nonlinear MPC, which exceeds the gains observed in the unidirectional flow case. The superior performance of KPC may be attributed to the weakness of nonlinear MPC in avoiding convergence to spurious sub-optimal local minima, particularly as problem complexity increases. The magnitude of this improvement is noteworthy given the relatively modest system modifications implemented, i.e., only a single bidirectional pipeline.

Despite the computational challenges of MIQP formulation, the results establish a clear pathway for advancing DHN control through systematic application of linear control theory and integer programming techniques. Successfully circumventing local minima results in significant performance gains that offset the additional computational cost inherent in the Koopman predictive control approach. Furthermore, the mentioned approximation strategies, besides other potential similar techniques, demonstrate clear scalability potential for large-scale industrial systems.

7.7. Conclusion

WE proposed a physics-guided Koopman operator approach for district heating networks exceeding an order-of-magnitude computational speedup while improving cost efficiency, operational performance, and constraint satisfaction. Our framework leverages structural graph properties and physical constraints to construct observables that accurately predict temperature and flow evolution, establishing a promising foundation for scalable linear MPC implementation. Following successful validation in controlled environments, the approach demonstrates significant potential for revolutionizing large-scale thermal network management in practical applications. Future research focusing on closed-loop stability analysis, observable set optimization, multi-step predictors, and real-world implementation will determine the broader viability of this Koopman-based approach for district heating systems.

8

Conclusions and Recommendations

8.1. Conclusions

DISTRICT heating networks are expected to play an increasingly important role in the energy transition, yet realizing their potential requires a departure from conventional control strategies. As networks grow in scale and incorporate bidirectional flow, distributed generation, and renewable sources, rule-based approaches can no longer achieve the required operational performance. Predictive control offers a principled solution, but obtaining tractable formulations for large-scale networks remains a significant challenge. This dissertation focused on addressing this issue, developing scalable model predictive control methods with tractability as the unifying theme across all contributions.

To achieve scalable predictive control for DHNs, several challenges had to be overcome. The first concerns modeling: while the governing physics of DHNs are well understood, they derive from multiscale nonlinear partial differential equations that are not directly amenable to optimization. Building on earlier discretization approaches, Chapter 2 developed an adaptable modeling framework that starts from a minimal network representation and allows the practitioner to increase resolution where needed, balancing tractability against accuracy. Crucially, the framework captures bidirectional flow through complementarity constraints rather than binary variables, avoiding the combinatorial complexity of mixed-integer formulations.

The continuous nonlinear formulation resulting from this modeling approach is amenable to gradient-based optimization, albeit limited to locally optimal solutions. However, scale remains a challenge: even the minimal model representation of a practical DHN can have hundreds or thousands of internal states, making the real-time solution of the underlying optimal control problem intractable with standard methods. A central contribution of this dissertation is, therefore, the development of decomposition and approximation strategies that scale to realistic problem sizes while preserving solution quality.

One source of computational difficulty lies in the nonlinear Kirchhoff loop constraints governing pressure and flow. By assuming sufficient valve and pump infrastructure throughout the network, we showed in Chapter 4 that these constraints can be convexified, yielding a tractable feasibility region for the flow variables without requiring explicit optimization over the hydraulic actuators. This reformulation eliminates a significant source of nonconvexity from the optimization problem without compromising the quality of the solution.

A second approach, introduced in Chapter 6, addresses the temporal dimension of the problem. Long prediction horizons are essential for capturing the slow thermal dynamics of DHNs, yet they dramatically increase problem size. Using a dual dynamic programming formulation with nested Benders' decomposition, we decomposed the horizon into a detailed near-term segment and an approximated long-term segment, substantially reducing computational burden while preserving the ability to account for long-horizon effects.

The most significant computational gains came from reformulating the control problem entirely in Chapter 7. The bilinear thermal dynamics, while continuous, still pose challenges for gradient-based solvers. By constructing a tailored basis of observable functions, we lifted the dynamics into a coordinate frame where they evolve

linearly, enabling the application of well-established linear MPC techniques. Beyond the order-of-magnitude speedups compared to nonlinear MPC, this Koopman-based formulation yields a convex optimization problem, guaranteeing global optimality within the lifted representation and eliminating sensitivity to initialization.

The simulation studies conducted in this dissertation provide evidence that these methods can substantially improve both computational tractability and operational performance. Since computational complexity scales with both the system dimension and the prediction horizon, the relevant metric is the total number of decision variables across the planning window. The networks and horizons considered in this work ranged up to approximately 5000 physical decision variables (temperatures and flows across all time steps), and in the case of the Koopman formulation, up to 10000 lifted decision variables. For problems of this scale, the Koopman-based MPC achieved solve times of two to five seconds, compared to 10-50 seconds for nonlinear MPC, while the dual dynamic programming approach enabled extended planning horizons without proportional increase in computational cost. Nevertheless, we note that computational time depends on many factors beyond optimization variable dimension, including the number and type of constraints, the demand and price data, and numerical properties such as scaling and matrix conditioning.

