Coordinated Optimization and Control for Smart Grids

THÈSE N^{O} 8814 (2018)

PRÉSENTÉE LE 29 AOÛT 2018

À LA FACULTÉ DES SCIENCES ET TECHNIQUES DE L'INGÉNIEUR
LABORATOIRE D'AUTOMATIQUE 3
PROGRAMME DOCTORAL EN GÉNIE ÉLECTRIQUE

ÉCOLE POLYTECHNIQUE FÉDÉRALE DE LAUSANNE

POUR L'OBTENTION DU GRADE DE DOCTEUR ÈS SCIENCES

PAR

Altuğ BITLISLIOĞLU

acceptée sur proposition du jury:

Prof. G. Ferrari Trecate, président du jury
Prof. C. N. Jones, directeur de thèse
Prof. T. Keviczky, rapporteur
Prof. I. Necoara, rapporteur
Dr A. Karimi, rapporteur

ÉCOLE POLYTECHNIQUE
FÉDÉRALE DE LAUSANNE

Suisse
2018

…seek out and identify structures of authority, hierarchy, and domination in every aspect of life, and challenge them; unless a justification for them can be given, they are illegitimate, and should be dismantled, to increase the scope of human freedom.
— Noam Chomsky

To my parents…

Acknowledgements

I would like to thank Colin (Prof. Jones), for putting his trust in me for carrying out my doctoral studies with him. With his positive, calm attitude and sharp mind, he was always a source of inspiration and he made my journey much more easier and bearable. I hope I had been able to match his expectations. He was always there with good advice and support, when I needed it.
I would like to thank the members of my thesis comittee Professors Keviczky, Necoara, Karimi and Ferrari-Trecate, for dedicating their time and mind on my work, improving the quality and the delivery of this thesis.
My gratitude and love goes to the LA (laboratoire d'automatique) family, with its unique and warm culture, under the protection of our Professors. I always felt at home and safe at LA with the limitless friendship and support. I abstain from listing names, fearing that I will miss out dear people. Special mention goes to Tomasz, my academic and climbing buddy, to the members of the Silent Bug; Giorgios and Luca, and to my chess nemesis, Faran. Furthermore, I would like to mention Christophe Salzmann, the corner stone of LA, who set an example for me with his dedication to the community and to his job. Big thanks to Truong, for helping me improve my programming with his kind supervision and to Yannis, for his friendship and support in times of need.
For the last few years that I had spent working on this thesis, I had dearly missed my friends from the joyful and often difficult times in Ankara and Zurich. I tried to contact them when I could and always found them with their warm and supporting embrace. I thank all, and hope to see them more in the future.
My deepest gratitude goes to my extended family in Ankara and Bursa for their unconditional love and support. I was mostly away from my father, mother and brother but I've always had them in my mind and felt their love and encouragement. My significant other, Gozde, there are no words to express my gratitude to you, so I abstain from a futile attempt here. And last but not least, Kiraz, thank you for the smile that takes away all the burden of life.

Lausanne, July 2018 A. B.

Abstract

In this thesis, we consider commercial buildings with available heating, ventilation and air conditioning (HVAC) systems, and develop methods to assess and exploit their energy storage and production potential to collectively offer ancillary services to the power grid. This demand response problem can be put in the generic framework of multi-agent optimization and control. In this setting, various agents interact through their objectives, constraints or dynamics over a network. In the example of demand response, individual buildings are connected to the power network and coupled via their common objective of providing service and the constraints of the power network. Within this generic multi-agent framework, we develop layers of abstractions that enables efficient coordination of the agents while making sure that the network constraints are satisfied, and the common goal of the agents is achieved.

The first approach is based on quantifying the tracking capability of a local system using robust optimization. Different from a standard robust optimization problem, we modify and optimize over the uncertainty set that represents the set of reference trajectories the system is required to track while rejecting external disturbances. The method facilitates hierarchical control by using reference sets for coordinating many agents.

In the second approach, we consider coordination of multiple agents by using local cost and constraint approximations. Specifically we consider decomposition of interior point methods in a multi-agent setting and analyze the computation and modeling task for the agents and the coordinator. We further consider decomposition of state of the art predictor-corrector type interior point methods and show that a naive implementation may result in excessive communication in a multi-agent setting. In order to remedy this issue, we propose a modification of the standard algorithm that uses decentralized predictions. We analyze convergence of the method and test the performance with numerical experiments.

Finally, we look into applying decomposition based interior point methods in a distributed model predictive control problem that includes dynamic coupling between the agents. Instead of solving the problem to optimality, adding barrier functions to the objective enhances numerical performance significantly, an approach that is well-known in model predictive control (MPC) literature. We consider applying this method in economic MPC problems with terminal equilibrium constraints, which is suitable for decomposition due to the simplicity of terminal constraints. However in this case standard results for MPC with barrier functions do not apply. We propose iterative re-centering of the barriers, which allows interpreting them as a regularizing cost in the problem that penalizes deviation from open-loop predictions. We show that regularizing barrier functions not only improve the numerical performance and facilitate

decomposition, but also enhance system theoretical properties.

Keywords : Control, Optimization, Model Predictive Control, Convex Optimization, Distributed Control, Distributed Optimization, Coordinated Optimization, Hierarchical Control, Interior-Point Method, Robust Optimization, Decomposition, Coordination, Smart Grids, Demand Response

Résumé

Dans cette thèse, nous considérons les bâtiments commerciaux avec des systèmes de chauffage, ventilation et climatisation (CVC) disponibles et développons des méthodes pour évaluer et exploiter leur potentiel de stockage et de production d'énergie afin d'offrir collectivement des services auxiliaires au réseau électrique. Ce problème de réponse à la demande peut être placé dans le cadre générique de l'optimisation et du contrôle multi-agents. Dans ce cadre, différents agents interagissent à travers leurs objectifs, contraintes ou dynamiques sur un réseau. Dans l'exemple de la réponse à la demande, les bâtiments individuels sont connectés au réseau électrique et couplés via leur objectif commun de fourniture de service et les contraintes du réseau électrique. Dans ce cadre générique multi-agent, nous développons des couches d'abstractions qui permettent une coordination efficace des agents tout en s'assurant que les contraintes réseau sont satisfaites et que l'objectif commun des agents est atteint.

La première approche est basée sur la quantification de la capacité de suivi d'un système local en utilisant une optimisation robuste. Différent d'un problème d'optimisation standard robuste, nous modifions et optimisons l'ensemble d'incertitudes qui représente l'ensemble des trajectoires de référence que le système doit suivre tout en rejetant les perturbations externes. La méthode facilite le contrôle hiérarchique en utilisant des ensembles de référence pour coordonner de nombreux agents.

Dans la seconde approche, nous considérons la coordination de plusieurs agents en utilisant des approximations de coût et de contrainte locales. Plus précisément, nous considérons la décomposition des méthodes de points intérieurs dans un environnement multi-agents et analysons la tâche de calcul et de modélisation pour les agents et le coordinateur. Nous considérons en outre la décomposition des méthodes de point intérieur de type prédicteur-correcteur de l'état de l'art et montrons qu'une implémentation naïve peut entraîner une communication excessive dans un environnement multi-agent. Afin de remédier à ce problème, nous proposons une modification de l'algorithme standard qui utilise des prédictions décentralisées. Nous analysons la convergence de la méthode et testons la performance avec des expériences numériques.

Enfin, nous examinons l'application de méthodes de points intérieurs basés sur la décomposition dans un problème de contrôle prédictif de modèle réparti qui inclut un couplage dynamique entre les agents. Au lieu de résoudre le problème à l'optimalité, l'ajout de fonctions de barrière à l'objectif améliore considérablement la performance numérique, une approche qui est bien connue dans la littérature de contrôle prédictif de modèle (MPC). Nous considérons l'application de cette méthode dans des problèmes de MPC économiques avec des contraintes d'équilibre terminales, ce qui convient à la décomposition en raison de la simplicité des contraintes termi-

nales. Cependant, dans ce cas, les résultats standard pour MPC avec des fonctions de barrière ne s'appliquent pas. Nous proposons un recentrage itératif des barrières, ce qui permet de les interpréter comme un coût de régularisation du problème qui pénalise l'écart par rapport aux prédictions en boucle ouverte. Nous montrons que les fonctions de barrière de régularisation améliorent non seulement la représentation numérique et facilitent la décomposition, mais améliorent également les propriétés théoriques de système.

Mots-clés : Contrôle, Optimisation, Contrôle prédictif de modèle, Optimisation convexe, Contrôle distribué, Optimisation distribuée, Optimisation coordonnée, Contrôle hiérarchique, Méthode des points intérieurs, Optimisation robuste, Décomposition, Coordination, Réseaux intelligents, Réponse à la demande.

Contents

Acknowledgements iv

Abstract (English/Français) v

1 Introduction 1
- 1.1 Motivation: Demand response with commercial building networks . . . 1
- 1.2 A multi agent coordination framework . . . 2
- 1.3 Summary and contributions . . . 2
- 1.4 Notation . . . 6

2 Hierarchical Control via Tracking Commitment 8
- 2.1 Introduction . . . 9
- 2.2 Preliminaries . . . 13
  - 2.2.1 Robust optimization . . . 13
  - 2.2.2 Model predictive control . . . 19
  - 2.2.3 Robust MPC . . . 20
- 2.3 Constrained tracking under uncertainty . . . 24
  - 2.3.1 Robust Tracking Commitment . . . 27
  - 2.3.2 Computationally tractable approximations . . . 32
- 2.4 Frequency regulation service with buildings . . . 39
  - 2.4.1 Building Models . . . 40
  - 2.4.2 Case study 1: Commitment with a single building under disturbance . 41
  - 2.4.3 Case study 2: Hierarchical tracking with a building fleet . . . 47
- 2.5 Conclusion . . . 50

3 Coordinated Control and Optimization via Interior Point Decomposition 54
- 3.1 Introduction . . . 55
- 3.2 Preliminaries . . . 57
  - 3.2.1 Multi-agent optimization . . . 57
  - 3.2.2 Decomposition of optimization problems . . . 58
  - 3.2.3 Primal decomposition . . . 59
  - 3.2.4 Dual decomposition . . . 59
  - 3.2.5 Interior-point methods . . . 60
- 3.3 Coordinated Interior-Point Method . . . 67

3.3.1 Coordinated Newton method . . . 69
3.3.2 Initialization with surrogate sets . . . 70
3.3.3 Line search . . . 70
3.3.4 Case study: Voltage support with commercial buildings . . . 72
3.4 Coordinated primal-dual interior point method . . . 77
3.4.1 Decomposition of Primal-Dual Interior point method . . . 78
3.4.2 Decentralized predictor-corrector . . . 81
3.4.3 Convergence with non-uniform perturbation . . . 82
3.4.4 Numerical experiments . . . 83
3.5 Conclusion . . . 84

4 Coordinated Economic MPC with Barrier Functions 88
4.1 Introduction . . . 89
4.2 Preliminaries . . . 90
4.2.1 Economic MPC . . . 90
4.2.2 Coordinated MPC . . . 92
4.2.3 Multi-objective optimization . . . 94
4.3 Economic MPC with barriers . . . 95
4.3.1 Regularized economic MPC . . . 95
4.3.2 Barrier functions as regularizing cost . . . 96
4.4 Closed loop properties . . . 97
4.4.1 Sub-optimality of the terminal cost . . . 99
4.4.2 Reachability of equilibrium points with improved terminal cost . . . 100
4.4.3 Convergence and average performance . . . 101
4.5 Case study: Power grid frequency control . . . 103
4.6 Conclusion . . . 106

5 Conclusion 109
5.1 Future directions . . . 111

A Appendix 112
A.1 Derivation of sensitivity equations . . . 112
A.2 Ellipsoid projection . . . 113
A.3 Convergence of PDIP with non-uniform complementarity perturbation . . . 114
A.4 Reachability of steady states . . . 119

Bibliography 129

Curriculum Vitae 131

List of Figures

2.1 Conceptual sketch of the relationships between uncertainty sets and applied functions. The initial uncertainty set $\hat{\Xi}$ is not necessarily a subset of the projection of $\mathcal{Q}$, therefore might not be admissible according to Lemma 1. However once we find a feasible lifting of this set into $\mathcal{Q}$, we can take its projection as an admissible uncertainty set, which is given by $\Xi = \boldsymbol{\nu}(\hat{\Xi})$. The corresponding admissible control policy can be obtained by letting $\boldsymbol{\pi} = \hat{\boldsymbol{\pi}} \circ \boldsymbol{\nu}^{-1}$. 29

2.2 Architecture of the control system for tracking service commitment and tracking. 39

2.3 Diagram of the HVAC system, which consists of a chiller, a chilled water (CW) loop, an air handling unit (AHU), duct system and variable air volume (VAV) boxes. $\dot{m}$ is mass flow rate, T temperature, P_{comp} the electrical power consumption of the chiller, q_{cool} the heat power extracted from the chilled water loop, v_{coil} the valve control parameter of the cooling coil inside the AHU. The controlled variables are $\dot{m}_{\text{VAV}}$. The input air temperature $T^{\text{in}}_{\text{AHU}}$ depends on both zone temperatures T_{zones} and ambient temperature T_{amb}, due to the mixing of internal and external air in the return duct. 41

2.4 Information structure for the example. $\boldsymbol{\mathcal{I}}$ shows that decisions at time step t can depend on the reference up to time t and disturbance up to time $t-1$. $\hat{\boldsymbol{\mathcal{I}}}$ is the resulting information structure for the modifier function. 43

2.5 Trajectories for different weather and reference scenarios in the optimized reference set. Shaded region is the tracking commitment period. Black lines show the 'nominal' scenario where the reference is zero and the weather takes its predicted value. From top to bottom: temperature in zones, total power consumption, AGC, tracking error, and weather scenarios. 45

2.6 Tracking capacity bid versus duration of participation for box and battery reference sets. 46

2.7 Hierarchical control structure for AGC tracking with a fleet of buildings. . . . 47

2.8 Block diagram of the building control structure. The upper level MPC controller provides thermal energy u set-points to the building automation system (BAS). The BAS consists of three PI control loops that control the cooling power output of the HVAC system u, by means of adjusting water temperature at the evaporator of the chiller T_{evap}, and the air supply temperature that exits the AHU, T_{AHU}. It is the responsibility of MPC to choose a suitable cooling power reference u_{ref} that results in the desired level of power consumption y. 49

2.9 Realized power inflows and state of charge for the virtual battery models of 30 buildings in the fleet. The trajectories of the building 10 is highlighted. 51
2.10 Aggregated AGC tracking performance with the outputs from 30 buildings. The output of the building 10 is highlighted. 52
2.11 Temperature values of the HVAC system and the thermal zones for building 10. The bottom plot shows the tracking of electrical power output and the resulting cooling power supply. 53

3.1 Predicted voltage magnitudes, without the demand response coordination (top), and with optimal coordination of 12 buildings in the grid (bottom). 74
3.2 The top figure shows the voltage magnitude at node 32, where the solar power infeed is placed. The figures below show the zone temperatures and total cooling power for a large office building placed at node 31. 75
3.3 Evolution of suboptimality measure $\|\boldsymbol{y} - \boldsymbol{y}^*\|_2$ for the problem (3.41), using the Dikin's ellipsoid feasibility search and backtracking with a reduction factor of $\kappa = 0.8$. Total Newton steps are 58 whereas 77 communication steps were carried out. Barrier shrinking factor μ is set to 0.4. 76
3.4 Comparison between ADMM and IPD applied to the problem (3.41). The deviation from optimal profiles is defined as the sum of absolute difference in MWh, with respect to the optimal power consumption profiles, for all buildings: $\sum_{i=1}^{n_a} \sum_{k=1}^{N} \delta t |y^i(k) - y^i(k)^*|$, with δt being the sampling time of 30 minutes. . 76
3.5 Newton step for the coordinated PDIP method 81
3.6 Evolution of iteration and communication rounds required to bring the suboptimality $\|y - y^*\|/n_y$ below 10^{-5}, for increasing number of agents for randomly generated coupled QPs with feasible starting points, using the mean of 10 sampled problems with $n_x = 20$, $n_y = 5$. 84
3.7 Comparison of various CPDIP mehtods for a realistic example of demand response [BPJ17a]. In this case 12 smart-buildings help a distribution grid for voltage support in case of a significant solar power in feed. The algorithm is started from a point that is locally centered but infeasible with respect to the coupling constraint. 85

4.1 Closed loop (solid) and predicted terminal state and input trajectories (dashed) for G-EMPC (left) and with regularization (right). 98
4.2 Closed loop stage cost of the state(solid-blue) and predicted terminal steady state (dashed-orange) for G-EMPC (left) and with regularization (right). 98
4.3 Comparison of controller and solver performance for different integration of barriers. From left to right, the cases with: no barriers, fixed barriers around the origin and iteratively re-centered barriers. For the case with barriers, the reduction factor for the complementarity target is set to 0.1. For all cases, an aggressive long-step approach is used, which takes 99% of the maximum feasible step-length. 106

4.4 Comparison of Closed loop input and state trajectories for individual areas for the case with (a) no barriers, (b) fixed barriers around the origin and (c) iteratively re-centered barriers. 107
4.5 Closed loop trajectories (a) and numerical performance (b) of regularized EMPC with re-centered barriers under disturbance, with $\beta = 1$ and early termination after four iterations of primal-dual Newton method. 108

A.1 Schematic representation of reachability of steady state $\tilde{x}$ from any state $x \in \mathcal{B}(\bar{x}, \delta)$ with a horizon length $N_1 + N_2$. 120

List of Tables

2.1 Optimization type for the robust counterpart formulation (2.69) of the tracking commitment problem, considering commonly used combinations of polytopes and ellipsoids as uncertainty sets. Note that the polytopic representation also covers 1 and ∞ norm balls. 35
2.2 Simulation parameters 50

3.1 Summary of buildings participating in the DR study 72

4.1 Normalized cost of power production per kW. 104

1 Introduction

1.1 Motivation: Demand response with commercial building networks

The social and economic cost of using fossil fuels is rising and will continue to rise due to looming climate change and finite resources [eur10]. While on the other hand environmental risks of nuclear fuel has forced some European countries to phase out their operational power plants [JK12]. The alternative to nuclear and fossil fuels is renewable energy resources, which will expose the electricity grid to their highly stochastic nature [RHF+08]. The question is: can we handle large scale renewable penetration and decentralized power generation while ensuring the safety of the electricity grid. In order to answer this question, we need to discover the potential of the today's grid, and develop methods to use it.

More uncertainty in power demand and generation means that we need more ancillary service providers in order to balance the unpredictable generation-demand mismatch in real time. If we don't want to burn more coal for balancing the uncertainty of renewable energy, we have to find energy storage solutions to absorb production uncertainty [MWM+14]. A viable option that is being considered by the power systems community is demand response, which means utilizing the flexibility of power demand to provide services to the power grid. Possible services might cover collective load shifting, real time power regulation for load balancing and capacity firming that will mitigate the uncertainty in the renewable power generation [BLR+10].

In this thesis, we will consider commercial buildings with available heating, ventilation and air conditioning (HVAC) systems, and develop methods to assess and exploit their energy storage and production potential to collectively offer ancillary services to the power grid. Power consumption of a single commercial building is in general only a fraction of one megawatt, and therefore a single building can not match the power scale that is used in the electricity markets. Therefore, providing ancillary services to the power grid will require collective operation of a network of buildings in order to aggregate their capacities.

1.2 A multi agent coordination framework

The problem of providing services to the power grid with a network of buildings can be put in the generic framework of multi-agent optimization and control. In this setting, various agents interact through their objectives, constraints or dynamics over a network. In the example of demand response, individual buildings are connected to the power network and coupled via their common objective of providing service and the constraints of the power network.

One key aspect here is the split of domain knowledge. The power network operators do not have the domain knowledge for modeling and controlling the operation of buildings directly. On the other hand the building operators are ready to lend their capabilities to the network, however they do not know the models and constraints of the power network. This split of knowledge and capability brings in the necessity of using a coordinated approach, that relies on a central agent which handles the constraints and objectives of the network, while coordinating the actions of the remaining agents.

The task of the coordinator is challenging however, as the number of agents might be overwhelming and each agent can be a very complex system in itself with local constraints, dynamics and objectives. In order to facilitate the coordination task, it is necessary to use simple layers of abstractions representing the capabilities and desires of the agents. A useful layer of abstraction would communicate 'just-enough' information to the coordinator, to ensure that the network constraints are satisfied, and the common goal of the agents is achieved.

The goal of this thesis is *to develop layers of abstractions to assess and exploit the capability of a network of agents, in order to solve control and optimization problems arising in smart grid applications.* In line with this goal, the key questions this thesis aims to answer are:

- *How to create layers of abstractions in a coordinated optimization and control framework, in order to simplify the task of the coordinator.*
- *How to ensure the safety of the coordinated agents such that the actions and demands of the coordinator does not put the agents in risk.*
- *How can the developed methods and algorithms be applied efficiently in practice, in terms of computation and communication effort.*

1.3 Summary and contributions

Hierarchical control with tracking commitment

In Chapter 2, we consider the generic problem of hierarchical control within a multi agent setting, where the agents track a reference signal that is broadcast by the coordinator. In order to create a layer of abstraction between the coordinator and the agent, we formulate the

robust tracking commitment problem. In this setting, the lower-level controller commits a set of possible reference trajectories over a finite horizon to the coordinator in exchange for a reward corresponding to the size of the reference set and the allowed margin of tracking error. If the commitment is accepted, the lower-level system is required to track any reference trajectory that can be sampled from the committed set. We present the framework of robust tracking commitment and a method to solve the optimal commitment problem for constrained linear systems subject to uncertain disturbance and reference signals. The proposed method allows tractable computations via convex optimization for conic representable uncertainty sets and lends itself to distributed solution methods. We demonstrate the proposed method of tracking commitment in simulation based case studies with commercial buildings that collectively offer frequency regulation service to the power grid.

Parts of the material in this chapter was published in [BGJ17] and the PhD thesis of Tomasz Gorecki [Gor17]. Some parts of [BGJ17] that are specifically due to the work of Tomasz Gorecki has been left out, related to information structures and nonlinear policies. The reformulation of linear adjustable robust programs to incorporate modification of uncertainty sets as explained in 2.2.1 is due to the author of this thesis, but the development of this idea and rest of the material was equally contributed by all authors. The material in Section 2.4 is partially composed from the papers [BGJ17] and [NBG+16]. The case study in 2.4.3 was not published previously.

Related publications:

[BGJ17] A. Bitlislioğlu, T.T. Gorecki, and C.N. Jones. Robust Tracking Commitment. *IEEE Transactions on Automatic Control*, 2017

[GBSJ15] T.T. Gorecki, A. Bitlislioğlu, G. Stathopoulos, and C.N. Jones. Guaranteeing input tracking for constrained systems: theory and application to demand response. In *American Control Conference (ACC), 2015*, pages 232–237. IEEE, 2015

[NBG+16] T.X. Nghiem, A. Bitlislioğlu, T.T. Gorecki, F.A. Qureshi, and C.N. Jones. OpenBuildNet framework for distributed co-simulation of smart energy systems. In *2016 14th International Conference on Control, Automation, Robotics and Vision (ICARCV)*, pages 1–6, November 2016

[FGQ+18] L. Fabietti, T.T. Gorecki, F.A. Qureshi, A. Bitlislioğlu, I. Lymperopoulos, and C.N. Jones. Experimental Implementation of Frequency Regulation Services Using Commercial Buildings. *IEEE Transactions on Smart Grid*, 9(3):1657–1666, May 2018

Coordinated Control via Interior Point Decomposition

In Chapter 3, we consider another layer of abstraction between the agents and the coordinator. Instead of using feasible sets as in the tracking commitment methodology, which can be compu-

tationally demanding in an iterative framework, we rely on quantifying local approximations of the cost and constraints of the agents in the output domain that is relevant to the coordination task. The approximations can be constructed using gradient and curvature information, which carries information related to both the cost and the constraints of the agents. Using barrier functions to replace constraints allows embedding information of constraints into the gradient and curvature information, while guaranteeing the satisfaction of the constraints. Therefore we consider decomposition of interior-point methods and analyze the arising algorithms in a multi-agent framework.

Coordinated Interior Point Method

In Section 3.3, first we present the application of the coordinated interior-point method (C-IP), which was originally formulated for stochastic programming, to optimization problems involving multiple agents that are coupled through constraints and objectives. C-IP eliminates the need to communicate local constraints and cost functions for all variables that relate to internal dynamics and objectives of the agents. Instead, by using embedded barrier functions, the problem is solved in the space of coupling variables, which are in general much lower in dimension compared to the internal variables of individual agents. Therefore, C-IP contributes to both problem size reduction as well as data hiding. The method is a distributed version of the primal barrier method, with locally and globally feasible iterations and faster convergence compared to first-order distributed optimization methods. Hence, C-IP is suitable for early termination in time-critical applications. We illustrate these attractive properties of the C-IP method with a distributed Model Predictive Control (MPC) application in the context of smart-grids, where a collection of commercial buildings provide voltage support to a distribution grid operator.

Coordinated primal-dual Interior point method

In Section 3.4, we consider the extension of the C-IP method to a primal-dual setting. In practical primal-dual interior point (PDIP) implementations, predictor-corrector variants are widely used, due to their very fast convergence. We show that in the coordinated case, a naive implementation of a coordinated (C)-PDIP with predictor-corrector scheme introduces extra communication steps between local and central agents. We propose a decentralized predictor-corrector scheme that uses a non-uniform perturbation on the complementary slackness condition, which is able to reduce the number of communication steps while preserving fast convergence of the original methods. We analyze the general framework of PDIP methods with non-uniform perturbations and provide convergence and complexity results, that directly apply to the proposed coordinated PDIP with decentralized predictor-corrector scheme.

The material in this chapter was published in papers [BPJ17b, BJ17a] and the technical report [BJ17b]. In [BPJ17b] the main contributor is the author of this thesis, with minor support from

Ivan Pejcic.

Related publications

[BPJ17b] A. Bitlislioğlu, I. Pejcic, and C.N. Jones. Interior Point Decomposition for Multi-Agent Optimization. In *20th IFAC World Congress*, 2017

[BJ17a] A. Bitlislioğlu and C.N. Jones. On Coordinated Primal-Dual Interior-Point Methods for Multi-Agent Optimization. In *56th IEEE Conference on Decision and Control*, 2017

[BJ17b] A. Bitlislioğlu and C.N. Jones. On generalized primal-dual interior-point methods with non-uniform complementarity perturbations for quadratic programming. Technical Report 231144, École Polytechnique Fédérale de Lausanne, September 2017

Coordinated Economic MPC with barrier functions

In Chapter 4, we consider the use of barrier functions in economic model predictive control (EMPC) for large scale problems in multi-agent settings. Well known benefits of using barrier functions are providing smoothing of the constrained problem, allowing the use of second-order methods and warm-starting, which reduce the iteration count significantly in closed loop operation, allowing efficient implementation of decomposed interior point methods for large scale MPC problems. Apart from the numerical benefits, barrier functions can be used as a regularizing cost in the EMPC problem for enhancing closed loop convergence properties. We show that in the case of EMPC with generalized terminal constraints, which allows the terminal state to be any equilibrium point, regularizing the problem provides (i) convergence of the predicted terminal state to a neighborhood of the overall optimal equilibrium point, (ii) asymptotic average performance guarantees for the closed-loop system and (iii) accelerated numerical solution of the optimal control problem. Specifically we use a proximal-like regularization, which penalizes the deviation from the previously predicted trajectories. We analyze system theoretical properties of the proposed scheme, and provide a simulation example illustrating the numerical and system theoretical benefits of using barriers, in the case of distributed frequency control for an interconnected wide-area power network.

The material in this chapter was submitted for publication as [BJ18].

Related publications

[BJ18] A. Bitlislioğlu and C.N. Jones. Regularized economic model predictive control with barrier functions. Technical report, École Polytechnique Fédérale de Lausanne, June 2018

1.4 Notation

$\mathbb{R}^n$ denotes the Euclidean space of dimension n, $\mathbb{Z}$ denotes the set of integers and $\mathbb{N}$ denotes the set of natural numbers. $\mathbb{R}^n_+$ represents the non-negative orthant of dimension n. For two integers $i \in \mathbb{Z}$ and $j \in \mathbb{Z}$ such that $i < j$, let $\mathbb{Z}_{[i,j]} := \{i, i+1, \dots, j\}$. I_n denotes the identity matrix of dimension n and $\otimes$ denotes the Kronecker product. For a matrix $\mathbf{M} \in \mathbb{R}^{n \times m}$, an integer $i \in \mathbb{Z}_{[1,n]}$ and a set $\mathcal{J} \subseteq \mathbb{Z}_{[1,m]}$, $\mathbf{M}(i, \mathcal{J})$ indicates the set of components that belong to the ith row and columns for which the indices belong to $\mathcal{J}$. For a set $\mathcal{Q} \subseteq \mathbb{R}^n \times \mathbb{R}^m$, the orthogonal projection operator is defined as $\text{Proj}_x(\mathcal{Q}) := \{x \in \mathbb{R}^n | \ \exists y \in \mathbb{R}^m, (x, y) \in \mathcal{Q}\}$. Given two functions $f : \mathbb{R}^n \to \mathbb{R}^m$ and $g : \mathbb{R}^m \to \mathbb{R}^l$, $f \circ g : \mathbb{R}^n \to \mathbb{R}^l$ denotes the composition of f and g , such that $f \circ g(x) = f(g(x))$. The operator diag(.) constructs a diagonal matrix from the input vector, for $x \in \mathbb{R}^n$ we have $\text{diag}(x) \in \mathbb{R}^{n \times n}$, $\text{diag}(x)(i, i) = x(i)$ and the remaining elements are equal to zero.

If a variable is subscripted or superscripted, as $y^i \in \mathbb{R}^{m_i}$, for $i \in \mathbb{Z}_{[1,n]}$, we define the stacked vector $\boldsymbol{y} \in \mathbb{R}^l$ as $\boldsymbol{y} := [{y^1}^T, \cdots, {y^n}^T]^T$, with $l = \sum_{i=1}^n m_i$ and T indicating the transposition. Similarly if a variable is indexed as $y(k) \in \mathbb{R}^{n/N}$ for $k \in [1, \dots, N]$, $y \in \mathbb{R}^n$ represents $[y(1)^T, \cdots, y(N)^T]^T$. Superscript indexing is used to specify local variables among agents, whereas parenthesis indexing is used to specify time steps in the MPC formulation.

For a function with two arguments, as in $f(x, y)$ with $f : \mathbb{R}^{n_x \times n_y} \to \mathbb{R}$, $\nabla_x f(\bar{x}, \bar{y})$ represents the gradient vector, $\partial_x f(\bar{x}, \bar{y})$ represents the sub-differential and $\nabla_{xx} f(\bar{x}, \bar{y})$ represents the Hessian matrix with respect to the variable in the subscript, evaluated at $(\bar{x}, \bar{y})$. For a function with a single argument, we omit the subscript. For a point $x \in \mathbb{R}^{n_x}$ and a positive scalar δ, $\mathcal{B}(x, \delta)$ denotes the 2-norm ball centered at x, $\mathcal{B}(x, \delta) := \{z \in \mathbb{R}^{n_x} \ : \ \|x - z\|_2 \leq \delta\}$.

The following definitions are adopted from [RW09]. A function $f : \mathbb{R}^n \to \mathbb{R} \cup \{\infty\}$ is *proper* if for at least one $x \in \mathbb{R}$, $f(x) < \infty$. f is *lower semi-continuous* (l.s.c) iff the level sets, $\{x \in \mathbb{R}^n | f(x) \leq a\}$, are all closed in $\mathbb{R}^n$. If the level sets are bounded for all $a \in \mathbb{R}$, then the function is also *level-bounded*. For a set $C \subset \mathbb{R}^n$, relint(C) denotes the relative interior of C which is the interior of C relative to its affine hull, that is the smallest affine set which includes C.

A directed graph is represented as $\mathcal{G} = (\mathcal{V}, \mathcal{E})$ where $\mathcal{V} \in \mathbb{N}$ is the set of vertices, describing each node, and $\mathcal{E} \in \mathbb{N} \times \mathbb{N}$ is the set of ordered pairs of vertices describing directed edges. An edge (i, j) is directed as $i \longleftarrow j$. For each node i, $\mathcal{N}^{-i} := \{j \in \mathbb{N} \mid (j, i) \in \mathcal{E}\}$ represents the set of neighbors that are affected by node i, and $\mathcal{N}^i := \{j \in \mathbb{N} \mid (i, j) \in \mathcal{E}\}$ represents the set of neighbors that affect the node i.

List of Abbreviations

HVAC ..... Heating Ventilation and Air Conditioning

BAS ..... Building Automation System

DR ..... Demand Response

AGC ..... Area Generation Control

PID ..... Proportional-Integral-Derivative Control

MPC ..... Model Predictive Control

EMPC ..... Economic MPC

G-EMPC ..... EMPC with terminal equilibrium constraints

IP ..... Interior-Point

PDIP ..... Primal-Dual Interior-Point

C-(PD)IP ..... Coordinated (PD)IP

PC ..... Predictor-Corrector

PF ..... Path Following

ARO ..... Adjustable Robust Optimization

LP ..... Linear Program

QP ..... Quadratic Program

SOCP ..... Second-Order Cone Program

2 Hierarchical Control via Tracking Commitment

2.1 Introduction

Regulation of large and complex systems that include many agents is usually handled with several control layers that interact in a hierarchical fashion for breaking down the complex control task into simpler sub-tasks. Frequency regulation of the power grid is a good example of such a setting, where an upper level controller, run by the grid operator, sends reference signals to the lower level subsystems, reserve providers, which are expected to track the reference within an acceptable error bound. For the upper level controller, it is crucial to know the tracking capability of the subsystems. In the power grid, this information is provided to the grid operator in terms of a reserve commitment, which represents the set of robustly trackable power generation/consumption trajectories by the lower level subsystem for a specific time window [RKTR07]. The assessment of the optimal reserve capacity of a large scale complex subsystem, such as a commercial building, is far from trivial. The difficulty comes from large dimensional models that limits available computational methods and the fact that the controlled system is restrained by its primary objectives and external disturbances such as weather conditions and occupancy levels.

In this part of the thesis, motivated by the reserve commitment problem for the power grid, we formulate the *robust tracking commitment problem* for constrained linear systems. The problem corresponds to the assessment of a subsystem's tracking capability for a finite horizon and optimal commitment of this capability to an upper control layer. Of course, the subsystem's overall cost of operation and feasibility have to be taken into account in this assessment. The tracking capability can be expressed in terms of a set of possible reference trajectories the system can track robustly. By robustness, we mean that the subsystem will not violate internal constraints and always stay close to the revealed reference trajectory within an allowed error set, while being subjected to external disturbances.

For a fixed reference set over a finite horizon, considering that both the reference and disturbance are external uncertainty sources, the ability of the system to track any reference from this set can be assessed by means of a standard robust optimization problem [GBSJ15, MRSVC14], given that the system is linear, constraint and uncertainty sets are polytopic. The peculiarity of the commitment problem comes from the freedom in the choice of uncertainty sets that are admissible for robust tracking, since it is up to the subsystem to determine the size of the uncertainty set it can handle. In other words, we are faced with a robust optimization problem, in which the size of the uncertainty set is also a decision variable.

Finite horizon robust control for linear systems is well established in the model predictive control (MPC) literature [Rak12]. However, the related work is mainly concerned with obtaining a control policy that guarantees robust feasibility and stability under a given uncertainty set. Available methods mainly rely on choosing a nominal trajectory and a control policy that will keep the system around the nominal trajectory under the effect of the uncertainty [MSR05, RKFC12, RKC+12, GKM06]. The use of a closed-loop control policy reduces the conservatism significantly compared to open-loop robust policies.

Among many available robust MPC methods, we build on the affine disturbance feedback method [GKM06] which is a translation of the work of [BTEGN09] to the MPC framework. The method can be directly applied to conic representable uncertainty sets over the whole prediction horizon, capturing the temporal correlation of the uncertain signal while eliminating the need of enumerating extreme trajectories. Affine disturbance feedback policies are shown to be equivalent to time-varying state feedback policies [GKM06, Rak12] and the method allows online optimization of the policy via convex optimization, rather than working with an a priori fixed policy. Other tube MPC methods that can supersede the disturbance feedback approach [RKFC12], [RKC+12], rely heavily on the assumption of uncorrelated realization of the uncertainty at each time step. This assumption restricts the family of trajectory sets that can be used to characterize the inherent tracking capability of the system. In our work, we take into account the temporal correlation of the uncertainty, especially for the reference signal, and hence consider a more expressive family of sets for characterizing the inherent capability of the system. We demonstrate the advantage of this approach in the simulation study in Chapter 2.4. Considering temporal correlation for uncertainty modeling is also found to be beneficial in the context of the multistage economic dispatch problem [LS15]. In addition, the machinery of affine, or equivalently linear, policies for treatment of uncertainty can be also used to obtain more generic nonlinear policies [GWK14]. The work of [GWK14] generalizes linear policies by considering lifted uncertainty sets under a nonlinear operator, a possibility also shown by [BTEGN09].

An illustrative example of robust model predictive control applied to the operation of the power grid can be found in [WGMM13], where the authors allocate reserves while considering temporal correlation of the demand-generation forecast and assuming the forecast error to belong to a polytopic set defined over a finite prediction horizon. However, the uncertainty set is fixed prior to allocation, therefore the authors do not address the problem of assessing disturbance rejection capabilities of a given reserve fleet. The *commitment problem*, on the other hand, requires searching over uncertainty sets that the system can accommodate rather than guaranteeing robustness against a fixed uncertainty set. The interest in this problem has peaked in recent years due to the prospect of demand response applications[OBB+13], where the reference set to be tracked that represents the reserve capacity of the system is not given, but should be computed by the service provider. In [MRSVC14] the problem is posed in the robust model predictive control context, and simple up-down flexibility of a single actuator is optimized. [VOAZ14] and [BOCT14] consider aggregation of several subsystems to track a reference signal and optimize maximum up-down limits on the reference, however the robust formulation is again limited either to single dedicated actuators or predetermined schemes that distribute the required change in the total power consumption among actuators. [ZKGL14] considers reference sets that are norm balls and optimizes over linear mappings to modify the uncertain reference set utilizing dual norm formulations. However, in all aforementioned works on demand response, the authors do not consider the temporal correlation that will be present in the uncertain reference signal. This problem is tackled in our work [GBSJ15], where we consider polytopic uncertainty sets defined over the whole prediction horizon and formulate

a method for modification of the uncertainty set via linear maps. A similar work is [VOA16], where the authors propose optimizing over a linear map to be applied to a polytopic reference set that represents energy constraints in frequency regulation signals.

Other works have tackled similar problems in more specific contexts. [HSPV14] and [ZZ16] propose aggregation methods for characterizing the power consumption flexibility of a collection of thermostatically controlled loads (TCLs). This can be considered as a particular case of the commitment problem and can be addressed with the methodology proposed in this thesis, that can also handle more general systems. Another related work is [NNPPV14] where the problem of serving a set of time constrained load requests (such as electric vehicle charging) is tackled. Sufficient conditions are given for a supply profile to be able to serve the loads and a dispatch strategy is proposed. This work differs from ours by considering a continuous time formulation and again using the specific structure of the problem to derive a solution.

A connected problem is the *output regulation*, which deals with the capability of the system to track a reference trajectory that is generated by an external dynamical system [Fra77]. In the finite horizon framework, the external system serves as a generator for the reference trajectory set. Most of the work in output regulation deals with asymptotic tracking guarantees [Fra77], [MM10]. Similarly, for systems under additive disturbance, the authors of [LAAC10] show robust convergence to a neighborhood of a fixed reference that is allowed to change occasionally. However, within the finite horizon tracking commitment context, it is necessary to guarantee tracking during the whole commitment period, which includes transients. Therefore, similar to [FM13], we consider tracking requirements during the whole horizon, including the transients, rather than showing convergence to a neighborhood of a constant reference. In this direction, the authors of [DCB13a] utilize robust invariant sets to guarantee tracking with specified error bounds during and after the finite prediction horizon. The guarantees are sought for a given reference generator under the assumption that there exists a feasible solution to the problem. However, none of the aforementioned works in guaranteed tracking or output regulation framework consider the problem of modulating the reference set while solving the control problem.

The idea of modulating initial sets for inner approximation of admissible sets is used in [RKFC12]. Namely the authors develop a method to partially design control policies online, by using approximations of state and input tubes, which overcomes the difficulty of designing fixed local controllers in the regular tube MPC [MSR05]. Even though both the framework and the methodology are significantly different than our case, we take a standpoint similar to that of [RKFC12], in the sense of applying modifier functions to fixed shaped simple sets for recovering inner approximations of more complex sets.

The robust commitment problem can be classified as a generalized semi-infinite program [GVRSS08] due to the variability of the uncertainty set. This kind of problems is in general more difficult to solve than semi-infinite programs, that is robust optimization problems with a fixed uncertainty set [Sti04]. The related literature is mostly concerned with generic nonlinear

programs and global optimization methods, that can be computationally demanding [Sti04, MT14, HB16, DHSS12]. Within the framework of linear systems and polytopic state and input constraints, we reformulate the robust tracking commitment problem such that it can be approximated by a simpler robust optimization problem. The novelty of our approach comes from the notion of modifier functions that are used to modulate uncertainty sets, which enables reformulation of the commitment problem into a standard multistage adjustable robust program that can be solved efficiently [BTEGN09]. The solution method is based on linear decision rules [BTEGN09, GKM06] together with linear modifier functions over conic representable sets and can be used for sizing of reference sets while simultaneously guaranteeing robust tracking.

In the following, we formulate a general robust tracking commitment problem for linear systems that allows practical computation of optimal reference sets to be committed for guaranteed tracking. The contributions are the formulation of the robust tracking commitment framework and enhancement of similar formulations presented in [GBSJ15, VOA16, ZKGL14, ZKG+17] in three ways. First we incorporate generic output tracking, where the output is a function of both input and the state of the system, rather than just the input. Second we consider additional disturbances acting on the system as another source of uncertainty and allow for inexact tracking with certified error bounds. Finally, we establish sufficient conditions for modifier functions and control policies on uncertainty sets, that ensure causal admissibility for robust tracking according to the available information on the uncertainty. Finally, we illustrate the results with simulation studies.

2.2 Preliminaries

2.2.1 Robust optimization

Consider the optimization problem

$$\begin{aligned} \min_{u} \quad & f(u) \\ \text{s. t.} \quad & g(u, \xi) \leq 0, \forall \xi \in \Xi \end{aligned} \tag{2.1}$$

where $f : \mathbb{R}^{n_u} \rightarrow \mathbb{R}$ is the cost function, $g : \mathbb{R}^{n_u \times n_\xi} \rightarrow \mathbb{R}$ is the constraint function, $u \in \mathbb{R}^{n_u}$ is the decision and $\xi \in \mathbb{R}^{n_\xi}$ represents the uncertainty in the problem. The specific value of ξ is not known when we solve the problem, but it is known that ξ can take any value within the set $\Xi \subset \mathbb{R}^{n_\xi}$, that is assumed to be bounded. The above optimization problem is called a 'robust' optimization problem, since any choice of u has to satisfy the constraints robustly, in the sense that the constraint satisfaction is guaranteed for any possible value of ξ. In general the constraint function g can be vector valued, representing a series of constraints. However, in robust programming each constraint should be robustified independently as the worst-case realization for each constrained value can be a different sample from the uncertainty set Ξ [BTEGN09]. Therefore, we will discuss the case with a single constraint without loss of generality.