Beyond computational performance, these methods resulted in operational cost reductions of approximately 5–10% compared to rule-based baselines, primarily through effective management of supply temperatures and exploitation of storage flexibility. Furthermore, for the bidirectional flow case, the Koopman formulation with binary flow direction variables achieved an additional 8% cost reduction compared to nonlinear MPC—a gap we attribute to the nonlinear solver converging to suboptimal local minima. Whether these gains extend to still larger networks and real-world operating conditions remains to be validated; factors such as model mismatch, measurement noise, and unforeseen disturbances will inevitably affect performance. Nevertheless, the results suggest that the fundamental computational barriers to predictive control of district heating networks can be addressed, and that further investigation toward practical deployment is warranted.

Tractable algorithms are a necessary but not sufficient condition for practical deployment. Among the further challenges that remain, this dissertation addressed two: guaranteeing stable closed-loop behavior, and obtaining the state estimates that MPC requires. While by no means exhaustive, these contributions represent steps toward a more complete framework.

Unless terminal conditions are imposed, model predictive control provides no inherent guarantee of closed-loop stability. This is particularly challenging in economic formulations, where the cost function does not directly penalize deviation from a setpoint. Drawing on dissipativity theory, we established in Chapter 5 the conditions under which the proposed controllers guarantee stable operation, providing the theoretical foundation necessary for practical adoption. Furthermore, we proposed a method based on sum of squares programming to determine whether a valid storage function exists for a given stage cost. This is a promising technique to determine strict dissipativity for the optimal control problem; however, limitations in scalability currently prevent the verification of dissipativity for large networks.

Finally, MPC requires an estimate of the system state, yet large-scale DHNs cannot be instrumented exhaustively to measure every single temperature state in the network. State estimation techniques such as the Kalman filter can reconstruct unmeasured states, but their accuracy depends critically on the placement of the available sensors. In Chapter 3, we addressed the sensor placement problem by formulating it as a combinatorial optimization problem that maximizes a metric of observability under budget constraints. We showed that this metric exhibits submodularity—a diminishing returns property—which allows efficient greedy selection algorithms to achieve near-optimal solutions with provable performance bounds. This approach enables accurate state reconstruction from a limited number of measurements.

Collectively, these contributions move scalable model predictive control for district heating networks closer to practical viability. While significant challenges remain before widespread deployment, the framework developed here establishes that the computational barriers to optimizing large-scale thermal networks can be overcome.

8.2. Recommendations

BASED on the findings of this dissertation, we offer the following recommendations for future research:

a — Further approaches to structural decomposition of DHN optimal control

Chapter 4 introduced structure-exploiting techniques that improve the tractability of optimal control problems. The specific structure of the nonlinear formulation reveals promising avenues for further research. When employing quadratic stage costs, the standard NLP formulation from Chapter 4, which consists of linear, bilinear, and quadratic constraints, represents a quadratically constrained quadratic program (QCQP). Since any QCQP can be reformulated as a semidefinite program, this connects our framework to extensive optimization literature that provides sophisticated relaxation techniques with rigorous tightness bounds [111–114]. This connection suggests that semidefinite relaxations could enable the solution of larger-scale district heating control problems within computational time constraints.

b — Non-uniform sampling MPC

Following the approach presented in [115]¹, a non-uniform sampling strategy offers an elegant solution to the bilinearity challenge. By adapting MPC time steps based on flow velocity, this method ensures that the traveled distance in each step equals exactly one spatial discretization unit. This transforms the optimization fundamentally: flows become functions computed a posteriori rather than decision variables, thereby fully linearizing the MPC formulation. While this approach naturally accommodates single-loop DHN topologies, extending it to branched networks remains challenging due to the heterogeneous flow velocities across different branches, which would require synchronized yet non-uniform time stepping across the network.

¹Inspired by discussions with the main author, Prof. J.M. Lemos.

c — Actuator scheduling using frequency-based Gramians

The sensor placement methodology from Chapter 3 extends naturally to thermal generator scheduling. Consider a district heating network with numerous prosumers, each equipped with heat pumps capable of injecting heat into the network. The scheduling problem seeks to determine which subset must operate at each time to meet demands while minimizing the number of active units. The frequential Gramian approach applies directly: controllability metrics in the frequency domain quantify each generator's contribution to satisfying the periodic demand patterns inherent to district heating systems. By decomposing demand into spectral components, the framework identifies minimal generator configurations that ensure sufficient controllability across dominant frequencies, naturally accommodating the cyclic structure of thermal loads. Such analysis becomes increasingly relevant as decentralized generation expands in modern thermal networks.

d — Learning terminal cost functions using input convex neural networks

In Chapter 6, we proposed an approach to split the optimization problem into two parts, linearize the tail end, and solve the decomposed problem using dual dynamic programming. While this approach saves significant computational time, it still requires online iterations to solve the two subproblems, with no guarantee regarding the number of iterations needed. Therefore, we suggest exploring a fully offline approach, where one learns the tail end of the optimization problem by generating many samples of initial states (terminal states for the first problem) and their corresponding optimal cost values, and subsequently training an input convex neural network on that data, creating a convex, differentiable mapping from terminal states to cost values [116–118]. This function can then be embedded as a terminal cost in the MPC formulation, allowing the use of substantially shorter prediction horizons and making the problem more tractable.