A possible approach of solving (2.2) is to incorporate samples of the constraint $g(u, \xi) \leq 0$ by enumerating all possible values of ξ. This is only possible if the cardinality of the set Ξ is finite. If Ξ has infinite cardinality, then the robust optimization problem falls into the generic category of semi-infinite programming, since there are infinite number of constraints to satisfy [GVRSS08].

One way of handling the semi-infinite nature of the problem is to solve an approximate problem that has a finite number of constraints. The most straightforward way is to construct a subset $\Xi_s \subset \Xi$ by generating samples or 'scenarios' from the set Ξ [SDR09]. This approach does not guarantee constraint satisfaction strictly. However, under convexity assumptions, the probability of constraint violation can be quantified as a function of the number of samples [CGP09].

A more rigorous way of handling the semi-infinite nature of the constraints is to consider the 'worst-case' scenario. The 'robust-counterpart' of the problem (2.2) can be written as

$$\begin{aligned} \min_{u} \quad & f(u) \\ \text{s. t.} \quad & \phi(u) \leq 0 \end{aligned} \tag{2.2}$$

with the constraint function

$$\phi(u) = \max_{\xi \in \Xi} g(u, \xi)$$

The main goal of robust programming research is to find cases when the function $\phi(u)$ can be treated efficiently [BTEGN09]. This is possible if $\phi(u)$ has a closed form or the maximization problem can be replaced with its dual to obtain an alternative constraint of the form

$$\min_{z \in \mathbb{L}} d(u, z) \leq 0$$

where z is the (possibly vector valued) dual variable of the constraint $\xi \in \Xi$ and $\mathbb{L}$ is the set of feasible dual variables, the formulation of which depends on the description of the set Ξ. In this case the minimization is not necessary for feasibility, since if $d(u, z) \leq 0$ for some feasible u and z then the minimum is guaranteed to be non-positive. With this 'trick' the semi-infinite problem can be formulated as a finite optimization problem

$$\begin{aligned} \min_{u,z} \quad & f(u) \\ \text{s. t.} \quad & d(u, z) \leq 0 \\ & z \in \mathbb{L} \end{aligned} \tag{2.3}$$

In the following section we will introduce some cases where it is possible to apply the dualization 'trick' to re-formulate a robust program so that the resulting problem can be solved efficiently.

Tractable robust programs

If certain properties are satisfied by the problem data, a semi-infinite robust program can be formulated into a standard optimization problem that can be solved efficiently. Specifically, if the cost function f and the constraint function g are convex functions, we obtain a convex robust program. However a convex problem with infinitely many constraints is still not tractable. Therefore we will rely on stricter assumptions on g and the uncertainty set Ξ. The results presented here are mostly adopted from [BTEGN09] and a comprehensive tutorial can be found in [GYdH15].

We will rely on the following assumptions on the problem data.

- The cost function f is convex
- The constraint function g is linear in x and ξ
- The uncertainty set is Ξ is representable as the intersection of convex cones

$$\hat{\Xi} = \left\{ \xi \mid F_i \xi + f_i \in \mathbf{K}_i, \ \ i \in \mathbb{Z}_{[1,m]} \right\} \tag{2.4}$$

  where the cone $\mathbf{K}_i$ is proper.

Note that the considered class of uncertainty sets is very extensive, as it allows the description of well known cones such as the non-negative orthant, the second-order cone and the positive

semi-definite cone as well as their intersections and products. A simple example is a polytopic constraint set that can be written as

$$Fw + f \in \mathbb{R}^n_+ .$$

for some $f \in \mathbb{R}^n$ and $F \in \mathbb{R}^{n_\xi \times n}$. For simplicity of development, we will consider a single constraint and a single cone for the uncertainty set. Under these assumptions we can write the optimization problem as

$$\begin{aligned} \min_{u} \quad & f(u) \\ \text{s. t.} \quad & \forall \xi |\ F\xi + f \in \mathbf{K} \\ & a^T u + b^T \xi + c \leq 0, \end{aligned} \tag{2.5}$$

We utilize strong duality of convex conic programs [BV04] to re-formulate the robust constraint as

$$\begin{aligned} \max_{\xi} \quad & a^T u + b^T \xi + c \\ \text{s. t.} \quad & F\xi + f \in \mathbf{K}, \end{aligned} \iff \begin{aligned} \min_{z} \quad & a^T u + z^T f + c \\ \text{s. t.} \quad & F^T z = -b, \quad z \in \mathbf{K}^* \end{aligned} \tag{2.6}$$

where $z \in \mathbb{R}^n$ is the vector of Lagrange multipliers for the constraint set and $\mathbf{K}^*$ is the dual cone for $\mathbf{K}$. Removing the 'min' operator as discussed before, the robust problem (2.37) can be re-formulated as a standard convex program

$$\begin{aligned} \min_{u,z} \quad & f(u) \\ \text{s. t.} \quad & a^T u + z^T f + c \leq 0, \qquad F^T z = -b \\ & z \in \mathbf{K}^* \end{aligned} \tag{2.7}$$

In case of multiple constraints, each constraint should be robustified independently even if they are affected by the same uncertain variable. For the case of polytopic constraints of type $Ax + B\xi + c \leq 0$, describing m constraints, we can write the robust counterpart as

$$\begin{aligned} \min_{u,z} \quad & f(u) \\ \text{s. t.} \quad & Au + Zf + c \leq 0 \\ & ZF = -B \\ & Z \in \mathbf{K}^* \end{aligned} \tag{2.8}$$

by stacking the dual variables for each constraint as rows of the matrix $Z \in \mathbb{R}^{m \times n}$ with $Z \in \mathbf{K}^*$ implying row-wise membership.

Adjustable robust programs

In some robust optimization problems a part of the decision can be taken after the realization of the uncertainty, introducing a causality link between the uncertain variable and the decision. These type of problems can be categorized under 'two-stage' or 'multi-stage' problems, where

we optimize over a sequence of decisions, as in the case of MPC. The generic adjustable robust optimization (ARO) problem [BTEGN09] can be written as

$$\begin{aligned} \min_{u,\pi(.)} \quad & f(u) \\ \text{s. t.} \quad & g(u, \pi(\xi), \xi) \leq 0, \quad \forall \xi \in \Xi \end{aligned} \tag{2.9}$$

where the decision $\pi(.)$ is a 'policy' that is a function of the realization of the uncertain variable. The adjustable nature of the decision brings in the possibility of compensating the effect of the uncertainty and results in a more realistic solution. However, solving the ARO is even more challenging, as we need to optimize over functions. As a work around, the policy can be formulated as a finite decision, by introducing a parameterization of the policy.

$$\begin{aligned} \min_{u,\theta} \quad & f(u) \\ \text{s. t.} \quad & g(u, \pi(\theta, \xi), \xi) \leq 0, \quad \forall \xi \in \Xi \end{aligned} \tag{2.10}$$

Again restricting the properties of the problem data will enable us to solve ARO efficiently, using the methodology from the previous section. Consider the adjustable problem with convex cost and linear robust constraints.

$$\begin{aligned} \min_{u} \quad & f(u) \\ \text{s. t.} \quad & \forall \xi |\ F\xi + f \in \mathbf{K} \\ & A(u + \pi(\xi)) + B\xi + c \leq 0, \end{aligned} \tag{2.11}$$

In this case, limiting the space of policies to linear functions as

$$\pi(\xi) = M\xi$$

and optimizing over the linear policy M enables us to reformulate the problem as a standard convex program [BTEGN09].

$$\begin{aligned} \min_{u,z} \quad & f(u) \\ \text{s. t.} \quad & Au + Zf + c \leq 0 \\ & ZF = -(B + AM) \\ & Z \in \mathbf{K}^* \end{aligned} \tag{2.12}$$

The formulation (2.12) is restricted to linear functions, but in certain cases it is possible to deal with nonlinear policies in a computationally tractable manner. The key principle, introduced in [BTEGN09] and studied in greater detail in [GWK14] is to consider a modified uncertainty set which is the image of the original uncertainty under a nonlinear lifting. If the lifted uncertainty set or its convex hull can be represented in the conic form of (2.64), the machinery of linear policies can be applied. The combination of the nonlinear lifting and linear policy results in a nonlinear policy. For details on this subject, see [GWK14]

Modifying uncertainty sets

The following material is one of the contributions of the thesis and is adapted from [GBSJ15, BGJ17]. We note here that similar methods were developed simultaneously by other groups, see [ZKGL14, VOA16, ZKG+17].

The basic robust program (2.2) can be considered as a special case of a more generic formulation, which is called a generalized semi-infinite program [GVRSS08]. Consider the optimization problem

$$\begin{aligned} \min_{u} \quad & f(u) \\ \text{s. t.} \quad & g(u, \xi) \leq 0, \forall \xi \in \Xi(u) \end{aligned} \tag{2.13}$$

The generalization arises due to the fact that the uncertainty set is characterized as a set valued map and therefore the uncertainty set takes shape according to the decision variable. Similar to the previous analysis, it is possible to obtain tractable versions of this problem if we limit the scope of parameterization of the uncertainty set.

The set valued map $\Xi(u)$ can be equivalently formulated by using modifier functions. We first define the *uncertainty modifier* function $\nu : \mathbb{R}^{Nn_\xi} \longrightarrow \mathbb{R}^{Nn_\xi}$, which is assumed to be bijective and used for reshaping a given uncertainty set.

$$\nu(\hat{\Xi}) = \{\xi = \nu(\hat{\xi}) : \hat{\xi} \in \hat{\Xi}\} \tag{2.14}$$

Similar to the adjustable case with policies (2.9), it is necessary to parameterize the modifier functions to obtain a finite dimensional formulation

$$\nu(\theta, \hat{\Xi}) = \{\xi = \nu(\theta, \hat{\xi}) : \hat{\xi} \in \hat{\Xi}\} \tag{2.15}$$

With the parameterization of the modifier function, the robust optimization problem with modified uncertainty sets can be written as

$$\begin{aligned} \min_{u,\theta} \quad & f(u, \theta) \\ \text{s. t.} \quad & g(u, \nu(\theta, \hat{\xi})) \leq 0, \forall \hat{\xi} \in \hat{\Xi} \end{aligned} \tag{2.16}$$

Computationally tractable cases of 2.16 rely on linearity assumptions as before. We define a set valued linear map as

$$L\hat{\Xi} = \{\xi = L\hat{\xi} \mid \hat{\xi} \in \hat{\Xi}\} \tag{2.17}$$

For a conic set $\hat{\Xi} = \left\{ \hat{\xi} \mid F\hat{\xi} + f \in \mathbf{K} \right\}$ we have

$$L\hat{\Xi} = \{\xi = L\hat{\xi} \mid F\hat{\xi} + f \in \mathbf{K}\} . \tag{2.18}$$

If the matrix L is invertible, we can simply write

$$L\hat{\Xi} = \{\xi \mid FL^{-1}\xi + f \in \mathbf{K}\} . \tag{2.19}$$

By using linear (or affine) maps on the uncertainty set, we can formulate a generalized robust program

$$\begin{aligned} \min_{u,L} \quad & f(u, L) \\ \text{s. t.} \quad & \forall \xi \in L\hat{\Xi} \\ & Au + B\xi + c \leq 0, \end{aligned} \tag{2.20}$$

for which the robust constraint can be re-formulated as

$$Au + BL\hat{\xi} + c \leq 0, \quad \forall \hat{\xi} \in \hat{\Xi} . \tag{2.21}$$

Note that this constraint has a form which can be dualized to obtain a standard convex program.

$$\begin{aligned} \min_{u,L,Z} \quad & f(u, L) \\ \text{s. t.} \quad & Au + Zf + c \leq 0 \\ & ZF = -BL \\ & Z \in \mathbf{K}^* \end{aligned} \tag{2.22}$$

We can also incorporate adjustable decisions by directly finding a policy on the unmodified uncertainty variable $y(\hat{\xi}) = \hat{M}\hat{\xi}$. The actual policy on the uncertain variable ξ can be recovered as

$$\pi(\xi) = \hat{M}L^{-1}\xi . \tag{2.23}$$

2.2.2 Model predictive control

We consider the discrete-time dynamical system with state transition map $f : \mathbb{R}^{n_x \times n_u} \longrightarrow \mathbb{R}^{n_x}$, state $x \in \mathbb{R}^{n_x}$ and input $u \in \mathbb{R}^{n_u}$, described by

$$x^+ = f(x, u) \tag{2.24}$$

where x^+ is the successor state, the state at the next time step.

In model predictive control (MPC) we use predicted trajectories of the system to find optimal input sequences, with respect to some objective. Following the notation in [Gr3], we define predicted trajectories as follows. Starting from an initial state $x \in \mathbb{X}$, for an input sequence of N steps $\boldsymbol{u} := \{u(0), u(1), \dots, u(N-1)\}$, N step predicted trajectory is given by the solution to the difference equations (2.24),

$$\boldsymbol{x_u}(x) := \{x_{\boldsymbol{u}}(0|x), x_{\boldsymbol{u}}(1|x), \dots, x_{\boldsymbol{u}}(N|x)\} \tag{2.25}$$

When the dependency on initial state x is evident from the context and we would like to emphasize dependency on time, we will simply write $x_{\boldsymbol{u}}(k|t) := x_{\boldsymbol{u}}(k|x(t))$.

The state and inputs are constrained within admissible compact sets,

$$x \in \mathbb{X} \subset \mathbb{R}^{n_x}, \quad u \in \mathbb{U} \subset \mathbb{R}^{n_u} . \tag{2.26}$$

Starting from a feasible initial state $x \in \mathbb{X}$, a control sequence $\boldsymbol{u}$ is admissible, if $\boldsymbol{u} \in \mathbb{U}^N$ and the predicted trajectory $\boldsymbol{x_u}(x)$ is feasible, $\boldsymbol{x_u}(x) \in \mathbb{X}^{N+1}$. Thus, the set of admissible input sequences are defined as

$$\mathcal{U}_N(x) = \{\boldsymbol{u} \in \mathbb{U}^N \mid \boldsymbol{x_u}(x) \in \mathbb{X}^{N+1}\} . \tag{2.27}$$

With the notion of admissible trajectories, it is important to characterize the capability of the system, in terms of reachability. To this end, we will use the following definitions

Definition 1 (Reachable set [BM08]). *Given set $\mathcal{P}$, the reachable set $\mathcal{R}_N(\mathcal{P})$ from $\mathcal{P}$ in time $N < \infty$ is the set of all x for which there exists $x_0 \in \mathcal{P}$ and $\boldsymbol{u} \in \mathcal{U}_N(x)$ such that $x_{\boldsymbol{u}}(N|x_0) = x$.*

Definition 2 (Steerable set [BM08]). *Given set $\mathcal{S}$, the steerable set $\mathcal{C}_N(\mathcal{S})$ to $\mathcal{S}$ in time $N < \infty$ is the set of all x_0 for which there exists $\boldsymbol{u} \in \mathcal{U}_N(x)$ such that $x_{\boldsymbol{u}}(N|x_0) \in \mathcal{S}$.*

For a finite horizon length N, $J_N : \mathbb{R}^{Nn_x} \times \mathbb{R}^{Nn_u} \longrightarrow \mathbb{R}$ denotes a cost function over input and state trajectories, which usually emanates from a stage cost

$$J_N(\boldsymbol{x}, \boldsymbol{u}) = \sum_{k=0}^{N-1} l(x(k), u(k)) \tag{2.28}$$

with $l : \mathbb{R}^{n_x} \times \mathbb{R}^{n_u} \longrightarrow \mathbb{R}$.

It is customary to add a terminal cost and terminal constraint to the finite horizon control problem, in order to ensure desired system properties such as stability and recursive feasibility. Finally, the finite horizon optimal control problem can be formulated as

$$\begin{aligned} \mathcal{P}(x) : \min_{\boldsymbol{u}} \quad & J_N(x_{\boldsymbol{u}}(x), \boldsymbol{u}) + \tau(x_{\boldsymbol{u}}(N|x)) \\ \text{s. t.} \quad & \boldsymbol{u} \in \mathcal{U}(x) \end{aligned} \tag{2.29}$$

where $\tau : \mathbb{R}^{n_x} \longrightarrow \mathbb{R}$ is the terminal cost with the optimal input sequence $\boldsymbol{u}^*(x)$ and the associated control policy $\mu(x) := u^*(0|x)$. The resulting closed loop system can be written as

$$x^+ = f(x, \mu(x)) \tag{2.30}$$

See [MRRS00] for a detailed survey on the stability of MPC with proper selection of the stage cost, terminal cost and constraints.

2.2.3 Robust MPC

Consider the linear system

$$x^+ = Ax + Bu + w \tag{2.31}$$

with constrained state, inputs $(x, u) \in \mathbb{X} \times \mathbb{U} \subset \mathbb{R}^{n_x} \times \mathbb{R}^{n_u}$, and disturbance $w \in \mathbb{R}^{n_x}$. The sets $\mathbb{X}$ and $\mathbb{U}$ are assumed to be convex, polytopic and compact. For a finite horizon N, starting from an initial state x at time 0, the input sequence $\boldsymbol{u}$ is applied, and the disturbance sequence $\boldsymbol{w} = (w(0), \dots, w(N-1))$ is observed, resulting in the state trajectory $\boldsymbol{x}_{\boldsymbol{u}}(x, \boldsymbol{w})$.

The idea behind robust MPC is to ensure that the system remains feasible for all possible realizations of the uncertainty. This description already implies that the uncertain variable w has a finite support. We will consider the case where the disturbance trajectory is contained within a convex set

$$\boldsymbol{w} \in \mathcal{W} \subset \mathbb{R}^{n_x} . \tag{2.32}$$

It is common to assume that the disturbance has no correlation in time, that is $\mathcal{W} = \mathbb{W}^N$ for some set $\mathbb{W} \subset \mathbb{R}^{n_x}$ that bounds the disturbance at every stage. However correlation of the disturbance trajectory over time can be captured by a more generic choice of $\mathcal{W}$. With the bounded disturbance assumption, we can define the set of input sequences that are admissible with respect to a initial state and a sequence of disturbance realizations as

$$\mathcal{U}_N(x, \boldsymbol{w}) = \{\boldsymbol{u} \in \mathbb{U}^N \mid \boldsymbol{x}_{\boldsymbol{u}}(x, \boldsymbol{w}) \in \mathbb{X}^{N+1}\} . \tag{2.33}$$

Considering a nominal disturbance scenario $\bar{\boldsymbol{w}} \in \mathcal{W}$ for the objective function, the robust MPC

problem can be formulated as

$$\begin{aligned} \mathcal{P}(x) \ : \min_{\boldsymbol{u}} \quad & J_N(x_{\boldsymbol{u}}(x, \bar{\boldsymbol{w}}), \boldsymbol{u}) + \tau(x_{\boldsymbol{u}}(N|x, \bar{\boldsymbol{w}})) \\ \text{s. t.} \quad & \boldsymbol{u} \in \mathcal{U}(x, \boldsymbol{w}) \quad \forall \boldsymbol{w} \in \mathcal{W} \end{aligned} \tag{2.34}$$

Note that it is possible to consider the worst-case scenario for the cost function, but for simplicity we will stick to the nominal cost - worst case feasibility formulation. This formulation is commonly regarded as 'open-loop' robust MPC, as we robustify a single input sequence, without considering the possibility of observing the disturbance and adjusting the input over-time in a 'closed-loop' manner. As such, this formulation is overly conservative as it does not capture the fact that the problem will be re-solved in the next time step, providing compensation for the realized disturbance. In order to capture this effect, we can consider feedback policies for the input. Some common options are [Jon16]

1. Pre-stabilization: $u(k) = Kx(k) + \bar{u}(k)$, for a pre-determined K, optimizing over the nominal input sequence $\bar{\boldsymbol{u}}$.

2. Linear state feedback: $u(k) = K(k)x(k) + \bar{u}(k)$, optimizing directly the controller gain K and the nominal sequence $\bar{\boldsymbol{u}}$.

3. Disturbance feedback: $u(k) = \sum_{i=0}^{k} M_{ki} w(k) + \bar{u}(k)$ optimizing over the disturbance feedback gains M_{ki} and the nominal sequence $\bar{\boldsymbol{u}}$.

4. Tube MPC: $u(k) = K(x(k) - \bar{x}(k)) + \bar{u}(k)$ optimizing over the setpoint for state feedback $\bar{\boldsymbol{x}}$ and sequence $\bar{\boldsymbol{u}}$ with a pre-determined stabilizing gain K.

Among the available options the most flexible is the linear state feedback, which could be combined with tube MPC, that optimizes over time-varying state feedback policies. This formulation however leads to a non-convex problem. It was shown in [GKM06] that the disturbance feedback is in fact equivalent to linear state feedback and allows a convex formulation since it can be formulated as an adjustable robust program with a linear policy (2.6). We will briefly show how this can be done in the case of a linear system with polytopic feasibility and disturbance sets.

Linear robust MPC with disturbance feedback

The dense form of the system equations (2.38), which describes the evolution of the system for N steps, is given by

$$\boldsymbol{x}_{\boldsymbol{u}}(x, \boldsymbol{w}) = \mathbf{A}x_0 + \mathbf{B}\boldsymbol{u} + \mathbf{E}\boldsymbol{\omega}$$

The matrices $\mathbf{A} \in \mathbb{R}^{(N+1)n_x \times n_x}$, $\mathbf{B} \in \mathbb{R}^{(N+1)n_x \times Nn_u}$ and $\mathbf{E} \in \mathbb{R}^{(N+1)n_x \times Nn_x}$ are defined as

$$\begin{aligned}\mathbf{A} &:= \begin{bmatrix} I_{n_x} \\ A \\ A^2 \\ \vdots \\ A^N \end{bmatrix}, \mathbf{E} := \begin{bmatrix} 0 & \cdots & \cdots & 0 \\ I_{n_x} & 0 & \cdots & \vdots \\ A & I_{n_x} & \cdots & \vdots \\ \vdots & \vdots & \ddots & \vdots \\ A^{N-1} & A^{N-2} & \cdots & I_{n_x} \end{bmatrix} \\ \mathbf{B} &:= \mathbf{E}(I_N \otimes B),\end{aligned}$$

The polytopic state, input constraints and the disturbance set can be described as

$$\begin{aligned}\mathbb{X}^{N+1} &:= \{\boldsymbol{x} \in \mathbb{R}^{(N+1)n_x} : \mathbf{F}_x \boldsymbol{x} \leqslant \mathbf{f}_x\} \\ \mathbb{U}^N &:= \{\boldsymbol{u} \in \mathbb{R}^{Nn_u} : \mathbf{F}_u \boldsymbol{u} \leqslant \mathbf{f}_u\} \\ \mathcal{W} &:= \{\boldsymbol{w} \in \mathbb{R}^{Nn_x} : \mathbf{F}_w \boldsymbol{w} \leqslant \mathbf{f}_w\}\end{aligned}$$

Given these definitions of the problem data, the set of admissible input sequences is decribed by a polytopic set

$$\mathcal{U}(x, \boldsymbol{w}) = \{(\boldsymbol{u}, \boldsymbol{\xi}) \,|\, T\boldsymbol{u} + V\boldsymbol{w} \leqslant h(x)\}$$

where matrices used in the definition are given by

$$T = \begin{bmatrix} [\mathbf{0} \;\; \mathbf{F}_x]\mathbf{B} \\ \mathbf{F}_u \end{bmatrix}, \; V = \begin{bmatrix} [\mathbf{0} \;\; \mathbf{F}_x]\mathbf{E} \\ \mathbf{0} \end{bmatrix}, \; h(x) = \begin{bmatrix} \mathbf{f}_x \\ \mathbf{f}_u \end{bmatrix} - \begin{bmatrix} [\mathbf{0} \;\; \mathbf{F}_x]\mathbf{A} \\ \mathbf{0} \end{bmatrix} x$$

In this setting, the robust MPC problem (2.34) can be written as

$$\begin{aligned}\min_{u} \quad & J_N(x_{\boldsymbol{u}}(x, \bar{\boldsymbol{w}}), \boldsymbol{u}) + \tau(x_{\boldsymbol{u}}(N|x, \bar{\boldsymbol{w}})) \\ \text{s. t.} \quad & \forall \boldsymbol{w} | \; -\mathbf{F}_w \boldsymbol{w} + \mathbf{f}_w \geq 0 \\ & T\boldsymbol{u} + G\boldsymbol{w} \leq h(x),\end{aligned} \tag{2.35}$$

This problem can be posed as an adjustable robust program considering the disturbance feedback formulation. The feedback policy can be written in a compact way as

$$\boldsymbol{u} = \boldsymbol{M}\boldsymbol{w} + \bar{\boldsymbol{u}} \,. \tag{2.36}$$

Note that the matrix $\boldsymbol{M}$ is structured in a block lower triangular manner, complying with the causality of the feedback policy over the disturbance vector $\boldsymbol{w}$. Finally, using the dualization of the adjustable robust program as in (2.12) the robust MPC with disturbance feedback problem

can be equivalently formulated as

$$\begin{aligned} \min_{\bar{\boldsymbol{u}},Z} \quad & J_N(x_{\boldsymbol{u}}(x,\bar{\boldsymbol{w}}),\bar{\boldsymbol{u}}) + \tau(x_{\bar{\boldsymbol{u}}}(N|x,\bar{\boldsymbol{w}})) \\ \text{s. t.} \quad & T\bar{\boldsymbol{u}} + Z\mathbf{f}_w \leq h(x) \\ & Z\mathbf{F}_w = G + T\boldsymbol{M} \\ & Z \leq 0\,. \end{aligned} \tag{2.37}$$

2.3 Constrained tracking under uncertainty

Consider the linear uncertain system

$$\begin{aligned} x^+ &= Ax + Bu + w \\ y &= Cx + Du \end{aligned} \tag{2.38}$$

with constrained state and inputs $(x, u) \in \mathbb{X} \times \mathbb{U} \subset \mathbb{R}^{n_x} \times \mathbb{R}^{n_u}$, disturbance $w \in \mathbb{R}^{n_x}$ and output $y \in \mathbb{R}^{n_y}$. The sets $\mathbb{X}$ and $\mathbb{U}$ are assumed to be convex, polytopic and compact. Similar to the state trajectory, the output of the system over the horizon is denoted by $\boldsymbol{y_u}(x, \boldsymbol{w})$. Since the system is linear, the output trajectory can be written as

$$\boldsymbol{y_u}(x, \boldsymbol{w}) = \mathbf{C}\boldsymbol{x_u}(x, \boldsymbol{w}) + \mathbf{D}\boldsymbol{u}$$

where, the matrices $\mathbf{C} \in \mathbb{R}^{Nn_y \times (N+1)n_x}$ and $\mathbf{D} \in \mathbb{R}^{Nn_y \times Nn_u}$ are defined as

$$\mathbf{C} := [I_N \otimes C \;\; \mathbf{0}], \;\; \mathbf{D} := I_N \otimes D \, .$$

Given a reference signal $r \in \mathbb{R}^{n_y}$ at time k, the tracking error is denoted by $e(k) = r(k) - y(k)$, whereas for a reference sequence $\boldsymbol{r} := (r(0), \dots, r(N-1))$ the corresponding error sequence over the finite horizon is given by

$$\boldsymbol{e} := \boldsymbol{r} - \boldsymbol{y_u}(x, \boldsymbol{w}) \tag{2.39}$$

In order to qualify tracking, we introduce the *tracking error set* $\mathcal{E} \subset \mathbb{R}^{Nn_y}$. The objective of the control action is to maintain the difference between the output of the system and the reference signal, that is $\boldsymbol{e}$, within the set $\mathcal{E}$, which is assumed to be compact. We can now define the set of input, reference and disturbance sequences that satisfy the system and tracking constraints over an N–step horizon

$$\mathcal{Q}(x) := \left\{ (\boldsymbol{u}, \boldsymbol{r}, \boldsymbol{w}) \;\middle|\; \boldsymbol{x_u}(x, \boldsymbol{w}) \in \mathbb{X}^{N+1}, \;\; \boldsymbol{u} \in \mathbb{U}^N, \;\; \boldsymbol{e} \in \mathcal{E} \right\} \tag{2.40}$$

As seen from the definition, the feasibility set $\mathcal{Q}$ is parameterized by the initial condition x of the system. In the following, we will drop the argument of $\mathcal{Q}$ for notational simplicity.

Remark 1. *The reference is only defined over the finite horizon and we do not discuss tracking guarantees for the infinite horizon case. Note that any terminal condition that would ensure robust invariance under a persistent disturbance after the finite horizon can easily be added to the set $\mathcal{Q}(x)$ without changing its structure. At this point, we acknowledge the difficulty of constructing invariant sets for self- correlated or state dependent uncertainties [RKML06]. Therefore we consider tracking obligation only over the finite horizon, and advise using standard robustly invariant sets generated for an uncorrelated disturbance after the finite horizon. The reader is referred to [RKMK07] for computation of positively robust invariant sets under time invariant disturbance sets, to be used as the terminal constraint.*

Our goal is to determine if, starting from a given initial state x, an admissible control policy exists such that the output y of system (2.38) can robustly track a reference signal r, that is all system and tracking constraints can be satisfied, for all possible realizations of the reference sequence $\boldsymbol{r}$ and disturbance sequence $\boldsymbol{w}$ such that

$$\boldsymbol{r} \in \mathcal{R}, \quad \boldsymbol{w} \in \mathcal{W} \tag{2.41}$$

where $\mathcal{R} \subset \mathbb{R}^{Nn_y}$ and $\mathcal{W} \subset \mathbb{R}^{Nn_x}$ represent the *reference set* and the *disturbance set* and are assumed to be compact. If such a control policy exists, then we can conclude that the system is qualified for robust tracking with respect to the bounded sets $\mathcal{R}$ and $\mathcal{W}$. Note that this definition of the reference and disturbance sets over entire sequences allows temporal correlation along the horizon, providing a critical flexibility that will be exploited in the application studied in Section 2.4.

Both the reference $\boldsymbol{r}$ and disturbance $\boldsymbol{w}$ are exogenous uncertain signals for the system. From the point of view of the controller, the main difference between these two is the time they are observed by the controller. This causality condition can be easily incorporated in the control policy to be applied over the horizon as we will show in the following sections. Therefore, we define $\boldsymbol{\xi} = (\boldsymbol{r}, \boldsymbol{w}) \in \mathbb{R}^{n_\xi}$ as the unified uncertain signal. We further define the general uncertainty set as

$$\Xi = \mathcal{R} \times \mathcal{W} \ .$$

Let the map $\pi_k : \mathbb{R}^{Nn_\xi} \longrightarrow \mathbb{R}^{n_u}$ be the control policy to be used at time step k. The control policy sequence for the finite horizon can be defined as

$$\boldsymbol{\pi}(\boldsymbol{\xi}) = (\pi_0(\boldsymbol{\xi}), \dots, \pi_{N-1}(\boldsymbol{\xi})) \tag{2.42}$$

We can now define the set of all admissible finite-horizon control policies mapping from disturbance and reference sequences to input sequences

$$\Delta(\Xi) := \left\{ \boldsymbol{\pi} : \Xi \longrightarrow \mathbb{U}^N \mid \forall \boldsymbol{\xi} \in \Xi, \ (\boldsymbol{\pi}(\boldsymbol{\xi}), \boldsymbol{\xi}) \in \mathcal{Q} \right\} \tag{2.43}$$

Given the feasibility and tracking conditions and the uncertainty set, a controller using a policy that belongs to the set $\Delta(\Xi)$, starting from state x, can keep the tracking error within the set $\mathcal{E}$ throughout the finite horizon for any realization of the disturbance $\boldsymbol{w}$ and the reference $\boldsymbol{r}$. The existence of such control policies is not guaranteed: if the system constraints and the tracking set is too restrictive or the uncertainty sets are too large, then it may not be possible for any controller to satisfy system feasibility and tracking requirements simultaneously.

Let us now characterize the uncertainty sets that allow existence of admissible control policies for tracking. This characterization will be instrumental in the following sections, when we optimize over reference sets that the system can track robustly.

Definition 3. *The set $\Xi \subset \mathbb{R}^{Nn_\xi}$ is admissible for tracking by system* (2.38) *in state x if*

$$\Delta(\Xi) \neq \emptyset \,. \tag{2.44}$$

The following lemma provides more insight into the geometry of the admissibility condition:

Lemma 1. *The set $\Xi \subset \mathbb{R}^{Nn_\xi}$ is admissible for tracking by system* (2.38) *in state x if and only if:*

$$\Xi \subseteq Proj_{\xi}(\mathcal{Q}) \tag{2.45}$$

where $Proj_{\xi}(\mathcal{Q})$ denotes the projection of the set $\mathcal{Q}$ onto the ξ subspace.

Proof. : The proof directly follows from the definition of the projection operator and the definition of $\Delta(\Xi)$. Suppose that $\Xi \subseteq \mathrm{Proj}_{\xi}(\mathcal{Q})$, we have that $\forall \boldsymbol{\xi} \in \Xi,\ \exists \boldsymbol{u} \,:\, (\boldsymbol{u}, \boldsymbol{\xi}) \in \mathcal{Q}$. This indicates the existence of a function $\boldsymbol{\pi}$, which maps every $\boldsymbol{\xi} \in \Xi$ to a feasible $\boldsymbol{u} = \boldsymbol{\pi}(\boldsymbol{\xi})$, such that $(\boldsymbol{\pi}(\boldsymbol{\xi}), \boldsymbol{\xi}) \in \mathcal{Q}$, and hence $\Delta(\Xi) \neq \emptyset$. Conversely, suppose that $\Delta(\Xi) \neq \emptyset$ and let $\boldsymbol{\pi} \in \Delta(\Xi)$. By definition of $\Delta(\Xi)$, $\forall \boldsymbol{\xi} \in \Xi$, $\exists \boldsymbol{u}$ such that $\boldsymbol{u} = \boldsymbol{\pi}(\boldsymbol{\xi})$ and $(\boldsymbol{u}, \boldsymbol{\xi}) \in \mathcal{Q}$, and hence $\Xi \subseteq \mathrm{Proj}_{\xi}(\mathcal{Q})$. □

Lemma 1 also illuminates a method of testing the admissibility of a given uncertainty set for robust tracking, by means of verifying set inclusion.

Information structure of control policies

The control policy $\boldsymbol{\pi}$ should account for the fact that the uncertain exogenous signals are revealed partially to the controller as time progresses. Generally speaking, any decision variable $u(k)$ might depend on a subset of the uncertainty vector $\boldsymbol{\xi}$ and only on this subset. To make this claim more precise, the concept of the information structure of a function f is introduced. The presentation follows concepts from Section 14.2 of [BTEGN09] but adopts a different formulation.

Definition 4. *Let $\mathcal{I}$ be a subset of $\{1, 2, \ldots, n\}$, and*

$$\mathcal{F}(\mathcal{I}) = \{f \,:\, \mathbb{R}^n \longrightarrow \mathbb{R} \mid x_{\mathcal{I}} = \hat{x}_{\mathcal{I}} \Longrightarrow f(x) = f(\hat{x})\} \tag{2.46}$$

where $x_{\mathcal{I}}$ denotes the entries of x defined by the indices of $\mathcal{I}$.
Let $\boldsymbol{\mathcal{I}} = (\mathcal{I}_k)_{k \in \mathbb{Z}_{[1,m]}}$ be a collection of index subsets and

$$\mathcal{F}(\boldsymbol{\mathcal{I}}) = \{f \,:\, \mathbb{R}^n \longrightarrow \mathbb{R}^m,\ f_k \in \mathcal{F}(\mathcal{I}_k)\ \forall k \in \mathbb{Z}_{[1,m]}\} \tag{2.47}$$

If $f \in \mathcal{F}(\boldsymbol{\mathcal{I}})$, then we refer to $\boldsymbol{\mathcal{I}}$ as the information structure of f.

$\mathcal{F}(\mathcal{I})$ denotes the set of real-valued functions that depend only on the input indexed in $\mathcal{I}$. For functions with multiple outputs, the information structure is defined output-wise. $\boldsymbol{\mathcal{I}}$ summarizes

the information structure of the function f: the k^{th} component of f depends only on inputs indexed in $\mathcal{I}_k$. For example, in the robust multi-stage control setting considered here, a typical requirement of the control policy will be causality (also called non-anticipativity) which states that the current control action can depend on observations made in the past; in our notation, this fact translates to: for each stage, every control action can depend on past measurements, so that $\pi_k \in \mathcal{F}(\mathcal{I}_k)$ with $\mathcal{I}_k = \{1, \dots, k-1\}$. Notice here a small abuse of notation in the sense that π_k is a function with values in $\mathbb{R}^{n_u}$, and by $\pi_k \in \mathcal{F}(\mathcal{I}_k)$ we mean that every component of π_k is in $\mathcal{F}(\mathcal{I}_k)$.

Definition 5. *The set $\Xi \subset \mathbb{R}^{N n_\xi}$ is causally admissible for tracking by system* (2.38) *in state x with respect to the information structure $\boldsymbol{\mathcal{I}}$ if*

$$\boldsymbol{\mathcal{F}}(\boldsymbol{\mathcal{I}}) \cap \Delta(\Xi) \neq \emptyset \,. \tag{2.48}$$

In contrast to Definition 3, we now require that the control policy satisfies a particular information structure. For example, the reference trajectory to track will usually be known at the current time step but not the disturbance. It is also possible that the reference is known either partially or totally in advance. In Section 2.4, we will see that our examples may display more complex information structures.

2.3.1 Robust Tracking Commitment

Consider the problem of finding a reference set, such that the combined uncertainty set $\Xi = \mathcal{R} \times \mathcal{W}$, composed of the disturbance and reference, is admissible for robust tracking for system (2.38). The admissible reference set $\mathcal{R}$ is to be committed to an external agent together with the guarantee of robust tracking for a finite horizon. We call this problem, a *robust tracking commitment* problem, the application of which can be found in ancillary service provision to the power grid, as we will demonstrate in Section 2.4, or similar hierarchical control settings.

Essentially, tracking commitment requires evaluation of the admissibility of uncertainty sets, in order to be able to find a suitable reference set to be committed. Based on Definition 5, we can write the family of causally admissible uncertainty sets for tracking with respect to a given information structure

$$\Omega = \left\{ \Xi \subset \mathbb{R}^{N n_\xi} \mid \exists \boldsymbol{\pi} \in \boldsymbol{\mathcal{F}}(\boldsymbol{\mathcal{I}}) \cap \Delta(\Xi) \right\} \tag{2.49}$$

First we tackle the problem of simply finding a causally admissible reference set, without attaching any cost function to the problem. The tracking commitment problem can be written as

$$\text{find } \mathcal{R} \,:\, \mathcal{R} \times \mathcal{W} \in \Omega \tag{2.50}$$

For a fixed uncertainty set, admissibility can be verified by searching over control policies. However it is not obvious how to search over possible admissible sets and corresponding control

policies simultaneously. In order to treat the problem with a unified methodology, we will characterize admissible sets as images of a modifier function applied to an initial uncertainty set $\hat{\Xi}$ as explained in section 2.2.1. The advantage of this approach will be evident in the following section 2.3.2, when we formulate computationally tractable methods for evaluating the admissibility of uncertainty sets for tracking.

Implicit modification of uncertainty sets

Note that we do not distinguish between the reference and the disturbance as they are just parts of the combined uncertainty signal, and all presented results apply to general robust programs. However, in the context of tracking commitment, only the reference set $\mathcal{R}$ can be modulated, whereas the disturbance set $\mathcal{W}$ is fixed. Also in general, the information structure of the reference and the disturbance are different. Nevertheless, without loss of generality, these particularities of the uncertain signal can be easily incorporated in the definition of the information structure and the mapping $\boldsymbol{\nu} : \mathbb{R}^n_{\xi} \longrightarrow \mathbb{R}^n_{\xi}$ by defining the *uncertainty modifier* function $\boldsymbol{\nu}$ as

$$\boldsymbol{\nu}(\xi) = \boldsymbol{\nu}(\boldsymbol{r},\ \boldsymbol{w}) := (\boldsymbol{\nu}_r(\boldsymbol{r}),\ \boldsymbol{w}). \tag{2.51}$$

In the following, we will show that we can evaluate the admissibility of the set $\Xi = \boldsymbol{\nu}(\hat{\Xi})$ via conditions on the composite function $\hat{\boldsymbol{\pi}} = \boldsymbol{\pi} \circ \boldsymbol{\nu}$ that is applied to the initial set $\hat{\Xi}$, as depicted in Figure 2.1. This allows us to fix an initial uncertainty set $\hat{\Xi}$, embed the modifier function into the control policy and implicitly modulate uncertainty sets and control policies simultaneously. To this end we introduce the following lemma,

Lemma 2. *Let $\boldsymbol{\nu} : \mathbb{R}^{Nn_{\xi}} \longrightarrow \mathbb{R}^{Nn_{\xi}}$, be a bijection and $\hat{\Xi}$ be a compact set with non-empty interior . The set $\Xi := \boldsymbol{\nu}(\hat{\Xi})$ is causally admissible for tracking by system* (2.38) *in state x with respect to the information structure $\boldsymbol{\mathcal{I}}$ if and only if*

$$\exists \hat{\boldsymbol{\pi}} \in \Delta_{\boldsymbol{\nu}}(\hat{\Xi}) : \hat{\boldsymbol{\pi}} \circ \boldsymbol{\nu}^{-1} \in \mathcal{F}(\boldsymbol{\mathcal{I}}) \tag{2.52}$$

where $\Delta_{\boldsymbol{\nu}}$ is defined as

$$\Delta_{\boldsymbol{\nu}}(\Xi) := \{\boldsymbol{\pi} : \forall \xi \in \Xi,\ (\boldsymbol{\pi}(\xi), \boldsymbol{\nu}(\xi)) \in \mathcal{Q}\} \tag{2.53}$$

Proof. : Suppose $\hat{\boldsymbol{\pi}} \in \Delta_{\boldsymbol{\nu}}(\hat{\Xi})$ and $\hat{\boldsymbol{\pi}} \circ \boldsymbol{\nu}^{-1} \in \mathcal{F}(\boldsymbol{\mathcal{I}})$. Then we have

$$\forall \hat{\xi} \in \hat{\Xi},\ (\hat{\boldsymbol{\pi}}(\hat{\xi}), \boldsymbol{\nu}(\hat{\xi})) \in \mathcal{Q} \tag{2.54}$$

Let $\xi := \nu(\hat{\xi})$. Since ν is bijective, we have $\hat{\xi} = \nu^{-1}(\xi)$. Therefore (2.54) is equivalent to

$$\begin{aligned} &\forall \nu^{-1}(\xi) \in \hat{\Xi},\ (\hat{\pi}\circ\nu^{-1}(\xi), \nu\circ\nu^{-1}(\xi)) \in \mathcal{Q} \\ \Longleftrightarrow &\forall \xi \in \nu(\hat{\Xi}),\ (\hat{\pi}\circ\nu^{-1}(\xi), \xi) \in \mathcal{Q} \\ \Longleftrightarrow &\forall \xi \in \Xi,\ (\pi(\xi), \xi) \in \mathcal{Q} \text{ with } \pi = \hat{\pi}\circ\nu^{-1} \\ \Longleftrightarrow &\pi \in \Delta(\Xi) \end{aligned} \tag{2.55}$$

Moreover, we have that $\pi = \hat{\pi}\circ\nu^{-1} \in \mathcal{F}(\mathcal{I})$. This concludes that Ξ is causally admissible for tracking according to Definition 5. The reverse direction follows from the equivalence of all steps.

□

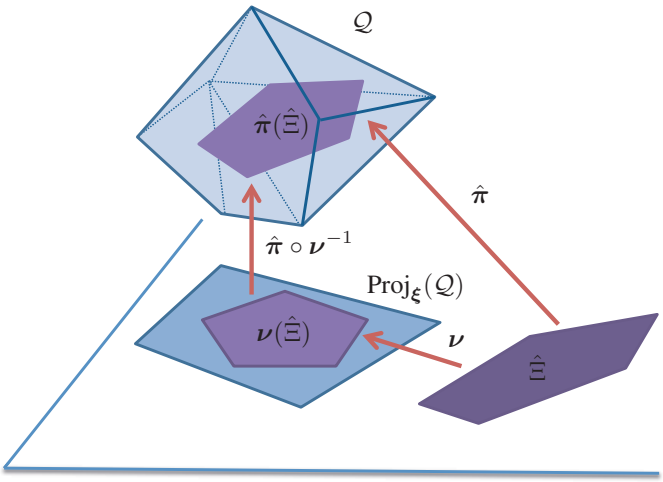

Figure 2.1 – Conceptual sketch of the relationships between uncertainty sets and applied functions. The initial uncertainty set $\hat{\Xi}$ is not necessarily a subset of the projection of $\mathcal{Q}$, therefore might not be admissible according to Lemma 1. However once we find a feasible lifting of this set into $\mathcal{Q}$, we can take its projection as an admissible uncertainty set, which is given by $\Xi = \nu(\hat{\Xi})$. The corresponding admissible control policy can be obtained by letting $\pi = \hat{\pi}\circ\nu^{-1}$.

Remark 2. *Since the modifier function ν is an arbitrary bijection, we do not lose generality when we consider uncertainty sets that can be characterized as the image of a given initial compact set $\hat{\Xi}$ with non-empty interior, under ν.*

According to Lemma 2 we can write an equivalent formulation of the family of admissible sets for a given initial set $\hat{\Xi}$

$$\Omega = \left\{ \Xi \subset \mathbb{R}^{Nn_\xi} \;\middle|\; \begin{array}{l} \exists \nu, \hat{\pi} \,:\, \Xi = \nu(\hat{\Xi}), \\ \hat{\pi} \,:\in \Delta_\nu(\hat{\Xi}), \hat{\pi}\circ\nu^{-1} \in \mathcal{F}(\mathcal{I}) \end{array} \right\} \tag{2.56}$$

When we look for a causally admissible set that belongs to Ω, the description (2.56) allows us to implicitly manipulate uncertainty sets and control policies simultaneously to verify admissibility, as will be seen in Section 2.3.2. However, while searching for a modifier function $\boldsymbol{\nu}$, the condition $\hat{\boldsymbol{\pi}} \circ \boldsymbol{\nu}^{-1} \in \mathcal{F}(\boldsymbol{\mathcal{I}})$ is difficult to evaluate since it is a condition on a composite function that involves the inverse of $\boldsymbol{\nu}^{-1}$. In the following, we will propose a simple sufficient condition directly on $\boldsymbol{\nu}$, that is easy to evaluate and ensures causal admissibility of the modified uncertainty set. We start by splitting the causality conditions of the composite function $\hat{\boldsymbol{\pi}} \circ \boldsymbol{\nu}^{-1}$.

Lemma 3. *Let $(\boldsymbol{\mathcal{I}}_k)_{k \in \mathbb{Z}_{[1,m]}}$ be a set of information structures and $f : \mathbb{R}^n \longrightarrow \mathbb{R}$. If for all k, $f \in \mathcal{F}(\boldsymbol{\mathcal{I}}_k)$ then $f \in \mathcal{F}(\bigcap_k \boldsymbol{\mathcal{I}}_k)$.*

The proof of Lemma 3, as well as other technical proofs in this section related to information structures can be found in [BGJ17]. Lemma 3 states an intuitive fact, that is if the output of a function f depends only on inputs indexed by $\boldsymbol{\mathcal{I}}_1$ and $\boldsymbol{\mathcal{I}}_2$, then it actually depends only on inputs indexed by their intersection. This directly motivates the next lemma.

Lemma 4. *Let $\boldsymbol{g} : \mathbb{R}^n \longrightarrow \mathbb{R}^n$, be a bijection. Given an information structure $\boldsymbol{\mathcal{I}}$, define $\hat{\boldsymbol{\mathcal{I}}}$ as*

$$\hat{\boldsymbol{\mathcal{I}}}_j = \bigcap_{\{i | j \in \boldsymbol{\mathcal{I}}_i\}} \boldsymbol{\mathcal{I}}_i \tag{2.57}$$

The following equivalence holds

$$\forall \boldsymbol{f} \in \mathcal{F}(\boldsymbol{\mathcal{I}}), \boldsymbol{f} \circ \boldsymbol{g} \in \mathcal{F}(\boldsymbol{\mathcal{I}}) \Longleftrightarrow \boldsymbol{g} \in \mathcal{F}(\hat{\boldsymbol{\mathcal{I}}}) \tag{2.58}$$

Equation (2.57) characterizes a set of functions which do not change the information structure of $\boldsymbol{f}$. Loosely speaking, it states that if $\boldsymbol{f}_i$ depends on x_j then $\boldsymbol{g}_j$ should not depend on anything that $\boldsymbol{f}_i$ does not depend on. Notice that $\hat{\boldsymbol{\mathcal{I}}}_j$ is always nonempty and in particular it contains j. This reflects the fact that a "diagonal" mapping (where $\boldsymbol{g}_j$ depends only on j for all j) does not change the information structure of any function it is composed with (for linear functions it means that multiplying by a diagonal matrix always preserves the sparsity pattern). For futher discussion on this subject, see [BGJ17].

Sufficient conditions for causal admissibility of modified uncertainty sets

In view of Lemma 2, simultaneous optimization over $\boldsymbol{\pi}$ and $\boldsymbol{\nu}$ would be beneficial for searching admissible uncertainty sets. Lemma 4 is instrumental in proving that from a control policy $\hat{\boldsymbol{\pi}}$ defined on $\hat{\Xi}$ and an invertible mapping $\boldsymbol{\nu}$, a control policy defined on $\boldsymbol{\nu}(\hat{\Xi})$ which has the desired information structure can be recovered. Indeed, $\hat{\boldsymbol{\pi}} \in \mathcal{F}(\boldsymbol{\mathcal{I}})$ and $\boldsymbol{\nu}^{-1} \in \mathcal{F}(\hat{\boldsymbol{\mathcal{I}}})$ ensures that $\hat{\boldsymbol{\pi}} \circ \boldsymbol{\nu}^{-1}$ defined on $\boldsymbol{\nu}(\hat{\Xi})$ belongs to $\mathcal{F}(\boldsymbol{\mathcal{I}})$ according to Lemma 4. However, conditions on $\boldsymbol{\nu}^{-1}$ are inconvenient since the aim is to optimize directly over $\boldsymbol{\nu}$. Sufficient conditions on $\boldsymbol{\nu}$ are sought

to replace the condition $\boldsymbol{\nu}^{-1} \in \mathcal{F}(\hat{\boldsymbol{\mathcal{I}}})$. Unfortunately a certain information structure for $\boldsymbol{\nu}^{-1}$ does not usually result in a specific information structure for $\boldsymbol{\nu}$. In particular, a sparse information structure for $\boldsymbol{\nu}^{-1}$ does not generally result in a sparse information structure for $\boldsymbol{\nu}$. For example, the inverse of a causal function is not generally causal. The following lemma gives sufficient conditions on $\boldsymbol{\nu}$.

Lemma 5. *Suppose* $\boldsymbol{\nu} : \mathbb{R}^n \longrightarrow \mathbb{R}^n$ *is a continuous bijection of* $\mathbb{R}^n$ *and* $\boldsymbol{\nu} \in \mathcal{F}(\hat{\boldsymbol{\mathcal{I}}})$ *as defined by equation* (2.57)*. Define* $\mathcal{G} = \{\boldsymbol{f} \circ \boldsymbol{\nu} \mid \boldsymbol{f} \in \mathcal{F}(\boldsymbol{\mathcal{I}})\}$*. We have*

$$\mathcal{G} = \mathcal{F}(\boldsymbol{\mathcal{I}})$$

Under mild assumptions, Lemma 5 states that composing $\boldsymbol{f} \in \mathcal{F}(\boldsymbol{\mathcal{I}})$ with $\boldsymbol{\nu}$ results in a function with the same information structure.

Corollary 1. *Given an information structure* $\boldsymbol{\mathcal{I}}$ *and* $\hat{\boldsymbol{\mathcal{I}}}$ *as defined in equation* (2.57)*, if* $\boldsymbol{\nu}$ *is a continuous bijection and* $\boldsymbol{\nu} \in \mathcal{F}(\hat{\boldsymbol{\mathcal{I}}})$*, then* $\boldsymbol{\nu}^{-1} \in \mathcal{F}(\hat{\boldsymbol{\mathcal{I}}})$*.*

Proof. We have $\mathcal{F}(\boldsymbol{\mathcal{I}}) = \{\boldsymbol{f} \mid \boldsymbol{f} \in \mathcal{F}(\boldsymbol{\mathcal{I}})\} = \{\boldsymbol{f} \circ \boldsymbol{\nu} \circ \boldsymbol{\nu}^{-1} \mid \boldsymbol{f} \in \mathcal{F}(\boldsymbol{\mathcal{I}})\}$ by bijectivity of $\boldsymbol{\nu}$. In turn, $\mathcal{F}(\boldsymbol{\mathcal{I}}) = \{\boldsymbol{g} \circ \boldsymbol{\nu}^{-1} \mid \boldsymbol{g} \in \mathcal{G}\}$ by definition of $\mathcal{G}$. By Lemma 5, using the assumptions on $\boldsymbol{\nu}$, we have that $\mathcal{G} = \mathcal{F}(\boldsymbol{\mathcal{I}})$ and therefore $\mathcal{F}(\boldsymbol{\mathcal{I}}) = \{\boldsymbol{g} \circ \boldsymbol{\nu}^{-1} \mid \boldsymbol{g} \in \mathcal{F}(\boldsymbol{\mathcal{I}})\}$. Hence, for any $\boldsymbol{g} \in \mathcal{F}(\boldsymbol{\mathcal{I}})$, it holds that $\boldsymbol{g} \circ \boldsymbol{\nu}^{-1} \in \mathcal{F}(\boldsymbol{\mathcal{I}})$. The fact that $\boldsymbol{\nu}^{-1} \in \mathcal{F}(\hat{\boldsymbol{\mathcal{I}}})$ then follows from Lemma 4. □

Theorem 1. *Let* $\boldsymbol{\nu} : \mathbb{R}^{Nn_\xi} \longrightarrow \mathbb{R}^{Nn_\xi}$*, be a continuous bijection and* $\boldsymbol{\mathcal{I}}$ *an information structure,* $\hat{\boldsymbol{\mathcal{I}}}$ *defined by equation* (2.57) *and* $\Delta_{\boldsymbol{\nu}}$ *in equation* (2.53)*.* $\boldsymbol{\nu}(\hat{\Xi})$ *is causally admissible for tracking with respect to the information structure* $\boldsymbol{\mathcal{I}}$ *if*

$$\begin{aligned} &\mathcal{F}(\boldsymbol{\mathcal{I}}) \cap \Delta_{\boldsymbol{\nu}}(\hat{\Xi}) \neq \emptyset \\ &\boldsymbol{\nu} \in \mathcal{F}(\hat{\boldsymbol{\mathcal{I}}}) \end{aligned} \tag{2.59}$$

Proof. Suppose there exists $\hat{\boldsymbol{\pi}} \in \mathcal{F}(\boldsymbol{\mathcal{I}}) \cap \Delta_{\boldsymbol{\nu}}(\hat{\Xi})$. Since $\boldsymbol{\nu}$ is a continuous bijection, $\boldsymbol{\nu} \in \mathcal{F}(\hat{\boldsymbol{\mathcal{I}}})$ implies that $\boldsymbol{\nu}^{-1} \in \mathcal{F}(\hat{\boldsymbol{\mathcal{I}}})$ by Corollary 1. Lemma 4 in turn ensures that $\hat{\boldsymbol{\pi}} \circ \boldsymbol{\nu}^{-1} \in \mathcal{F}(\boldsymbol{\mathcal{I}})$. Finally, application of Lemma 2 concludes the proof. □

Theorem 1 provides sufficient conditions for causal admissibility of an uncertainty set for tracking. We can define the family of admissible sets that comply with these sufficient conditions as

$$\tilde{\Omega}(\hat{\Xi}) = \left\{ \Xi \subset \mathbb{R}^{Nn_\xi} \;\middle|\; \begin{array}{l} \exists \boldsymbol{\nu}, \hat{\boldsymbol{\pi}} \\ \Xi = \boldsymbol{\nu}(\hat{\Xi}),\ \ \boldsymbol{\nu} \in \mathcal{F}(\hat{\boldsymbol{\mathcal{I}}}) \\ \hat{\boldsymbol{\pi}} \in \mathcal{F}(\boldsymbol{\mathcal{I}}) \cap \Delta_{\boldsymbol{\nu}}(\hat{\Xi}) \end{array} \right\} \tag{2.60}$$

For the definition of $\tilde{\Omega}$ we have replaced the condition $\hat{\boldsymbol{\pi}} \circ \boldsymbol{\nu}^{-1} \in \mathcal{F}(\boldsymbol{\mathcal{I}})$ with the sufficient but simpler conditions $\hat{\boldsymbol{\pi}} \in \mathcal{F}(\boldsymbol{\mathcal{I}})$ and $\boldsymbol{\nu} \in \mathcal{F}(\hat{\boldsymbol{\mathcal{I}}})$. Therefore $\tilde{\Omega}$ is a restriction of the original family of admissible sets Ω.

$$\tilde{\Omega}(\hat{\Xi}) \subseteq \Omega \tag{2.61}$$

The restriction will depend on the initial set $\hat{\Xi}$ and thus the argument of $\tilde{\Omega}$ is added to reflect this fact. However, this restriction leads to tractable formulations based on the available robust programming literature, as we will show in the next section.

Finally, we write the *modified robust tracking commitment* problem that is based on sufficient conditions (2.59) as

$$\text{find } \mathcal{R} \ : \ \mathcal{R} \times \mathcal{W} \in \tilde{\Omega}(\hat{\Xi}) \tag{2.62}$$

2.3.2 Computationally tractable approximations

The problem formulation (2.62) allows us to search over uncertainty sets implicitly by means of modifier functions. However the problem is still difficult in its general form, due to the infinite dimension of the search space and the infinite number of constraints. Therefore we will look for finite dimensional and tractable approximations of the tracking commitment problem in order to solve it efficiently.

Using the definitions of $\tilde{\Omega}$ and $\Delta_{\boldsymbol{\nu}}$, we can rewrite the modified robust tracking commitment problem as

$$\begin{aligned} \text{find} \quad & \hat{\boldsymbol{\pi}}, \boldsymbol{\nu} \\ \text{subject to} \quad & \forall \hat{\boldsymbol{\xi}} \in \hat{\Xi} \\ & (\hat{\boldsymbol{\pi}}(\hat{\boldsymbol{\xi}}), \boldsymbol{\nu}(\hat{\boldsymbol{\xi}})) \in \mathcal{Q} \\ & \hat{\boldsymbol{\pi}} \in \mathcal{F}(\boldsymbol{\mathcal{I}}) \\ & \boldsymbol{\nu} \in \mathcal{F}(\hat{\boldsymbol{\mathcal{I}}}) \ . \end{aligned} \tag{2.63}$$

Note that (2.63) is an *adjustable robust optimization* (ARO) with a modified uncertainty set problem as described in sections 2.2.1 and 2.2.1. Therefore, we can utilize the results presented in section 2.2.1 to obtain a computationally tractable reformulation, under the following assumption.

Assumption 1. *The system under consideration is linear and described by* (2.38)*. State and inputs of the system are subject to polytopic constraints. The allowed error set $\mathcal{E}$ is polytopic.*

Following Assumption 1, we define the polytopic tracking error set as

$$\mathcal{E} := \{\boldsymbol{e} \mid G\boldsymbol{e} \leqslant g\}$$

The feasibility set $\mathcal{Q}$ also becomes polytopic and can be written as

$$\mathcal{Q} = \{(\boldsymbol{u}, \boldsymbol{\xi}) \mid H\boldsymbol{u} + Q\boldsymbol{\xi} \leqslant q\}$$

The matrices Q, H, and q are defined as

$$P := \mathbf{CB} + \mathbf{D}, \quad S = \mathbf{CE}, \quad \bar{\boldsymbol{y}} = \mathbf{CA}x_0$$

$$H := \begin{bmatrix} T \\ -P \end{bmatrix}, \quad Q = \begin{bmatrix} \mathbf{0} & V \\ -S & G \end{bmatrix}, \quad q = \begin{bmatrix} h \\ g + \bar{\boldsymbol{y}} \end{bmatrix}$$

where $\mathbf{0}$'s are matrices of zeros with proper dimensions and the remaining definitions are provided in Section 2.2.3.

Note that q is an affine function of the initial condition x and the vector g that determines the size of the tracking set. This polytopic description of the feasibility set allows treatment of the modified robust tracking commitment problem (2.63) in the uncertain linear optimization framework. Until this point, we have not made any strong assumptions on the families of uncertainty sets, policy and modifier functions. The presented theory in section 2.3.1 applies to generic functions and sets. Therefore, the sufficient conditions in (2.60) can be used to verify causal admissibility of any uncertainty set, using generic policies and modifier functions. In the following we restrict ourselves to sets for which computationally tractable methods can be derived for verification of admissibility conditions (2.60). According to the results presented in Section 2.2.1, we introduce the following assumption on the uncertainty set.

Assumption 2. *The uncertainty sets under consideration are representable by intersections of convex cones as*

$$\hat{\Xi} = \left\{ \boldsymbol{\xi} \mid F_i\boldsymbol{\xi} + f_i \in \mathbf{K}_i, \ i \in \mathbb{Z}_{[1,m]} \right\} \tag{2.64}$$

where the cone $\mathbf{K}_i$ *is proper.*

Linear policy and modifier functions

In this section we build our formulation based on the results presented in Section 2.2.1, which shows that restricting the search space of policies to linear (or affine) functions leads to finite dimensional and tractable formulations of adjustable robust programs. Let us define the linear versions of the control policy and the uncertainty modifier

$$\boldsymbol{\pi}_{lin}(\boldsymbol{\xi}) := \mathbf{M}\boldsymbol{\xi}, \quad \boldsymbol{\nu}_{lin}(\boldsymbol{\xi}) := \mathbf{L}\boldsymbol{\xi} \tag{2.65}$$

where $\mathbf{M} \in \mathbb{R}^{Nn_u \times Nn_\xi}$ and $\mathbf{L} \in \mathbb{R}^{Nn_\xi \times Nn_\xi}$. We can describe the causality conditions by constraints on $\mathbf{M}$ and $\mathbf{L}$

$$\begin{aligned} &\mathbf{M}(k, \mathbb{Z}_{[1,Nn_\xi]} \backslash \boldsymbol{\mathcal{I}}_k) = 0, \quad k \in \mathbb{Z}_{[1,N]} \quad \iff \boldsymbol{\pi} \in \mathcal{F}(\boldsymbol{\mathcal{I}}) \\ &\mathbf{L}(k, \mathbb{Z}_{[1,Nn_\xi]} \backslash \hat{\boldsymbol{\mathcal{I}}}_k) = 0, \quad k \in \mathbb{Z}_{[1,N]} \quad \iff \boldsymbol{\nu} \in \mathcal{F}(\hat{\boldsymbol{\mathcal{I}}}) \end{aligned} \tag{2.66}$$

Note that the constraints (2.66) impose that the elements of $\mathbf{M}$ and $\mathbf{L}$ multiplying the elements of the uncertain variable which are not included in the information structure at step k to be zero, thus enforcing causality of the linear functions $\boldsymbol{\pi}_{lin}$ and $\boldsymbol{\nu}_{lin}$.

Let us now formulate the robust tracking commitment problem (2.63) with linear policies given in (2.65) and conic uncertainty sets described by (2.64)

$$\begin{aligned} \text{find} \quad & \hat{\mathbf{M}}, \mathbf{L} \\ \text{subject to} \quad & \forall \hat{\boldsymbol{\xi}} \,:\, F_i \hat{\boldsymbol{\xi}} + f_i \in \mathbf{K}_i, \quad i \in \mathbb{Z}_{[1,m]} \\ & H \hat{\mathbf{M}} \hat{\boldsymbol{\xi}} + Q \mathbf{L} \hat{\boldsymbol{\xi}} \leqslant q \\ & (\hat{\mathbf{M}}, \mathbf{L}) \text{ satisfies (2.66)} \end{aligned} \tag{2.67}$$

Once the problem is solved, a feasible solution $\hat{\mathbf{M}}^*$ and $\mathbf{L}^*$ can be used to construct the uncertainty set that is causally admissible for tracking and the corresponding control policy

$$\Xi = \mathbf{L}^* \hat{\Xi}, \quad \boldsymbol{\pi}(\xi) = \mathbf{M} \xi, \quad \mathbf{M} = \hat{\mathbf{M}}^* \mathbf{L}^{*-1} \tag{2.68}$$

Applying conic duality as in Section 2.2.1 and stacking dual variables in matrix $\mathbf{Z}$, the robust counterpart can be formulated as

$$\begin{aligned} \text{find} \quad & \mathbf{Z}, \hat{\mathbf{M}}, \mathbf{L} \\ \text{subject to} \quad & \mathbf{Z}_i^T \in \mathbf{K}_i^*, \;\; i \in \mathbb{Z}_{[1,m]} \\ & \sum_{i=1}^{m} \mathbf{Z}_i f_i \leqslant q \\ & \sum_{i=1}^{m} \mathbf{Z}_i F_i = -\left(H \hat{\mathbf{M}} + Q \mathbf{L} \right) \\ & (\hat{\mathbf{M}}, \mathbf{L}) \text{ satisfies (2.66)} \end{aligned} \tag{2.69}$$

The robust counterpart for the tracking commitment problem (2.67) is convex in linear control policies parameterized by $\hat{\mathbf{M}}$ and linear uncertainty modifiers parameterized by $\mathbf{L}$. Therefore, when sets $\mathbf{K}_i$ are polyhedral, second order or semi-definite cones, the problem formulation (2.69) allows tractable computations of feasible reference sets admissible for tracking with respect to the information structure $\boldsymbol{\mathcal{I}}$ and the tracking error set $\mathcal{E}$, by system (2.38). Table 2.1 gives a summary of problem complexity in case of most common uncertainty sets for the reference and disturbance.

Remark 3. *Any additional decision variable that is independent of the uncertainty, $\bar{\boldsymbol{u}}$, as well as non zero nominal values for the disturbance and reference, $\bar{\boldsymbol{\xi}}$, can be easily incorporated in the above formulation as*

$$H \bar{\boldsymbol{u}} + Q \bar{\boldsymbol{\xi}} + \sum_{i=1}^{m} \mathbf{Z}_i f_i \leqslant q$$

$\mathcal{R}$	$\mathcal{W}$	Robust Counterpart
$F_r \boldsymbol{r} \leqslant f_r$	$F_w \boldsymbol{w} \leqslant f_w$	LP
$F_r \boldsymbol{r} + f_r \,:\, \|\boldsymbol{r}\|_2 \leqslant 1$	$F_w \boldsymbol{w} \leqslant f_w$	SOCP
$F_r \boldsymbol{r} \leqslant f_r$	$F_w \boldsymbol{w} + f_w \,:\, \|\boldsymbol{w}\|_2 \leqslant 1$	SOCP
$F_r \boldsymbol{r} + f_r \,:\, \|\boldsymbol{r}\|_2 \leqslant 1$	$F_w \boldsymbol{w} + f_w \,:\, \|\boldsymbol{w}\|_2 \leqslant 1$	SOCP

Table 2.1 – Optimization type for the robust counterpart formulation (2.69) of the tracking commitment problem, considering commonly used combinations of polytopes and ellipsoids as uncertainty sets. Note that the polytopic representation also covers 1 and ∞ norm balls.

Modulating the tracking error set

In Section 2.3.2, we have formulated a tractable version of the robust tracking commitment problem (2.50) which seeks a causally admissible reference set with respect to a fixed tracking error set $\mathcal{E}$. However for some applications, it might be preferable to modulate the reference set together with the associated tracking error set, since the relative sizes of the two sets will indicate the tracking performance. For example, for frequency regulation service to the power grid in Switzerland, the service providers are allowed to deviate from the reference up to a certain percentage of the total service capacity committed, therefore a service provider who is committing a larger reference set is allowed to have a larger error set [Swi15b].

Notice that the problem (2.69) is convex in q which parameterizes the polytopic feasibility set $\mathcal{Q}$. Therefore one can freely optimize over modifications of the feasibility set. For clarity let us write the system feasibility and tracking constraints separately as

$$\mathcal{Q} = \left\{ (\boldsymbol{u}, \boldsymbol{r}, \boldsymbol{w}) \middle| \begin{matrix} T\boldsymbol{u} + V\boldsymbol{w} \leqslant h \\ G(\boldsymbol{r} - (P\boldsymbol{u} + S\boldsymbol{w} + \bar{\boldsymbol{y}})) \leqslant g \end{matrix} \right\} \tag{2.70}$$

where $\bar{\boldsymbol{y}}$ is the nominal output of the system without control action and the derivation of matrices T, V, P, S can be defined as

$$P := \mathbf{CB} + \mathbf{D}, \quad S = \mathbf{CE}, \quad \bar{\boldsymbol{y}} = \mathbf{CA}x_0$$

From (2.70), we can immediately observe that the problem (2.69) is also convex in g which parameterizes the tracking error set $\mathcal{E}$, therefore allows modulation in a tractable manner. The sizes of the tracking error set and the reference set can be related by enforcing a joint constraint on the uncertainty modifier function parameter $\mathbf{L}$ and the error set parameter g.

Optimal tracking commitment

As mentioned earlier, the robust tracking commitment (2.50) is a feasibility problem. On the other hand, the optimal robust tracking commitment problem is finding the control policy,

reference and tracking error sets, that minimize a cost function.

$$\underset{\pi \in \mathcal{F}(\boldsymbol{I}) \cap \Delta(\mathcal{R} \times \mathcal{W})}{\text{minimize}} \; J(\boldsymbol{\pi}, \mathcal{R}, \mathcal{E}) \tag{2.71}$$

Relying on the tractable formulation with linear control policies and uncertainty modifiers (2.67), we can solve the tractable version of the optimal robust tracking commitment.

$$\begin{aligned} \text{minimize} \quad & J(\boldsymbol{u}, \mathbf{L}_r, g) \\ \text{subject to} \quad & \forall \hat{\boldsymbol{r}} \; : \; F_r \hat{\boldsymbol{r}} + f_r \in \mathbf{K}_r, \\ & \forall \boldsymbol{w} \; : \; F_r \boldsymbol{w} + f_w \in \mathbf{K}_w, \\ & T\boldsymbol{u} + V\boldsymbol{w} \leqslant h \\ & G(\mathbf{L}_r \hat{\boldsymbol{r}} - (P\boldsymbol{u} + S\boldsymbol{w} + \bar{\boldsymbol{y}})) \leqslant g \\ & \boldsymbol{u} = \hat{\mathbf{M}}_r \hat{\boldsymbol{r}} + \mathbf{M}_w \boldsymbol{w} + \bar{\boldsymbol{u}} \\ & (\hat{\mathbf{M}}, \mathbf{L}) \text{ satisfies (2.66)} \end{aligned} \tag{2.72}$$

where $\hat{\mathbf{M}}$ and $\mathbf{L}$ are defined as

$$\hat{\mathbf{M}} = \begin{bmatrix} \hat{\mathbf{M}}_r & \\ & \mathbf{M}_w \end{bmatrix}, \quad \mathbf{L} = \begin{bmatrix} \mathbf{L}_r & \\ & \mathbf{I}_{Nn_x} \end{bmatrix} \tag{2.73}$$

For notational simplicity, the reference and disturbance sets are represented as single conic sets, but they can also be defined as intersection of several conic sets as in (2.64).

With a suitable cost function, the optimal commitment problem (2.72) can be solved as explained in section 2.3.2. The cost function usually depends on the uncertain realization of the reference and disturbance. However this dependence can be qualified out by either considering the worst-case or expectation of the possible cost realizations [SDR09]. Furthermore, as long as J is bilinear in the uncertain variables and decision variables, the tractable robust formulation for the minimization of an upper bound, that constitutes an equivalent problem with certain cost function, can be obtained [BTEGN09].

After solving (2.72), the optimal control policy, reference and tracking error sets can be obtained as

$$\begin{aligned} \mathcal{R}^* &= \{\boldsymbol{r} = \mathbf{L}_r^* \hat{\mathbf{r}} \mid F_r \hat{\boldsymbol{r}} + f_r \in \mathbf{K}_r\} \\ \boldsymbol{\pi}^*(\boldsymbol{r}, \boldsymbol{w}) &= \hat{\mathbf{M}}_r^* \mathbf{L}_r^{*-1} \boldsymbol{r} + \mathbf{M}_w^* \boldsymbol{w} + \bar{\boldsymbol{u}}^* \\ \mathcal{E}^* &= \{\boldsymbol{e} \mid G\boldsymbol{e} \leqslant g^*\} \end{aligned} \tag{2.74}$$

where * indicates that the variable is an optimizer of (2.72).

Collective tracking

In this section, we will consider the collective tracking commitment problem, where the reference signal is to be tracked by the summation (or average) of the outputs of several subsystems that are not coupled via constraints or dynamics. In this case, the aggregate tracking error can be written as

$$\boldsymbol{e} = \boldsymbol{r} - \sum_{j=1}^{n} \boldsymbol{y}^j \tag{2.75}$$

where superscript j indicates that the associated variable belongs to subsystem j.

In order to compute a causally admissible reference set for the collection of subsystems, it is necessary to compute the aggregated tracking capability simultaneously. An obvious option is to treat the collective system as a single system with block diagonal system matrices, and solve the problem (2.67) centrally. However this requires collecting the knowledge of detailed subsystem models by an aggregator and will possibly result in a very large number of decision variables. Therefore it is desirable to distribute the problem, such that it can be solved without central knowledge of the models and using limited communication between the agents.

Let us briefly show that the problem (2.67) in fact easily lends itself to distributed solution methods. The subsystems are supposed to track a single reference by collective action. This is equivalent to saying that the subsystems are tracking separate reference signals, which sum up to the central reference. Using a common nominal reference set and a linear uncertainty modifier function for each subsystem, the reference can be split as

$$\boldsymbol{r}^j = \mathbf{L}^j \hat{\boldsymbol{r}}, \quad \sum_{j=1}^{n} \mathbf{L}^j = \mathbf{L} \tag{2.76}$$

where $\mathbf{L}$ parameterizes the global linear modifier function. Therefore the aggregate reference set can be described as

$$\mathcal{R} = \sum_{j=1}^{n} \mathbf{L}^j \hat{\mathcal{R}} \tag{2.77}$$

Many distribution schemes are possible, given the cost-reward framework of the collective tracking task. As an example, we consider the case where the error set is fixed, the reward is split between agents according to their contributions to tracking characterized by parameter $\mathbf{L}^j$ and the objective is to minimize the total cost

$$\sum_{j=1}^{n} J^j(\boldsymbol{\pi}^j, \mathbf{L}^j) \tag{2.78}$$

Subsystem feasibility constraints $T^j \boldsymbol{u}^j + V^j \boldsymbol{w}^j \leqslant h^j$ and causality conditions on the local control policies can be treated separately by each subsystem. However (2.76) and the collective tracking constraint (2.75) introduces a constraint coupling among all subsystems. The collective tracking

constraint can be written as

$$\begin{aligned}
&\forall \hat{\boldsymbol{r}} \,:\, F_r \hat{\boldsymbol{r}} + f_r \in \mathbf{K}_r, \quad \forall \boldsymbol{w}^j \,:\, F_w^j \boldsymbol{w}^j + f_w^j \in \mathbf{K}_w^j \\
&\sum_{j=1}^{n} \left(G \left(\mathbf{L}^j \hat{\boldsymbol{r}} - (P^j \boldsymbol{u}^j + S^j \boldsymbol{w}^j + \bar{\boldsymbol{y}}^j) \right) \right) \leqslant g \\
&\boldsymbol{u}^j = \hat{\mathbf{M}}_r^j \hat{\boldsymbol{r}} + \mathbf{M}_w^j \boldsymbol{w}^j + \bar{\boldsymbol{u}}^j
\end{aligned} \tag{2.79}$$

Utilizing linear control policies $\boldsymbol{u}^j = \hat{\mathbf{M}}_r^j \hat{\boldsymbol{r}} + \mathbf{M}_w^j \boldsymbol{w}^j + \bar{\boldsymbol{u}}^j$ and formulating the dual of the robust counterpart as in (2.69), we obtain a tractable formulation for the collective tracking constraint

$$\mathbf{Z}_r^T \in \mathbf{K}_r^* , \quad {\mathbf{Z}_w^j}^T \in {\mathbf{K}^j}_w^* , \qquad \forall j \in \mathbb{Z}_{[1,n]} \tag{2.80a}$$

$$\mathbf{Z}_r f_r + \sum_{j=1}^{n} \left(\mathbf{Z}_w^j f_w^j - P^j \bar{\boldsymbol{u}}^j - \bar{\boldsymbol{y}}^j \right) \leqslant g \tag{2.80b}$$

$$\mathbf{Z}_r F_r = \sum_{j=1}^{n} G(P^j \hat{\mathbf{M}}_r^j - \mathbf{L}^j) , \tag{2.80c}$$

$$\mathbf{Z}_w^j F_w^j = G(P^j \mathbf{M}_w^j + S^j) , \qquad \forall j \in \mathbb{Z}_{[1,n]} \tag{2.80d}$$

We observe that the coupling constraints are (2.80b) and (2.80c). The partial Lagrangian with the coupling constraints can be written as

$$\sum_{j=1}^{n} J^j(\pi^j, \mathbf{L}^j) + \lambda^T \mathrm{vec}(\mathbf{Z}_r F_r - \sum_{j=1}^{n} G(P^j \hat{\mathbf{M}}_r^j - \mathbf{L}^j)) + \mu^T (\mathbf{Z}_r f_r + \sum_{j=1}^{n} \left(\mathbf{Z}_w^j f_w^j - P^j \bar{\boldsymbol{u}}^j - \bar{\boldsymbol{y}}^j \right) - g) \tag{2.81}$$

which is separable given the variables λ, μ and $\mathbf{Z}_r$. Therefore the collective tracking commitment problem can be solved in a distributed manner with global updates of these variables or enforcing consensus on their local copies [Boy10]. Note that, after the commitment of the aggregate reference set, it is not necessary to further communicate for guaranteed tracking if all subsystems have access to the aggregate reference $\boldsymbol{r}$, since collective tracking is robustly guaranteed by independent local control policies. However, the performance can be improved by repetitively solving the problem (2.80) online for redistributing the tracking task among subsystems according to the available information on the disturbance and reference as time progresses. Adjusting the error set parameter g is also possible with a suitable constraint on the aggregate modifier $\mathbf{L}$ and g. In case of dynamic couplings, inter-system constraints or common disturbances that couples the state and inputs of several subsystems, the same methodology can be applied to derive the tractable robust counterpart formulation and distribute the computation.

2.4 Frequency regulation service with buildings

In this section, we will focus on the provision of secondary frequency control service with commercial buildings, using the methodology of robust tracking commitment. In secondary frequency regulation, the service providers, also known as 'reserves', are asked to track a certain power injection reference, called the *Automatic Generation Control (AGC)* signal, over the period of service commitment [Swi15a]. As the AGC signal compensates the demand and generation uncertainty in the system, the value of the reference is not known a-priori. Therefore, the service providers are required to provide a 'baseline' power injection profile and a flexibility window, which describes their ability to track AGC reference signals on top of their baseline injection. The reserve flexibility is described by a parameter γ called the *capacity bid*, that is used as the scaling factor for the AGC reference. For example if γ = 10kW the normalized AGC reference $-1 \leq \hat{r} \leq 1$ is scaled by 10kW and allowed to take values in the range $-\gamma \leq r \leq \gamma$. The power injection(consumption) of the reserve provider is allowed to deviate from the reference within an error margin proportional to the size of the commitment.

The frequency regulation service provision can be split into two main tasks [FGQ+18]

- Commitment: Offline phase when the reserve provider determines the baseline output $\bar{y}$ and the commitment size γ,
- Tracking: Online phase when the reserve provider tracks the received AGC signal.

For the case of a building providing the service, these two phases are illustrated in in Figure 2.2. Notice that the robust tracking commitment is particularly suited for this application. Indeed,

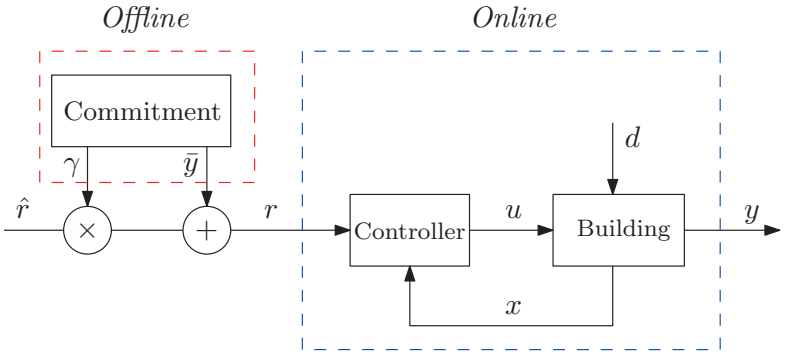

Figure 2.2 – Architecture of the control system for tracking service commitment and tracking.

all possible values of the normalized AGC can be represented by a reference set $\hat{\mathcal{R}}$, that the building operator can modify using the scaling parameter γ to obtain the admissible reference set for AGC around the baseline output $\bar{y}$.

$$\mathcal{R} = \gamma\hat{\mathcal{R}} + \bar{y}$$

An optimal set size can be computed provided that an appropriate cost function is chosen as will be demonstrated next. By choosing $\boldsymbol{y}$ defined in equation (2.38) to be the total power consumption of the building and $\mathbb{X}$ and $\mathbb{U}$ to represent the operating constraints of the building, we can see that the robust tracking commitment method directly allows us to address this problem.

2.4.1 Building Models

Providing ancillary service to the power grid with a building requires modelling and controlling its total power consumption. The modeling task can be split into two parts as (i) thermal dynamics of the building structure considering thermal power as input and zone temperatures as output, (ii) dynamics of the HVAC system, considering thermal power and electrical power consumption as outputs.

The buildings are originally modeled in EnergyPlus and are too complex to use for MPC [UR]. We use OpenBuild [GQJ15a] to automatically extract a reduced-order linear thermodynamic model of each building from its EnergyPlus data files. OpenBuild also outputs for each building the disturbance sequences that capture solar gains, internal gains and the effect of the ambient temperature. The thermal dynamics of the building can be written as

$$\begin{aligned} x^+ &= Ax + B_u u + B_w w \\ \theta &= Cx \end{aligned} \tag{2.82}$$

where $x \in \mathbb{R}^{n_x}$ is the state, $u \in \mathbb{R}^{n_z}$ is the thermal energy input to each thermal zone, $w \in \mathbb{R}^3$ is the disturbance signal consisting of internal heat gain, solar radiation, and the ambient temperature. The output $\theta \in \mathbb{R}^{n_z}$ describes the mean temperature of each thermal zone in the building. A typical model generation with OpenBuild for a large office building results in state dimension $n_x = 57$ and input-output dimension $n_z = 19$, which is the same as the number of zones.

The time scale of the thermal dynamics of the building is in general much slower compared to the dynamics of the HVAC system. Assuming that the HVAC system achieves steady state quickly, we can describe the power consumption with a static relation as

$$y(t) = \mathrm{COP}(t) \sum_{i=1}^{n_z} u(t) \tag{2.83}$$

The coefficient of performance (COP) is a nonlinear term that depends on the internal state of the building and ambient conditions. For simplifying the control design, it can be represented by a time varying prediction that is computed by the expected values of the internal and ambient temperatures. This allows formulating an MPC problem that controls the power consumption relying on a time-varying linear model.

Even though the high level control of the dynamics might rely on a simplified linear model, the

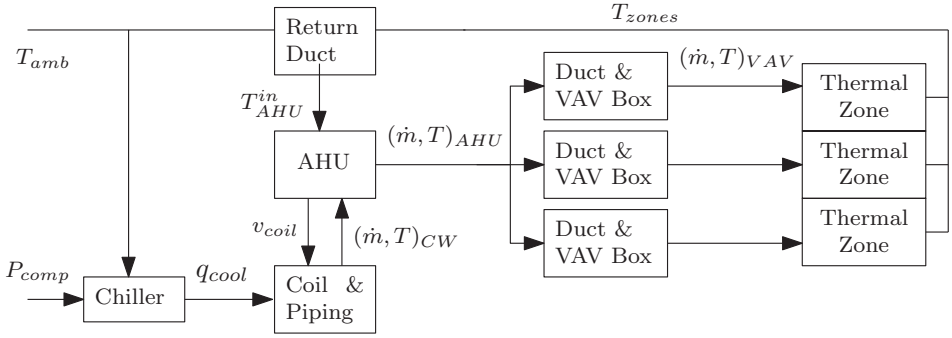

Figure 2.3 – Diagram of the HVAC system, which consists of a chiller, a chilled water (CW) loop, an air handling unit (AHU), duct system and variable air volume (VAV) boxes. $\dot{m}$ is mass flow rate, T temperature, P_{comp} the electrical power consumption of the chiller, q_{cool} the heat power extracted from the chilled water loop, v_{coil} the valve control parameter of the cooling coil inside the AHU. The controlled variables are $\dot{m}_{\text{VAV}}$. The input air temperature $T^{\text{in}}_{\text{AHU}}$ depends on both zone temperatures T_{zones} and ambient temperature T_{amb}, due to the mixing of internal and external air in the return duct.

simulation will be more realistic with a model of the underlying HVAC system. The nonlinear dynamics for the HVAC can be written as

$$\begin{aligned} \psi^+ &= f(\psi, v, \theta) \\ u &= g(\psi, v, \theta) \\ y &= h(\psi, v, \theta) \end{aligned} \tag{2.84}$$

where $\psi \in \mathbb{R}^{n_\psi}$ is the internal state , $v \in \mathbb{R}^{n_v}$ is the input of the HVAC and θ is the zone temperatures that is the output of the thermal model (2.82). The total electrical power consumption is represented by $y \in \mathbb{R}$ where as u is the thermal power input to each zone as in (2.82).

We choose to model a forced-air system as described in Figure 2.3, and size it according to the maximum heating/cooling demand of the building model. In this study we consider the cooling period.