e — Extending Koopman predictive control for DHNs

Several directions for future work emerge from the framework developed in Chapter 7. Firstly, the physics-guided dictionary construction yields accurate models, though the choice of observables was not exhaustive. Systematic basis selection methods could improve the trade-off between lifted dimension and approximation quality. Secondly, the linear structure admits distributed MPC implementations that scale with network size. One can either learn subsystem dynamics independently and coordinate them afterward [119], or partition the full linear model directly to exploit network topology. Both approaches enable parallel computation while managing coupling through coordination mechanisms. Thirdly, a particularly promising extension combines the Koopman framework with the dual dynamic programming method from Chapter 6. That approach requires linearization of the second-stage problem, but finding a suitable linearization is nontrivial when the system operates far from equilibrium across diverse conditions. The Koopman embedding circumvents this issue by providing a globally valid linear representation, eliminating the need for successive local approximations. Combined with horizon decomposition, this controller synthesis method directly

addresses scalability limitations for large-scale networks while preserving model fidelity across the operating envelope.

Bibliography

- [1] Heat Roadmap Europe. *Heating and Cooling: Facts and Figures – The Transformation Towards a Low-carbon Heating & Cooling Sector*. 2017. URL: www.heatroadmap.eu.
- [2] L. Pelkmans, B. Kulišić, M. Georgiadou and M. Buffi. *Implementation of bioenergy in the European Union – 2024 update*. Country Report. IEA Bioenergy, 2024.
- [3] ISO New England. *Final 2024 Heating Electrification Forecast*. 2024. URL: www.iso-ne.com/static-assets/documents/100010/final-2024-heating-electrification-forecast.pdf.
- [4] A. Battaglini, N. Komendantova, P. Brtnik and A. Patt. ‘Perception of barriers for expansion of electricity grids in the European Union’. In: *Energy Policy* 47 (2012), pp. 254–259.
- [5] S. Watson. ‘Predicting the additional GB electricity demand resulting from a widespread uptake of domestic heat pumps’. PhD thesis. Loughborough University, 2019.
- [6] S. Frederiksen and S. Werner. *District Heating and Cooling*. Lund: Studentlitteratur, 2013. ISBN: 978-91-44-08530-2.
- [7] A. Vandermeulen, B. van der Heijde and L. Helsen. ‘Controlling district heating and cooling networks to unlock flexibility: A review’. In: *Energy* 151 (2018), pp. 103–115.
- [8] H. Lund, S. Werner, R. Wiltshire, S. Svendsen, J. E. Thorsen, F. Hvelplund and B. V. Mathiesen. ‘4th Generation District Heating (4GDH): Integrating smart thermal grids into future sustainable energy systems’. In: *Energy* 68 (2014), pp. 1–11.
- [9] N. N. Novitsky, Z. I. Shalaginova, A. A. Alekseev, V. V. Tokarev, O. A. Grebneva, A. V. Lutsenko, E. A. Mikhailovsky, O. V. Vanteeva, R. Pop, P. Vorobev and M. Chertkov. ‘Smarter smart district heating’. In: *Proceedings of the IEEE* 108.9 (2020), pp. 1596–1611.
- [10] S. Buffa, M. Cozzini, M. D’Antoni, M. Baratieri and R. Fedrizzi. ‘5th generation district heating and cooling systems: A review of existing cases in Europe’. In: *Renewable and Sustainable Energy Reviews* 104 (2019), pp. 504–522.
- [11] S. Boesten, W. Ivens, S. C. Dekker and H. Eijndems. ‘5th generation district heating and cooling systems as a solution for renewable urban thermal energy supply’. In: *Advances in Geosciences* 49 (2019), pp. 129–136.