To simplify the model, we neglect the mass transport dynamics and only consider thermal dynamics of the transport media: water between the heat pump / chiller and AHU, and air in the duct system. The model is described by a continuous time nonlinear ordinary differential equation of type (2.84), which is solved numerically during the simulations.

2.4.2 Case study 1: Commitment with a single building under disturbance

In this case study, we solve the tracking commitment problem for a single building. We consider a small office building with three controlled zones served by individual air handling units, described by the linear model (2.82). The model of the building has state dimension $n_x = 10$ and input dimension $n_u = 3$. The disturbance captures the effect of internal gains, solar radiation

and outdoor temperature representing a typical summer weather for the city of Chicago. In this study, we assume a constant COP value to compute the electrical power from thermal power input to the building and neglect HVAC dynamics.

The peak thermal cooling load of the building is 45kW for the summer period. The input constraint set $\mathcal{U}$ specifies maximum and minimum cooling levels in the rooms so that $u_{i,\min} \leq u_i \leq u_{i,\max} = 0$ for each thermal zone input, reflecting the sizing of the equipment. The state constraints $\mathcal{X}$ specifies temperature zones in the constraints so that the temperature is maintained between 20°C and 25°C.

Details of the bidding are as follows: at time $t_0 = 0$, the building starts in initial condition x_0. The tracking period starts at time t_1 and ends at time t_2, therefore leaving a 'preparation' period for the building controller. The capacity bid consists of the commitment of a baseline consumption during the tracking period $\mathbf{u}_{\text{nom}}$ and up-down regulation limits around that baseline for power tracking. Up-down regulation bids result in a "box" uncertainty set. We therefore fix the basic uncertainty shape as the unit box

$$\hat{\mathcal{R}}_{box} = \{\boldsymbol{r} \mid \|\boldsymbol{r}\|_\infty \leq 1\} \tag{2.85}$$

For the external disturbance from weather and internal gains, the disturbance set is defined as follows

$$\mathcal{W} = \{\boldsymbol{w}_{\text{nom}} + \boldsymbol{w}_{\text{stoch}} \mid \boldsymbol{w}_{\text{stoch,i}}^T Q_i \boldsymbol{w}_{\text{stoch,i}} \leq 1,\ i = 1, 2, 3\} \tag{2.86}$$

As such, $\mathcal{W}$ is the direct product of three "uncorrelated" ellipsoidal uncertainty sets so that $\mathcal{W} = \mathcal{W}_{\text{sun}} \times \mathcal{W}_{\text{gains}} \times \mathcal{W}_{\text{temp}}$. $\boldsymbol{w}_{\text{nom}}$ is the nominal prediction of the uncertainty over the prediction horizon and $\boldsymbol{w}_{stoch}$ is the stochastic part of the uncertainty around the prediction. The three ellipsoids represent confidence sets that should cover a reasonable part of the possible outcomes for the disturbance. The choice of Q_i determines the size of the set $\mathcal{W}$ and should be done so that $\mathcal{W}$ contains the actual weather forecast with a high confidence (see , *e.g.* [MGL14]). Generally speaking, the selection of good uncertainty sets in classical robust optimization are a subject of active research [BGK14] and fall outside the scope of the present work, but notice that rather than fixing the uncertainty $\mathcal{W}$, the method proposed in this work could also be used to optimize for $\mathcal{W}$ as well and by doing so, evaluate how much prediction error in the weather and the internal gains can be accommodated. Finally we have $\hat{\Xi} = \hat{\mathcal{R}} \times \mathcal{W}$.

We consider here an affine controller and modifier function as in (2.65). Following the rules of the Swiss ancillary market, the bid is a fixed up/down capacity over the tracking period. This means the allowable modifier function is a uniform scaling of the uncertainty set (that is, time-varying capacity is not allowed). For clarity we keep the description of the uncertainty split between the reference to track and the external disturbance, so that: $\xi = (\boldsymbol{r}, \boldsymbol{w})$ and $\nu = (\nu_{\mathcal{R}}, \nu_{\mathcal{W}})$. We assume the disturbance is unknown at the time of the decision whereas the reference is revealed as it needs to be tracked: this results in an information structure that is depicted in Figure 2.4. We see that the modifier function could theoretically modify the

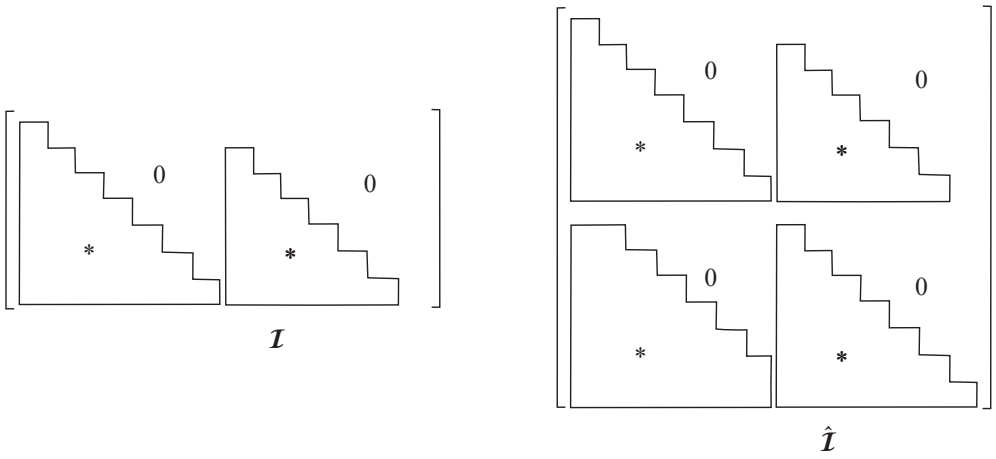

Figure 2.4 – Information structure for the example. $\mathcal{I}$ shows that decisions at time step t can depend on the reference up to time t and disturbance up to time $t-1$. $\hat{\mathcal{I}}$ is the resulting information structure for the modifier function.

uncertainty set so as to "mix" the external disturbance and the reference. In this application, it would not have physical sense so it is preferable to keep a block diagonal structure for the modifier's information structure. The disturbance uncertainty set is fixed a priori while the reference set can be modified. Furthermore, in the case that the reference set is a fixed up/down box along the horizon then the reference tracking set can only be scaled uniformly so that the modifier function will reduce to the simpler form:

$$\mathbf{L} = \begin{pmatrix} \gamma \mathbb{I}_N & 0_{N,Nn_u} \\ 0_{Nn_w,N} & \mathbb{I}_{Nn_w} \end{pmatrix} \tag{2.87}$$

To maximize the up/down capacity bid, it suffices to maximize the scaling factor γ. Notice that enforcing (2.87) implicitly enforces the requirement that $\boldsymbol{\nu} \in \hat{\mathcal{I}}$. The description of the uncertainty set $\hat{\Xi} = \hat{\mathcal{R}} \times \mathcal{W}$ can easily be put under the form of equation (2.64) since it is the direct product of a polyhedron with three ellipsoids.

The aim of the building operator is to maximize the financial reward gained from participating in the reserve market. In accordance with the Swiss market regulation, a payment proportional to the bid is made to the reserve provider, while the energy is bought at the day-ahead price, yielding the cost function

$$J = \mathbf{c_e}^T \bar{\mathbf{y}} - c_{\text{comm}} \gamma$$

where $\mathbf{c_e}$ is the vector of time-varying prices of electricity, $\bar{\mathbf{y}}$ is the baseline consumption and c_{comm} is the unit reward price of the power tracking commitment (hence promising to track $\pm$1kW for the commitment period is rewarded at the price c_{comm}). Notice that according to the Swiss market rule, the deviation in power consumption resulting from the tracking still needs to be paid for, therefore adding a stochastic component to the cost. It has however been shown

in [LQN+15] that this part is minor and can therefore be neglected.

The tracking error is sized proportionally to the tracking requirement so that tracking errors amounting up to 10% of the maximum tracking requirement are allowed. This yields:

$$\mathcal{E} := \{ \boldsymbol{e} \mid \|\boldsymbol{e}\|_\infty \leqslant 0.1\gamma \}$$

which results in a tractable reformulation as suggested in section 2.3.2.

A horizon of one day with a time step of one hour is considered. For the sake of illustration, we take $\|\mathbf{c_e}\| << \|c_{\text{comm}}\|$ in order to favor participation in the tracking commitment. The problem solved is a second-order cone problem with 200,000 non-zero variables and 900 second-order cone constraints. Solving time on a 2.7GHz i-Core 7 platform was 7 seconds. The optimal value of γ is 5.4, meaning that the building can offer a 5.4kW up/down power tracking capacity for a period of 10 hours. This represents 8% of the peak cooling power and 36% of the average power consumption for that day.

Figure 2.5 shows the trajectories generated in response to randomly generated weather and reference signals inside the reference sets. In each of the plots, the shaded band shows the reference tracking times. The different plots shows the average temperature in the building as well as in individual zones, the total power consumption in the building, the requested power consumption to be tracked on top of the nominal consumption, and the tracking error. It can be observed that in the nominal case, the power consumption increases during the day to compensate for higher solar radiation and outside temperature as shown in the bottom plot. Therefore, the baseline consumption varies during the day to maintain temperature at the nominal value of 23 ^{o}C. In addition, it is seen how the temperature and power consumption changes in response to varying tracking requests (depicted in the middle plot). As a result of the requested increase or decrease of the power consumption, the temperature respectively drops or rises in the rooms, within the prescribed comfort constraints. Notice that outside the tracking period the building is only subject to uncertainty in the weather and internal gains, while during the tracking period, it has to modify its power consumption in response to the tracking signal, which causes a larger variability of the inputs and outputs during that period.

Influence of the integral constraint in the reference set

Frequency control bids theoretically impose the providers to be able to offer up or down regulation for long periods of time, which can be limiting for loads. In this section, we propose a way of mitigating this issue by using time-correlated (meaning the constraint describing the set couple different time stages) reference sets with integral constraints that capture more

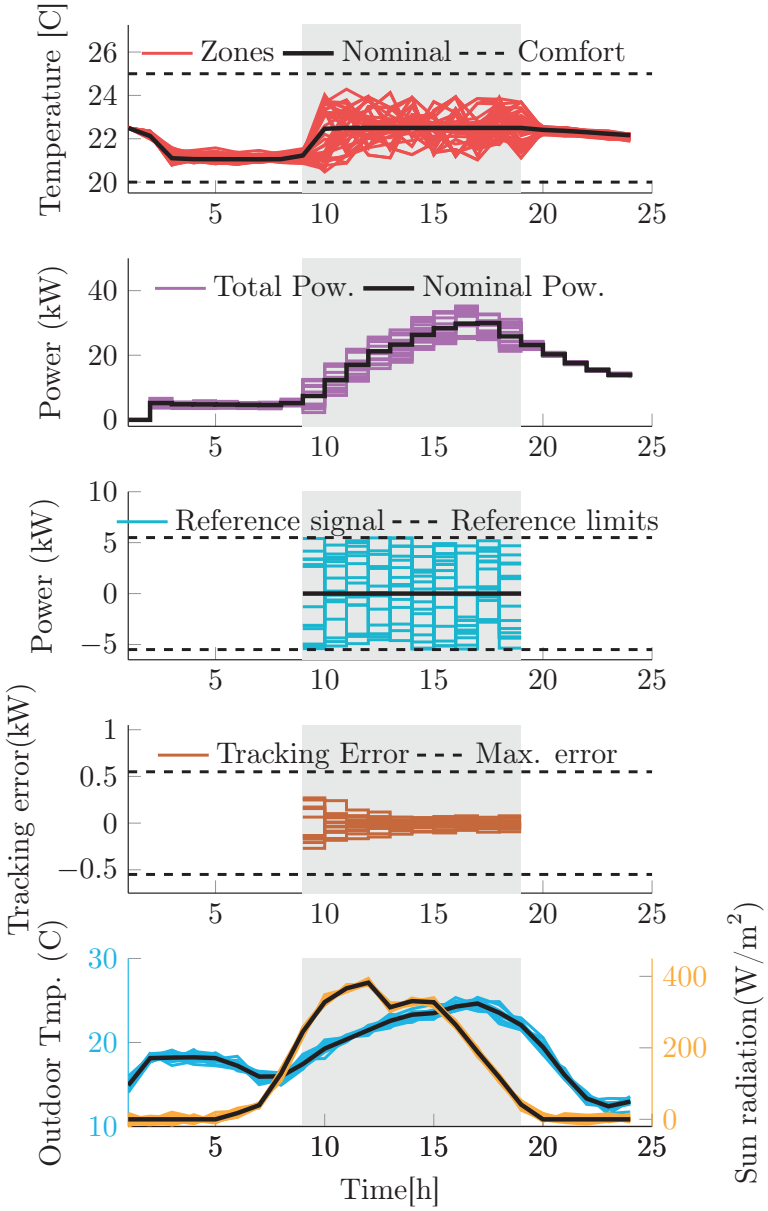

Figure 2.5 – Trajectories for different weather and reference scenarios in the optimized reference set. Shaded region is the tracking commitment period. Black lines show the 'nominal' scenario where the reference is zero and the weather takes its predicted value. From top to bottom: temperature in zones, total power consumption, AGC, tracking error, and weather scenarios.

accurately the capabilities of the load. Let us consider an uncertainty set of the form

$$\hat{\mathcal{R}}_{batt}(s_0) = \left\{ \mathbf{r} \;\middle|\; \begin{array}{ll} s(0) = \bar{s}_0 & \\ 0 \le s(t) \le s_{max}, & \forall t \in \mathbb{Z}_{[1,n]} \\ s(t+1) = s(t) + r(t) & \forall t \in \mathbb{Z}_{[1,n]} \\ -1 \le r_t \le 1 & \forall t \in \mathbb{Z}_{[1,n]} \end{array} \right\} \tag{2.88}$$

This set is analogous to the feasible set of a simplified battery model where s_t represents the state of charge of that battery. Therefore, we will refer to this uncertainty set as the "battery" reference set. Notice that contrary to the box reference set, the battery reference set is time-correlated, and our approach directly allows to account for that.

We respectively consider a box reference set (2.85) and a battery reference set (2.88) and compute the maximum bid that can be offered. Notice that the design of battery reference set requires the choice of a value for the integral constraint limit $s_{\max}$. This value was chosen as the worst case value of the integral of actual realization of the AGC over the year 2013. Therefore, for both reference sets, the system would at least be robust to any occurrence of the AGC that happened in 2013. The value computed is given by $s_{\max} = 5.6$. This suggests that the tracking signal is relatively well-behaved in the sense that it does not typically ask for maximum positive or negative power tracking for long periods of time.

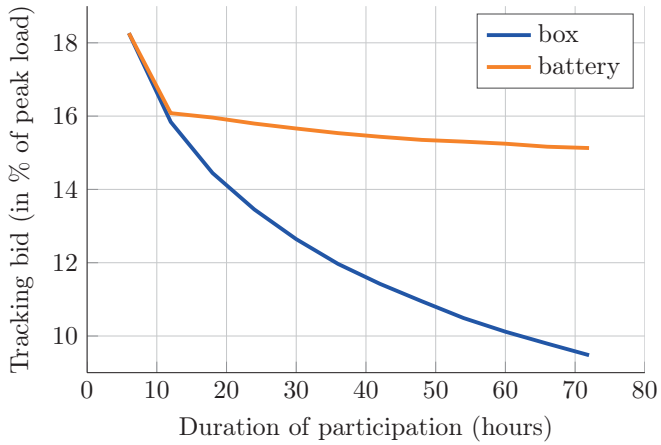

Figure 2.6 – Tracking capacity bid versus duration of participation for box and battery reference sets.

To study the influence of the integral limit in the reference set, the tracking bid is evaluated as a function of the duration of the tracking commitment. the preparation time is kept at 8 hours. The weather is considered known perfectly in advance in this case to rule out other factors of uncertainty in the computation. Beyond 66 hours, the computational burden becomes prohibitive. The maximum bids for the battery and the box uncertainty sets are reported on Figure 2.6. We can observe that, beyond 12h of consecutive participation, introducing an integral limit for the tracking commitment allows to increase the tracking bid, and more so as the duration of participation time increases. Thanks to the integral constraints, situations of long lasting positive or negative tracking requests are ruled out, thus relieving the tracking requirements on the building, and leading to less conservative solutions.

2.4.3 Case study 2: Hierarchical tracking with a building fleet

In this section we will present a simulation study that replicates a practical implementation of robust tracking commitment, for frequency regulation service with a fleet of 30 office buildings over a 24 hour period. We consider the case where an aggregator coordinates the buildings for providing the AGC tracking. Before the tracking service begins, each building characterizes its flexibility as a battery using robust tracking commitment, as explained in the previous Section 2.4.2. The battery models are collected by the aggregator, which allocates the AGC signal according to the estimated 'state of charge' of the battery abstractions. On the lower level, the buildings track the allocated power reference, using a simple MPC implementation. The schematic representation of the control structure of the building fleet is given in Figure 2.7

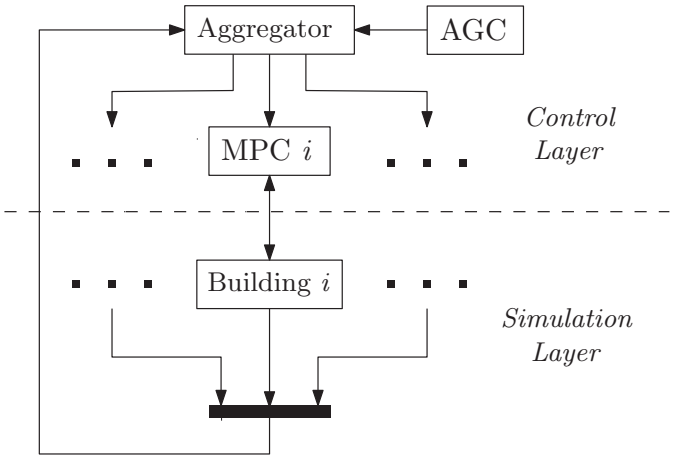

Figure 2.7 – Hierarchical control structure for AGC tracking with a fleet of buildings.

The flexibility model for each building represents a virtual battery, with state of charge and power capacity constraints. The batteries are simple integrators without leakage or efficiency terms

$$s^+ = s + \Delta_t p \tag{2.89}$$

where s is the state of charge, p is the power input and Δ_t is the sampling rate. The set of feasible power injection trajectories for the virtual battery is described by the set $\mathcal{R}_{\text{batt}} = \gamma\hat{\mathcal{R}}_{\text{batt}}$ as described in Section 2.4.2. Having access to the flexibility models $\mathcal{R}_{\text{batt}}$, the aggregator computes in real-time, the allocation of the current reference, the Automatic Generation Control (AGC) signal, and the predicted upcoming reference signals. The building MPC regulates the power consumption of the HVAC system, according to the COP prediction and the power consumption reference trajectory received from the aggregator.

The battery representations are computed for 24 hours with a crude sampling of 1 hour. Strictly speaking, the battery representations are only valid if the sampling rate is kept the same and a robust MPC is implemented for the closed loop operation. In the simulation study, we sample the virtual batteries at a faster rate of 5 minutes and operate a simple tracking MPC with a horizon of 2 hours. Due to the conservative flexibility computation, this practical implementation is

sufficient for keeping the building variables within constraints. Given a total power consumption reference y_{ref} and a desired temperature level θ_{ref}, the tracking MPC formulation is given by

$$\begin{aligned} \min_{\boldsymbol{u}} \quad & \sum_{i=0}^{N} \|y_{\boldsymbol{u}}(i) - r(i)\|_2^2 + \rho \|\theta_{\boldsymbol{u}}(i) - \theta_{\text{ref}}\|_2^2 \\ \text{s. t.} \quad & \boldsymbol{u} \in \mathcal{U}(x) \\ & y(i) = \text{COP}(i) \sum_{i=1}^{n_z} u(i) \end{aligned} \tag{2.90}$$

where ρ is the weighting factor for the temperature regulating term. Offset-free tracking of the power consumption is achieved by estimating the value of the actual COP with a simple observer, which compensates for the model mismatch between the MPC layer and the nonlinear HVAC system.

The aggregator also solves an MPC problem for determining the allocation of the AGC signal. The model for the aggregator represents a collection of single state simple integrator battery models provided by the buildings at the beginning of the day. The aggregator measures the power consumption realization of the buildings and updates the virtual battery states according to predefined baselines $\bar{\boldsymbol{y}}$. The AGC signal is allocated among the batteries while respecting their storage and power capacity constraints. At a given virtual state of charge level s, the allocation problem is given by

$$\begin{aligned} \min_{\boldsymbol{p}} \quad & \sum_{i=0}^{N} \| \sum_{k=1}^{m} p^k(i) - r_{\text{AGC}}(i)\|_1 + \nu \|p(i)\|_1 \\ \text{s. t.} \quad & \boldsymbol{p}^k \in \mathcal{R}^k(s^k), \quad \forall k \in \mathbb{Z}_{[1:m]} \end{aligned} \tag{2.91}$$

where r_{AGC} is the total AGC signal reference, $\boldsymbol{p}^k$ is the power trajectory of the kth virtual battery, representing the kth building and ν is the weighing factor for the regularization. The battery power consumption $\boldsymbol{p}^k$ is send to the building k as its output reference $\boldsymbol{r}^k := p^k$. We use 1-norm penalties for the tracking and regularization terms in the allocation problem to encourage a sparse activation of the batteries, which is a desired feature in general, when there is activation compensation for the participants or switched actuators. The sampling period of the allocation problem is the same as of building predictive controllers, that is 5 minutes, with a horizon of 2 hours. The resulting power trajectories are broadcast to building predictive controllers for tracking as depicted in Figure 2.7.

Implementation

The simulations are run with the openBuildNet (OBN) toolbox [NBG+16]. OBN is a distributed simulation platform that is designed and developed by Truong Nghiem, with Python interface and simulation support from the author of this thesis. OBN enables integrating simulation platforms that run on separate hardware, by means of a synchronized communication and input-output framework. The advantage of OBN lies in its scalability, that it is straightforward

to run a simulation including hundreds of agents, given that one has access to enough parallel or centralized computational power. For a detailed overview of the toolbox, see [NBG+16].

To obtain a more realistic simulation, we consider the nonlinear dynamics of the HVAC system, at a faster time scale than the thermal dynamics. Each building is controlled by a two-level hierarchical controller. In the upper layer, an MPC controller optimizes the electrical power consumption cost over a finite horizon and specifies the thermal power input to the zones in the building. The lower level controller consists of multiple proportional-integral (PI) controllers, which control the internal temperatures of the HVAC system as well as the thermal cooling/heating power input to the zones, as shown in Figure 2.8.

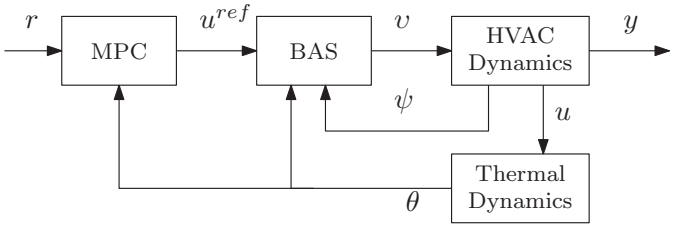

Figure 2.8 – Block diagram of the building control structure. The upper level MPC controller provides thermal energy u set-points to the building automation system (BAS). The BAS consists of three PI control loops that control the cooling power output of the HVAC system u, by means of adjusting water temperature at the evaporator of the chiller T_{evap}, and the air supply temperature that exits the AHU, T_{AHU}. It is the responsibility of MPC to choose a suitable cooling power reference u_{ref} that results in the desired level of power consumption y.

In order to have an estimate of the electrical energy used by the HVAC system, we use a time varying coefficient of performance (COP) estimation. For the cooling case it can be approximated as $\text{COP} = \eta \frac{T_{\text{cold}}}{(T_{\text{hot}} - T_{\text{cold}})}$, where η describes the overall efficiency of the heat pump. T_{cold} can be taken as equal to the water temperature at the evaporator T_{evap}, and T_{hot} can be taken as the ambient temperature T_{amb}. We use the disturbance values provided by OpenBuild for the ambient temperature prediction, whereas T_{evap} is regulated by the BAS at a known fixed value.

Results

The results of a simulation with a fleet of 30 office buildings over a 24 hours period is shown in Figures 2.9, 2.10 and 2.11. Each building is a small office building with identical dynamics, but randomized disturbance values and initial states. On average, the scaled virtual batteries describe state of charge limit of 6.3kWh and a bidirectional maximum power inflow of 2.1kW. Given the power inflow capacity of each battery, the fleet is able to commit a regulation service of 63kW. Table 2.2 summarizes important simulation parameters.

The results of the collective tracking is shown in Figure 2.10. The sampled AGC signal is taken from the data set provided by Swissgrid for the year 2013. We can observe that the fleet is able to maintain tracking with acceptable accuracy. Small deviations from the reference power

Table 2.2 – Simulation parameters

# Buildings	30
Building Type	Small Office
# Zones	5
# Internal State	15
Cooling Capacity [kW]	28
COP	2-3
Virtual Battery Power [kW]	2.1
Virtual Battery Storage [kWh]	6.3

trajectories can be observed, due to the mismatch of the linearized models and the nonlinear HVAC dynamics.

Evolution of the virtual battery variables are shown in Figure 2.9. We see that thanks to the 1-norm penalty, the allocation is sparse and chooses different batteries for different time periods, rather than demanding tracking from all batteries simultaneously. In Figure 2.11 we observe the effect of service provision on a single building and its HVAC system. We can observe that sparse activation creates rather large transients in the HVAC system and might not be desirable from this point of view. Another point is that the temperature values briefly exceed maximum allowed values, due to the simplistic implementation of the tracking MPC and unmodeled nonlinear dynamics in the MPC level. However, the violation is kept small and the situation is anticipated by the depleted virtual battery state of charge level, highlighted in Figure 2.10.

2.5 Conclusion

In this part of the thesis, we have formulated the problem of optimal robust tracking commitment and proposed a computationally tractable solution method. By implicit modification of uncertainty sets, the set of possible reference trajectories that can be tracked under additive disturbance with a guaranteed error bound can be efficiently computed over a finite prediction horizon. The presented tracking commitment framework is representative of many practical problems encountered in the hierarchical control of complex systems, that requires communication of tracking capability of subsystems to an upper level control layer. We have illustrated the description capability of the framework and the solution method with a practical example that investigates ancillary service provision to the power grid by commercial buildings.

The simulation studies show that the robust tracking commitment framework is suitable for complex hierarchical control tasks, as it allows building simple abstractions of agents. Utilizing these simple abstractions, a coordinator can easily decide on the allocation of the tracking task amongst the agents. Thanks to the robust optimization framework that incorporates disturbance uncertainty as well as reference uncertainty, the agents are guaranteed to comply with the tracking requests from the coordinator.

Our results have been also successfully demonstrated in experiments [FGQ+18] with a collaborative effort. The experiments were not included to this thesis as they were presented in detail in the PhD thesis of Tomasz Gorecki [Gor17].

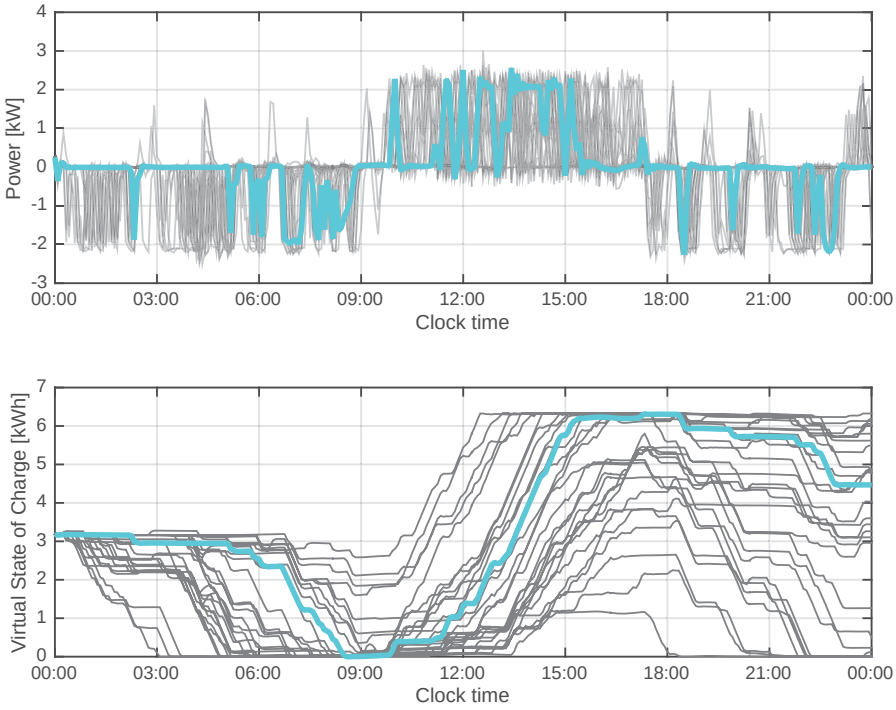

Figure 2.9 – Realized power inflows and state of charge for the virtual battery models of 30 buildings in the fleet. The trajectories of the building 10 is highlighted.

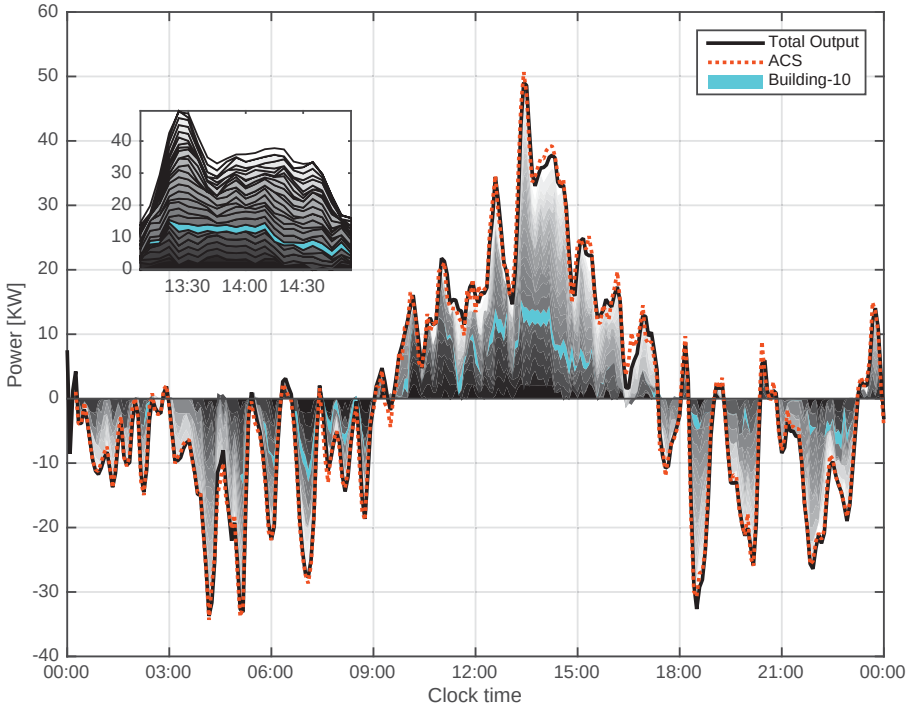

Figure 2.10 – Aggregated AGC tracking performance with the outputs from 30 buildings. The output of the building 10 is highlighted.

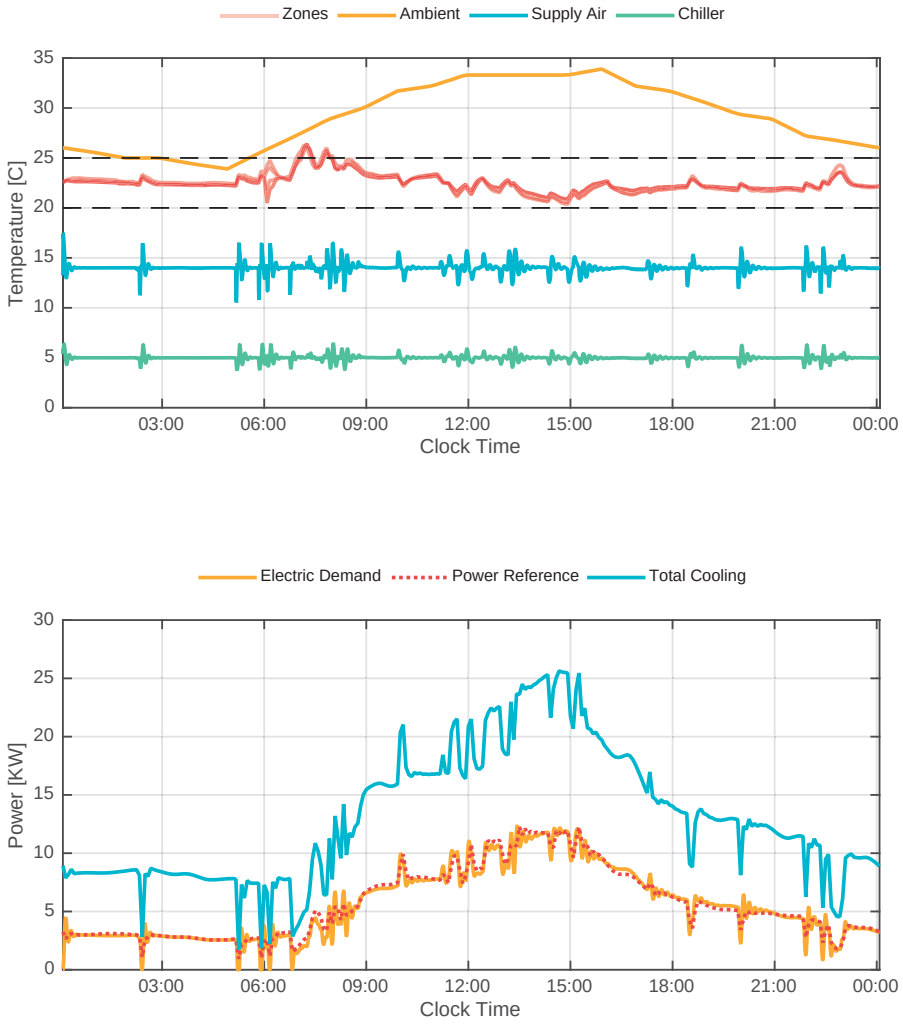

Figure 2.11 – Temperature values of the HVAC system and the thermal zones for building 10. The bottom plot shows the tracking of electrical power output and the resulting cooling power supply.

3 Coordinated Control and Optimization via Interior Point Decomposition

3.1 Introduction

Multi-agent optimization is receiving increased attention from the academy and industry with the rising prospect of smart grid applications. 'Smart' operation of the future's electricity grid will involve coordination of various large and small players such as generators, consumers and grid operators. Efficient coordination of these complex agents with various dynamics, constraints and interests is an inherently difficult task and introduces many challenges.

For example in a distribution grid, one can consider a demand response scheme where buildings, factories, and local generation facilities such as solar panels and generators are coordinated by a central agent, that is responsible for optimizing the aggregate power consumption/generation output and security of the grid. In this setting, it is desirable for the central agent to establish a consistent interface with local agents and consider only variables which affect the grid and aggregate objectives directly, whereas for commercial and residential agents, it is important to protect private data, such as internal states, constraints and objectives. Another requirement arises when optimization of the aggregate response is carried out in real-time. Due to time constraints for taking a decision, an algorithm should either terminate in a predetermined final time, or be able to terminate prematurely while having a feasible sub-optimal decision at hand. Another consideration is the communication effort for solving the multi-agent problem. Even though some distributed algorithms can achieve the same number of iterations as a central method, the number of communication steps per iteration can be significant.

Distributed and coordinated methods for solving optimization problems involving multiple agents is an active area of research. We distinguish here 'coordinated' schemes from 'distributed' schemes by the existence of a central agent that can collect information from all agents, inducing a 'star' topology on the information network. Many engineering applications, such as power grids and chemical plants, demand a central point of information, e.g. supervisory control and data acquisition (SCADA) systems, for enabling fault surveillance, intervention by human operators and data storage. Furthermore, without a central agent, the responsibility of modelling and managing the coupling network will fall on the agents, who might not have the domain knowledge. Another motivation for having a central agent is that the convergence of the optimization method will typically be much faster compared to fully distributed schemes since information can be shared efficiently. Distributed methods can be applied as a coordinated algorithm by choosing the splitting of the problem accordingly. On the other hand coordinated methods can often be made fully distributed if it is possible to allocate the task of the central node among all agents.

Coordinated solution methods are strongly linked to decomposition methods where the problem is split into a master problem and many sub-problems. Many decomposition based distributed optimization methods that are developed in the optimization research [BT97] are being extended and applied in multi-agent contexts [NND11, NÖ09, NOP10]. Decomposition methods can be collected under two main categories; primal and dual decomposition. In both approaches the optimization problem is split into a master problem that manages coupling data and many

sub-problems that manage local data, although the master problem can often also be distributed. Dual decomposition methods are based on Lagrangian relaxation and therefore cannot guarantee primal feasibility until convergence. Primal decomposition methods on the other hand decompose the problem directly. In the constrained setting, the master problem becomes non-differentiable and is usually solved with suitable first order methods such as sub-gradient or cutting plane methods [Ber08]. Applying second order methods is still possible by smoothing the problem with barrier functions, either in the dual setting [Zha05, NS09, KG14] or in the primal setting [Zha01, MO09b, BPJ17a]. Using second order curvature information, these interior-point decomposition methods become robust against scaling and can converge more quickly to high accuracy solutions compared to first order methods [NND11, BPJ17a].

In this part of the thesis, we first consider the requirements of local data hiding, and anytime termination of optimization algorithms in time critical applications. We show that the interior-point decomposition (IPD) method [Zha01, MO09b], or as we name it, the coordinated interior point method (C-IP) in a multi agent framework, satisfies both of these requirements since the implementation of the method only necessitates sharing information related to coupling variables, which connects agents through common constraints or objectives, and the method always outputs a feasible solution candidate at each iteration and achieves fast convergence. Next, we will consider the coordinated primal dual interior point (C-PDIP) method that enjoys improved convergence speed, while preserving the possibility of primal feasible iterations. For an algorithm used in the multi-agent setting, not only the number of iterations, but also the number of communications is of utmost importance for practical applicability. We will show that applying Mehrotra's predictor-corrector scheme [Meh92] in a multi-agent setting, results in excessive communication effort. To tackle this issue we introduce a decentralized predictor-corrector scheme in the C-PDIP setting, which enjoys the high performance of widely used predictor-corrector (PC) methods [Meh92, Gon96], without introducing too many communication steps. We further analyze the algorithm from a general point of view of a primal-dual interior point method with non-uniform complementarity perturbations, or barrier parameters. This analysis provides safeguards that can be used to ensure convergence of the decentralized PC scheme.

3.2 Preliminaries

3.2.1 Multi-agent optimization

We consider the following convex multi-agent optimization problem:

$$\begin{aligned} \min_{x^1,\ldots,x^m} \quad & \sum_{i=1}^{m} f^i(x^i) \\ \text{s. t.} \quad & \forall i \in \mathbb{Z}_{[1,m]} \,, \\ & g^i(x^i) \leq 0 \,, \\ & C^i x^i = y^i \,, \\ & \sum_{i=1}^{m} D^i y^i = d \,, \end{aligned} \tag{3.1}$$

where superscript i indicates that the function or variable belongs to agent i. The cost function and the constraint set are separable, except for the last equality constraint which introduces coupling between agents. We write (3.1) in the 'output coupling' format, in order to distinguish the local variables $x^i \in \mathbb{R}^{n_{x^i}}$, which determine feasibility and the cost of the agent, and the output variables $y^i \in \mathbb{R}^{n_{y^i}}$, which contribute to the coupling between agents. In practical applications, the dimension of the output n_y is often much smaller compared to the dimension of the local variables n_x, which makes the problem (3.1) suitable for decomposition methods. The coupling between the agents are captured by the equality constraint involving the coupling matrices $D^i \in \mathbb{R}^{n_d \times n_{y^i}}$ and the global parameter $d \in \mathbb{R}^{n_d}$.

Problem (3.1) captures a wide area of interest in estimation and control in networks [NND11]. Distributed control problems with dynamic or output coupling between agents can be formulated in the form of (3.1). Coupling in the cost and inequality constraints can also be incorporated into (3.1) by reformulations with auxiliary variables. See [NND11] for a survey on distributed problems on estimation and control and categorization of formulations and solution methods.

In the development of coordinated interior-point methods, we will rely on the following assumptions regarding the problem data.

Assumption 1. *The local cost functions f^i's are convex.*

Assumption 2. *The local constraint sets $\mathcal{G}^i := \{x^i : g^i(x^i) \leq 0\}$ are convex and compact.*

Assumption 3. *There exists a strictly feasible point for problem* (3.1).

Assumption 4. *The local output matrices, C^i's, have full row rank.*

The first two assumptions ensure convexity of the optimization problem. The third assumption is necessary for using interior-point methods, whereas the last assumption will facilitate the decomposition of the problem with respect to the output variables. Specifically, Assumption 4 is required to ensure that the variable y^i is a 'well defined' output, in the sense that for every

$y^i \in \mathbb{R}^{n_{y^i}}$ the relation $C^i x^i = y^i$ has a solution. In case Assumption 4 is not satisfied, one should introduce an intermediate output variable that facilitates decomposition. When C^i does not have full row rank, we can apply singular value decomposition (SVD) to obtain an intermediate output for which the output multiplier has full row rank

$$\mathcal{U}\mathcal{S}\mathcal{V} = C^i$$

where we remove rows and columns from $\mathcal{U}, \mathcal{S}, \mathcal{V}$ that correspond to zero singular values. The output and coupling matrices can be redefined as

$$\hat{C}^i = \mathcal{S}\mathcal{V}$$
$$\hat{D}^i = D^i\mathcal{U} \ .$$

where $\hat{C}^i$ has full-row rank and $\hat{y}^i = \hat{C}^i x^i$ is a valid output formulation.

3.2.2 Decomposition of optimization problems

In this section we will overview possible ways of solving (3.1) by means of decomposition methods, which distribute the computation effort among the agents. We define the local problem for an agent as

$$\begin{aligned} p^i(y^i) = \min_{x^i} \quad & f^i(x^i) \\ \text{s. t.} \quad & C^i x^i = y^i \\ & g^i(x^i) \leq 0 \ . \end{aligned} \tag{3.2}$$

$p^i(y^i)$ is called the 'primal' function [Ber08] and represents the optimal value of the problem (3.2) for a given y^i that parameterizes the constraints of (3.2). The primal function $p^i(y^i)$ of a convex program is convex, and non-differentiable in general [BV04].

In decomposition methods, the problem (3.1) is split into sub-problems that solve (3.2) and a master problem that coordinates them as we will see in the following sections. Dual decomposition methods use Lagrangian relaxation and do not provide feasibility until convergence, whereas primal decomposition methods decompose the problem directly and preserve feasibility. In the following section, we will present both primal and dual decomposition methods.

3.2.3 Primal decomposition

Primal decomposition operates directly on the original formulation of the problem. Using the local primal function definitions, the master problem is formulated as

$$\begin{aligned} \min_{y} \quad & p(y) \\ \text{s. t.} \quad & \mathbf{D}y = d \end{aligned} \quad \Longleftrightarrow \quad \begin{aligned} \min_{y^1,\dots,y^m} \quad & \sum_{i=1}^{m} p^i(y^i) \\ \text{s. t.} \quad & \sum_{i=1}^{m} D^i y^i = d \end{aligned} \tag{3.3}$$

with $y := [{y^1}^T, \dots, {y^m}^T]^T$ and $\mathbf{D} := [D^1, \dots, D^m]$. In general, due to the non-smoothness of the value functions $p^i(y^i)$ the master problem is solved with a first-order method. A standard method to solve such problems is the projected sub-gradient method [Ber08]

$$y_{k+1} = y_k - [t_k \theta(y_k)]^+ \tag{3.4}$$

where $\theta \in \partial p(y)$ is a subgradient of p at y_k and $[.]^+$ is projection on the local and global coupling constraints.

$$\sum_{i=1}^{m} D^i y^i = d, C^i x^i = y^i, g_i(x^i) \leq 0 \tag{3.5}$$

A global sub-gradient $\theta \in \partial p(y)$ can be constructed by stacking local sub-gradients as

$$\theta = \left[{\theta^1}^T, \quad \cdots, \quad {\theta^m}^T \right]^T . \tag{3.6}$$

The local sub-gradient θ^i can be obtained locally, by solving the problem (3.2). Since the local problem is parameterized by y^i in the constraints, the sensitivity theorem [Ber08] indicates that the negative of the dual variable $\lambda^{*,i}$ of the constraint $C^i x^i = y^i$ is a sub-gradient of $p(y^i)$.

$$-\lambda^{*,i} \in \partial p^i(y^i) \tag{3.7}$$

3.2.4 Dual decomposition

In dual decomposition we first dualize the problem with respect to the coupling constraint $\sum_{i=1}^{m} D^i y^i = d$. The dual problem takes the form

$$\max_{\gamma} \left\{ \mathrm{d}(\gamma) = \min_{y} \sum_{i=1}^{m} p^i(y^i) + \gamma^T \left(\sum_{i=1}^{m} D^i y^i - d \right) \right\} \tag{3.8}$$

The global dual function $\mathrm{d}(\gamma) : \mathbb{R}^{n_d} \longrightarrow \mathbb{R}$ can be further split among agents as

$$\max_{\gamma} \left\{ \mathrm{d}(\gamma) = \sum_{i=1}^{m} \mathrm{d}^i(\gamma) \right\} \tag{3.9}$$

where the local dual function $\mathrm{d}^i(\gamma)$ is the value function of the local problem (3.2) with a modified cost

$$\begin{aligned} \mathrm{d}^i(\gamma) = \min_{x^i} \quad & f^i(x^i) + \gamma^T(D^i y^i - \frac{d}{m}) \\ \text{s. t.} \quad & C^i x^i = y^i \\ & g^i(x^i) \leq 0 \,. \end{aligned} \tag{3.10}$$

In dual decomposition, the sub-gradient of the dual function is directly given by the infeasibility of the coupling constraint [Ber08]

$$d - \sum_{i=1}^{M} D^i y^i(\gamma) \in \partial \mathrm{d}(\gamma) \,. \tag{3.11}$$

In contrast to primal decomposition, the local output variables are not directly manipulated by the master problem, therefore feasibility of the y variables is ensured for any value of the dual variable γ, see for example [DKDS11]. This enables one to take a simple sub-gradient step without any projections to solve the master problem.

$$\gamma_{k+1} = \gamma_k + t(\sum_{i=1}^{m} D^i y^i(\gamma_k) - d) \tag{3.12}$$

At each iteration the primal output variables $y^i(\gamma)$ are determined by solving the local problems in parallel.

3.2.5 Interior-point methods

Barrier functions

The logarithmic barrier function, or simply the log-barrier, for an inequality $g(x) \leq 0$ with continuous $g : \mathbb{R}^{n_x} \longrightarrow \mathbb{R}$ is defined as

$$\phi^g(x) = -\log(-g(x)) \,, \tag{3.13}$$

where $\log : \mathbb{R}_+ \longrightarrow \mathbb{R} \cup \{\infty\}$ is the natural logarithm with the extension $\log(u) = -\infty$, for $u \leq 0$. The log-barrier has the properties that it is continuous over the set of strictly feasible points, $\{x \mid g(x) < 0\}$ and $\lim_{k\to\infty} \phi^g(x(k)) = \infty$ for a sequence $x(k)$ that converges to some $\bar{x}$ with $g(\bar{x}) = 0$.

The logarithmic barrier function for a set $\mathcal{G} = \{x \mid g_i(x) \leq 0, \quad i \in \mathbb{Z}_{[1,n_g]}\}$ can be defined as

$$\Phi^{\mathcal{G}}(x) = \sum_{i=1}^{n_g} \phi^{g_i} = -\sum_{i=1}^{n_g} \log(-g_i(x)) \,, \tag{3.14}$$

Lemma 6. *If the constraint set $\mathcal{G}$ has non-empty interior and is bounded, then the corresponding barrier function $\Phi^{\mathcal{G}}$ is proper, lower semi continuous and level-bounded over $\mathbb{R}^{n_g}$.*

Proof. If the constraint set $\mathcal{G}$ has non-empty interior and compact, then the barrier function is proper since there are $x : -g_i(x) > \epsilon, \forall i$ for some $\epsilon > 0$ and $\Phi^{\mathcal{G}}(x) < \infty$. Furthermore, note that the level sets $\{x \mid \Phi^{\mathcal{G}}(x) \leq a\}$ are subsets of relint($\mathcal{G}$). Since $\Phi^{\mathcal{G}}(x)$ is continuous over the set relint($\mathcal{G}$) then the level sets are closed, providing lower-semi continuity over $\mathbb{R}^{n_g}$. Finally, the level sets are bounded since $\mathcal{G}$ is itself bounded. □

The gradient and the Hessian of $\Phi^{\mathcal{G}}$ are given by

$$\nabla\Phi^{\mathcal{G}}(x) = \sum_{i=1}^{n_g} \frac{1}{-g_i(x)} \nabla g_i(x)$$
$$\nabla^2\Phi^{\mathcal{G}}(x) = \sum_{i=1}^{n_g} \frac{1}{g_i(x)^2} \nabla g_i(x) \nabla g_i(x)^T + \frac{1}{-g_i(x)} \nabla^2 g_i(x) \ .$$

Similarly for a polytopic set $\mathcal{C} := \{x \mid Gx \leq h\}$ we have

$$\Phi^{\mathcal{C}}(x) = -\sum_{i=1}^{n_g} \log(-G(i)x + h(i)) \ , \tag{3.15}$$

where $G(i) \in \mathbb{R}^{n_x}$ is the ith row of $G \in \mathbb{R}^{n_g \times n_x}$ and $h \in \mathbb{R}^{n_g}$. The barrier function for a set of constraints can be manipulated further by choosing a weighting of the barriers of individual constraints. We define such a weighted barrier function as

$$\Phi^{\mathcal{C}}(\omega, x) = -\sum_{i=1}^{n_g} \omega(i) \log(-G(i)x + h(i)) \ , \tag{3.16}$$

For a polytopic set, the weighted barrier function can be re-centered around any strictly feasible point, by choosing the weights $\omega \in \mathbb{R}_+^{n_g}$ such that the gradient is zero at the chosen center point [FE15].

$$\nabla\Phi^{\mathcal{C}}(\omega, x) = \sum_{i=1}^{n_g} \frac{\omega(i)}{-G(i)x + h(i)} G(i) \ . \tag{3.17}$$

For any strictly feasible point $x \in \text{relint}(\mathcal{C})$, there exists $\omega > 0$ such that $\nabla\Phi^{\mathcal{C}}(\omega, x) = 0$ [FE15]. In general ω is not unique, and various constraints or criteria could be applied for the selection, such as enforcing a minimum curvature condition on the Hessian. In the following, we will write $\Phi_y(.)$, to indicate that Φ is centered at some feasible point y, such that $\nabla\Phi_y^{\mathcal{C}}(y) = 0$.

The barrier method

Consider the following convex optimization problem.

$$\begin{aligned} \min_x \quad & f(x) \\ \text{s. t.} \quad & x \in \mathcal{G} \\ & Ax = b \end{aligned}$$

with $\mathcal{G} := \{x \,:\, g(x) \leq 0\}$. Instead of solving the inequality constrained problem above, we instead formulate a *barrier augmented* version with only equality constraints

$$\min_x f(x) + \beta\Phi^{\mathcal{G}}(x) \quad : \quad Ax = b \tag{3.18}$$

where $\beta > 0$ is the barrier weighting parameter that determines the level of approximation in the problem. For a large β we obtain a smooth problem that is easier to solve with second order methods, but the solution is a cruder approximation of the original problem. In the standard barrier method, the equality constrained augmented problem (3.18) is solved with the primal-dual Newton method [BV04] for a decreasing sequence $\{\beta_k\}$, with $\beta_k \longrightarrow 0$. The solution of the barrier problem converges to the solution of the original problem as $\beta_k \longrightarrow 0$ [RW09].

The barrier method is sketched below by Algorithm 1.

Algorithm 1 Barrier method

1: **given:** $\beta, \sigma \,:\, \beta > 0, 0 < \sigma < 1$
2: **while** $\beta < \epsilon$ **do**
3: $\quad x = \arg\min f(x) + \beta\Phi^{\mathcal{G}}(x) \quad : \quad Ax = b$
4: $\quad \beta = \sigma\beta$
5: **return** (x) ▷

Remark 4. *All iterations of the barrier method are feasible with respect to the constraint set $\mathcal{G}$. Furthermore when the primal-dual Newton method is used for solving* (3.18)*, if an iteration achieves feasibility with respect to the equality constraint $Ax = b$, then all further iterations remain feasible for all constraints [BV04].*

Primal-dual interior point method

This section provides a summary of standard primal-dual interior point (PDIP) framework. See [PW00, NW00] for a more in depth treatment and [Gon12] for an up to date survey. This type of methods are preferred over the simple barrier method, as they are known to converge much faster [BV04].

Consider the equivalent formulation of the problem (3.2.5), using slack variables

$$\begin{aligned} \min_x \quad & f(x) \\ \text{s. t.} \quad & g(x) + s = 0 \\ & Ax = b \qquad (\mu) \\ & s \geq 0 \qquad (\lambda) \end{aligned} \tag{3.19}$$

where $s \in \mathbb{R}^{n_g}$ is the slack variable for n_g constraints and the dual variables $\mu \in \mathbb{R}^{n_b}, \lambda \in \mathbb{R}^{n_g}$ are shown next to their corresponding constraints. The first order optimality conditions for

problem (3.19), or the so called KKT system, can be written as

$$r(x, s, \lambda, \mu) = 0. \tag{3.20}$$

Where $r = (r_{\text{dual}}^T, r_{\text{prim}}^T)^T$, $r_{\text{primal}} = (r_{\text{prim,in}}^T, r_{\text{prim,eq}})^T$ and the following definitions are used:

$$\begin{aligned} r_{\text{dual}}(x, \mu, \lambda) &= \nabla f(x) + A^T \mu + \mathcal{J}_g(x)^T \lambda \ , \\ r_{\text{prim,in}}(x, s) &= g(x) + s \ , \\ r_{\text{prim,eq}}(x, y) &= Ax - b \ , \\ r_{\text{comp}}(\lambda, s) &= \Lambda S \mathbf{1} \ . \end{aligned} \tag{3.21}$$

Here $\Lambda = \text{diag}(\lambda)$, $S = \text{diag}(s)$ and the *Jacobian* matrix for the constraints $\mathcal{J}_g(x)$ is defined as:

$$\mathcal{J}_g(x) = \left[\nabla g_1(x), \dots, \nabla g_n(x)\right]^T$$

PDIP methods operate by applying Newton's method on a smoothed version of the nonlinear set of equations (3.20). The smoothing is obtained by perturbing the complementary slackness condition

$$r_{\text{cent}}(\beta, \lambda, s) = r_{\text{comp}}(\lambda, s) - \boldsymbol{\beta} = \Lambda S \mathbf{1} - \boldsymbol{\beta}, \quad \boldsymbol{\beta} > 0 \tag{3.22}$$

with $\boldsymbol{\beta} \in \mathbb{R}_+^{n_g}$. The modified residual is defined by

$$\hat{r}(\boldsymbol{\beta}) = (r_{\text{dual}}^T, r_{\text{prim}}^T, r_{\text{cent}}^T(\boldsymbol{\beta}))^T.$$

Note that for a given perturbation parameter $\boldsymbol{\beta}$, when $r_{\text{prim,in}} = 0$, $r_{\text{cent}} = 0$, we have that

$$\lambda(i) = -\frac{1}{g_i(x)}, \quad \forall i \in \mathbb{Z}_{[1,n_g]} \tag{3.23}$$

and the KKT system (3.44) can be seen as the KKT system for the barrier augmented problem (3.18). This equivalence reveals the fact that the perturbation parameter $\boldsymbol{\beta}$ is the weighting factor for the barrier functions, when we eliminate λ and s from the problem. Therefore primal-dual interior point methods are a natural extension of the barrier method with extra degrees of freedom during the solution process [BV04].

In general the perturbation vector $\boldsymbol{\beta}$ is uniform, meaning that it is obtained by multiplying a scalar with a vector of ones as in $\boldsymbol{\beta} = \beta \mathbf{1}$. However we will let it attain different values for each constraint, similar to [CG08], which will allow us to decentralize the computation in Section 3.4.2. In order to retain the optimizer of the original KKT system, the perturbation parameter is decreased progressively. As $\boldsymbol{\beta} \longrightarrow 0$ for each component, the primal-dual pairs converge to the optimizer of (3.2.5) following the so called 'central path' with $\hat{r}(\boldsymbol{\beta}) = 0$ for some uniform $\boldsymbol{\beta}$. The Newton direction $(\Delta x, \Delta s, \Delta \lambda, \Delta \mu)$, for the perturbed KKT system can be found

by solving the following linear system of equations:

$$\begin{bmatrix} H(x,\lambda) & \mathbf{C}^T & \mathcal{J}_g(x)^T & 0 \\ \mathbf{C} & 0 & 0 & \mathbf{I}_s \\ \mathcal{J}_g(x) & 0 & 0 & 0 \\ 0 & S & 0 & \Lambda \end{bmatrix} \begin{bmatrix} \Delta x \\ \Delta\mu \\ \Delta\lambda \\ \Delta s \end{bmatrix} = - \begin{bmatrix} \nabla f(x) + A^T\mu + \mathcal{J}_g(x)^T\lambda \\ g(x) + s \\ Ax - b \\ \Lambda S\mathbf{1} - \boldsymbol{\beta} \end{bmatrix} \tag{3.24}$$

where we use the definitions :

$$H(x,\lambda) = \nabla^2 f(x) + \mathcal{H}_g(x,\lambda) \,,$$
$$\mathcal{H}_g(x,\lambda) = \sum_{i=1}^{n} \lambda_i^1 \nabla^2 g_i(x) \,.$$

The linear system of equations for the Newton step can be written in a compact form as

$$M\Delta z = - \begin{bmatrix} r_{\text{primal,dual}} \\ \Lambda S\mathbf{1} - \sigma\hat{\eta}\mathbf{1} \end{bmatrix}$$

with $\Delta z = (\Delta x, \Delta s, \Delta\lambda, \Delta\mu)$. After computing the Newton direction, it is necessary to choose an appropriate step-size t. 'Long step' methods use the maximum step-size that preserves strict positivity of the pair λ, s such that $(s + t\Delta s, \lambda + t\Delta\lambda) > 0$ holds, whereas 'short step' variants further restrict the step size to preserve 'centrality' and/or to ensure sufficient decrease in the residuals.

Remark 5. *If the primal residual is zero and all constraints are linear, the PDIP method will take primal feasible iterations [BV04] since equation* (3.24) *preserves linear equalities.*

The primal dual interior point method will terminate when the duality gap measure

$$\eta = \lambda^T s$$

and the primal dual residuals are reduced below desired thresholds. η becomes equal to the actual duality gap, when the iterates are primal and dual feasible, $r_{\text{dual}} = 0$, $r_{\text{primal}} = 0$ [BV04]. For a uniform perturbation vector $\boldsymbol{\beta}$, when the perturbed KKT system is solved, that is we obtain a point on the central path, we have that

$$\hat{\eta} := \frac{\eta}{n_g} = \frac{\lambda^T s}{n_g}, \quad \hat{\eta}\mathbf{1} = \boldsymbol{\beta} \,,$$

where the average complementarity value $\hat{\eta}$ provides a measure for the duality gap. In practice it is not necessary to wait for convergence to the central path, before reducing the perturbation parameter. In the standard 'feasible path following' method, for which convergence and complexity results are relatively easy to provide [NW00], the perturbation vector after each Newton step is chosen as

$$\boldsymbol{\beta} = \sigma\hat{\eta}\mathbf{1}, \quad \sigma \in (0,1) \,.$$

To ensure that proximity to the central path is preserved, the iterates are constrained to follow a shrinking region around the central path. Existing options in the literature [NW00, CG08] are

$$\mathcal{N}_2(\theta) = \{(x, s, \lambda) \in \mathcal{F}^0 \ : \ \|\Lambda S - \hat{\eta}\|_2 \leq \theta\hat{\eta}\} \tag{3.25}$$

$$\mathcal{N}_{-\infty}(\theta) = \{(x, s, \lambda) \in \mathcal{F}^0 \ : \ \lambda_j s_j \geq \theta\hat{\eta} \quad , \forall j \in \mathbb{Z}_{[1,n_g]}\} \tag{3.26}$$

$$\mathcal{N}_s = \{(x, s, \lambda) \in \mathcal{F}^0 \ : \ \frac{1}{\theta}\hat{\eta} \geq \lambda_j s_j \geq \theta\hat{\eta} \ , \forall j \in \mathbb{Z}_{[1,n_g]}\} \tag{3.27}$$

where the strictly feasible set $\mathcal{F}^0$ is defined as

$$\mathcal{F}^0 = \{(x, s, \lambda) \ : \ (r_{\text{prim}}, r_{\text{dual}}) = 0, \ (s, \lambda) > 0\}.$$

The PDIP method is sketched below as Algorithm 2.

Algorithm 2 Primal-dual interior point method

initialize: $z, \boldsymbol{\beta} = \beta\mathbf{1}, \sigma \ : \ \beta > 0, 0 < \sigma < 1$
while $\hat{\eta} > \epsilon_\eta$ or $\|r\| > \epsilon_r$ **do**
  $\hat{\eta} \longleftarrow \lambda^T s / n_g$
  $\boldsymbol{\beta} \longleftarrow \sigma\hat{\eta}\mathbf{1}$
  $\Delta z \longleftarrow -M^{-1} \begin{bmatrix} r_{\text{primal,dual}}(z) \\ \Lambda S\mathbf{1} - \sigma\hat{\eta}\mathbf{1} \end{bmatrix}$
  $t \longleftarrow \text{linesearch}(z, \Delta z)$
  $z \longleftarrow z + t\Delta z$
return $(x, \lambda, \mu, s) \longleftarrow z$ ▷

Predictor-corrector method

Practical implementations of PDIP methods involve many heuristic modifications for accelerating the algorithm. Mehrotra's predictor corrector [Meh92] together with Gondzia's multiple corrections [Gon96] are widely used in commercial solvers [BCC+10] and results in much faster convergence. In this section we will briefly explain the predictor-corrector approach.

We first define the 'affine scaling direction' as

$$M\Delta z^{\text{aff}} = -\begin{bmatrix} r_{\text{primal,dual}} \\ \Lambda S\mathbf{1} \end{bmatrix}$$

which gives a Newton step for the original KKT system. This direction is named after the *affine scaling method*, which was an early interior-point method(*e.g.* [Ber08]). The duality gap reduction parameter σ is chosen according to the performance of the affine scaling direction. The duality gap measure that would be achieved by the affine scaling direction with the maximum feasible step-size t^{aff} is

$$\hat{\eta}^{\text{aff}} = (\lambda + t^{\text{aff}}\Delta\lambda^{\text{aff}})^T(s + t^{\text{aff}}\Delta s^{\text{aff}})/n_g$$

The reduction parameter is chosen as $\sigma = \left(\hat{\eta}^{\text{aff}}/\hat{\eta}\right)^3$, which takes a small value if the affine scaling direction is preferable, that is $\hat{\eta}^{\text{aff}}$ is small, or takes on a large value, aiming for the central path. Further information is extracted from the affine scaling step by computing a corrector term, which aims to correct the error in the linearization by adding a second order approximation. The composite predictor-corrector step is computed by solving

$$H\Delta z^{\text{pc}} = -\begin{bmatrix} r_{\text{primal,dual}} \\ \Lambda S\mathbf{1} + \upsilon - \sigma\hat{\eta}\mathbf{1} \end{bmatrix}$$

where the corrector υ is defined as

$$\upsilon = \rho\Delta\Lambda^{\text{aff}}\Delta S^{\text{aff}}\mathbf{1}$$

and ρ is a weighing factor for the corrector (for example, $\rho = t^{\text{aff}}$ [SPT07] or $\rho = (t^{\text{aff}})^2$ [Car05]).

The PDIP method with prediction and correction with adaptive reduction factor, or the so called Mehrotra's predictor-corrector method, is sketched below as Algorithm 3.

Algorithm 3 Mehrotra's predictor–corrector method

initialize: $z, \boldsymbol{\beta} = \beta\mathbf{1}, \sigma : \beta > 0, 0 < \sigma < 1$
while $\hat{\eta} > \epsilon_\eta$ or $\|r\| > \epsilon_r$ **do**
  $\hat{\eta} \leftarrow \lambda^T s/n_g$
  $\boldsymbol{\beta} \leftarrow \sigma\hat{\eta}\mathbf{1}$
  $\Delta z^{\text{aff}} \leftarrow -M^{-1}\begin{bmatrix} r_{\text{primal,dual}}(z) \\ \Lambda S\mathbf{1} - \sigma\hat{\eta}\mathbf{1} \end{bmatrix}$ ▷ predictor
  $t^{\text{aff}} \leftarrow \text{maxstep}(z, \Delta z^{\text{aff}})$
  $\hat{\eta}^{\text{aff}} \leftarrow (\lambda + t^{\text{aff}}\Delta\lambda^{\text{aff}})^T(s + t^{\text{aff}}\Delta s^{\text{aff}})/n_g$
  $\sigma \leftarrow \left(\hat{\eta}^{\text{aff}}/\hat{\eta}\right)^3$ ▷ adaptive reduction factor
  $\upsilon \leftarrow \rho\Delta\Lambda^{\text{aff}}\Delta S^{\text{aff}}\mathbf{1}$ ▷ corrector
  $\Delta z^{\text{pc}} \leftarrow -M^{-1}\begin{bmatrix} r_{\text{primal,dual}}(z) \\ \Lambda S\mathbf{1} + \upsilon - \sigma\hat{\eta}\mathbf{1} \end{bmatrix}$
  $t \leftarrow \text{linesearch}(z, \Delta z^{\text{pc}})$
  $z \leftarrow z + t\Delta z^{\text{pc}}$
return $(x, \lambda, \mu, s) \leftarrow z$ ▷

3.3 Coordinated Interior-Point Method

In this chapter, we analyze the decomposition of the primal barrier method in a multi-agent setting, which we name the coordinated interior-point (C-IP) method. The multi-agent problem is smoothed by using local and global barrier functions, allowing Newton's method to be applied on the master problem, achieving fast convergence as well as feasible iterations.

Primal decomposition of the barrier method was first formulated by [Zha01] for solving two stage linear stochastic programs and later applied to more generic convex stochastic programs [MO09b, CM11]. The method is based on the primal barrier method [Ber08, BV04] which converts a constrained optimization problem into an unconstrained one by using barrier functions. The method is mostly used in the stochastic programming context [SMSS09] where the optimization is decomposed among scenarios that represent the possible realizations of an uncertain variable in the problem. Instead, we use C-IP in the context of multi-agent optimization and decompose the problem among agents. Using the barrier method for solving distributed MPC problems was previously proposed by [CS11]. However the authors of [CS11] consider a Gauss-Seidel type distributed framework for solving the barrier problem, where each agent solves a local problem and communicates coupling variables with its neighbor in a sequential manner. Instead, we will consider a 'coordinated' approach with a central agent for solving the barrier problem.

Recently, several second order distributed optimization methods based on Newton's method have been proposed [PHA15, NS09, WO12]. These methods are similar to C-IP as they apply the barrier method to their problem and compute Newton steps in a distributed fashion. The method in [PHA15] is applicable to problems where agents are 'loosely' coupled and the descent direction is computed in a distributed manner without any central agent. On the other hand, IPD can handle more generic coupling constraints and cost functions by relying on a central agent. Similarly, the method proposed in [WO12] is applicable to a specific structure arising from network utility maximization.

A similar framework is applying dual decomposition to the barrier method, as proposed in [Zha05, NS09, KG14]. Since these methods are based on dual decomposition, they cannot maintain primal feasible iterations for both the local and the network constraints. On the other hand the primal decomposition applied in C-IP maintains global primal feasibility of the iterations and enables early termination and safe operation.

In the following, we will briefly introduce the method, and discuss its attractive properties for solving a multi-agent optimization problem: fast convergence, primal feasible iterates for both local and coupling constraints, data hiding for local agents and problem size reduction for the central agent. For finding a feasible initial point and feasibility assessment in line-search, we propose novel methods based on simple surrogate sets [BGJ17, ZdH15]. We also present an application of the method to an example from smart-grids: voltage support in a distribution grid with active buildings, together with comparison of performance with a popular first order

distributed optimization algorithm, *alternating direction method of multipliers* (ADMM)(*e.g.* [Boy10]).

The C-IP method consists of applying the barrier method for solving the multi-agent problem (3.1) with primal decomposition. The *barrier augmented* multi-agent problem can be written as

$$\begin{array}{ll} \min\limits_{y} & \rho(\beta, y) \\ \text{s. t.} & \mathbf{D}y = d \end{array} \quad \Longleftrightarrow \quad \begin{array}{ll} \min\limits_{y^1,\ldots,y^m} & \sum\limits_{i=1}^{m} \rho^i(\beta, y^i) \\ \text{s. t.} & \sum\limits_{i=1}^{m} D^i y^i = d \end{array} \tag{3.28}$$

for some $\beta \in \mathbb{R}_+$ where the barrier augmented primal functions $\rho^i(\beta, y^i)$ are defined as

$$\begin{array}{rll} \rho^i(\beta, y^i) = \min\limits_{x^i} & f^i(x^i) + \beta\Phi^{\mathcal{G}^i}(x^i) \\ \text{s. t.} & C^i x^i = y^i \end{array} \tag{3.29}$$

Here, $\Phi^{\mathcal{G}^i}(\beta, x^i)$ is the barrier function, as defined in Section 3.2.5, for the feasible set defined by the local constraints $\mathcal{G}^i = \{x^i \mid g^i(x^i) \leq 0\}$. It was shown by [CM11] that, when logarithmic barrier functions are used, $\rho^i(\beta, .)$ is a self-concordant barrier function for the feasible set of the output variable y^i, defined as

$$\mathcal{F}^i := \{y^i \mid \exists x^i \ , \ y^i = C^i x^i \ , \ x^i \in \mathcal{G}^i\} \ . \tag{3.30}$$

In general, especially when $n^{y^i} << n^{x^i}$, it is difficult to obtain an explicit representation of $\mathcal{F}^i$ with functional constraints directly on y^i. Therefore instead of constructing a barrier function for the set $\mathcal{F}^i$, we can use ρ^i as a barrier function. Similarly, for ρ^i, an explicit description is not available, however at a given point y^i, the value of the function as well as its gradient and Hessian can be obtained by solving the equality constrained convex optimization problem (3.29). The gradient and the Hessian of ρ^i is given by

$$\begin{aligned} \nabla_{y^i}\rho^i(\beta, y^i) &= -\lambda^{*,i} \\ \nabla^2_{y^i y^i}\rho^i(\beta, y^i) &= \left(C^i(\nabla^2_{x^i x^i}L^i(x^{*,i}, \lambda^{*,i})^{-1}{C^i}^T\right)^{-1} \end{aligned} \tag{3.31}$$

where L^i is the Lagrangian for the local problem (3.29)

$$L^i(x^i, \lambda^i) = f^i(x^i) + \Phi^{\mathcal{G}^i}(\beta, x^i) + {\lambda^i}^T(C^i x^i - y^i) \tag{3.32}$$

and the optimal primal dual pair $(x^{*,i}, \lambda^{*,i})$ satisfy the second order sufficient conditions for optimality [Ber08]. The relations (3.31) describe the first and second order sensitivity of the optimal value with respect to the variable y^i and can be derived by applying the implicit function theorem to the optimality conditions, see [CM11] and the alternative proof in Appendix A.1.

3.3.1 Coordinated Newton method

Thanks to the smoothing with barrier functions, the problem (3.28) can be solved with second-order methods, instead of sub-gradient methods as in Section 3.2.3. The unconstrained minimization step in the primal-barrier method is carried out with Newton's method, which iteratively minimizes a second order approximation of the barrier-augmented cost function using the gradient and the Hessian

$$\begin{aligned} \nabla\rho(\beta, \boldsymbol{y}) &= \left[\nabla\rho^1(\beta, y^1)^T \quad \cdots \quad \nabla\rho^m(\beta, y^m)^T\right]^T \\ \nabla^2\rho(\beta, \boldsymbol{y}) &= \text{blkdiag}(\nabla^2\rho^1(\beta, y^1), \ldots, \nabla^2\rho^m(\beta, y^m)) \end{aligned} \tag{3.33}$$

As seen from the above equations, one needs to collect $\nabla\rho^i(\beta, y^i), \nabla^2\rho^i(\beta, y^i)$ from all agents at each iteration. Each agent can compute these values in parallel, by solving (3.29) and using relations (3.31), and send them to the central agent. The central agent then computes the descent direction Δy, by computing a primal-dual Newton step [BV04] for the problem 3.28,

$$\begin{bmatrix} \nabla^2\rho(y, \beta) & \mathbf{D}^T \\ \mathbf{D} & 0 \end{bmatrix} \begin{bmatrix} \Delta y \\ \Delta\gamma \end{bmatrix} = - \begin{bmatrix} \nabla\rho(y, \beta) + \mathbf{D}^T\gamma \\ \mathbf{D}y - d \end{bmatrix} = - \begin{bmatrix} r_{\text{dual}} \\ r_{\text{prim}} \end{bmatrix} \tag{3.34}$$

where γ is the dual variable for the coupling constraint $\mathbf{D}y = d$ and r_{dual} and r_{prim} are the dual and primal residuals respectively. The solution to the above linear system is given by

$$\begin{aligned} \Delta y &= -(\nabla^2\rho)^{-1}(r_{\text{dual}} + \mathbf{D}^T\Delta\gamma) \\ \Delta\gamma &= \left(\mathbf{D}(\nabla^2\rho)^{-1}\mathbf{D}^T\right)^{-1}\left(r_{\text{prim}} - \mathbf{D}(\nabla^2\rho)^{-1}r_{\text{dual}}\right) \end{aligned} \tag{3.35}$$

Thanks to the block diagonal structure of the Hessian $\nabla^2\rho$ and the stacked structure of $\mathbf{D}$ the solution can also be written as

$$\begin{aligned} \Delta y^{i,*} &= -(\nabla^2\rho^i)^{-1}r^i_{\text{dual}} - (\nabla^2\rho^i)^{-1}(D^i)^T\Delta\gamma^* \\ \Delta\gamma^* &= \left(\sum_{i=1}^m D^i(\nabla^2\rho^i)^{-1}(D^i)^T\right)^{-1}\left(\sum_{i=1}^m D^i(\nabla^2\rho^i)^{-1}r^i_{\text{dual}} + d - \sum_{i=1}^m D^i y^i\right) \end{aligned} \tag{3.36}$$

Note that the computation of the Newton direction requires summation over the data of all the agents and computing the descent direction for the dual variable γ. After this step, which can be done by a central coordinator or in a distributed fashion, the descent direction of the agents can be computed at the local level. However if the model of the coupling constraints D^i is not known by the local agent, as is the case in power networks, the descent directions can be directly computed at the central node and broadcast to the agents. We will describe the algorithm considering the latter case. The step-size for the Newton step is computed according to a suitable line search strategy, which we will discuss in the following section. The algorithm proceeds until the Newton decrement δ is reduced under a specified threshold. A description of the method is given in Algorithm 4.

Algorithm 4 Coordinated Newton method

given: ϵ, β
while $\delta > \epsilon$ **do**
  for $i \in [1, \dots, m]$ **do** ▷ parallel
    solve $\min_{x^i} f^i(x^i) + \Phi^{\mathcal{G}^i}(\beta, x) \mid C^i x^i = y^i$
    return $\rho^i(\beta, y^i), \nabla\rho^i(\beta, y^i), \nabla^2\rho^i(\beta, y^i)$
  $\Delta y \leftarrow$ solve (3.35)
  $t \leftarrow \text{linesearch}(y, \Delta y)$
  $y \leftarrow y + t\Delta y$
  $\delta \leftarrow \sqrt{\Delta y^T \nabla^2 \rho(\beta, y) \Delta y}$ ▷ Newton decrement
return y

3.3.2 Initialization with surrogate sets

If strictly feasible iterations are required, the primal barrier method requires an initial point that is feasible for both local and coupling constraints. If such a point is not available, one can solve a 'Phase 1' problem for finding one; for various options see [BV04]. Another possibility, also suggested by [MO09b], is to add slack variables with large penalties to the problem. Due to the distributed nature of IPD, these methods will require additional communication and optimization steps.

We propose using simple inner approximations of the feasible sets $\mathcal{F}^i$'s of local outputs. Note that an explicit description of $\mathcal{F}^i$ is in general not available, even to the local agent itself. However, one can use the methods proposed in Section 2.3.1, [BGJ17, ZdH15] to recover a simple surrogate set

$$\mathcal{S}^i \subset \mathcal{F}^i, \tag{3.37}$$

with a relatively large volume [ZdH15], using robust programming. For initialization, each agent computes the set $\mathcal{S}^i$, which can be an ellipsoidal or a polytopic set, and passes it to the central agent. The central-agent then solves a surrogate feasibility problem to find a feasible point that satisfies coupling constraints within the aggregate set. The feasibility problem can be formulated in different ways. One option is to minimize the maximum constraint violation within the set $S^1 \times \cdots \times S^m$. If the minimizer is not feasible, the procedure can be repeated by re-computing the surrogate sets around the new point. This method is likely to generate a feasible point with very few communication steps between central and local agents, if the volumes of surrogate sets are substantial. In our application example in Section 3.3.4, a single step of the method was sufficient to find a strictly feasible point.

3.3.3 Line search

One of the most important ingredients of Newton method is the line-search, for determining the step-size in the descent direction. The step-size should be selected so as to preserve feasibility

and to guarantee a sufficient decrease at each iteration [Ber08]. We can split the line-search into two steps; first, a feasibility search for determining the maximum step size that can maintain feasibility and second, a minimizing search for finding the step-size with a sufficient decrease in the cost function.

Feasibility search with Dikin's ellipsoid:

For a given descent direction, each agent needs to solve an additional feasibility problem in order to determine the maximum feasible step-size. However this can be avoided by limiting the search within a surrogate set that is communicated to the central agent together with gradient and Hessian information. The surrogate sets can be computed as mentioned in the previous section. However, the computational cost of finding $\mathcal{S}^i$ at every step might be prohibitive depending on the dimensionality. For this reason we propose using a feasible ellipsoid that can be constructed using the Hessian of the barrier function which is already at hand.

Lemma 7. *At any strictly feasible local variable $\hat{x}$ the ellipsoid defined by*

$$\mathcal{E}(\hat{x}) = \{y \,:\, (y - C\hat{x})^T (C\nabla^2 \Phi^{\mathcal{G}^{-1}} C^T)^{-1} (y - C\hat{x}) < 1\} \tag{3.38}$$

is contained in the feasible set for the output variable y; $\mathcal{E}(\hat{x}) \subset \mathcal{F}$.

Proof. At any strictly feasible $\hat{x} \in \mathcal{G}$, the Dikin's ellipsoid is defined by

$$\mathcal{D}(\hat{x}) := \{x \,:\, (x - \hat{x})^T \nabla^2 \Phi^{\mathcal{G}}(\hat{x})(x - \hat{x}) < 1\} \ . \tag{3.39}$$

For a self concordant barrier function $\Phi^{\mathcal{G}}$ of set $\mathcal{G}$, $\mathcal{D}(\hat{x})$ is contained within the set $\mathcal{G}$; $\mathcal{D} \subset \mathcal{G}$ [NN94]. It can be shown that $\mathcal{E}(\hat{x})$ is the image of the Dikin's ellipsoid $\mathcal{D}(\hat{x})$, under the linear operator C, see Appendix A.2. Therefore, $\mathcal{E}(\hat{x})$ is contained in the feasible set $\mathcal{F}$ for the output variable y, which is the image of the set $\mathcal{G}$ under C. □

After solving the local problem (3.29) the agents can easily construct the ellipsoid $\mathcal{E}^i(x^{*,i})$ and pass its parameters to the central agent, avoiding any further communication for a feasibility check. The resulting step-size can be significantly smaller compared to an approach that uses the actual maximum step-size, however the burden of extra communication and optimization by all agents is eliminated.

Minimizing search

Minimizing line search is carried out until a sufficient descent condition [BV04] is satisfied. For $0 < \alpha < 0.5$ the condition is given by

$$\rho(y + t\Delta y) < \rho(x) + \alpha t \nabla \rho(y)^T \Delta y \ . \tag{3.40}$$

The standard method for finding a point satisfying (3.40) is backtracking, which consists of decreasing the step-size by a predetermined reduction factor $0 < \kappa < 1$ until (3.40) is satisfied. This method will require communicating the step-size to individual agents, solution of the problems (3.29) and re-communication of the value to the central-agent. Depending on whether the burden of communication or optimization dominate, one can take different approaches, such as; aggressive backtracking, adaptive step-size that relates to the Newton decrement [Zha01, MO09a, WO12] or using pre-computed values collected from agents.

3.3.4 Case study: Voltage support with commercial buildings

In this section we present an application of C-IP to a distributed MPC problem related to smart grids. We consider a case where a distribution grid operator coordinates a group of commercial buildings that joins a Demand Response (DR) program, in order to balance the voltage surge caused by a large solar plant.

The grid under consideration is a simplified version of the IEEE 123 test feeder, taken from [BZ16], which is modeled as a single-phase system, and consists of 56 nodes. The power flow equations are linearized with the method of [BZ16] that is suitable for distribution grids. The linear model is accurate enough for our demonstration purposes and all following results are computed with it. The grid is populated by a mix of commercial buildings, the properties of which are given below in Table 3.1.

Table 3.1 – Summary of buildings participating in the DR study

Type	Cooling Cap.[kW]	Thermal Zones	# in DR
Large Office	250	18	2
Small Office	30	5	7
Warehouse	150	3	3

The data for the building models are obtained from [UR] and [sta16]. We consider a summer scenario, where the cooling systems of buildings that participate in DR are controlled by MPC, which relies on a thermal model that is generated with the OpenBuild Toolbox [GQJ15b] as explained in Section 2.4. The grid is mostly populated by small offices representing static loads, whereas 12 buildings: two large offices, three warehouses and seven small offices contribute to the DR program for voltage support.

The grid operator is responsible for maintaining voltage magnitudes of all nodes within ±5% of the nominal value. Starting from an initial grid state, it predicts the solar power in-feed and

power injections from all nodes and solves the following multi-agent optimization problem.

$$\begin{aligned} \min \quad & \sum_{i=1}^{m} p^i(y^i) \\ \text{s. t.} \quad & v(k) = \sum_{i=1}^{m} G^i y^i(k) + \bar{v}(k) \quad \forall k \in \mathbb{Z}_{[1,N]}, \\ & 0.95 \le v(k) \le 1.05 \qquad\qquad \forall k \in \mathbb{Z}_{[1,N]}, \end{aligned} \tag{3.41}$$

where N is the time horizon, $\bar{v}(k) \in \mathbb{R}^l$ represents the predicted voltage magnitudes (in p.u.) of all l nodes at step k, $y^i \in \mathbb{R}^{n_{y,i}}$ represents the active power injections from the buildings that participate in DR. $G^i \in \mathbb{R}^{l \times n_{y,i}}$ models the effect of the injections from the ith node on the voltage magnitudes of all nodes. We assume that the power factor of the buildings are constant, and therefore reactive power computations can be embedded inside G^i. The voltage constraints are respected while the power injections are allocated fairly according to the local cost functions p^i, defined as

$$\begin{aligned} p^i(y^i) = \min_{u^i} \quad & \sum_{k=1}^{N} \|\theta^i(k) - \theta_{i,\text{ref}}\|_Q^2 + \pi(k) y^i(k) \\ \text{s. t.} \quad & \forall k \in \mathbb{Z}_{[1,N]}, \\ & x^i(k+1) = A^i x^i(k) + B^i u^i(k) + E^i d^i(k) \\ & \theta^i(k+1) = C^i x^i(k) + D^i u^i(k) \\ & y^i(k) = F^i(k) u^i(k) \\ & \theta^i_{\min} \le \theta^i(k) \le \theta^i_{\max}, \quad u^i_{\min} \le u^i(k) \le u^i_{\max} \end{aligned} \tag{3.42}$$

where N is the horizon length, the variables x, u, d and (θ, y) describe state, input, disturbance and outputs for the linear thermal dynamics. The input to thermal dynamics u is the thermal energy over the sampling period. The output θ describes the mean zone temperatures, which are constrained to lie inside the comfort limits. The electrical power consumption y^i is modeled as a time varying linear function of the thermal cooling power input to the building, which is limited by actuation constraints. The cost function represents the economic cost of electricity determined by the price π and penalizes deviation from the reference temperature values that represent ideal comfort conditions.

The problem (3.41) can be formulated as the standard multi-agent problem (3.1) and thus can be solved with C-IP. In this setting, each building can compute the barrier augmented cost function related to its power consumption and communicate with the grid operator that collects data from all-agents and implements the primal barrier method, given in Algorithm 1, with Newton's method, given in Algorithm 4.