- [12] J. Jansen, A. Vandermeulen, D. Patteeuw and L. Helsen. 'Optimal control of a fourth generation district heating network using an integrated non-linear model predictive controller'. In: *Applied Thermal Engineering* 223 (2023), p. 120030.
- [13] F. Verrilli, S. Srinivasan, G. Gambino, M. Canelli, M. Himanka, C. Del Vecchio, M. Sasso and L. Glielmo. 'Model Predictive Control-Based Optimal Operations of District Heating System With Thermal Energy Storage and Flexible Loads'. In: *IEEE Transactions on Automation Science and Engineering* 14.2 (2017), pp. 547–557.
- [14] M. Schwenzer, M. Ay, T. Bergs and D. Abel. 'Review on model predictive control: An engineering perspective'. In: *The International Journal of Advanced Manufacturing Technology* 117.5 (2021), pp. 1327–1349.
- [15] D. Quaggiotto, J. Vivian and A. Zarrella. 'Management of a district heating network using model predictive control with and without thermal storage'. In: *Optimization and Engineering* 22.3 (2021), pp. 1897–1919.
- [16] B. Talebi, P. A. Mirzaei, A. Bastani and F. Haghghat. 'A Review of District Heating Systems: Modeling and Optimization'. In: *Frontiers in Built Environment* 2 (2016), p. 22.
- [17] S.-A. Hauschild, N. Marheineke, V. Mehrmann, J. Mohring, A. M. Badlyan, M. Rein and M. Schmidt. 'Port-Hamiltonian modeling of district heating networks'. In: *Progress in Differential-Algebraic Equations II*. Springer International Publishing, 2020, pp. 333–355.
- [18] R. Krug, V. Mehrmann and M. Schmidt. 'Nonlinear optimization of district heating networks'. In: *Optimization and Engineering* 22 (2021), pp. 783–819.
- [19] S. Kuntarova, T. Lickleder, T. Huynh, D. Zinsmeister, T. Hamacher and V. Perić. 'Design and simulation of district heating networks: A review of modeling approaches and tools'. In: *Energy* 305 (2024), p. 132189.
- [20] A. Benonysson, B. Bøhm and H. F. Ravn. 'Operational optimization in a district heating system'. In: *Energy Conversion and Management* 36.5 (1995), pp. 297–314.
- [21] P. Kundur *et al.* 'Power system stability'. In: *Power System Stability and Control* 10.1 (2007).
- [22] J. Simonsson, K. T. Atta and W. Birk. 'A Graph Theoretical Approach to Modeling of District Energy Networks'. In: *IEEE Transactions on Control Systems Technology* 32.5 (2024), pp. 1616–1627.
- [23] H. Dänschel, V. Mehrmann, M. Roland and M. Schmidt. 'Adaptive nonlinear optimization of district heating networks based on model and discretization catalogs'. In: *SeMA Journal* (2023), pp. 1–32.
- [24] J. Bjørnskov, L. K. Mortensen, K. Filonenko, H. R. Shaker, M. Jradi and C. Veje. 'Optimization of district heating production with thermal storage using mixed-integer nonlinear programming with a new initialization approach'. In: *Energy Informatics* 4.Suppl 2 (2021), p. 34.
- [25] J. Jansen, F. Jorissen and L. Helsen. 'Mixed-integer non-linear model predictive control of district heating networks'. In: *Applied Energy* 361 (2024), p. 122874.

-
- [26] S. Zhang, W. Gu, H. Qiu, S. Yao, G. Pan and X. Chen. 'State estimation models of district heating networks for integrated energy system considering incomplete measurements'. In: *Applied Energy* 282 (2021), p. 116105.
- [27] J. Mankad, B. Natarajan and B. Srinivasan. 'Integrated approach for optimal sensor placement and state estimation: A case study on water distribution networks'. In: *ISA Transactions* 123 (2022), pp. 272–285.
- [28] L. Reichle. 'Optimization of District Heating Networks: Differential Algebraic Equations, Hyperbolic Systems and Receding Horizon Control'. In: (2025).
- [29] J. E. Machado, M. Cucuzzella and J. M. A. Scherpen. 'Modeling and passivity properties of multi-producer district heating systems'. In: *Automatica* 142 (2022), p. 110397.
- [30] L. T. Biegler. 'A perspective on nonlinear model predictive control'. In: *Korean Journal of Chemical Engineering* 38.7 (2021), pp. 1317–1332.
- [31] J. B. Rawlings, D. Angeli and C. N. Bates. 'Fundamentals of economic model predictive control'. In: *Proceedings of the 51st IEEE Conference on Decision and Control*. IEEE. 2012, pp. 3851–3861.
- [32] L. Grüne and J. Pannek. *Nonlinear Model Predictive Control*. Springer, 2017.
- [33] T. Faulwasser, L. Grüne and M. A. Müller. 'Economic nonlinear model predictive control'. In: *Foundations and Trends® in Systems and Control* 5.1 (2018), pp. 1–98.
- [34] M. Sibeijn, S. Ahmed, M. Khosravi and T. Keviczky. 'Economic Nonlinear Model Predictive Control of Prosumer District Heating Networks'. In: *IEEE Transactions on Control Systems Technology* 33.5 (2025), pp. 1879–1894. DOI: [10.1109/TCST.2025.3561501](https://doi.org/10.1109/TCST.2025.3561501).
- [35] M. Sibeijn, S. Pequito, D. Boskos and M. Khosravi. 'Sensor Placement in District Heating Networks Using Frequency-Domain Gramians'. In: *IEEE Control Systems Letters* 9 (2025), pp. 1634–1639. DOI: [10.1109/LCSYS.2025.3580027](https://doi.org/10.1109/LCSYS.2025.3580027).
- [36] M. Sibeijn, S. Ahmed, M. Khosravi and T. Keviczky. 'Dissipativity analysis for economic nonlinear MPC of district heating networks'. In: *Proceedings of the European Control Conference*. 2024, pp. 1111–1118.
- [37] M. Sibeijn, M. Khosravi and T. Keviczky. 'Mitigating short-sightedness of MPC for district heating networks using dual dynamic programming'. In: *Proceedings of the IEEE Conference on Decision and Control*. 2024, pp. 4608–4614.
- [38] M. Sibeijn, M. Khosravi, S. Pequito and T. Keviczky. 'Scalable Predictive Control for District Heating Networks: A Physics-Guided Koopman Operator Approach'. In: *Under review for IEEE Transactions on Control Systems Technology* (2025).
- [39] A. J. Neale, R. F. Babus'Haq, S. D. Probert and M. J. Shilston. 'Thermal design of district-heating distribution networks'. In: *Applied Energy* 28.4 (1987), pp. 269–282.
- [40] S. El Mrabet, B. Lamrani, M. Abd-Lefdil and T. Kousksou. 'A brief overview of district heating pipe network progress'. In: *Energy Conversion and Management: X* 23 (2024), p. 100641.