The results of the optimization are shown in Figures 3.1-3.4. Figure 3.1 shows the voltage profiles of all nodes without and with coordination. The control problem is solved for the time of day between 8:00 and 15:00, sampled every half-hour, resulting in a horizon length $N = 14$. For initialization the method explained in Section 3.3.2 is used, where the central agent collects feasible ellipsoids from agents and minimizes the maximum voltage magnitude among

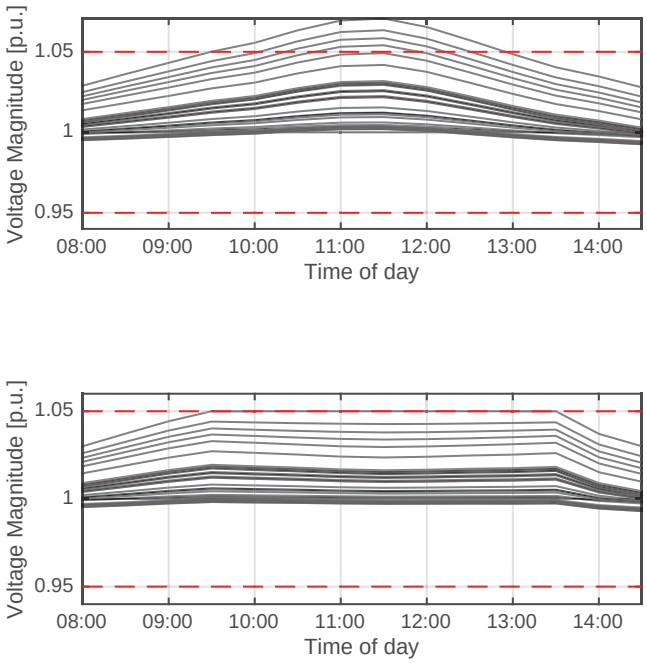

Figure 3.1 – Predicted voltage magnitudes, without the demand response coordination (top), and with optimal coordination of 12 buildings in the grid (bottom).

all nodes. After this step, IPD proceeds from a feasible point, and maintains local and global feasibility as shown in Figure 3.2. For the line-search, Dikin's ellipsoids around the current point for feasibility and regular backtracking for sufficient decrease search are used, resulting in a reasonable number of optimization and communication steps, as shown in Figure 3.3.

Finally, we provide a comparison of C-IP with a first order distributed method: ADMM. In order to preserve the same communication structure, we introduce global copies of local outputs as $\boldsymbol{y}_{\text{glob}} = \boldsymbol{y}$ and apply ADMM to the partial Lagrangian with respect to this constraint. By doing this, one can split the primal updates into a central and parallel local problems similar to the C-IP method. The penalty parameter for ADMM is adaptive as described in [Boy10]. We can observe from Figure 3.4 that ADMM quickly reduces the deviation from optimal energy profiles, however the solution candidates remain infeasible for many iterations. On the other hand C-IP always takes feasible steps by keeping the maximum voltage values below limits and achieves much higher accuracy in a smaller number of communication steps, compared to ADMM.

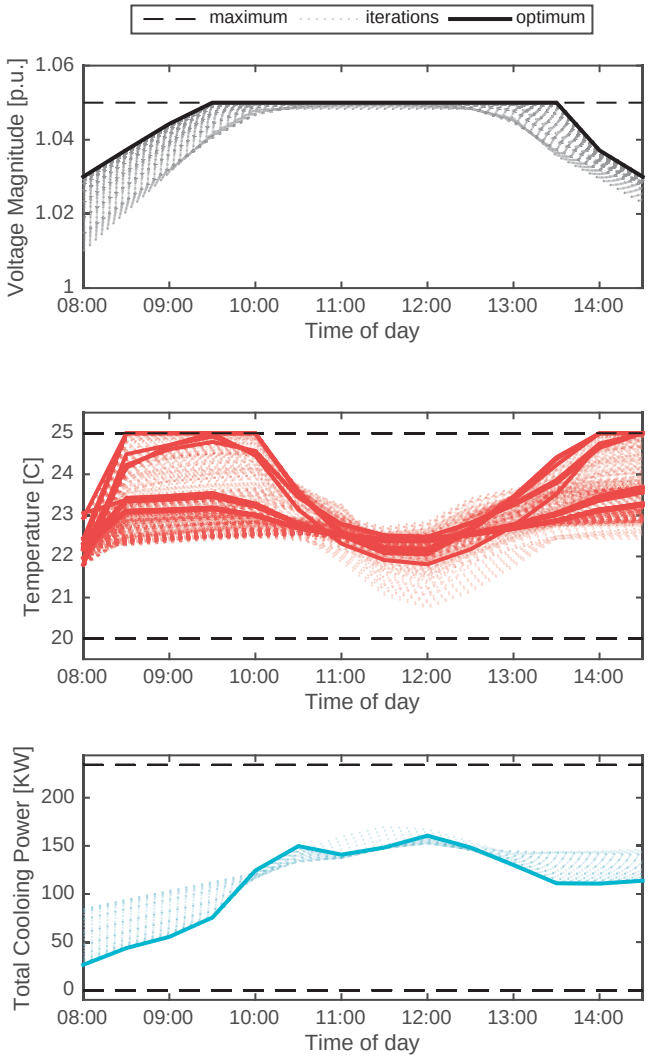

Figure 3.2 – The top figure shows the voltage magnitude at node 32, where the solar power infeed is placed. The figures below show the zone temperatures and total cooling power for a large office building placed at node 31.

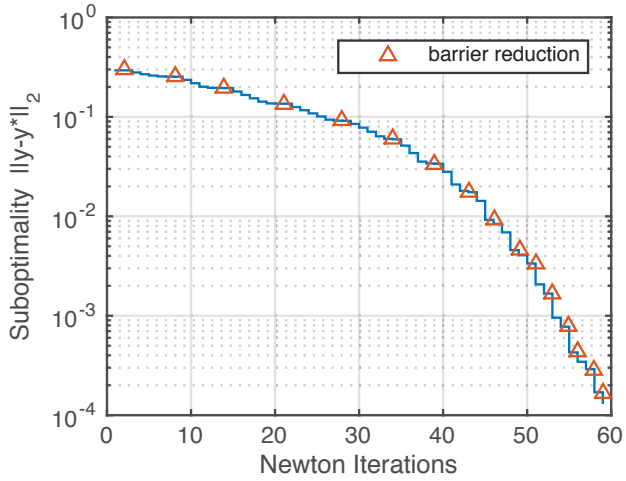

Figure 3.3 – Evolution of suboptimality measure $\|y - y^*\|_2$ for the problem (3.41), using the Dikin's ellipsoid feasibility search and backtracking with a reduction factor of $\kappa = 0.8$. Total Newton steps are 58 whereas 77 communication steps were carried out. Barrier shrinking factor μ is set to 0.4.

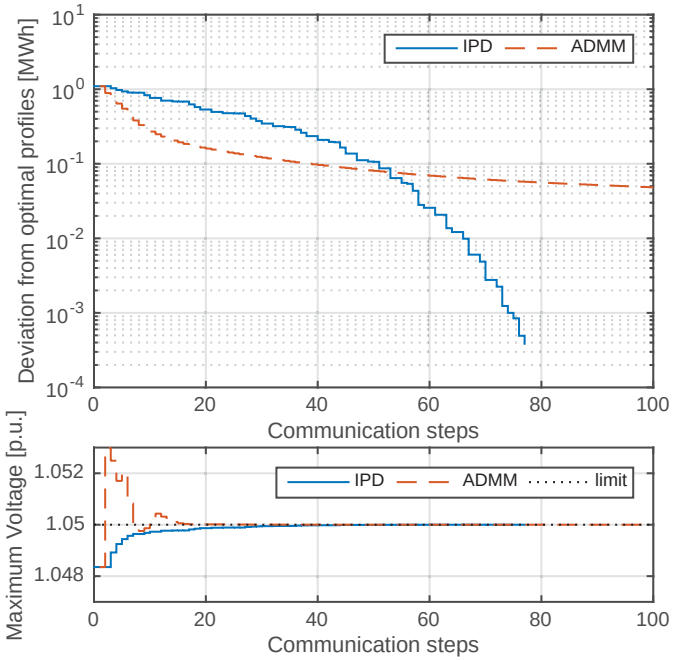

Figure 3.4 – Comparison between ADMM and IPD applied to the problem (3.41). The deviation from optimal profiles is defined as the sum of absolute difference in MWh, with respect to the optimal power consumption profiles, for all buildings: $\sum_{i=1}^{n_a} \sum_{k=1}^{N} \delta t |y^i(k) - y^i(k)^*|$, with δt being the sampling time of 30 minutes.

3.4 Coordinated primal-dual interior point method

Apart from primal and dual interior-point decomposition methods, there is a third option: applying a primal-dual interior point method (PDIP) [Wri97] directly to the problem and decomposing the Newton step, which provides a coordinated (C-)PDIP method. This primal-dual scheme was proposed in [PHAN17, APHW14], where the authors focus on distributed solutions of the Newton step to eliminate the need for a central agent, using proximal splitting [APHW14] and message passing [PHAN17] methods.

The advantage of the primal-dual interior-point setting is three fold: Firstly there is no distinction between local and global iterates and iterations for solving local problems which only generate intermediate Hessian and gradient information [BPJ17b, NS09] are eliminated. Secondly, in certain cases, it is still possible to enforce primal feasible iterations while enjoying increased convergence speed. Finally, by applying predictor-corrector schemes [Meh92, Gon96], PDIP methods can become very efficient and therefore are commonly used in implementations aiming for high-performance [BCC$^+$10, DCB13b, MB12]. However, in the interior-point decomposition framework, none of the aforementioned works apply a predictor-corrector scheme for enhancing convergence speed, with the exception of [KG14] which apply predictors in the dual decomposition setting, with primal-dual iterations in the local agents. Note that the method of [KG14] does not fall into the primal-dual framework as we use here.

Primal-dual interior methods are widely used for convex programming, due to their polynomial worst-case complexity as well as practical fast convergence properties [Wri97]. For an algorithm used in the multi-agent setting, not only the number of iterations, but also the number of communications is of utmost importance for practical applicability. Even though coordinated algorithms can achieve the same number of iterations as a central method, the number of communication steps per iteration can be significant. We will show that applying Mehrotra's predictor-corrector scheme [Meh92] in the C-PDIP setting, results in four communication steps per iteration. If the communication delay is significant, this might nullify the benefit of the reduced number of iterations. To tackle this issue we introduce a decentralized predictor-corrector scheme in the C-PDIP setting, which enjoys the high performance of widely used predictor-corrector (PC) methods [Meh92, Gon96], without introducing too many communication steps. We further analyze the algorithm from a general point of view of a primal-dual interior point method with non-uniform complementarity perturbations, or barrier parameters. This analysis provides safeguards that can be used to ensure convergence of the decentralized PC scheme.

3.4.1 Decomposition of Primal-Dual Interior point method

We will use the multi-agent problem (3.1) for describing the method. For brevity we re-write (3.1) in a compact form with slack variables

$$\begin{aligned} \min_{x} \quad & f(x) \\ \text{s. t.} \quad & g(x) + s = 0 \\ & \mathbf{C}x - y = 0 \qquad (\mu) \\ & \mathbf{D}y = d \qquad (\gamma) \\ & s \geq 0 \qquad (\lambda) \end{aligned} \tag{3.43}$$

where the dual variables μ, γ, λ are shown next to their corresponding constraints. The first order optimality conditions for problem (3.43), or the KKT system, can be written as

$$\begin{aligned} r_{\text{dual},x}(x, \mu, \lambda) &= \nabla f(x) + \mathbf{C}^T \mu + \mathcal{J}_g(x)^T \lambda \\ r_{\text{dual},y}(\gamma, \mu) &= \mathbf{D}^T \gamma - \mu \\ r_{\text{prim,in},x}(x, s) &= g(x) + s \\ r_{\text{prim,eq},x}(x, y) &= \mathbf{C}x - y \\ r_{\text{prim,y}}(y) &= \mathbf{D}y - d \\ r_{\text{compl}}(\lambda, s) &= \Lambda S \mathbf{1} \end{aligned} \tag{3.44}$$

or in a compact form;

$$r(x, y, s, \mu, \gamma, \lambda) = 0. \tag{3.45}$$

Where $r = (r_{\text{dual}}^T, r_{\text{comp}}^T, r_{\text{prim}}^T)^T$, $r_{\text{dual}} = (r_{\text{dual},x}^T, r_{\text{dual},y}^T)^T$ and r_{prim} defined similarly. We replace the complementarity residual, as explained in Section 3.2.5, and obtain the modified residual $\hat{r}(\boldsymbol{\beta})$. The Newton direction $\Delta z = (\Delta x, \Delta y, \Delta s, \Delta\lambda, \Delta\mu, \Delta\gamma)$, for the perturbed KKT system can be found by solving the following linear system of equations:

$$\begin{bmatrix} H(x,\lambda) & 0 & 0 & \mathcal{J}_g(x)^T & \mathbf{C}^T & 0 \\ 0 & 0 & 0 & 0 & -\mathbf{I}_\lambda & \mathbf{D}^T \\ 0 & 0 & \Lambda & S & 0 & 0 \\ \mathcal{J}_g(x) & 0 & \mathbf{I}_s & 0 & 0 & 0 \\ \mathbf{C} & -\mathbf{I}_y & 0 & 0 & 0 & 0 \\ 0 & \mathbf{D} & 0 & 0 & 0 & 0 \end{bmatrix} \begin{bmatrix} \Delta x \\ \Delta y \\ \Delta s \\ \Delta\lambda \\ \Delta\mu \\ \Delta\gamma \end{bmatrix} = -\hat{r}(\boldsymbol{\beta}) \tag{3.46}$$

where we use the definitions ;

$$H(x,\lambda) = \nabla^2 f(x) + \mathcal{H}_g(x,\lambda)\,, \tag{3.47}$$
$$\nabla^2 f(x) = \text{blkdiag}(\nabla^2 f^1(x), \dots, \nabla^2 f^m(x))\,, \tag{3.48}$$
$$\mathcal{H}_g(x,\lambda) = \text{blkdiag}(\sum_{j=1}^{n^1} \lambda_j^1 \nabla^2 g_j^1(x), \dots, \sum_{j=1}^{n^m} \lambda_j^m \nabla^2 g_j^m(x))\,, \tag{3.49}$$
$$\mathcal{J}_g(x) = \text{blkdiag}(\mathcal{J}_g^1(x^1), \dots, \mathcal{J}_g^m(x^m))\,, \tag{3.50}$$
$$\mathcal{J}_g^i(x^i) = \left[\nabla g_1^i(x^i), \dots, \nabla g_{n^i}^i(x^i)\right]^T\,, \tag{3.51}$$
$$\Lambda = \text{blkdiag}(\Lambda^1, \dots, \Lambda^m),\ \Lambda^i = \text{diag}(\lambda^1, \dots, \lambda^m)\,, \tag{3.52}$$
$$S = \text{blkdiag}(S^1, \dots, S^m),\ S^i = \text{diag}(s^1, \dots, s^m)\,. \tag{3.53}$$

In this section we will show how to exploit the structure of the linear system (3.46), such that a Newton direction can be computed by communicating a minimum amount of information between the agents and the central coordinator.

We start by eliminating $\Delta s, \Delta\lambda, \Delta x, \Delta\mu$ from the system (3.46), in the given order, to obtain

$$\Delta s = -\Lambda^{-1} S \Delta\lambda - \Lambda^{-1} r_{\text{cent}}\,, \tag{3.54}$$
$$\Delta\lambda = S^{-1}\Lambda \tilde{r}_{\text{pr,in,x}} + S^{-1}\Lambda \mathcal{J}_g(x)\Delta x\,, \tag{3.55}$$
$$\Delta x = -\mathbf{R}^{-1}\tilde{r}_{\text{dual},x} + -\mathbf{R}^{-1}\mathbf{C}^T\Delta\mu\,, \tag{3.56}$$
$$\Delta\mu = -\mathbf{P}\Delta y + \mathbf{P}\tilde{r}_x\,, \tag{3.57}$$

where we use the definitions

$$\tilde{r}_{\text{prim,in},x} = r_{\text{prim,in},x} - \Lambda^{-1} r_{\text{cent}}\,, \tag{3.58}$$
$$\tilde{r}_{\text{dual,x}} = r_{\text{dual,x}} + \mathcal{J}_g(x)^T S^{-1}\Lambda \tilde{r}_{\text{prim,in,x}}\,, \tag{3.59}$$
$$\mathbf{R} = H(x,\lambda) + \mathcal{J}_g(x)^T S^{-1}\Lambda \mathcal{J}_g(x)\,, \tag{3.60}$$
$$\tilde{r}_x = r_{\text{prim,eq,x}} - \mathbf{C}\mathbf{R}^{-1}\tilde{r}_{\text{dual},x}\,, \tag{3.61}$$
$$\mathbf{P} = (\mathbf{C}\mathbf{R}^{-1}\mathbf{C}^T)^{-1}\,, \tag{3.62}$$
$$q = (r_{\text{dual},y} - \mathbf{P}\tilde{r}_x)\,. \tag{3.63}$$

Finally we are left with two equations

$$\mathbf{P}\Delta y + \mathbf{D}^T\Delta\gamma = -q\,, \tag{3.64}$$
$$\mathbf{D}\Delta y = -r_{\text{prim},y}\,, \tag{3.65}$$

which can be written in a compact form

$$\begin{bmatrix} \mathbf{P} & \mathbf{D}^T \\ \mathbf{D} & 0\,, \end{bmatrix} \begin{bmatrix} \Delta y \\ \Delta\gamma \end{bmatrix} = - \begin{bmatrix} q \\ r_{\text{prim},y} \end{bmatrix}\,, \tag{3.66}$$

where

$$\mathbf{P} = \text{blkdiag}(P^1, \dots, P^m),\ q = ((q^1)^T, \dots, (q^m)^T)^T. \tag{3.67}$$

The linear system (3.66) can also be written as

$$P^i \Delta y + (D^i)^T \Delta \gamma = -q^i, \quad \forall i \in \mathbb{Z}_{[1\,m]} \tag{3.68}$$

$$\sum_{i=1}^{m} D^i \Delta y^i = -r_{\text{prim},y}\,, \tag{3.69}$$

revealing the underlying structure. The matrices $\mathbf{R}$ and $\mathbf{P}$ are guaranteed to be invertible by virtue of the assumptions made in Section 3.2.1. Note that the linear system (3.66) is very similar to relation (3.34) that was obtained from the decomposition of the primal barrier method. The only difference is that the matrix P^i is not directly the Hessian of the local primal function at local optimality, but a projection of the Hessian of the local Lagrangian on the output y^i space at some sub-optimal primal-dual point. Likewise the vector q^i is not directly the dual residual for the local problem but it is a combination of the local primal and dual residuals. Similar to (3.36), the solution can be computed as

$$\begin{aligned} \Delta y^{i,*} &= -(P^i)^{-1} q^i - (P^i)^{-1} (D^i)^T \Delta \gamma^* \\ \Delta \gamma^* &= \left(\sum_{i=1}^{m} D^i (P^i)^{-1} (D^{i\,T}) \right)^{-1} \left(\sum_{i=1}^{m} D^i (P^i)^{-1} q^i + d - \sum_{i=1}^{m} D^i y^i \right) \end{aligned} \tag{3.70}$$

The system of equations (3.66) can be solved in a coordinated manner as illustrated in Figure 3.5. The coordinated step distributes the linear algebra operations required to compute P^i and q^i to the agents, which can operate in parallel, and therefore reduces the computational effort at the central agent. Furthermore, local data of the agents is partially hidden, since the central agent does not have direct access to the information regarding the values of local variables $x^i, s^i, \mu^i, \lambda^i$ or functions $g^i(x^i), f^i(x^i)$.

Starting the iteration step k of the coordinated PDIP method, the agents receive the global step-size t, the perturbation (or barrier) parameter β and either the global dual variable γ or directly the product $(D^i)^T \gamma$, which will be used in the computation of $r_{\text{dual},y}$ and q^i. Firstly the agent updates its local variables and residuals using the step-size and the direction computed in the previous step. Using these values the agent then constructs the data P^i, q^i. The data (P^i, q^i) is sent to the central agent together with the norms of local residuals. The central agent constructs the overall residual norm, block diagonal $\mathbf{P}$ and stacked q and solves (3.66) for obtaining the direction in output variables Δy. Upon receiving it's local direction Δy^i the agent evaluates $\Delta x^i, \Delta \lambda^i, \Delta s^i$, the maximum feasible step-size and constructs the one-dimensional quadratic function $\eta^i(t)$ which outputs the sum of the local complementarity products for a given step-size t. Finally the central agent collects local maximum step-sizes t^i and the local 'duality gap' functions $\eta^i(t)$ for computing the global step-size t, evaluating the duality gap measure $\hat{\eta}$ and the perturbation vector β according to the reduction parameter σ. Having access to $\hat{\eta}$ and the norm of the primal-dual residuals of the previous step, the central agent can check if the specified termination criterion is satisfied.

Inspection of Figure 3.5 reveals that two communication rounds between the central agent and local agents are necessary to complete a single Newton step; one for computing parameters $\mathbf{P}, q$ and another for coordinating the step-sizes. In case of a short step approach, where the step-size is selected as to guarantee centrality and/or sufficient decrease of the residuals, additional communication steps might be necessary.

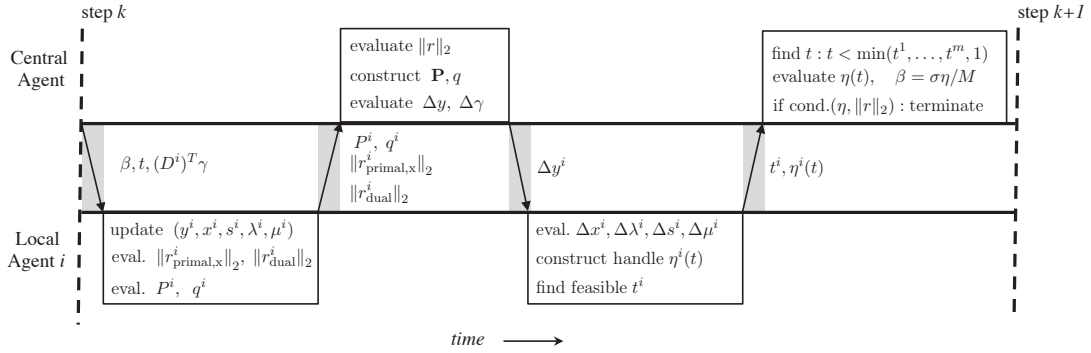

Figure 3.5 – Newton step for the coordinated PDIP method

3.4.2 Decentralized predictor-corrector

In this section we will investigate how the predictor-corrector methodology can be applied to the coordinated PDIP setting, without introducing extra communication burden. The predictor-corrector approach requires computing two Newton steps per iteration, using the same Hessian with pre-factorization. In practice, the reduction in the number of iterations greatly compensates the extra time spent in additional back-solves. However in the coordinated setting, the cost of communication becomes considerable, as now four communication rounds are necessary per iteration. In order to remedy this effect, we propose a decentralized predictor-corrector approach, which does not require additional communication rounds. The predictor step is essentially carried out for getting an idea of how much decrease could be achieved with an affine scaling step, and what would be the second order term in the linearization of the complementary slackness condition in that case. This collected information is used for determining the corrector terms and the duality gap reduction parameter σ.

We break the communication in the predictor step and accept a cruder prediction; all agents solve a local affine scaling problem, disregarding the global equality constraint. This decentralized affine scaling step is used to determine the local duality gap reduction parameters σ^i and corrector vectors. The actual Newton step is then determined in a coordinated manner, respecting the coupling constraints. The main modification to the standard coordinated Newton step is in the first computational block shown in Figure 3.5, and changes the way local agents compute q^i using the local centrality residual r^i_{cent} as

$$\Delta(y^i)^{\text{aff}} = -(P^i)^{-1} q^{i,\text{aff}} \tag{3.71}$$

$$(t^i)^{\text{aff}} \;:\; (s^i + (t^i)^{\text{aff}}\Delta s^i, \lambda^i + (t^i)^{\text{aff}}\Delta\lambda^i) > 0 \tag{3.72}$$

$$\sigma^i = \left(\hat{\eta}^{i,\text{aff}}/\hat{\eta}^i\right)^3 \tag{3.73}$$

$$v^i = (\Delta\Lambda^i)^{\text{aff}}(\Delta S^i)^{\text{aff}}\mathbf{1} \tag{3.74}$$

$$r^i_{\text{cent}} = \Lambda^i S^i \mathbf{1} + v^i - \sigma^i\hat{\eta}^i\mathbf{1}, \tag{3.75}$$

where q^{aff} is computed using the local complementarity residual r^i_{comp} as explained in Section 3.2.5. The rest of the process is identical to the standard coordinated Newton step as shown in Figure 3.5, except that the global reduction parameter σ and the perturbation vector β are not needed.

Note that the local affine scaling direction could be significantly different from the original affine scaling direction, since the coupling constraint $\mathbf{D}\Delta y = -r_{\text{primal},y}$ is ignored. This diminishes the quality of the

corrector and the reduction parameter σ^i in exchange for reduced communication effort. In addition, by allowing different σ^i and corrector terms for each agent, we break the uniformity in the perturbation vector β. This last effect is not necessarily undesired, as discussed by [CG08] and may even enhance the robustness of the method. For guaranteed convergence in this setting, certain conditions that preserves centrality should be enforced, as explained in the next section.

3.4.3 Convergence with non-uniform perturbation

Consider the modified Newton step with a non-uniform perturbation vector on the complementarity condition:

$$\mathcal{M}\Delta z = -\begin{bmatrix} r_{\text{prim,dual}} \\ \Lambda S\mathbf{1} - \Lambda S\xi \end{bmatrix} \tag{3.76}$$

where the vector ξ specifies how much change in each complementarity pair is desired, with $\xi \in \mathbb{R}^M$ and $\xi > 0$. For a given $\sigma > 0$, ξ can be selected as $\xi_i = \sigma \frac{\hat{\eta}}{\lambda_i s_i}$ to recover the standard uniform perturbation vector with σ. This more generic definition captures the effect of the decentralized predictor-corrector described in the previous chapter. Note that any corrector applied together with the uniform-perturbation can be considered as a non-uniform perturbation vector.

Using different perturbation values for each complementarity pair was previously proposed by [Gon96, CG08]. The authors propose a 'multiple-corrector' scheme, where instead of targeting centrality with a uniform perturbation vector, they target the 'symmetric-neighborhood' $\mathcal{N}_s$ (3.27) and try to correct only the complementarity pairs that are outside of this neighborhood. The algorithm works well in practice [CG08] but the authors provide convergence results only for the standard path-following PDIP method within the symmetric neighborhood, in order to justify the use of $\mathcal{N}_s$ as a target for the complementarity pairs. Instead, we will provide a full convergence analysis and convergence safeguards for using non-uniform complementarity perturbations. The analysis and safeguards can be applied to any PDIP method, whether it uses multiple correctors or the standard Mehrotra's corrector, providing an alternative to the existing convergence safeguards in the literature [Car05, ML05, SPT07]. We will state the convergence conditions for quadratic programs (QP) in the feasible case, where the primal-dual residuals are equal to zero, to make the analysis simpler as is common in the interior point literature [Wri97, Gon13]. The results can be extended to the non-feasible case, and for LPs and QPs, a feasible point can be obtained with reformulations of the problem [Gon13, Gon12].

We restrict the iterations to lie inside the 'symmetric neighborhood' (3.27) as in [Gon13]. Furthermore we restrict the reduction vector ξ to satisfy the following criteria,

$$\exists c \;:\; 0 < c < 1, \;\; \frac{\sum_j \xi_j \lambda_j s_j}{M} = c\hat{\eta} \tag{3.77}$$

$$\frac{1}{\phi}\frac{c\hat{\eta}}{\lambda_j s_j} \geq \xi_j \geq \phi \frac{c\hat{\eta}}{\lambda_j s_j} \tag{3.78}$$

for some $1 \gg \phi > \psi$. The constraints (3.78), (3.77) enforce that the next point we would like to arrive with the Newton step, does not have to be on the central path as in the uniform case, but it should be at least strictly within the $\mathcal{N}_s$ and aim to reduce the overall duality gap. By applying these conditions, we will ensure that the linearization error in the Newton step is bounded, and there is always some finite step-size that can reduce the duality-gap and keep the system within $\mathcal{N}_s$. The following theorem states the convergence and complexity result.

Theorem 2. *If problem* (3.1) *is a convex QP, given* $\epsilon > 0$*, suppose that a starting point* $(x_0, y_0, s_0, \lambda_0) \in \mathcal{N}_s$ *satisfies*

$$\eta_0 \leq \frac{1}{\epsilon^\kappa} \tag{3.79}$$

for some positive constant κ*. Let* $\{\eta_k\}$ *be a sequence generated by the iteration scheme, which takes steps using the Newton relation* (3.76) *with admissible parameters* ϕ, ψ *and reduction vectors* ξ_k *that satisfy* (3.77), (3.78)*, with some* $\alpha, \gamma \in (0, 1)$, $\gamma < 1 - \alpha$ *and* $c_k \leq \gamma$*. Then, there exists an admissible step-size* $t_k \sim \mathcal{O}(\frac{1}{M})$ *and an index* $K \sim \mathcal{O}(M \log(\frac{1}{\epsilon}))$*, such that*

$$\eta_k \leq \epsilon, \quad \forall k \geq K. \tag{3.80}$$

Proof. see Appendix A.3. □

Integrating convergence conditions to the coordinated setting is not straightforward as one should balance extra communication burden and performance deterioration with the robustness of the algorithm. Here we state one possible way that does not require extra communication, but a detailed comparison and performance analysis is out of the scope of this paper. Constraints (3.78), (3.77) on ξ can be satisfied by enforcing

$$\frac{1}{\phi} c_{\min} \hat{\eta} \geq \sigma^i \hat{\eta}^i \mathbf{1} - w^i v^i \geq \phi c_{\max} \hat{\eta} \tag{3.81}$$

$$c_{\max} \geq (\sum_j \xi_j^i \lambda_j^i s_j^i)/M^i \geq c_{\min} \tag{3.82}$$

in the decentralized predictor-corrector computation step using pre-determined values $0 < c_{\min} < c_{\max} < 1$ and a corrector weighing term w^i that is selected by the agent. The central agent collects $\max(|\xi_i - 1|)$ from agents and constructs $\theta = 2^{-3/2}\|\xi - \mathbf{1}\|_\infty^2$ and bounding functions on the duality gap measure as

$$\hat{\eta}_{\max}(t) := \hat{\eta}((1 - t) + tc + t^2\theta) \tag{3.83}$$

$$\hat{\eta}_{\min}(t) := \hat{\eta}((1 - t) + tc) \tag{3.84}$$

which can be sent to the local agents together with Δy^i. The agents then compute their maximum allowable step-size that satisfy the neighborhood constraints as;

$$t^i_{\max} \ : \ (1/\psi)\hat{\eta}_{\min}(t) \geq (s^i + t\Delta s^i)^T(\lambda^i + t^i\Delta\lambda^i) \geq \psi\hat{\eta}_{\max}(t) \tag{3.85}$$

Finally, the sufficient decrease condition can be verified at the central node using the collected local duality gap functions $\eta^i(t)$.

3.4.4 Numerical experiments

We apply three different coordinated PDIP methods, PF: path following with a fixed reduction parameter σ and $\beta = \sigma\hat{\eta}\mathbf{1}$, PC: predictor-corrector, DPC: decentralized predictor corrector, on randomly generated problems as well as on the demand response problem from section 3.3.4. In order to compare maximum performance, we omit convergence safeguards and use the 'long-step' approach for step-size selection, which is a common practice [NW00].

We first apply the algorithms on a problem of the form

$$\begin{aligned} \min_{x} \quad & \sum_{i=1}^{m} x^{i\prime} Q^i x^i + c^{i\prime} x^i \\ \text{s. t.} \quad & A^i x^i \leq b^i, \quad \forall i \in \mathbb{Z}_{[1,m]} \\ & \sum_{i=1}^{m} C^i x^i = d \end{aligned} \tag{3.86}$$

with randomly generated data. On randomly generated problems, the benefit of the decentralized predictors in terms of reduced communication burden is clear as shown in Figure 3.6. However this reduction is not deterministic, as can be seen in Figure 3.7, where the performance of PF method and DPC method are similar for a more realistic problem data from the voltage support DR problem in section 3.3.4. In terms of iterations, PC method prevails, but suffers from increased communication rounds. The DPC method is effective in reducing the communication burden, however the reduction depends on the problem data and the starting point. The standard path-following method is also attractive due to small communication per iteration requirement and simplicity of the algorithm.

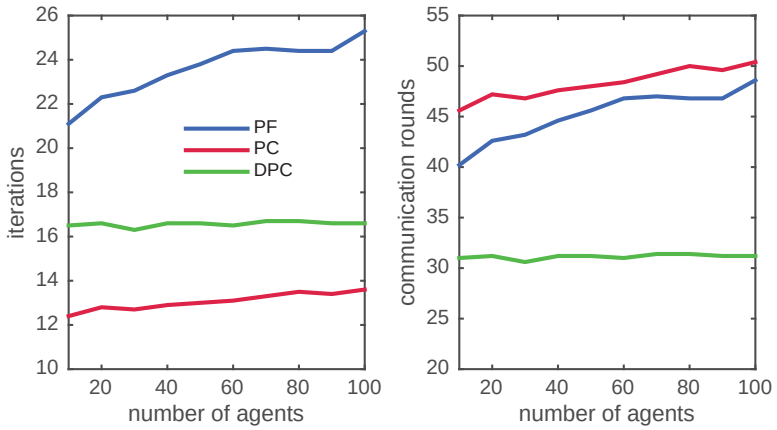

Figure 3.6 – Evolution of iteration and communication rounds required to bring the suboptimality $\|y - y^*\|/n_y$ below 10^{-5}, for increasing number of agents for randomly generated coupled QPs with feasible starting points, using the mean of 10 sampled problems with $n_x = 20$, $n_y = 5$.

3.5 Conclusion

In this part of the thesis, we have discussed two different coordinated interior-point methods. The first method C-IP is a decomposition of the barrier method, which carries attractive properties of the method to the multi-agent setting. These properties are

- Started from a globally feasible point, C-IP preserves local and global feasibility during the iterations.
- C-IP enjoys the fast convergence properties of the Newton's method, thanks to the smoothing of the problem with barrier functions.
- The coordination of the sub-problems only requires communication of relevant output variables with a central agent. The dimension of the output variables are much less compared to the internal

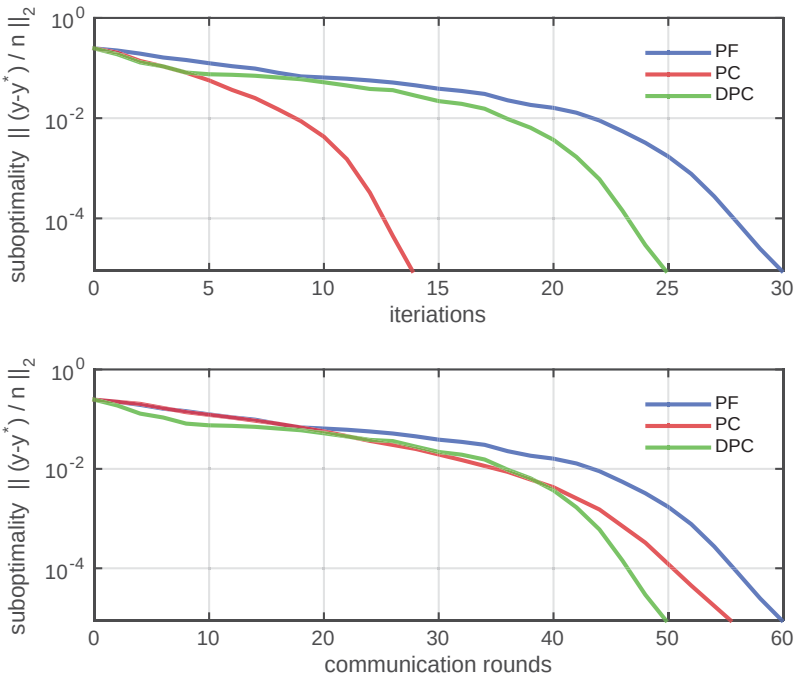

Figure 3.7 – Comparison of various CPDIP mehtods for a realistic example of demand response [BPJ17a]. In this case 12 smart-buildings help a distribution grid for voltage support in case of a significant solar power in feed. The algorithm is started from a point that is locally centered but infeasible with respect to the coupling constraint.

local variables in practical applications, resulting in efficient coordination.

- The task of the coordinator can be further distributed among the agents, due to the structure of the Newton step computation.

All these properties makes C-IP an attractive option for optimization and control in multi-agent settings. Comparison against the popular ADMM method shows that it is possible to obtain more accurate solutions in a shorter period, as expected from a second-order method.

C-IP method can be further enhanced by considering a decomposition of the PDIP method, which is the method of preference for commercial solvers. We have discussed three different coordinated PDIP methods: standard path following, Mehrotra's predictor-corrector and the novel decentralized predictor-corrector method, for solving the multi-agent problem (3.1). We provide convergence results for PDIP methods with non-uniform complementarity perturbations, as the decentralized PC scheme falls into this category. This convergence result can be applied to any PDIP method that utilizes corrector terms and breaks the uniformity of the complementarity perturbation. Finally we have tested the coordinated PDIP methods on numerical examples. We can draw the following conclusions.

- Naive application of predictor-corrector methods in coordinated settings might result in excessive communications.
- By relying on local predictions, it is possible to reduce the communication effort, in exchange for

reduced performance in terms of the number of iterations.

- Depending on the problem data, local predictions can reduce communication effort significantly and result in a more scalable method.
- In practical applications, it is necessary to test different options and choose the best performing method in terms of iteration number or communication steps. According to the results, depending on the expense of iterations and communications, all three C-PDIP methods are viable options.

4 Coordinated Economic MPC with Barrier Functions

4.1 Introduction

In MPC, the cost function is usually formulated so as to stabilize a target point. Several ingredients are added to ensure stability such as terminal cost and constraints [MRRS00]. The recent trend in MPC literature is to incorporate economic objectives directly in the dynamic MPC layer, named as Economic MPC (EMPC). The standard stability results do not apply, as the cost function might favor transient states over equilibrium points. Several stability results have been provided in the literature, depending on the additional constraints and assumptions enforced on the problem formulation [DAR11, RAB12, Gr3, FGM18].

A recently proposed variant, EMPC with generalized terminal constraints (G-EMPC) [FT13], is specifically attractive for large scale distributed implementations. The algorithm relies on a generic terminal constraint which only requires the terminal state to be a feasible steady state with an additional constraint that ensures monotonic decrease of the terminal stage cost. This approach eliminates the need of computing terminal sets and controllers in a distributed setting, which is a demanding task within itself, while ensuring an extended region of attraction compared to MPC with fixed terminal constraint [DAR11]. Furthermore, the optimal equilibrium point need not be specified and the system finds this point with appropriate tuning [FT13]. This property is specifically desired in large scale and distributed implementations. In such cases finding the optimal equilibrium point would amount to solving an expensive large scale problem each time the economic condition and therefore the cost function changes. Some formulations have been put forward to remedy this issue, such as solving in parallel a target optimization problem [MA14, KMA16]. However solving an additional target finding problem increases the communication and computation burden in a distributed framework. For this reason G-EMPC is suitable for large scale distributed implementations as it does not require an optimal target specification.

An important consideration in MPC is the numerical algorithms used in the solution of the finite horizon optimal control problem. Faster solvers and limited iteration numbers are desirable, especially for fast and large scale systems. To achieve desired numerical properties, it is common practice to solve the optimal control problem approximately, either modifying the cost and constraints, or by terminating the numerical solvers before convergence. An interesting case is the use of barrier functions that approximate the hard constraints in the problem. Instead of reducing the effect of the barriers progressively as in interior point solvers, the barriers are preserved in the optimal control problem resulting in an approximation. The result is a smoothed optimization problem that can be solved with second order methods and warm started effectively [WH04, FE15, FE16a, FE16b, WB10, ZQJD17].

The use of barriers in EMPC formulations is an attractive option, but mostly unexplored in the literature. In standard MPC with tracking cost, the added barrier functions are centered, so as to have zero gradient at the target point. Otherwise, addition of the barriers would shift the target point [WH04, FE15]. For EMPC problems with known optimal equilibrium point, the implementation would only require re-centering the barrier functions around the optimal equilibrium point, which adds a 'tracking cost' to the problem. This additional tracking term can be used to ensure convergence to the optimal equilibrium by providing strong duality [DAR11], or to a neighborhood similar to the dual objective formulation of [AAJ16]. For G-EMPC, the optimal equilibrium point is not known a-priori and the system finds it in closed loop, which brings a problem to the case of using barrier functions, as we do not know where to center the barriers.

In this work, we will investigate the idea of re-centering the barriers iteratively, updating the centering at every time step. There are two options at hand, first re-centering around the sub-optimal terminal

steady state that is available to the algorithm from the previous step, or re-centering around the predicted trajectory from the previous time step. In both cases, the barrier functions will work as a regularizing cost, penalizing deviations from previous solutions. The latter choice of trajectory regularization cost translates into a very intuitive idea, that is to penalize deviations from open-loop predicted trajectories to reduce the mismatch between predicted and closed loop behaviour, but it is mainly unexplored in the MPC literature. In both cases the analysis becomes challenging, due to the time varying nature of the cost. In this work we will focus on the trajectory regularization case, and show that it is relatively easy to analyze using multi-objective optimization results. We will show that under certain reachability assumptions, the addition of barriers and iterative re-centering not only preserves convergence properties of G-EMPC, but improves them by ensuring (i) convergence of the predicted terminal state to a neighborhood of the overall optimal equilibrium point and (ii) asymptotic average performance guarantees for the closed-loop system, without the need of implementing a switching algorithm as in [FT13] or terminal sets as in [MAA14]. Therefore, apart from the numerical benefits gained by the addition of barrier functions, using them as a regularizing cost brings system theoretic benefits to EMPC formulations.

4.2 Preliminaries

4.2.1 Economic MPC

Economic MPC (EMPC) is a branch of MPC methodology that focuses on using objective functions that reflect directly the economics of the process to be controlled. In mathematical terms, any objective function that does not only penalize the distance from a specific target steady state or trajectory, as done in tracking MPC, can be regarded as 'economic'.

Because of this arbitrary objective selection that might not promote system stability, the standard stability results does not apply, as the objective function might favor transient states over equilibrium points. Several stability results are provided in the literature, depending on the additional constraints and assumptions enforced on the problem formulation [DAR11, RAB12, Gr3, FGM18].

In this section we will introduce some of the technical notions used in EMPC. A feasible equilibrium point for the system is defined as

$$(x, u) \in \mathbb{X} \times \mathbb{U} \,:\, x = f(x, u). \tag{4.1}$$

We restrict the admissible equilibrium points to those which are strictly feasible. Therefore, the set of admissible equilibrium points is defined as

$$\mathcal{E} := \{(x, u) \in \lambda(\mathbb{X} \times \mathbb{U}),\ x = f(x, u)\} \tag{4.2}$$

for some $\lambda \in [0, 1)$. The purpose of this restriction is to ensure reachability of the set $\mathcal{E}$ as discussed in Appendix A.4. In the following, optimality of equilibrium points are discussed with respect to the admissible set $\mathcal{E}$. From an initial state x, the optimal reachable equilibrium point is given by

$$(z^R(x), v^R(x)) = \arg\min_{z,v}\ l(z, v) \ :\ z \in \mathcal{R}_N(x), (z, v) \in \mathcal{E}, \tag{4.3}$$

with the corresponding stage cost $l^R(x) = l(z^R(x), v^R(x))$. We can also define the optimal admissible

equilibrium point for the system

$$(z^*, v^*) = \arg\min_{z,v} \; l(z, v) \;\; : \;\; (z, v) \in \mathcal{E}. \tag{4.4}$$

Another useful notion in EMPC framework is the asymptotic average performance [AAR12] of the system, defined as

$$\bar{J}_\infty(\boldsymbol{x}_\mu) = \limsup_{T\to\infty} \frac{1}{T} \sum_{t=0}^{T-1} l(x_\mu(t), \mu(t)) \tag{4.5}$$

for a closed loop trajectory $\boldsymbol{x}_\mu$ with control policy μ. Due to the economic cost, system trajectories that do not stabilize or convergence to a specific equilibrium point might be preferable. Therefore $\bar{J}_\infty$ provides an alternative criteria to assess the validity of proposed EMPC schemes in the literature.