- [41] A. La Bella, L. Nigro and R. Scattolini. 'Predictive Control and Benefit Sharing in Multi-Energy Systems'. In: *IEEE Transactions on Control Systems Technology* (2023).
- [42] M. Rein, J. Mohring, T. Damm and A. Klar. 'Optimal control of district heating networks using a reduced order model'. In: *Optimal Control Applications and Methods* 41.4 (2020), pp. 1352–1370.
- [43] C. De Persis and C. S. Kallés. 'Pressure regulation in nonlinear hydraulic networks by positive and quantized controls'. In: *IEEE Transactions on Control Systems Technology* 19.6 (2011), pp. 1371–1383.
- [44] D. B. West. *Introduction to Graph Theory*. 2nd. Prentice Hall, 2001. ISBN: 978-0-13-014400-3.
- [45] R. J. LeVeque and R. J. LeVeque. *Numerical methods for conservation laws*. Vol. 132. Springer, 1992.
- [46] B. van der Heijde. 'Optimal integration of thermal energy storage and conversion in fourth generation thermal networks'. PhD thesis. KU Leuven, 2019.
- [47] A. Matei, A. Bott, L. Rehlich, F. Steinke and S. Ulbrich. 'Optimal sensor placement in district heating networks for Bayesian inference of uncertain demands'. In: *Proceedings of the 4th International Conference on Uncertainty Quantification in Computational Sciences and Engineering*. 2021.
- [48] T. H. Summers and J. Lygeros. 'Optimal sensor and actuator placement in complex dynamical networks'. In: *Proceedings of the 19th IFAC World Congress*. Vol. 47. 3. 2014, pp. 3784–3789.
- [49] F. Pasqualetti, S. Zampieri and F. Bullo. 'Controllability Metrics, Limitations and Algorithms for Complex Networks'. In: *IEEE Transactions on Control of Network Systems* 1.1 (2014), pp. 40–52.
- [50] K. Morris. 'Linear-Quadratic Optimal Actuator Location'. In: *IEEE Transactions on Automatic Control* 56.1 (2011), pp. 113–124.
- [51] L. M. Dang, S. Lee, Y. Li, C. Oh, T. N. Nguyen, H.-K. Song and H. Moon. 'Daily and seasonal heat usage patterns analysis in heat networks'. In: *Scientific Reports* 12.1 (2022), p. 9165.
- [52] A. J. Krener and K. Ide. 'Measures of unobservability'. In: *Proceedings of the IEEE Conference on Decision and Control*. 2009, pp. 6401–6406.
- [53] K. Manohar, B. W. Brunton, J. N. Kutz and S. L. Brunton. 'Data-driven sparse sensor placement for reconstruction: Demonstrating the benefits of exploiting known patterns'. In: *IEEE Control Systems Magazine* 38.3 (2018), pp. 63–86.
- [54] X. Wu, B. Jacob and H. Elbern. 'Optimal control and observation locations for time-varying systems on a finite-time horizon'. In: *SIAM Journal on Control and Optimization* 54.1 (2016), pp. 291–316.
- [55] A. Padovan and C. W. Rowley. 'Continuous-time balanced truncation for time-periodic fluid flows using frequential Gramians'. In: *Journal of Computational Physics* 496 (2024), p. 112597.

-
- [56] V. Tzoumas, M. A. Rahimian, G. J. Pappas and A. Jadbabaie. ‘Minimal Actuator Placement With Bounds on Control Effort’. In: *IEEE Transactions on Control of Network Systems* 3.1 (2016), pp. 67–78.
- [57] V. Tzoumas, M. A. Rahimian, G. J. Pappas and A. Jadbabaie. *Minimal Actuator Placement with Bounds on Control Effort*. arXiv preprint. Dec. 2016.
- [58] T. H. Summers, F. L. Cortesi and J. Lygeros. ‘Corrections to “On submodularity and controllability in complex dynamical networks”’. In: *IEEE Transactions on Control of Network Systems* 5.3 (2018), pp. 1503–1503.
- [59] F. Strehle, J. E. Machado, M. Cucuzzella, A. J. Malan, S. Hohmann and J. M. A. Scherpen. ‘A unifying passivity-based framework for pressure and volume flow rate control in district heating networks’. In: *IEEE Transactions on Control Systems Technology* (2024).
- [60] J. Simonsson, K. T. Atta and W. Birk. ‘Semi-decentralized temperature control in district heating systems’. In: *Journal of Process Control* 140 (2024), p. 103251.
- [61] S. Ahmed, J. E. Machado, M. Cucuzzella and J. M. A. Scherpen. ‘Control-oriented modeling and passivity analysis of thermal dynamics in a multi-producer district heating system’. In: *Proceedings of the 22nd IFAC World Congress*. Vol. 56. 1. 2023, pp. 175–180.
- [62] T. Scholten, C. De Persis and P. Tesi. ‘Modeling and control of heat networks with storage: The single-producer multiple-consumer case’. In: *IEEE Transactions on Control Systems Technology* 25.2 (2016), pp. 414–428.
- [63] J. E. Machado, J. Ferguson, M. Cucuzzella and J. M. A. Scherpen. ‘Decentralized Temperature and Storage Volume Control in Multiproducer District Heating’. In: *IEEE Control Systems Letters* 7 (2023), pp. 413–418.
- [64] H. Zhao, J. Holst and L. Arvastson. ‘Optimal operation of coproduction with storage’. In: *Energy* 23.10 (1998), pp. 859–866.
- [65] F. M. Hante and M. Schmidt. ‘Complementarity-based nonlinear programming techniques for optimal mixing in gas networks’. In: *EURO Journal on Computational Optimization* 7.3 (2019), pp. 299–323.
- [66] G. Sandou, S. Font, S. Tebbani, A. Hired and C. Mondon. ‘Predictive control of a complex district heating network’. In: *Proceedings of the 44th IEEE Conference on Decision and Control, and the European Control Conference*. 2005, pp. 7372–7377.
- [67] S. S. Farahani, Z. Lukszo, T. Keviczky, B. De Schutter and R. M. Murray. ‘Robust model predictive control for an uncertain smart thermal grid’. In: *Proceedings of the European Control Conference*. 2016, pp. 1195–1200.
- [68] L. Frison, M. Kollmar, A. Oliva, A. Bürger and M. Diehl. ‘Model predictive control of bidirectional heat transfer in prosumer-based solar district heating networks’. In: *Applied Energy* 358 (2024), p. 122617.
- [69] M. Rose, H. Gernandt, J. E. Machado and J. Schiffer. *Model Predictive Control of District Heating Grids Using Stabilizing Terminal Ingredients*. arXiv preprint arXiv:2404.01820. 2024.

- [70] F. Agner, P. Kergus, R. Pates and A. Rantzer. ‘Combating district heating bottlenecks using load control’. In: *Smart Energy* 6 (2022), p. 100067.
- [71] S. van der Zwan and I. Pothof. ‘Operational optimization of district heating systems with temperature limited sources’. In: *Energy and Buildings* 226 (2020), p. 110347.
- [72] R. Gondhalekar and J.-i. Imura. ‘Least-restrictive move-blocking model predictive control’. In: *Automatica* 46.7 (2010), pp. 1234–1240.
- [73] M. Lubin, O. Dowson, J. Dias Garcia, J. Huchette, B. Legat and J. P. Vielma. ‘JuMP 1.0: Recent improvements to a modeling language for mathematical optimization’. In: *Mathematical Programming Computation* (2023).
- [74] A. Wächter and L. T. Biegler. ‘On the implementation of an interior-point filter line-search algorithm for large-scale nonlinear programming’. In: *Mathematical Programming* 106 (2006), pp. 25–57.
- [75] C. Rackauckas and Q. Nie. ‘DifferentialEquations.jl—A performant and feature-rich ecosystem for solving differential equations in Julia’. In: *Journal of Open Research Software* 5.1 (May 2017), p. 15.
- [76] Ember. *European wholesale electricity price data*. Ember, Sandbag Climate Campaign CIC. 2024. URL: <https://ember-climate.org/data-catalogue/european-wholesale-electricity-price-data/>.
- [77] L. Grüne and S. Pirkelmann. ‘Economic model predictive control for time-varying system: Performance and stability results’. In: *Optimal Control Applications and Methods* 41.1 (2020), pp. 42–64.
- [78] L. Grüne. *Why does strict dissipativity help in model predictive control?* Tech. rep. Department of Mathematics, University of Bayreuth, 2020.
- [79] L. Grüne, S. Pirkelmann and M. Stieler. ‘Strict dissipativity implies turnpike behavior for time-varying discrete time optimal control problems’. In: Springer International Publishing, 2018, pp. 195–218.
- [80] T. Faulwasser, M. Korda, C. N. Jones and D. Bonvin. ‘Turnpike and dissipativity properties in dynamic real-time optimization and economic MPC’. In: *Proceedings of the 53rd IEEE Conference on Decision and Control*. IEEE, Dec. 2014, pp. 2734–2739.
- [81] J. Berberich, J. Köhler, F. Allgöwer and M. A. Müller. ‘Dissipativity properties in constrained optimal control: A computational approach’. In: *Automatica* 114 (2020).
- [82] S. Pirkelmann, D. Angeli and L. Grüne. ‘Approximate computation of storage functions for discrete-time systems using sum-of-squares techniques’. In: *Proceedings of the 21st IFAC World Congress*. Vol. 52. 16. 2019, pp. 508–513.
- [83] B. Flamm, A. Eichler, J. Warrington and J. Lygeros. ‘Dual dynamic programming for nonlinear control problems over long horizons’. In: *Proceedings of the European Control Conference*. 2018, pp. 471–476.