There are many alternative formulations of EMPC that incorporate various forms of terminal cost, constraints and auxiliary cost functions to ensure desired properties such as recursive feasibility and stability [FGM18]. In this thesis we will focus on the EMPC with generalized terminal constraints (G-EMPC) [FT13]. In the G-EMPC formulation of [FT13] a steady state constraint that enforces the terminal state-input pair to be any admissible equilibrium point for the system is used. This generic terminal constraint enlarges the feasible region significantly and eliminates the need of designing a terminal region and controller. As the terminal cost, the stage cost is used with a selected weight. Finally a terminal constraint that puts an upper bound on this terminal cost is added. From an initial state x the control problem is given by:

$$\begin{aligned} \mathcal{P}^e(x, \bar{l}) \; : \min_{\boldsymbol{u}} \quad & J_N(x_{\boldsymbol{u}}(x), \boldsymbol{u}) + \alpha l(\boldsymbol{x}_{\boldsymbol{u}}(N|x), u(N)) \\ \text{s. t.} \quad & \boldsymbol{u} \in \mathcal{U}(x) \\ & (\boldsymbol{x}_{\boldsymbol{u}}(N|x), u(N)) \in \mathcal{E} \\ & l(\boldsymbol{x}_{\boldsymbol{u}}(N|x), u(N)) \leq \bar{l} \end{aligned} \tag{4.6}$$

with the optimal input sequence $\boldsymbol{u}^e(x, \bar{l})$ and the control policy $\mu^e(x, \bar{l}) := u^e(0|x, \bar{l})$. The upper bound on the terminal stage cost is selected as the cost value at the previous time step in order to ensure that the terminal cost is non-increasing. This does not guarantee however, that the terminal cost will decrease and converge to the optimal steady state. Defining an augmented state as $\xi = (x, \bar{l})$, the resulting closed loop system can be written as

$$\begin{aligned} x^+ &= f(x, \mu^e(\xi)) \\ \bar{l}^+ &= l(x_{\boldsymbol{u}^e}(N|x), u^e(N|\xi)) \end{aligned} \tag{4.7}$$

The upper bound $\bar{l}$ can take values between the minimum and maximum possible stage cost for a feasible equilibrium point and therefore is contained within the set

$$\mathbb{L} = \{\bar{l} \mid \bar{l} \geq \{\min l(x, u) \; : \; x, u \in \mathcal{E}\} \tag{4.8}$$

$$\bar{l} \leq \{\max l(x, u) \; : \; x, u \in \mathcal{E}\}\} \; . \tag{4.9}$$

Finally, the set of admissible initial states and terminal cost bounds are given by

$$\mathbb{F} = \{(x, \bar{l}) \in \mathbb{X} \times \mathbb{L} \mid \mathcal{P}^e(x, \bar{l}) \text{ is feasible}\} \; . \tag{4.10}$$

4.2.2 Coordinated MPC

In this section, we consider a large scale implementation of the EMPC problem involving multiple systems. We will show how the problem can be formulated such that the control task can be posed as a generic multi-agent problem (3.1) and the decomposition methods explained in the previous Section 3.2.3 can be applied thanks to the simple terminal constraints.

We first decompose the global system into subsystems according to the allocation of computational resources and access to information. The formulation is adapted from [NND11] and [FBGS14]. We consider M linear subsystems that interact over a network that is described as a standard directed graph. We introduce the dynamic interconnection variables w, y to model the interaction between subsystems. The model for each subsystem can be written as

$$\begin{aligned} x^{i+} &= A^i x + B^i u^i + w^i, \quad \forall i \in \mathbb{Z}_{[1,m]} \\ w^i &= \sum_{j \in \mathcal{N}^i} E^{ij} y^{ij} \\ y^{ji} &= C^{ji} x^i + D^{ji} u^i, \quad \forall j \in \mathcal{N}^{-i} \end{aligned} \tag{4.11}$$

Here the variable $w^i \in \mathbb{R}^{n_x}$ is the exogenous input on the system caused by the outputs of neighboring systems' output with $j \in \mathcal{N}^i$ and y^{ji} is the output of the system that contributes to the exogenous input w^j of a system with $j \in \mathcal{N}^{-i}$. We further decompose the constraint sets into local sets

$$X = (\mathbb{X}^1 \times \cdots \times \mathbb{X}^M), \quad U = (\mathbb{U}^1 \times \cdots \times \mathbb{U}^M) \tag{4.12}$$

The decomposition of the dynamical system and constraints is without loss of generality, since any coupling constraint between the subsystems can be re-written as the local constraint of a trivial dynamical system which represents the coupling variables. This might however require a change of variables and re-formulation of the constraint sets. With the same approach, by introducing agents that keep track of non-separable terms, the cost function can also be decomposed

$$J_N(x_{\boldsymbol{u}}, \boldsymbol{u}) = \sum_{i=1}^{M} J_N^i(x^i_{\boldsymbol{u}^i}(x^i, \boldsymbol{w}^i), \boldsymbol{u}^i) \tag{4.13}$$

Using the compact notation from Section 2.2.3 for systems under exogenous inputs, we define the local optimization problem and the corresponding value function as

$$\begin{aligned} p^i(x^i, \boldsymbol{w}^i, \boldsymbol{y}^i) = \min_{\boldsymbol{u}^i} \quad & J_N^i(\boldsymbol{x}^i_{\boldsymbol{u}^i}(x^i, \boldsymbol{w}^i), \boldsymbol{u}^i) \\ \text{s. t.} \quad & \boldsymbol{y}^i = C^i \boldsymbol{x}^i_{\boldsymbol{u}^i}(x^i, \boldsymbol{w}^i) + D^i \boldsymbol{u}^i \\ & \boldsymbol{u}^i \in \mathcal{U}(x^i, \boldsymbol{w}^i) \end{aligned} \tag{4.14}$$

where $\boldsymbol{y}^i$, C^i and D^i are constructed by stacking the variables $\boldsymbol{y}^{ji}$ and the corresponding matrices for all $j \in \mathcal{N}^{-i}$. Further terminal cost or local constraints can be added to the formulation without changing the structure of the local problem.

Using the local value functions, the global MPC problem can be re-written as

$$\begin{aligned} \min_{\boldsymbol{w},\boldsymbol{y}} \quad & \sum_{i=1}^{m} p^i(x^i, \boldsymbol{w}^i, \boldsymbol{y}^i) \\ \text{s. t.} \quad & \boldsymbol{w}^i + \sum_{j \in \mathcal{N}^i} \mathbf{E}^{ij} \boldsymbol{y}^{ij} = 0, \quad \forall i \end{aligned} \tag{4.15}$$

By defining the output of an agent as $\boldsymbol{z}^i = [\boldsymbol{w}^{i\,T}, \boldsymbol{y}^{i\,T}]$, the distributed MPC problem can be posed as the generic multi-agent problem (3.1) allowing the use of decomposition methods. Note that reformulation requires the elimination of extra equality constraints and introducing intermediate output variables in order to satisfy the assumptions in Section 3.2.1 as explained there.

Terminal constraints

In MPC formulations a standard way to ensure recursive feasibility and stability is to enforce a terminal constraint on the system. The terminal constraint is commonly computed as an invariant set for the system using a simple stabilizing controller. In large scale problems however, it becomes challenging to compute terminal controllers and invariant sets. Decentralized controller and terminal region design is a research area by itself, and there are many variants [NM14, Bak08]. Even though it is possible to use such an approach in a large scale implementation, it makes the design process demanding. When a new subsystem is added in the dynamical network, either the whole terminal ingredients have to be re-computed or a local update should be made allowing plug-n-play [RFFT13].

A very simple terminal constraint that trivially ensures recursive feasibility is the terminal equilibrium condition [FT13, LAAC08, FRLC10]. In that case we simply enforce that the terminal state of the system should be a feasible steady state. This constraint can be introduced as a simple local constraint for each agent as

$$(x^i_{\boldsymbol{u}^i}(N|x^i, \boldsymbol{w}^i), u^i(N)) \in \mathcal{E}^i \tag{4.16}$$

without changing the problem structure.

In closed loop operation, the terminal equilibrium constraint brings in another advantage. Since at each step, one-step shifted input sequences from the previous step will preserve the predicted trajectories, the exogenous inputs $\boldsymbol{w}^i$ can be predicted locally by simple shifting as well, assuming that all other agents will apply the same protocol. Therefore, each agent can apply the shift operator to the internal input and exogenous inputs $\boldsymbol{u}^i$, $\boldsymbol{w}^i$, without further communication. This shifted solution provides a globally feasible solution and thus can be used to warm start the problem without any communication effort.

In G-EMPC formulation (4.6), there is an additional terminal constraint that puts an upper bound on the terminal cost. This constraint can be decomposed by introducing local bounds as additional outputs, and a network constraint that determines the sum of local bounds.

$$l^i(\boldsymbol{x}^i_{\boldsymbol{u}^i}(N|x^i, \boldsymbol{w}^i), u(N)) \leq \bar{l}^i, \qquad \sum_{i=1}^{m} \bar{l}^i = \bar{l} \tag{4.17}$$

4.2.3 Multi-objective optimization

We will see in the subsequent sections that, in the analysis of G-EMPC, it is important to quantify the sub-optimality of the terminal cost and the cost on the remaining trajectory, according to their relative weights. To this end, we provide a theorem that relates the weighting parameter in a parameterized multi-objective optimization problem, to upper bounds on the values of the resulting cost functions. The proof is based on the results of [FT13], however the theorem and the proof is extended so as to cover extended valued functions, such as barrier functions which tend to ∞ on the boundary of the feasible set.

Theorem 3. *Let $C(\theta) \subset \mathbb{R}^{n_x}$ be a compact set for any $\theta \in \Theta \subset \mathbb{R}^{n_\theta}$. Consider two parametric functions $f(\theta, x), g(\theta, x) : C(\theta) \longrightarrow \mathbb{R} \cup \infty$ that are proper, lower semi-continuous and level bounded within* relint$(C(\theta))$ *in x for all θ . For a given $\alpha \in \mathbb{R}^+$, let $x^*(\alpha, \theta)$ be the optimizer of the multi-objective optimization problem*

$$v(\alpha, \theta) = \min_{x \in C(\theta)} f(\theta, x) + \alpha g(\theta, x) \,. \tag{4.18}$$

Let $x^\dagger(\theta)$ be the optimizer of only g with the optimal value

$$g^\dagger(\theta) = \min_{x \in C(\theta)} g(\theta, x) \,. \tag{4.19}$$

If, for each $\delta > 0$ there exists a finite bound $\bar{f}(\delta) < \infty$ such that

$$\bar{f}(\delta) > \max_{\theta \in \Theta} \min_{x \in C(\theta)} f(\theta, x) \,:\, g(\theta, x) \leq g^\dagger(\theta) + \delta \tag{4.20}$$

then, for any $\epsilon > 0$, there exists an $\bar{\alpha}(\epsilon) < \infty$ such that for any $\theta \in \Theta$ if $\alpha \geq \bar{\alpha}(\epsilon)$ the following relation is satisfied.

$$g(\theta, x^*(\alpha, \theta)) \leq g^\dagger(\theta) + \epsilon \,. \tag{4.21}$$

Proof. Since $f(\theta, x)$ is proper, l.s.c and level bounded within sets $C(\theta)$, it attains a minimum with a bounded value [RW09]. Let $\underline{f}$ be the minimum value of $f(\theta, x)$ for all $\theta \in \Theta, x \in C(\theta)$. For any ϵ choose $\bar{\alpha} > 2\eta/\epsilon$ with $\eta = \bar{f}(\epsilon/2) - \underline{f}$. For the purpose of contradiction, assume that $g(\theta, x^*(\bar{\alpha}, \theta)) > g^\dagger(\theta) + \epsilon$. Define

$$\tilde{x} := \arg\min_{x \in C(\theta)} f(\theta, x) \,:\, g(\theta, x) \leq g^\dagger(\theta) + \epsilon/2 \,, \tag{4.22}$$

and consider the difference

$$v(\bar{\alpha}, \theta) - (f(\theta, \tilde{x}) + \bar{\alpha} g(\theta, \tilde{x})) \,. \tag{4.23}$$

From the definitions of $\bar{f}(\epsilon/2)$ and $\underline{f}$ and the assumptions on the values of $g(\theta, x^*(\bar{\alpha}, \theta)), f(\theta, \tilde{x}), g(\theta, \tilde{x})$, we have the following relations

$$\begin{aligned}
&v(\alpha, \theta) - f(\theta, \tilde{x}) - \bar{\alpha} g(\theta, \tilde{x}) \\
&\geq v(\alpha, \theta) - \bar{f}(\epsilon/2) - \bar{\alpha} g(\theta, \tilde{x}) \\
&= f(\theta, x^*(\bar{\alpha}, \theta)) + \bar{\alpha} g(\theta, x^*(\bar{\alpha}, \theta)) - \bar{f}(\epsilon/2) - \bar{\alpha} g(\theta, \tilde{x}) \\
&\geq \underline{f} - \bar{f}(\epsilon/2) + \bar{\alpha}(g(\theta, x^*(\bar{\alpha}, \theta)) - g(\theta, \tilde{x})) \\
&> \underline{f} - \bar{f}(\epsilon/2) + \bar{\alpha} \frac{\epsilon}{2} = -\eta + \frac{\epsilon}{2} \bar{\alpha} > 0
\end{aligned} \tag{4.24}$$

The positivity of (4.23) is in contradiction with the optimality of $\nu(\bar{\alpha}, \theta)$ under the assumption that $g(\theta, x^*(\bar{\alpha}, \theta)) > g^\dagger(\theta) + \epsilon$ and $g(\theta, \tilde{x}) \leq g^\dagger(\theta) + \epsilon/2$. Therefore it must be that $g(\theta, x^*(\alpha, \theta)) \leq g^\dagger(\theta) + \epsilon$ for any $\alpha > 2\eta/\epsilon$. □

4.3 Economic MPC with barriers

In this section, we will introduce a modification of the EMPC formulation by including fixed barrier functions in the problem. The main advantage of using fixed barrier functions is the smoothing of the optimization problem for improved numerical performance and the possibility of enforcing strictly feasible iterations [FE16b]. Since in G-EMPC formulation, the optimal equilibrium point is not known a-priori, adding barrier functions that are centered around an arbitrary point would result in sub-optimal behaviour in transients, and even worse at steady state. Therefore we will investigate an iteratively re-centered barrier function and the main question we will try to answer is whether including iteratively re-centered barrier functions to G-EMPC preserves the closed loop properties of the original scheme or not. If convergence and feasibility properties are not inhibited by barrier functions, then one is justified to use them for enhanced numerical performance. It turns out that using barriers as trajectory regularization in the problem, not only preserves the closed loop convergence properties, but also improves them by guaranteeing convergence of the terminal state to a neighborhood of the optimal steady state. We will first introduce the general framework of regularized MPC, and then explain how barrier functions can be used as a regularizing penalty.

4.3.1 Regularized economic MPC

In this section, we will introduce the novel idea in this chapter, that is the use of barrier functions as a trajectory regularizing cost. We will first provide definitions for the proposed variant for G-EMPC that utilizes a regularizing cost, and then analyze its closed loop properties.

Given a feasible input-state sequence $\bar{\boldsymbol{u}}$ and $\bar{\boldsymbol{x}}$, the regularizing trajectory penalty is defined as

$$P_{\bar{\boldsymbol{u}},\bar{\boldsymbol{x}}}(\boldsymbol{u}, \boldsymbol{x}) = \sum_{k=0}^{N-1} h_{\bar{x}(i)}(x(i)) + h_{\bar{u}(i)}(u(i)) \tag{4.25}$$

where $h_z : R^{n_z} \longrightarrow \mathbb{R} \cup \{\infty\}$ is lower semi-continuous, level bounded within relint($\mathcal{U}(x)$), and chosen to fulfill the property

$$h_z(y) \geq a(\|y - z\|), \quad h_z(z) = 0 \tag{4.26}$$

for some strictly increasing continuous function $a : [0, \infty) \longrightarrow [0, \infty]$ with $a(0) = 0$ and $\lim_{t\to\infty} a(t) = \infty$. From an initial state x, with some feasible $\bar{\boldsymbol{u}}$ and by letting $\bar{\boldsymbol{x}} = \boldsymbol{x}_{\bar{\boldsymbol{u}}}(x)$, the regularized EMPC problem can be written as

$$\begin{aligned} \mathcal{P}^r(x, \bar{\boldsymbol{u}}, \bar{l}) : \min_{\boldsymbol{u}} \quad & J_N(\boldsymbol{u}, \boldsymbol{x}_{\boldsymbol{u}}(x)) + \beta P_{\bar{\boldsymbol{u}},\bar{\boldsymbol{x}}}(\boldsymbol{u}, \boldsymbol{x}_{\boldsymbol{u}}(x)) + \alpha l(\boldsymbol{x}_{\boldsymbol{u}}(N|x), u(N)) \\ \text{s. t.} \quad & \boldsymbol{u} \in \mathcal{U}(x) \\ & (\boldsymbol{x}_{\boldsymbol{u}}(N|x), u(N)) \in \mathcal{E} \\ & l(\boldsymbol{x}_{\boldsymbol{u}}(N|x), u(N)) \leq \bar{l} \end{aligned} \tag{4.27}$$

for some $\alpha, \beta > 0$. The optimal input sequence is denoted by $\boldsymbol{u}^r(x, \bar{l}, \bar{\boldsymbol{u}})$, and the corresponding control policy by $\mu^r(x) := u^r(0|x, \bar{l}, \bar{\boldsymbol{u}})$, choosing the first element of the optimal sequence. The specific regularization we will investigate penalizes the deviation from the previous solution, as used for example in proximal step algorithms. Defining an augmented state as $\xi = (x, \bar{l}, \bar{\boldsymbol{u}})$, the closed loop system can be described as

$$\begin{aligned} x^+ &= f(x, \mu^r(\xi)) \\ \bar{l}^+ &= l(x_{\boldsymbol{u}^r}(N|x), u^r(N, \xi)) \\ \bar{\boldsymbol{u}}^+ &= \vec{\boldsymbol{u}}^r(\xi) \end{aligned} \tag{4.28}$$

where $\vec{\boldsymbol{u}}$ denotes the shifted input sequence

$$\vec{\boldsymbol{u}} := \{u(1), \dots, u(N-1), u(N-1)\} \tag{4.29}$$

Note that according to the above relations, the upper bound $\bar{l}$ is the terminal cost attained by applying the nominal input trajectory $\bar{\boldsymbol{u}}$. The feasible set of state, upper bound and nominal input trajectories is given by

$$\mathbb{F}^r = \{(x, \bar{l}, \bar{\boldsymbol{u}}) \in \mathbb{X} \times \mathbb{L} \times \mathcal{U}^r(x) \mid \bar{l} = l(\boldsymbol{x}_{\bar{\boldsymbol{u}}}(N|x), \bar{u}(N)), \mathcal{P}^r(x, \bar{\boldsymbol{u}}, \bar{l}) \text{ feasible.}\} \tag{4.30}$$

4.3.2 Barrier functions as regularizing cost

With the possibility of re-centering around arbitrary feasible points and being positive definite functions, barrier functions can be used as a regularizing cost in the regularized EMPC formulation (4.27). Using the notation from Section 3.2.5 for re-centered barrier functions, and the definitions of the sets $\mathbb{X}, \mathbb{U}, \mathcal{E}$, we can define the trajectory regularizing barrier cost as

$$B_{\bar{\boldsymbol{u}},\bar{\boldsymbol{x}}}(\boldsymbol{u}, \boldsymbol{x}) = \sum_{k=0}^{N} \Phi^{\mathbb{X}}_{\bar{x}(i)}(x(i)) + \Phi^{\mathbb{U}}_{\bar{u}(i)}(u(i)) + \Phi^{\lambda\mathbb{X}}_{\bar{x}(N)}(x(N)) + \Phi^{\lambda\mathbb{U}}_{\bar{u}(N)}(u(N)) \,. \tag{4.31}$$

Note that the terminal cost is the corresponding barrier functions for $\mathcal{E}$. If a full smoothing of the optimal control problem is desired, with the state and input constraints replaced by their corresponding barrier functions, there is one extra barrier that should be incorporated into the formulation, that is the terminal cost upper bound constraint.

$$\phi_{\bar{l}}(l, z) = -\log(\bar{l} - l(\boldsymbol{x}_{\boldsymbol{u}}(N|x), u(N))) \tag{4.32}$$

Note that the gradient of this barrier function always has the same direction as the gradient of the cost function l. Therefore, adding the barrier for the terminal constraint will have the same effect of increasing the weight α of the terminal cost and therefore does not change the convergence analysis fundamentally. Still, the barrier cannot be enforced directly without adding a relaxation factor, as there could be a case where decreasing the terminal cost is not feasible and the barrier function attains the value ∞. To eliminate this issue, one can add an infinitesimal relaxation ρ to the barrier as $\phi_l(l, z) = -\log(\bar{l} + \rho - l(z, v))$.

Finally, EMPC with barrier functions and terminal equilibrium constraints can be formulated as

$$\mathcal{P}^b(x, \bar{\boldsymbol{u}}, \bar{l}) : \min_{\boldsymbol{u}} \quad J_N(\boldsymbol{u}, \boldsymbol{x}_{\boldsymbol{u}}(x)) + \beta B_{\bar{\boldsymbol{u}},\bar{\boldsymbol{x}}}(\boldsymbol{u}, \boldsymbol{x}) + \alpha l(x_{\boldsymbol{u}}(N|x), u(N)) + \gamma \phi_{\bar{l}}(l(x_{\boldsymbol{u}}(N|x), u(N))) \,. \tag{4.33}$$

4.4 Closed loop properties

A drawback of G-EMPC is the difficulty of ensuring that the terminal state converges to a neighborhood of the optimal steady state for the system, so that the system finds a proper equilibrium point for the stage cost at hand. Even setting the weight for the terminal cost to a very high value or increasing it progressively does not rule out the possibility of the terminal cost getting 'stuck' at an arbitrary sub-optimal point, due to reachability issues [MAA13, FT13]. A solution proposed by [FT13] is switching to a shrinking horizon formulation if the terminal cost is not improving, whereas [MAA14] proposes using a terminal set and controller around equilibrium points. This formulation however takes away the simplicity of the terminal constraint formulation of G-EMPC, which makes it appealing for large scale distributed implementations.

Even though the theoretical convergence issue of stuck terminal cost is difficult to eliminate, in practice it remains as a pathological case. Until now, no examples of such a case has been presented in the literature. In practice, tuning the terminal weight and the horizon length often provides a convergent controller [FT13]. In this section, we provide a mild case of the undesired behaviour of G-EMPC, that can be solved with proper tuning of the parameters. It is still interesting however, that a sub-optimal periodic orbit arises with a linear system and a convex tracking cost. The aim is to see the effect of adding a regularizing cost to the problem.

Consider the following linear system

$$\begin{bmatrix} x_1(t+1) \\ x_2(t+1) \end{bmatrix} = \begin{bmatrix} \cos(\theta) & -\sin(\theta) \\ \sin(\theta) & \cos(\theta) \end{bmatrix} \begin{bmatrix} x_1(t) \\ x_2(t) \end{bmatrix} + \begin{bmatrix} 1 & 0 \\ 0 & 1 \end{bmatrix} \begin{bmatrix} u_1(t) \\ u_2(t) \end{bmatrix} \tag{4.34}$$

with $\theta = \pi/6$. The state and input are subject to constraints $-4 \le x \le 4$ and $-1 \le u \le 1$. The stage cost is selected as

$$l(x, u) = 0.01\|x\|_2^2 + \|u\|_1 \, . \tag{4.35}$$

We apply G-EMPC to this problem with a horizon length of $N = 20$ and terminal cost weight $\alpha = 1$. We then add a regularizing cost that penalizes deviations from the predicted trajectories as in (4.27) using the barrier functions for the state and input constraints, re-centered around the predicted trajectories at each step, with the weight $\beta = 1$. For a better comparison, we do not apply regularization for the initial solution. The closed loop trajectories and the evolution of the stage cost for the initial state, that is the *running* cost, and the terminal cost are shown in Figures 4.1 and 4.2. We can observe that adding regularization keeps the closed loop and open loop trajectories close to each other and prevents the system from getting stuck in a periodic orbit. We will formally analyze this effect in the next section.

Remark 6. *In the example provided above, the original G-MPC converges to the origin if the terminal weight is increased. Therefore, the example does not demonstrate the theoretically possible case of getting stuck in a sub-optimal periodic orbit for any selection of the terminal weight, but only illustrates the effect of regularization on the problem.*

We proceed with the analysis of the closed loop properties of the regularized G-EMPC formulation. The following assumptions are made on the problem data.

Assumption 1. *The feasible set $\mathbb{F}^r$ is bounded and non-empty.*

Assumption 2. *The functions for the dynamical system f, the terminal cost l and the trajectory cost function J_N are continuous.*

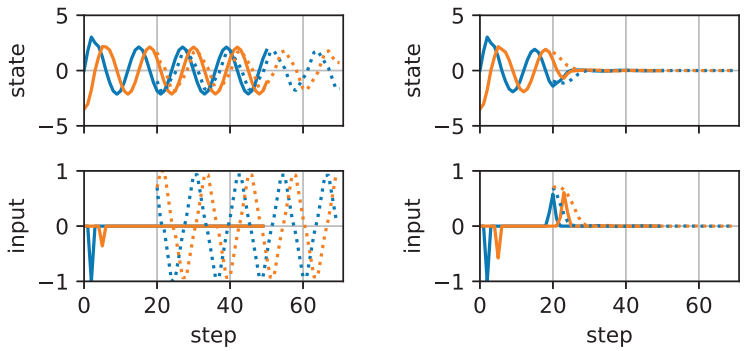

Figure 4.1 – Closed loop (solid) and predicted terminal state and input trajectories (dashed) for G-EMPC (left) and with regularization (right).

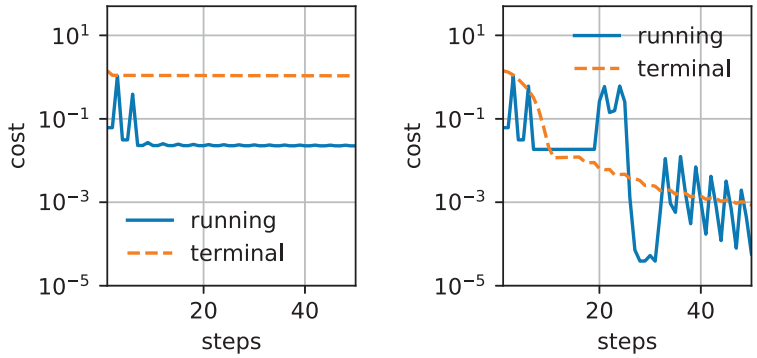

Figure 4.2 – Closed loop stage cost of the state(solid-blue) and predicted terminal steady state (dashed-orange) for G-EMPC (left) and with regularization (right).

Assumption 3. *$\mathcal{P}^r$ is solved to global optimality.*

Reachability between equilibrium points provides an important handle in the convergence analysis. The following assumption is similar to the reachability assumption in [FT13] with the additional requirement that equilibrium points become reachable from a sufficiently small distance.

Assumption 4 (Decreasing terminal stage cost)**.** *For some $\bar{\epsilon} > 0$, there exists a prediction horizon N and a $\delta > 0$, such that for any equilibrium point $(\bar{x}, \bar{u}) \in \mathcal{E}$, there exists at least one state input pair $(\tilde{x}, \tilde{u}) \in \mathcal{E}$ with*

- $\tilde{x} \in \mathcal{R}_N(x), \quad \forall x \in \mathcal{B}(\bar{x}, \delta)$
- $l(\tilde{x}, \tilde{u}) \leq \max(l(z^*, v^*), l(\bar{x}, \bar{u}) - \bar{\epsilon})$

The reachability Assumption 4 is a 'mild' assumption and relies on local controllability of equilibrium points. Specifically, it holds for linear systems considering strictly feasible equilibrium points as we show in the Appendix A.4. Under this assumption, by ensuring convergence to a neighborhood of a steady state, one can guarantee that the system will be able to discover better terminal states in terms of the stage cost. Therefore the undesired situation of having the terminal cost stuck is relieved at most

in N steps, ensuring convergence towards the overall optimal steady state z^*. The following theorem formalizes this claim.

Theorem 4. *Let $\boldsymbol{x}_\mu$ be the trajectory of the system* (4.28) *with an initial state $(x, \bar{l}, \bar{\boldsymbol{u}}) \in \text{relint}(\mathbb{F}^r)$. If Assumptions 1-4 hold, then there exists a finite selection of α, β and horizon length N such that*

- *The predicted terminal state $x_{\boldsymbol{u}^r}(N|x_\mu(t))$ converges to the set*

$$\mathcal{M} := \{x : \exists u, (x, u) \in \mathcal{E}, l(x, u) \leq l(z^*, v^*) + \epsilon\} \tag{4.36}$$

- *The asymptotic average performance is bounded by*

$$\bar{J}_\infty(\boldsymbol{x}_\mu) \leq l(z^*, v^*) + \epsilon \tag{4.37}$$

for some $\epsilon < \bar{\epsilon}$. Furthermore ϵ can be made arbitrarily small by proper selection of N, α and β.

Theorem 4 states that if properly tuned, the terminal cost of the regularized G-MPC converges to a neighborhood of the optimal steady state. Furthermore the asymptotic average performance of the system can be made arbitrarily close to the performance of the optimal equilibrium point (z^*, v^*). We will provide the proof of Theorem 4 at the end of the chapter, after developing the required preliminary results. We will also show that tuning the weights can be done in a simple manner by increasing them until convergent behaviour is obtained.

4.4.1 Sub-optimality of the terminal cost

Since the predicted terminal state is only constrained to be a steady state, it is not obvious how it will change according to problem data and in closed loop operation. Some properties of the terminal state input pair can be deduced from the problem formulation. The following lemma deals with the suboptimality of the terminal stage cost.

Lemma 8 (Terminal cost suboptimality[FT13])**.** *Let assumptions 1-3 hold. Let $\boldsymbol{u}^r$ be the solution of the problem $\mathcal{P}^r$. Given $\epsilon > 0$, there exists a finite value $\alpha(\epsilon)$ such that for any $\alpha > \alpha(\epsilon)$, $(x, \bar{l}) \in \mathbb{F}$,*

$$l(x_{\boldsymbol{u}^r}(N|x), u^r(N)) \leq l^R(x) + \epsilon\ . \tag{4.38}$$

Proof. The proof is an application of Theorem 3. □

Lemma 8 implies that, starting from an initial state x, one can increase the weight α in order to push the terminal state towards the optimal reachable steady state $z^R(x)$, and the distance can be made arbitrarily small. If the overall optimal steady state z^* is within the N step reachable set, that is $z^* \in \mathcal{R}_N(x)$ and $z^R(x) = z^*$, this will recover the closed loop properties of the standard EMPC with fixed terminal points [FT13]. However, this might not be the case and even though one can attain the optimal *reachable* steady state $z^R(x)$, the question of convergence to the optimal steady state z^* remains un-answered. In theory, there is nothing that prevents the system from getting stuck at a certain level of terminal cost, if the reachable sets are not expanding in favorable directions.

4.4.2 Reachability of equilibrium points with improved terminal cost

The reachability assumption (4) implies that in case the terminal cost is stuck at a sub-optimal level, one should get closer to an equilibrium point with the available terminal cost level. From a sufficiently small distance, one can expand the reachable set so as to discover equilibrium points that decreases the terminal cost. Considering the sub-optimality properties of the terminal cost established with Lemma 8, if the weight α is sufficiently large, reachability of sufficiently better terminal cost values will result in a decrease in the terminal cost of the solution trajectory.

The remaining issue is to ensure that the system will converge towards an equilibrium state with the currently available terminal cost, if the terminal cost is stuck. Instead of having a switched algorithm as in [FT13], we will show that the regularizing trajectory penalty will drive the system to a neighborhood of the currently available reachable steady state, the size of which depends on the weighting of the regularizing penalty. Based on the reachability assumption (4), we know that from an equilibrium point that is not optimal, the system will find a new terminal state with improved cost, with sufficiently large weight on the terminal cost. Therefore regularizing the trajectory cost will ensure convergence to the optimal steady state, without further modifications on the algorithm.

Due to the recursive upper bound constraint on the terminal cost there are two possibilities regarding the evolution of the terminal cost. The terminal cost function can either decrease or remain constant. We will analyze the case where the cost function is stuck, *i.e.* the system is not able to discover equilibrium points with sufficiently lower cost value. Note that while the terminal cost remains constant the terminal state may still change on the level set of the terminal cost. In such situations, the upper bound constraint on the terminal cost remains active. Since active inequalities can be replaced by equality constraints without affecting the solution, we will analyze the equality constrained problem, denoted by $\mathcal{P}^q(x, \bar{\boldsymbol{u}}, \bar{l})$, with

$$l(x_{\boldsymbol{u}}(N|x), u(N)) = \bar{l}. \tag{4.39}$$

In this case the terminal cost in the objective function has no effect due to the equality constraint. Therefore one can choose the regularizing weight β, independent of α, in order to enforce a strict level of regularization. The following lemma formalizes this idea.

Lemma 9. *Let Assumptions 1-3 hold. Let $\boldsymbol{u}^q$ be the solution to problem $\mathcal{P}^q$. Given $\delta > 0$ there exists a finite weight $\beta(\delta)$ such that for any $(x, \bar{l}, \bar{\boldsymbol{u}}) \in \mathbb{F}^r$,*

$$P_{\bar{\boldsymbol{u}},\bar{\boldsymbol{x}}}(\boldsymbol{u}^q, \boldsymbol{x}_{\boldsymbol{u}^q}(x)) \leq \delta \tag{4.40}$$

$$\|x_{\boldsymbol{u}^q}(i|x) - \bar{x}(i)\| \leq a^{-1}(\delta), \quad \forall i \in [0, \dots, N]\,. \tag{4.41}$$

Proof. Note that for a feasible $\bar{\boldsymbol{u}} \in (\mathcal{U}(x))$ and $\bar{l} = l(x_{\bar{u}}(N), \bar{u}(N))$, the minimum attainable value for the regularizing cost is 0. The proof follows directly from Theorem 3 and the definition of $P_{\bar{\boldsymbol{u}},\bar{\boldsymbol{x}}}$. □

The purpose of enforcing a strict level of proximity to the previous solution is to ensure convergence to a neighborhood of the N step previously available terminal state, with the same terminal cost as the current $\bar{l}$.

Lemma 10. *Let $\boldsymbol{x}_\mu$ be the trajectory of the system* (4.28)*. For any finite $\sigma > 0$ there exists a regularization weight $\bar{\beta}(\sigma)$ such that if $\beta > \bar{\beta}(\sigma)$ and $\bar{l}(t)$ remains constant for N steps starting from time step t, then*

$$x_\mu(t+N) \in \mathcal{B}(x_{\boldsymbol{u}}(N|x_\mu(t)), \sigma) \tag{4.42}$$

Proof. Since the nominal trajectory is defined as $\bar{\boldsymbol{x}}(t) = \vec{\boldsymbol{x}}_{\boldsymbol{u}^r}(x_\mu(t-1))$, from Lemma 9, we have that $\forall k \in [1, \dots, N-1]$,

$$\|x_{\boldsymbol{u}^r}(k|t+1) - x_{\boldsymbol{u}^r}(k+1|t)\| \leq \kappa, \tag{4.43}$$

for some $\kappa > 0$, which implies that the closed loop state remains close to the predicted one, due to the regularizing cost. By induction, after $N-1$ steps, we have that

$$\|x_{\boldsymbol{u}^r}(1|t+N-1) - x_{\boldsymbol{u}^r}(N|t)\| \leq (N-1)\kappa \tag{4.44}$$

Since without any disturbance $x_\mu(t+N) = x_{\boldsymbol{u}^r}(1|t+N-1)$, we have that

$$\|x_\mu(t+N) - x_{\boldsymbol{u}^r}(N|t)\| \leq (N-1)\kappa \tag{4.45}$$

By virtue of Theorem 3, κ can be made arbitrarily small by increasing the value of β and therefore there exists a $\bar{\beta}$ such that $(N-1)\kappa = \sigma$. □

4.4.3 Convergence and average performance

Equipped with the previous development, we are now in a state to prove Theorem 4, which claims guaranteed convergence of terminal steady state and asymptotic average performance.

Proof. Firstly, the recursive feasibility of the trajectory is trivially guaranteed by virtue of the terminal equilibrium conditions. Thanks to the terminal steady state, shifted solutions always provide a feasible solution, guaranteeing $(x, \bar{l}, \bar{\boldsymbol{u}})^+ \in \mathbb{F}^r$. Furthermore, in the case where the penalty functions go to infinity on the boundaries of the feasible set, such as the case with logarithmic barrier functions, the system always chooses trajectories within relint($\mathbb{F}^r$), since the trajectories on the boundaries are prevented by the regularization penalty term. Next, we consider the main claim on the convergence of predicted terminal states. At each time step t there are two possibilities with respect to the evolution of the terminal cost bound $\bar{l}(t)$. It either decreases compared to the previous step, or remains constant. When $\bar{l}(t)$ decreases, it contributes to the convergence of the terminal state towards the optimal steady state. Therefore we will analyze the undesirable case where it remains constant. From Lemma 10, in the case where the terminal cost remains constant for N steps, we know that it is possible to ensure that the state $x_\mu(t)$ reaches a σ-ball, $\mathcal{B}(0, \sigma)$, around the N step previously predicted terminal state $x_{\boldsymbol{u}^r}(N|x_\mu(t-N))$, by choosing a sufficiently large β value. Therefore, given $\bar{\epsilon} > 0$, with a sufficiently long horizon N, one can choose β such that $\sigma \leq \delta(\bar{\epsilon})$ as stated in Assumption 4 and Lemma 10. This guarantees that after N steps of the terminal cost remaining constant, there is an equilibrium point $(\tilde{x}, \tilde{u})$ within the reachable set $\mathcal{R}_N(x_\mu(t))$ such that

$$l^R(t) \leq l(\tilde{x}, \tilde{u}) \leq \max(l(z^*, v^*), \bar{l}(t) - \bar{\epsilon}) \,. \tag{4.46}$$

for a sufficiently long horizon N, as stated in Assumption 4. In this case, either the optimal feasible

steady state z^* becomes reachable with $z^R(t) = z^*$ or we have that

$$l^R(t) \leq \bar{l}(t) - \bar{\epsilon} \tag{4.47}$$

According to Lemma 8 there exists a sufficiently large α such that

$$\bar{l}(t+1) \leq \bar{l}(t) - \theta \tag{4.48}$$

with $\theta := \bar{\epsilon} - \epsilon$, for some $\epsilon < \bar{\epsilon}$, ensuring the reduction in the terminal cost by an amount no less than $\theta > 0$. Therefore, the solution is guaranteed to attain a strict improvement of the terminal cost in every N steps until $\bar{l}(t) - \bar{\epsilon} \leq l(z^*, v^*)$. At this point a final improvement is guaranteed such that the terminal cost is reduced by θ, leading to the inequality

$$\bar{l}(t) \leq l(z^*, v^*) + \epsilon \tag{4.49}$$

which ensures that $x_{\boldsymbol{u}^r}(N|x_\mu(t)) \in \mathcal{M}$, since $\bar{l}$ is an upper bound for the terminal cost.

The upper bound on the asymptotic performance of the system $\bar{J}_\infty$ can be obtained by analyzing the evolution of the cost values for the system. We first define the part of the cost function without the regularizing term

$$\begin{aligned} V_N(\boldsymbol{u}, x) &= J_N(\boldsymbol{u}, \boldsymbol{x}_{\boldsymbol{u}}(x)) + \alpha l(x_{\boldsymbol{u}}(N|x), u(N)) \\ &= \sum_{k=0}^{N-1} l(x_{\boldsymbol{u}}(k|x), u(k)) + \alpha l(x_{\boldsymbol{u}}(N|x), u(N)) \end{aligned} \tag{4.50}$$

and define

$$V_N^r(t) = V_N(\boldsymbol{u}^r(t), x_\mu(t)) \tag{4.51}$$

Note that the regularizing term penalizes deviation from shifted previous solutions, $\bar{\boldsymbol{u}}$, as defined in (4.28). Due to the optimality of $\boldsymbol{u}^r$, the definitions of the regularizing term with $P_{\bar{\boldsymbol{u}},\bar{\boldsymbol{x}}}(\bar{\boldsymbol{u}}, \bar{\boldsymbol{x}}) = 0, P_{\bar{\boldsymbol{u}},\bar{\boldsymbol{x}}}(.) \geq 0$ and the nominal trajectories $\bar{\boldsymbol{u}}(t) = \vec{\boldsymbol{u}}^r(t-1)$, $\bar{\boldsymbol{x}}(t) = \vec{\boldsymbol{x}}_{\boldsymbol{u}^r}(t-1)$, we have that

$$\begin{aligned} V_N^r(t+1) + \beta P_{\bar{\boldsymbol{u}}(t+1),\bar{\boldsymbol{x}}(t+1)}(\boldsymbol{u}^r(t+1), \boldsymbol{x}_{\boldsymbol{u}^r}(t+1)) &\leq V_N(\vec{\boldsymbol{u}}^r(t), x_\mu(t+1)) + \beta P_{\bar{\boldsymbol{u}}(t+1),\bar{\boldsymbol{x}}(t+1)}(\vec{\boldsymbol{u}}^r(t), \vec{\boldsymbol{x}}_{\boldsymbol{u}^r}(t)) \\ V_N^r(t+1) + \beta P_{\bar{\boldsymbol{u}}(t+1),\bar{\boldsymbol{x}}(t+1)}(\boldsymbol{u}^r(t+1), \boldsymbol{x}_{\boldsymbol{u}^r}(t+1)) &\leq \sum_{k=1}^{N-1} l(x_{\boldsymbol{u}^r}(k|t), u^r(k|t)) + (\alpha+1) l(x_{\boldsymbol{u}^r}(N|t), u^r(N|t)) \\ V_N^r(t+1) &\leq V_N^r(t) - l(x_\mu(t), \mu^r(t)) + l(x_{\boldsymbol{u}^r}(N|t), u^r(N|t)) \\ &\Downarrow \\ V_N^r(t+1) - V_N^r(t) &\leq -l(x_\mu(t), \mu^r(t)) + \bar{l}(t+1) \end{aligned} \tag{4.52}$$

We know that $\bar{l}$ converges to a value that is less than $l(z^*, v^*) + \epsilon$. Stated otherwise $\varepsilon(t)$ converges to a value less than ϵ, where $\varepsilon(t) := \bar{l}(t+1) - l(z^*, v^*)$. Plugging this definition into above relation, we get

$$V_N^r(t+1) - V_N^r(t) \leq -l(x_\mu(t), \mu^r(t)) + l(z^*, v^*) + \varepsilon(t) \tag{4.53}$$

Following the similar proofs in [MAA13, FT13], we take averages for both sides

$$\begin{aligned}\liminf_{T\to\infty}\frac{1}{T}\sum_{t=0}^{T-1}\left(V_N^r(t+1)-V_N^r(t)\right) &\le \liminf_{T\to\infty}\frac{1}{T}\sum_{t=0}^{T-1}\left(-l(x_\mu(t),\mu^r(t))+l(z^*,v^*)+\varepsilon(t)\right)\\ &\le l(z^*,v^*)+\limsup_{T\to\infty}\sum_{t=0}^{T-1}\frac{1}{T}\varepsilon(t)+\liminf_{T\to\infty}\frac{1}{T}\sum_{t=0}^{T-1}\left(-l(x_\mu(t),\mu^r(t))\right)\\ &\le l(z^*,v^*)+\epsilon-\limsup_{T\to\infty}\frac{1}{T}\sum_{t=0}^{T-1}\left(l(x_\mu(t),\mu^r(t))\right)\end{aligned} \tag{4.54}$$

In the derivation above, we have used the facts that, $\liminf(a(n)+b(n)) \le \liminf(a(n))+\limsup(b(n))$ for finite a_n, b_n and both $l(.)$ and ε are bounded. Since the cost function V_N is continuous and the set $\mathbb{F}^r$ is bounded, there exists $\bar{V}$ such that $V_N^r(t) \ge \bar{V}$ for all t. Therefore, we can further simplify the left hand side by

$$\begin{aligned}\liminf_{T\to\infty}\frac{1}{T}\sum_{t=0}^{T-1}\left(V_N^r(t+1)-V_N^r(t)\right) &= \liminf_{T\to\infty}\frac{1}{T}\left(V_N^r(T)-V_N^r(0)\right)\\ &\ge \liminf_{T\to\infty}\frac{1}{T}V_N^r(T)+\liminf_{T\to\infty}\frac{1}{T}(-V_N^r(0))\\ &\ge \liminf_{T\to\infty}\frac{1}{T}\bar{V}+\liminf_{T\to\infty}\frac{1}{T}(-V_N^r(0))\\ &= 0\end{aligned} \tag{4.55}$$

Combining the inequalities, we arrive at

$$\begin{aligned}0 &\le l(z^*,v^*)+\epsilon-\limsup_{T\to\infty}\frac{1}{T}\sum_{t=0}^{T-1}\left(l(x_\mu(t),\mu^r(t))\right)\\ \limsup_{T\to\infty}\frac{1}{T}\sum_{t=0}^{T-1}\left(l(x_\mu(t),\mu^r(t))\right) &\le l(z^*,v^*)+\epsilon\end{aligned} \tag{4.56}$$

Finally, by virtue of Lemma 8, for a fixed horizon length that satisfies Assumption 4, and a regularization weight β, one can increase the weight α such that θ is as large, thus ϵ is as small, as possible. □

Theorem 4 provides guidelines for tuning the weights on the terminal cost and the regularizing penalty. Firstly, it is necessary to have a sufficiently large regularizing weight β independent of α to guarantee convergence to a neighborhood of available terminal states. Given this β value, α should be chosen sufficiently large to guarantee progressive reachability. A simple tuning strategy would be increasing β while increasing α with an exponential factor when the terminal cost gets stuck.