-
- [84] P. M. Castro, I. E. Grossmann and Q. Zhang. 'Expanding scope and computational challenges in process scheduling'. In: *Computers & Chemical Engineering* 114 (2018), pp. 14–42.
- [85] B. Brunaud and I. E. Grossmann. 'Perspectives in multilevel decision-making in the process industry'. In: *Frontiers of Engineering Management* 4.3 (2017), pp. 256–270.
- [86] R. M. Lima, I. E. Grossmann and Y. Jiao. 'Long-term scheduling of a single-unit multi-product continuous process to manufacture high-performance glass'. In: *Computers & Chemical Engineering* 35.3 (2011), pp. 554–574.
- [87] D. I. Mendoza-Serrano and D. J. Chmielewski. 'Demand response for chemical manufacturing using economic MPC'. In: *Proceedings of the American Control Conference*. IEEE. June 2013, pp. 6655–6660.
- [88] N. V. Sahinidis and I. E. Grossmann. 'Convergence properties of generalized Benders decomposition'. In: *Computers & Chemical Engineering* 15.7 (1991), pp. 481–491.
- [89] I. Gabriellaitienė, B. Bøhm and B. Sundén. 'Modelling Temperature Dynamics of a District Heating System in Naestved, Denmark: A Case Study'. In: *Energy Conversion and Management* 48.1 (2007), pp. 78–86.
- [90] L. Nigro, A. La Bella, F. Casella and R. Scattolini. 'Control-Oriented Modeling, Simulation, and Predictive Control of District Heating Networks'. In: *IEEE Transactions on Automation Science and Engineering* (2024).
- [91] C. A. Floudas and A. R. Ciric. 'Strategies for overcoming uncertainties in heat exchanger network synthesis'. In: *Computers & Chemical Engineering* 13.10 (1989), pp. 1133–1152.
- [92] J. Maurer, J. Illerhaus, P. J. Soneira and S. Hohmann. 'Distributed optimization of district heating networks using optimality condition decomposition'. In: *Energies* 15.18 (2022), p. 6605.
- [93] B. O. Koopman. 'Hamiltonian systems and transformation in Hilbert space'. In: *Proceedings of the National Academy of Sciences* 17.5 (1931), pp. 315–318.
- [94] M. Korda and I. Mezić. 'Linear Predictors for Nonlinear Dynamical Systems: Koopman Operator Meets Model Predictive Control'. In: *Automatica* 93 (2018), pp. 149–160.
- [95] M. Khosravi. 'Representer theorem for learning Koopman operators'. In: *IEEE Transactions on Automatic Control* 68.5 (2023), pp. 2995–3010.
- [96] F. Nüske, S. Peitz, F. Philipp, M. Schaller and K. Worthmann. 'Finite-data error bounds for Koopman-based prediction and control'. In: *Journal of Nonlinear Science* 33.1 (2023), p. 14.
- [97] J. S. Kim, Y. S. Quan and C. C. Chung. 'Koopman operator-based model identification and control for automated driving vehicle'. In: *International Journal of Control, Automation and Systems* 21.8 (2023), pp. 2431–2443.

- [98] L. Do, M. Korda and Z. Hurák. ‘Controlled synchronization of coupled pendulums by Koopman Model Predictive Control’. In: *Control Engineering Practice* 139 (2023), p. 105629.
- [99] Y. Shi, X. Hu, Z. Zhang, Q. Chen, L. Xie and H. Su. ‘Data-driven identification and fast model predictive control of the ORC waste heat recovery system by using Koopman operator’. In: *Control Engineering Practice* 141 (2023), p. 105679.
- [100] Y. Chen, K. H. Kwak, D. D. Jung and Y. Kim. ‘Data-driven Koopman model predictive control for the integrated thermal management of electric vehicles’. In: *Control Engineering Practice* 160 (2025), p. 106323.
- [101] C. M. Jensen, M. C. Frederiksen, C. S. Kallesøe, J. N. Jensen, L. H. Andersen and R. Izadi-Zamanabadi. ‘HAVOK model predictive control for time-delay systems with applications to district heating’. In: *Proceedings of the 22nd IFAC World Congress*. Vol. 56. 2. 2023, pp. 2238–2243.
- [102] S. J. Cox, D. Kim, H. Cho and P. Mago. ‘Real time optimal control of district cooling system with thermal energy storage using neural networks’. In: *Applied Energy* 238 (2019), pp. 466–480.
- [103] L. B. de Giuli, A. La Bella and R. Scattolini. ‘Physics-informed neural network modeling and predictive control of district heating systems’. In: *IEEE Transactions on Control Systems Technology* 32.4 (2024), pp. 1182–1195.
- [104] T. Boussaid, F. Rousset, V.-M. Scuturici and M. Clause. ‘Enabling fast prediction of district heating networks transients via a physics-guided graph neural network’. In: *Applied Energy* 370 (2024), p. 123634.
- [105] S. Buffa, M. H. Fouladfar, G. Franchini, I. Lozano Gabarre and M. Andres Chicote. ‘Advanced control and fault detection strategies for district heating and cooling systems—A review’. In: *Applied Sciences* 11.1 (2021), p. 455.
- [106] Gurobi Optimization, LLC. *Gurobi Optimizer Reference Manual*. 2024.
- [107] B. Stellato, G. Banjac, P. Goulart, A. Bemporad and S. Boyd. ‘OSQP: An operator splitting solver for quadratic programs’. In: *Mathematical Programming Computation* 12.4 (2020), pp. 637–672.
- [108] B. Stellato, V. V. Naik, A. Bemporad, P. Goulart and S. Boyd. ‘Embedded Mixed-Integer Quadratic Optimization Using the OSQP Solver’. In: *Proceedings of the European Control Conference*. July 2018.
- [109] M. Haseli and J. Cortés. ‘Data-driven approximation of Koopman-invariant subspaces with tunable accuracy’. In: *Proceedings of the American Control Conference*. IEEE. 2021, pp. 470–475.
- [110] M. Haseli and J. Cortés. ‘Generalizing dynamic mode decomposition: Balancing accuracy and expressiveness in Koopman approximations’. In: *Automatica* 153 (2023), p. 111001.
- [111] X. Bao, N. V. Sahinidis and M. Tawarmalani. ‘Semidefinite relaxations for quadratically constrained quadratic programming: A review and comparisons’. In: *Mathematical Programming* 129.1 (2011), pp. 129–157.

-
- [112] J. Park and S. Boyd. ‘General heuristics for nonconvex quadratically constrained quadratic programming’. In: *arXiv preprint arXiv:1703.07870* (2017).
 - [113] H. Luo, X. Bai and J. Peng. ‘Enhancing semidefinite relaxation for quadratically constrained quadratic programming via penalty methods’. In: *Journal of Optimization Theory and Applications* 180.3 (2019), pp. 964–992.
 - [114] K. M. Anstreicher. ‘On convex relaxations for quadratically constrained quadratic programming’. In: *Mathematical Programming* 136.2 (2012), pp. 233–251.
 - [115] J. M. Lemos, R. Neves-Silva and J. M. Igreja. *Adaptive Control of Solar Energy Collector Systems*. Springer, 2014.
 - [116] B. Amos, L. Xu and J. Z. Kolter. ‘Input convex neural networks’. In: *Proceedings of the International Conference on Machine Learning*. PMLR, 2017, pp. 146–155.
 - [117] F. Büning, A. Schalbetter, A. Aboudonia, M. H. de Badyn, P. Heer and J. Lygeros. ‘Input convex neural networks for building MPC’. In: *Proceedings of Learning for Dynamics and Control*. PMLR, 2021, pp. 251–262.
 - [118] K. Seel, A. B. Kordabad, S. Gros and J. T. Gravdahl. ‘Convex neural network-based cost modifications for learning model predictive control’. In: *IEEE Open Journal of Control Systems* 1 (2022), pp. 366–379.
 - [119] D. Tellez-Castro, C. Garcia-Tenorio, E. Mojica-Nava, J. Sofrony and A. Vande Wouwer. ‘Data-driven predictive control of interconnected systems using the Koopman operator’. In: *Actuators* 11.6 (2022), p. 151.
 - [120] M. Sibeijn and S. Pequito. ‘Discrete-time Linear Time-invariant Distributed Minimum Energy Estimator’. In: *Proceedings of the 22nd IFAC World Congress*. 2023, pp. 3856–3861.

Curriculum Vitæ

Max Willem SIBEIJN

23-12-1996 Born in The Hague, The Netherlands.

Education

2009–2015 High School
The British School Muscat (2009–2012)
Haags Montessorri Lyceum (2012–2015)

2015–2019 BSc in Mechanical Engineering
Delft University of Technology

2019–2021 MSc in Systems & Control
Delft University of Technology

2022-2025 PhD
Delft University of Technology
Thesis: Scalable Predictive Control for District Heating Networks
Promotors: Prof. dr. T. Keviczky and dr. S.D. Gonçalves Melo Pequito
Co-promotor: Dr. M. Khosravi

List of Publications

6. M. Sibeijn, M. Khosravi, S. Pequito and T. Keviczky. 'Scalable Predictive Control for District Heating Networks: A Physics-Guided Koopman Operator Approach'. In: *Under review for IEEE Transactions on Control Systems Technology* (2025)
5. M. Sibeijn, S. Pequito, D. Boskos and M. Khosravi. 'Sensor Placement in District Heating Networks Using Frequency-Domain Gramians'. In: *IEEE Control Systems Letters* 9 (2025), pp. 1634–1639. DOI: [10.1109/LCSYS.2025.3580027](https://doi.org/10.1109/LCSYS.2025.3580027)
4. M. Sibeijn, S. Ahmed, M. Khosravi and T. Keviczky. 'Economic Nonlinear Model Predictive Control of Prosumer District Heating Networks'. In: *IEEE Transactions on Control Systems Technology* 33.5 (2025), pp. 1879–1894. DOI: [10.1109/TCST.2025.3561501](https://doi.org/10.1109/TCST.2025.3561501)
3. M. Sibeijn, M. Khosravi and T. Keviczky. 'Mitigating short-sightedness of MPC for district heating networks using dual dynamic programming'. In: *Proceedings of the IEEE Conference on Decision and Control*. 2024, pp. 4608–4614
2. M. Sibeijn, S. Ahmed, M. Khosravi and T. Keviczky. 'Dissipativity analysis for economic nonlinear MPC of district heating networks'. In: *Proceedings of the European Control Conference*. 2024, pp. 1111–1118
1. M. Sibeijn and S. Pequito. 'Discrete-time Linear Time-invariant Distributed Minimum Energy Estimator'. In: *Proceedings of the 22nd IFAC World Congress*. 2023, pp. 3856–3861