4.5 Case study: Power grid frequency control

We have shown in the previous section that if barrier functions are added as a trajectory regularizing cost to G-EMPC, convergence properties of the closed-loop system are preserved and improved. In this section, we will illustrate the practical benefits of using barrier smoothing with a case study. The problem in consideration is the automatic generation control (AGC) of power grids and the case study is taken from [RFT12]. We consider the case with four interconnected areas that are connected in a chain formation, corresponding to Scenario 1 in [RFT12]. The areas regulate their frequencies, under the effect of tie-line power flows from neighboring areas and the internal demand change. The dynamics of the

interconnected system in continuous time can be described as

$$\begin{aligned} \dot{x}^i &= A^i x + B^i u^i + L^i \Delta P_L^i + w^i, \quad \forall i \in \mathbb{Z}_{[1,m]} \\ w^i &= \sum_{j \in \mathcal{N}^i} E^{ij} \Delta\theta^j \\ \Delta\theta^i &= C^i x^i \end{aligned} \tag{4.57}$$

$$\begin{aligned} A^i &= \begin{bmatrix} 0 & 1 & 0 & 1 \\ -\frac{\sum_{j\in\mathcal{N}^i} P^{ij}}{2H^i} & \frac{D^i}{2H^i} & \frac{1}{2H^i} & 0 \\ 0 & 0 & -\frac{1}{T_t^i} & \frac{1}{T_t^i} \\ 0 & \frac{1}{R^i T_g^i} & 0 & -\frac{1}{T_g^i} \end{bmatrix} \quad B^i = \begin{bmatrix} 0 \\ 0 \\ 0 \\ \frac{1}{T_g^i} \end{bmatrix} \\ E^{ij} &= \begin{bmatrix} 0 \\ \frac{P^{ij}}{2H^i} \\ 0 \\ 0 \end{bmatrix} \quad C^{ji} = \begin{bmatrix} 1 & 0 & 0 & 0 \end{bmatrix} \quad L^i = \begin{bmatrix} 0 \\ -\frac{1}{2H^i} \\ 0 \\ 0 \end{bmatrix} \end{aligned} \tag{4.58}$$

We refer the reader to [RFT12] for the definition and values of the parameters $P^{ij}, H^i, R^i, D^i, T_t^i, T_g^i$. The state of a sub-system i is given by $x^i = (\Delta\theta^i, \Delta\omega^i, \Delta P_m^i, \Delta P_v^i)$ which represents deviation from nominal values of angular displacement of the rotor, speed of the rotor or the 'frequency' of the area, mechanical power and steam valve positions. The input u^i represents the power reference ΔP_{ref}^i and ΔP_L^i is the load deviation at area i acting as a disturbance. The state and input are constrained as $-0.1 \leq (\Delta\theta, \Delta\omega) \leq 0.1$ and $-0.55 \leq (\Delta P_m, \Delta P_v, \Delta P_{\text{ref}}) \leq 0.55$.

We modify the cost function to be an 'economic' one that reflects the material cost of power generation. For area-i, the stage cost is given by

$$l^i(x, u) = \|\Delta\theta^i\|^2 + \|\Delta\omega^i\|^2 + \|\Delta P_{\text{ref}}^i\|^2 + c^i \Delta P_m^i \tag{4.59}$$

where the the material price term c^i is chosen to create an 'economic' imbalance between the areas as shown below in Table 4.1. The remaining quadratic penalty terms penalize deviations in frequency,

Table 4.1 – Normalized cost of power production per kW.

Area	1	2	3	4
c^i	50	10	30	100

phase angle and regularize the input. The aggregate stage cost is the mean of the local area cost terms.

$$l(x, u) = \frac{1}{m} \sum_{i=1}^{m} l^i(x^i, u^i) \tag{4.60}$$

We apply the original G-MPC [FT13] with a sampling rate of 1 second, horizon length $N = 7$ and a terminal weight $\alpha = 1$. We further add barrier weights as a trajectory regularization cost with weight $\beta = 0.3$. The barrier weight is chosen to balance warm starting with improved numerical performance obtained from smoothing and the close loop performance. Finally we add simple fixed barriers centered around the origin to the G-MPC in the third case. The barriers are also added to the upper bound constraint on the terminal cost function with a weight of 10^{-4}. This small weight does not interfere with the closed loop behaviour of the system while providing smoothing.

The problem is solved with a custom solver that implements the C-PDIP method with path following, as explained in Chapter 3. For the case with no barriers, interior point methods can not be warm started efficiently, so the algorithm is restarted from a central point. For the case with barriers, we can use the shifted solutions directly for warm starting the algorithm.

A comparison of numerical solver and closed loop performance is provided in Figure 4.3. We can observe that G-MPC performs quite well with limited horizon and a small terminal weight. It converges to a very small neighborhood of the optimal steady state. Adding fixed barriers centered around the origin improves the numerical performance of the primal-dual interior point solver considerably, thanks to the warm starting effect of using shifted solutions. Another advantage is that warm starting is obtained with a primal feasible solution, which brings in the possibility of early termination of the algorithm. However, the down-side is that the cost function has been distorted due to barriers and the solutions stay far away from the constraints, diminishing the closed loop performance significantly. Finally re-centering the barriers at every step around the previous trajectories restores the closed-loop performance of the system and ensures convergence to a neighborhood of the optimal steady state. In return, the warm starting is less effective compared to fixed barriers, however the numerical performance is still significantly improved compared to the original G-MPC formulation with 45% reduction in worst-case iteration number.

We can observe from Figure 4.3 that the regularized G-MPC does not back away from the constraints while providing improved performance to the primal-dual solver with primal-feasible starting points. Note that the improved numerical performance is due to the smoothing effect of the barriers. However when the system converges to a steady state on the constraints, the barriers for the active constraints will have reduced weights in order to center the barriers cost around the current solution, also reducing the smoothing effect with respect to these 'active' constraints. In this case, if the system receives a large disturbance or a change in the cost function, it might face numerical difficulties [DFH09]. This issue can be mitigated by re-setting the barrier weights when there is a change in the cost, or limiting the reduction of the barrier weights in the re-centering, that will result in approximate re-centering.

In order to illustrate the practical aspects of implementing regularized EMPC with barrier functions, we carry out a simulation where the system is subjected to both persistent and intermittent disturbances, as shown in Figure 4.5. The disturbance in this case is a variation in the load of each power grid area. The load disturbance is sampled from a Gaussian distribution with zero mean and 20kW standard deviation. Furthermore the load of Area 3 is increased by 200kW, 30 seconds into the simulation.

As discussed above, the disturbance necessitates using a lower bound on the barrier weights of the constraints, for preserving the smoothing effect and the numerical performance of the solver. We set the regularization coefficient $\beta = 1$ and choose the lower bound as $\beta w \geq 0.05$, where w is the weight used for re-centering the barrier functions. The lower bound is enforced by simply saturating the combined barrier weights after re-centering. The upper bound constraint on the terminal cost was removed, as it always remained inactive. Without this constraint, all remaining constraints in the problem are linear, in which case the primal-dual Newton method is guaranteed to preserve primal feasibility [BJ17a]. When there is no disturbance present, even only one iteration suffices to preserve recursive feasibility of the system. However when the system is subjected to disturbance, the primal-feasibility levels of the shifted solutions will be reduced, depending on the amount of disturbance. We have found that with the given amount of disturbance, the chosen levels of regularization weight and lower bound on the barrier weights, a maximum of four iterations result in satisfactory performance, as shown in Figure 4.5. These values should be chosen judiciously according to the application at hand and the severity of disturbance values.

4.6 Conclusion

In this chapter we have presented a novel EMPC formulation that incorporates regularization in the cost function. The regularization can be implemented with barrier functions that are re-centered around previously predicted trajectories. The use of barrier functions as a regularizing cost brings two benefits: improved numerical performance due to smoothing, and guaranteed convergence to a neighborhood of the optimal steady state. The approach is based on the EMPC with generalized terminal constraints [FT13], which is suitable for decomposition and large scale distributed implementations following our work on interior point decomposition methods [BJ17a, BPJ17b]. Simulation studies indicate that it is possible to reduce iteration number significantly without giving away closed loop performance and convergence guarantees.

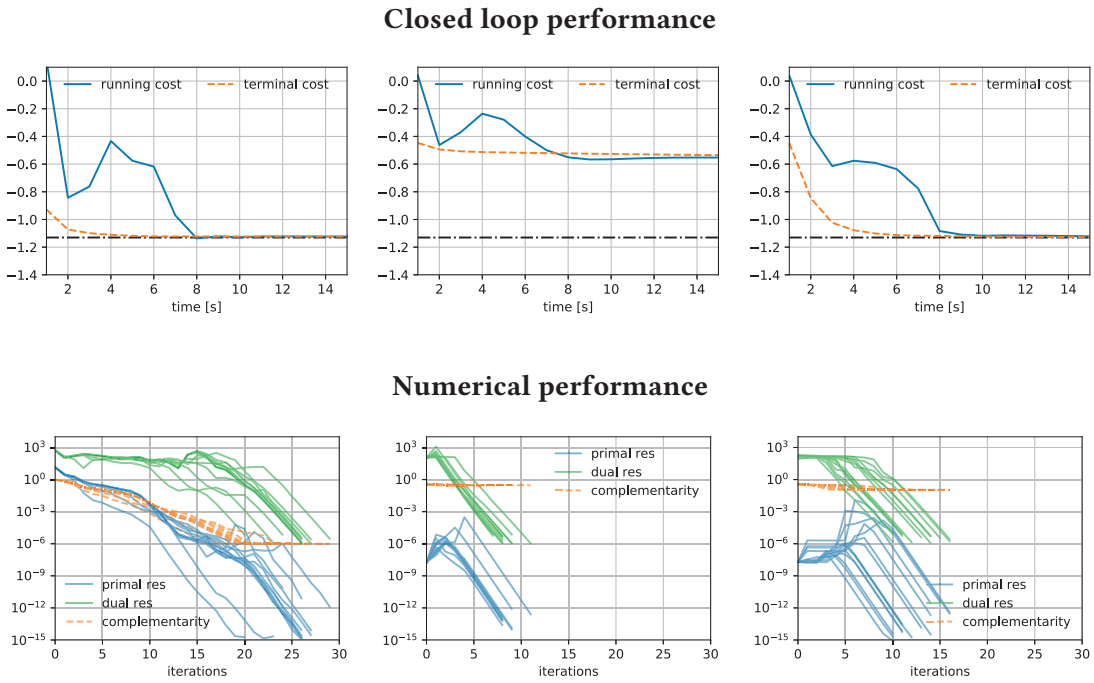

Figure 4.3 – Comparison of controller and solver performance for different integration of barriers. From left to right, the cases with: no barriers, fixed barriers around the origin and iteratively re-centered barriers. For the case with barriers, the reduction factor for the complementarity target is set to 0.1. For all cases, an aggressive long-step approach is used, which takes 99% of the maximum feasible step-length.

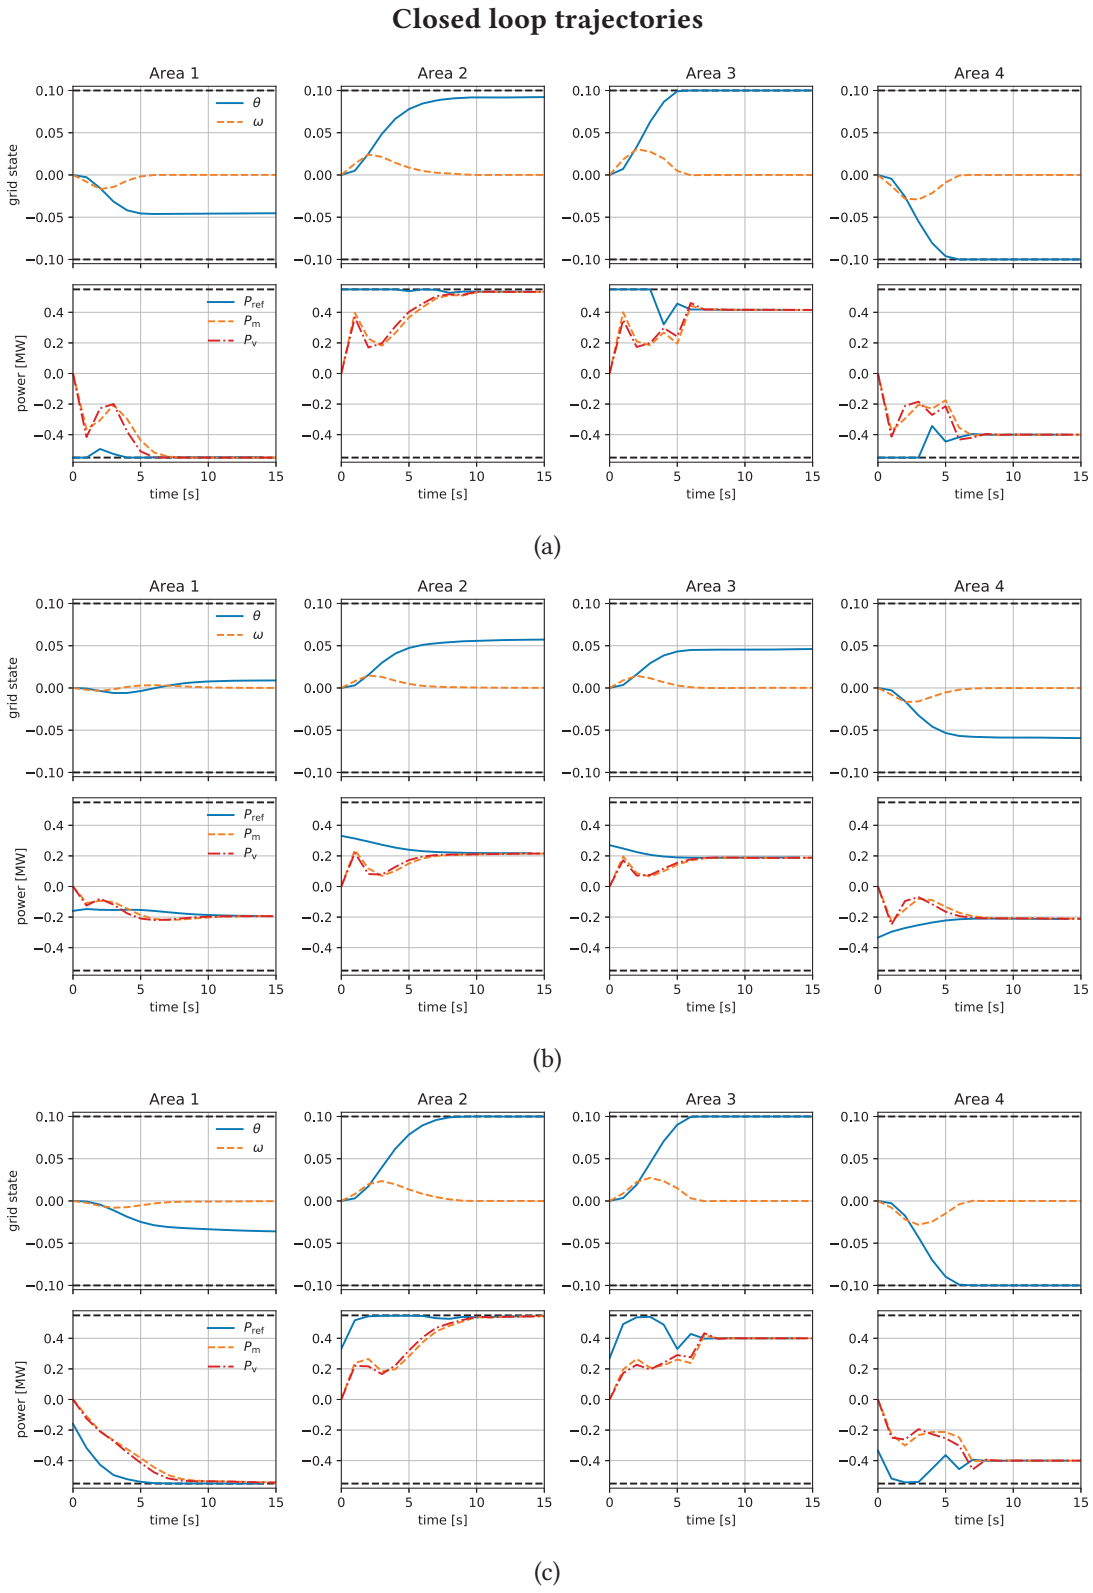

Figure 4.4 – Comparison of Closed loop input and state trajectories for individual areas for the case with (a) no barriers, (b) fixed barriers around the origin and (c) iteratively re-centered barriers.

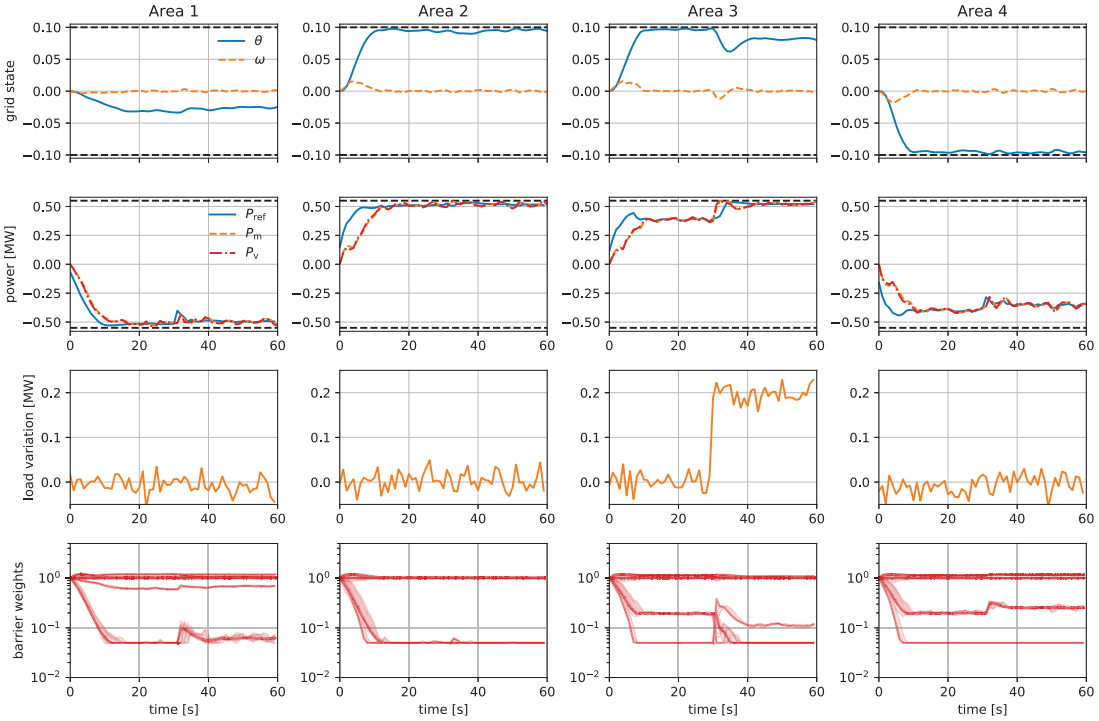

(a)

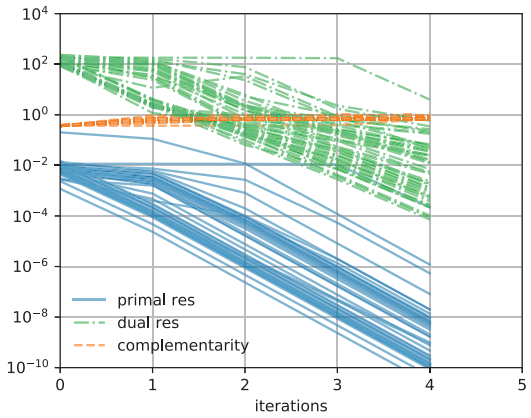

(b)

Figure 4.5 – Closed loop trajectories (a) and numerical performance (b) of regularized EMPC with re-centered barriers under disturbance, with $\beta = 1$ and early termination after four iterations of primal-dual Newton method.

5 Conclusion

In this thesis we have considered solution methods for multi-agent optimization and control problems that arise in smart-grid applications. The motivation is the prospect of large scale demand response using the available storage capacity of commercial-buildings. Once this capacity is facilitated, there is significant potential to support the power grid in applications such as frequency regulation and voltage support.

The problem of large scale demand response provision embeds the most challenging aspect of smart grids, coordination of complex agents in a safety critical environment. To tackle this challenge, we have put forward two approaches that facilitate safe and efficient coordination; hierarchical control with tracking commitment, and coordinated interior-point methods based on decomposition in the output domain. Furthermore, we have demonstrated that coordinated interior-point methods can be efficiently implemented in model predictive control to solve large scale problems, while enhancing numerical and system theoretic properties.

Hierarchical control with tracking commitment is based on characterizing a subsystem's capability of tracking reference signals from a prescribed set. The reference set is used by a coordinator which sends reference signals to many agents, in order to achieve a common goal. The formulation was motivated by and naturally applies to the frequency regulation problem in smart grids. Apart from the existing literature, we have formulated the problem in a way that enables the subsystem itself to compute the set of feasible set of references. The methodology is based on convex robust programming and allows efficient solution of problems that involve manipulation of the uncertainty set, which was not addressed in the adjustable robust optimization framework. We have demonstrated with simulation studies that the tracking commitment formulation can be used to characterize the power consumption flexibility of a building, allowing aggregation in large scale.

Next, we have analyzed the properties of coordinated interior-point methods that can be used to solve large scale multi-agent problems, while preserving network and subsystem safety. We formulated the multi-agent problem in output coupling form, which fits in naturally to smart grid applications. In the example of buildings providing demand response, each building is a complex system with internal states constraints and dynamics, however the coupling with other buildings arises due to a low dimensional output, total-power consumption. When the complex internal characteristics of the agent is projected down to the output domain, a coordinator can easily handle numerous agents, due to the lower dimensionality of the output variable. In technical terms, this is achieved by computing gradient and Hessian

information in the output domain of the agent, and using this information to solve the multi-agent problem at hand. Using interior-point methods allows including constraint information to the cost function, allowing safe iterations that preserve network and agent feasibility.

Interior point methods can be very efficient in terms of convergence, when applied in the primal-dual domain and using predictor corrector methods. We have analyzed the properties of the well-known Mehrotra's predictor corrector in the coordinated primal-dual interior point framework. Having observed that the communication effort may become excessive in a naive implementation, we have proposed a decentralized prediction strategy. The proposed method can be guaranteed to converge to the optimal solution, if certain safeguards are applied that are similar to the standard convergence safeguards in the interior-point literature. We have demonstrated with numerical examples that depending on the problem data, decentralized predictors can reduce the communication effort significantly. However in real problems it is not clear cut as to which type of coordinated interior-point method is more suitable. The path following method has the advantage of simple implementation and acceptable performance. The predictor-corrector method has very fast convergence properties in terms of iteration count, whereas the decentralized predictors result in a trade-off between performance in terms of iterations and communication rounds.

Finally we have considered the application of C-PDIP methods in large-scale distributed MPC problems that incorporate not only coupling in constraints and objectives, but also the dynamics between subsystems. Such problems can be easily formulated in our generic multi-agent optimization framework and solved with C-PDIP methods. In this setting, a very attractive option to guarantee recursive feasibility of the networked systems is using terminal equilibrium conditions that allows the problem to choose any admissible steady state for the system. Such constraints can be incorporated in the multi-agent problem as local constraints, without introducing extra variables and ensuring global steady state operation. While using interior point methods in MPC applications, a popular method for speeding up numerical performance, and facilitate warm starting is to fix the weight of barrier functions in the problem and consider them as an additional cost term. While improving the numerical performance, this additional cost term changes the closed-loop properties of the system and necessitates the analysis of these properties. We have shown that in case of EMPC with terminal equilibrium conditions, it is possible to incorporate barrier functions, in a way to also enhance system-theoretic properties. Specifically, we use barrier functions with iterative re-centering as a regularization cost on the problem. At each step, the barriers are re-centered such that they penalize deviations from previously predicted trajectories. We have shown that this type of trajectory regularization can be used to enforce the system to approach its predicted terminal state, which facilitates convergence of the terminal cost in the EMPC with terminal equilibrium conditions. Furthermore, the regularizing cost does not deteriorate the asymptotic average performance of the system. In light of these results, the use of barrier functions for enhanced numerical performance is not only justified but also encouraged, considering the closed-loop properties of the controlled system. We demonstrate the results with a popular smart-grid problem, frequency regulation of interconnected areas. The problem is solved with the C-PDIP method that incorporates re-centered barriers. We show that using barriers reduces the iteration count significantly while preserving the closed loop performance of the system.

In all our methods, we have taken a stance that promotes coordination using a central agent. A disadvantage of such an approach is that the system reliability and performance hinges on the central agent, creating a point of failure for the whole system. However, in complex engineering environments such as the smart grids, there is a split of domain knowledge and responsibility. The grid operators know how to operate power grids, and building operators know how to operate buildings. Implementing

fully distributed algorithms would mean that the responsibility of maintaining the grid safety is also distributed among the agents. Without sufficient domain knowledge and expertise, this responsibility can not be upheld by local agents, unless the domain knowledge is also distributed so that each local agent becomes a power-grid expert. Considering the difficulty and the cost of distributing responsibility and expertise in a safety critical engineering application and the efficiency of information exchange, coordinated methods based on a central agent is well justified in smart-grids applications. If a fully distributed implementation is necessary, it is often possible to distribute the task of the central agent among the agents as we have noted in Chapter 3.

5.1 Future directions

Most of the methodologies presented in this thesis rely on the key assumption of convexity of the problems solved. In model predictive control, this translates to linear system models. It is therefore a natural, but challenging, extension to incorporate nonlinear systems and non-convex problems into the presented methods.

In the case of robust tracking commitment, the key idea is to modulate uncertainty sets in a linear robust program with convex programming, allowing efficient computation. However the framework of modulating uncertainty sets can be formulated as a generalized semi-infinite program [GVRSS08] that consider generic non-linear, non-convex problems. There are numerical methods to tackle these problems, but there is ample room for improvement [GVRSS08]. A possible approach here would be using a sequential method, by approximating the problem at hand by a sequence of linear robust programs with adjustable sets. The practical application in terms of robust tracking commitment would be immediate, as one can incorporate nonlinear dynamics of the buildings, as well as the nonlinearities of the grid, in order to characterize flexibility of whole distribution networks. We note here the work of [Qur17] that solves the tracking commitment problem for a building considering the nonlinear HVAC system, by using a hierarchical approach. Another application would be characterizing the flexibility of industrial systems, such as chemical production plants, that represent a big share of electricity consumption coupled with complex nonlinear process dynamics.

Regarding the coordinated interior-point methods, the extension to non-convex problems and non-linear systems is very natural, in the sense that interior-point methods can incorporate non-convex problems directly. Of course primal feasible iteration guarantees and the efficiency of predictor-corrector methods are usually lost, however most of the presented methodology in this thesis can be still applied. Especially for the case of EMPC with barrier functions, the warm starting and smoothing effect of using barriers in the optimal control problems is a well known method for also nonlinear systems [Oht04, DFH09] and all of the results presented in Chapter 4 directly apply to nonlinear systems.

Finally, experimental validation of the proposed methods will be the final validation for their practicality and engineering value. For the case of robust tracking commitment, the methodology is already validated by experimental studies [FGQ$^+$18]. The next step would be actual participation in the frequency regulation service using this method. In case of coordinated interior point methods, the framework can be applied and tested in any multi-agent control and optimization framework, such as coordination of drones or chemical production plants.

A Appendix

A.1 Derivation of sensitivity equations

For brevity, we remove the superscript i from the variables. The gradient being equal to the Lagrange multiplier vector is a well known result from sensitivity theory and we refer the reader to [Ber08] for the proof. The Hessian expression can be obtained considering the first order optimality conditions,

$$\begin{aligned}
&\nabla_x L(x^*, \lambda^*) = 0 \\
&\nabla_\lambda L(x^*, \lambda^*) = 0 \\
&v^T \nabla^2_{xx} L(x^*, \lambda^*) v > 0, \quad \forall v \neq 0, C^T v = 0 \ .
\end{aligned}$$

Taking the optimality conditions as a function $r(y, x, \lambda) = 0$

$$r(y, x(y), \lambda(y)) = \begin{bmatrix} \nabla_u \phi(x) + C\lambda \\ Cx - y \end{bmatrix} = 0$$

we can observe that the gradient matrix of this function with respect to (u, λ) is given by:

$$\nabla_{(y,\lambda)} r = \begin{bmatrix} \nabla^2_{xx} L(x^*, \lambda^*) & C^T \\ C & 0 \end{bmatrix}$$

Since the Jacobian $\nabla_{(y,\lambda)} r^T$ is continuous and non-singular (due to $\nabla^2_{xx} L(x^*, \lambda^*)$ being non-singular and C full rank) Implicit Function Theorem [Ber08] holds. Therefore one can obtain the gradient of $(x(y), \lambda(y))$ with respect to y as

$$\nabla_y (x(y), \lambda(y)) = -\nabla_y r(y, x(y), \lambda(y)) (\nabla_{(x,\lambda)} r(y, x(y), \lambda(y)))^{-1}$$

Equivalently

$$\begin{bmatrix} \nabla_y x(y) & \nabla_y \lambda(y) \end{bmatrix} = \begin{bmatrix} 0 & \mathcal{I}_{n_y} \end{bmatrix} \begin{bmatrix} \nabla^2_{xx} L(x(y), \lambda(y)) & C^T \\ C & 0 \end{bmatrix}^{-1}$$

For ease of notation let $H = \nabla^2_{xx} L(x(y), \lambda(y))$. By using blockwise inversion

$$\begin{bmatrix} H & C^T \\ C & 0 \end{bmatrix}^{-1} = \begin{bmatrix} A & B \\ C & D \end{bmatrix}$$

where

$$\begin{aligned} A &:= H^{-1} + H^{-1}C^T(CH^{-1}C^T)^{-1}CH^{-1} \\ B &:= (CH^{-1}C^T)^{-1}CH^{-1} \\ C &:= H^{-1}C^T(CH^{-1}C^T)^{-1} \\ D &:= -(CH^{-1}C^T)^{-1} \end{aligned}$$

we arrive at

$$\begin{bmatrix} \nabla_y x(y) & \nabla_y \lambda(y) \end{bmatrix} = \begin{bmatrix} H^{-1}C^T(CH^{-1}C^T)^{-1} & -(CH^{-1}C^T)^{-1} \end{bmatrix}$$

leading finally to

$$-\nabla_y \lambda(y) = \nabla^2 \rho(y) = (C(\nabla^2_{xx} L(x(y), \lambda(y))^{-1} C^T)^{-1} .$$

A.2 Ellipsoid projection

We take an ellipsoid in $\mathbb{R}^{n_x}$

$$\mathcal{D} = \{x \in \mathbb{R}^{n_x} | x^T Q x \leq 1\} \tag{A.1}$$

with $Q \succ 0$.

The image of this ellipsoid under a non-injective linear map $C \in \mathbb{R}^{n_y \times n_x}$, $n_y < n_x$ is given by

$$C\mathcal{D} = \{y \in \mathbb{R}^{n_y} \mid \exists x \in \mathbb{R}^{n_u} \ : \ y = Cx, \ x^T Q x \leq 1\} \tag{A.2}$$

or equivalently we can write

$$C\mathcal{D} = \{y \in \mathbb{R}^{n_y} \mid \{\min x^T Q x \,|Cx = y\} \leq 1\} \tag{A.3}$$

The optimality conditions for the quadratic minimization problem reads as

$$\begin{bmatrix} Q & C^T \\ C & \end{bmatrix} \begin{bmatrix} x \\ \lambda \end{bmatrix} = \begin{bmatrix} 0 \\ y \end{bmatrix} \tag{A.4}$$

By block elimination we arrive at

$$\begin{aligned} x &= -Q^{-1}C^T\lambda \\ \lambda &= -(CQ^{-1}C^T)^{-1}y \\ x^T Q u &= \lambda^T CQ^{-1}QQ^{-1}C^T\lambda \\ &= \lambda^T CQ^{-1}C^T\lambda \\ &= y^T(CQ^{-1}C^T)^{-1}CQ^{-1}C^T(CQ^{-1}C^T)^{-1}y \\ &= y^T(CQ^{-1}C^T)^{-1}y \end{aligned}$$

Therefore the projected ellipsoid can be described with a direct quadratic function over the lower

dimensional variable y

$$C\mathcal{D} = \{x \in \mathbb{R}^{n_u} | y^T (CQ^{-1}C^T)^{-1} y \leq 1\} \tag{A.5}$$

A.3 Convergence of PDIP with non-uniform complementarity perturbation

The analysis and the proof of Theorem 2 will be similar to that of [Gon13] by showing that starting from a point within $\mathcal{N}_s$, it is possible to keep the linearization error small and obtain an admissible step-size of complexity $\mathcal{O}(\frac{1}{M})$ that reduces the duality gap sufficiently and keep the iterates within $\mathcal{N}_s$. Allowing non-uniformity necessitates additional constraints (3.78), (3.77) on the perturbation vector and results in different bounding terms compared to [Gon13], however the complexity result is identical.

We start the analysis by deriving bounds on the linearization error.

Lemma 11. *For $(x, y, s, \lambda) \in \mathcal{N}_s$, the Newton direction resulting from* (3.76) *satisfies the following relations;*

$$0 \leq \Delta\lambda^T \Delta s \leq \theta \lambda^T s \tag{A.6}$$

$$-\theta\lambda^T s \leq \Delta\lambda_j \Delta s_j \leq \sqrt{2}\theta\lambda_j s_j \tag{A.7}$$

where θ is defined as

$$\theta := 2^{-3/2} max\,(\frac{c}{\phi\psi} - 1, 1 - \phi\psi c). \tag{A.8}$$

Proof. Using the Newton step equations (3.24) in the feasible setting, it can be shown that

$$\Delta\lambda^T \Delta s = \Delta x^T H(x, \lambda) \Delta x$$

Since f and g are convex functions and $\lambda \geq 0$, we have $H(x, \lambda) \succeq 0$, which validates the left-hand side inequality in (A.6).

For the right-hand side of (A.6) we need a technical result ([Gon12], Lemma 3.3.). Given two vectors $u, v \in \mathbb{R}^n$, if $u^T v \geq 0$ then,

$$\|UV\mathbf{1}\| \leq 2^{-3/2} \|u + v\|^2 \tag{A.9}$$

where $U := \text{diag}(u)$, $V := \text{diag}(v)$. With relation (A.9) at hand, we start with the non-uniformly perturbed centrality equation

$$S\Delta\lambda + \Lambda\Delta s = \Lambda S(\xi - \mathbf{1}). \tag{A.10}$$

Multiply both sides with $(\Lambda S)^{-1/2}$;

$$\Lambda^{-1/2} S^{1/2} \Delta\lambda + \Lambda^{1/2} S^{-1/2} \Delta s = (\Lambda S)^{1/2} (\xi - \mathbf{1}). \tag{A.11}$$

Since $\Delta\Lambda, \Lambda, S, \Delta S$ are all diagonal matrices we have that

$$\|\Delta\Lambda\Delta S\mathbf{1}\| = \|\Lambda^{-1/2} S^{1/2} \Delta\lambda \Lambda^{1/2} S^{-1/2} \Delta s\|$$

Next we use (A.9) with $u := \Lambda^{-1/2} S^{1/2} \Delta\lambda$ and $v := \Lambda^{1/2} S^{-1/2} \Delta s$. Note that the condition $u^T v \geq 0$ is

satisfied since $\Delta\lambda^T\Delta s \geq 0$.

$$\begin{aligned}\|\Delta\Lambda\Delta S\mathbf{1}\| &= \|\Lambda^{-1/2}S^{1/2}\Delta\lambda\Lambda^{1/2}S^{-1/2}\Delta s\| \\ &\leq 2^{-3/2}\|\Lambda^{-1/2}S^{1/2}\Delta\lambda + \Lambda^{1/2}S^{-1/2}\Delta s\|^2 \\ &= 2^{-3/2}\|(\Lambda S)^{1/2}(\xi - \mathbf{1})\|^2 \\ &= 2^{-3/2}\sum_j(\lambda_j s_j)(\xi_j - 1)^2 \\ &\leq 2^{-3/2}\|\xi - \mathbf{1}\|_\infty^2\lambda^T s \qquad \text{due to}(\lambda_j s_j > 0,\ \forall j)\end{aligned}$$

Relying on the constraints (3.78), (3.77) imposed on the non-uniform perturbation vector ξ, we can obtain a parametric expression for $\|\xi - \mathbf{1}\|_\infty$

$$\|\xi - \mathbf{1}\|_\infty^2 = \max\left(\frac{c}{\phi\psi} - 1, 1 - \phi\psi c\right)$$

which is equal to $\theta/(2^{-3/2})$, by definition (A.8). Therefore, we can write the bound on $\|\Delta\Lambda\Delta S\mathbf{1}\|$ as

$$\|\Delta\Lambda\Delta S\mathbf{1}\| \leq \theta\lambda^T s$$

Finally using Cauchy-Schwarz inequality we arrive at the upper bound on $\Delta\lambda^T\Delta s$;

$$\Delta\lambda^T\Delta s = (\Delta\Lambda\Delta S\mathbf{1})^T\mathbf{1} \leq \|\Delta\Lambda\Delta S\mathbf{1}\|\|\mathbf{1}\| = \theta\lambda^T s$$

which proves the right-hand side of (A.6).

Next, we prove the relation (A.7) that provide bounds on $\Delta\lambda_j\Delta s_j$. We write equation (A.11) element-wise and square both sides

$$\begin{aligned}&\lambda_j^{-1/2}s_j^{1/2}\Delta\lambda_j + \lambda_j^{1/2}s_j^{-1/2}\Delta s_j = \lambda_j^{1/2}s_j^{1/2}(\xi_j - 1) \\ \implies\ &\lambda_j^{-1}s_j\Delta\lambda_j^2 + \lambda_j s_j^{-1}\Delta s_j^2 + 2\Delta\lambda_j\Delta s_j = (\xi_j - 1)^2\lambda_j s_j\end{aligned}$$

Since $\lambda, s > 0$, we have that

$$\begin{aligned}&2\Delta\lambda_j\Delta s_j \leq (\xi_j - 1)^2\lambda_j s_j \\ \Longrightarrow&\Delta\lambda_j\Delta s_j \leq 0.5\|(\xi_j - \mathbf{1})\|_\infty^2\lambda_j s_j\end{aligned}$$

Using the relation $|\Delta\lambda_j\Delta s_j| \leq \|\Delta\Lambda\Delta S\mathbf{1}\|_\infty \leq \|\Delta\Lambda\Delta S\mathbf{1}\|$, and the definition of θ, we arrive at

$$-\theta\lambda^T s \leq \Delta\lambda_j\Delta s_j \leq \sqrt{2}\theta\lambda_j s_j$$

□

In the following, we will show that a finite step-size can be obtained such that the neighborhood constraints are preserved and a sufficient decrease in the duality gap is achieved.

For a given direction, we describe the values after a step with a step-size $t \in (0, 1]$ as

$$\lambda_j(t) = \lambda_j + t\Delta\lambda_j, \quad s_j(t) = s_j + t\Delta s_j \ .$$

Lemma 12. *For a given point $(x, y, s, \lambda) \in \mathcal{N}_s$, and the Newton direction resulting from (3.76), if the step-size $t \in (0, 1]$ satisfies the following conditions;*

$$t \leq \frac{c(\phi - \psi)}{(\psi + M)\theta} \tag{A.12}$$

$$t \leq \frac{c(\phi - \psi)}{\sqrt{2}\theta\phi} \tag{A.13}$$

then the new point attained by taking a step in the Newton direction with step-size t stays in the neighborhood $\mathcal{N}_s$, that is

$$(x(t), y(t), s(t), \lambda(t)) \in \mathcal{N}_s. \tag{A.14}$$

Proof. Symmetric neighborhood $\mathcal{N}_s$ implies lower and upper bound constraints on complementarity pairs and feasibility of the iterates. We start with the lower bound constraint. Using Lemma 11, the centrality equation (A.10) and neighborhood constraints on the current point $\lambda_j s_j \geq \psi\hat{\eta}$ and the perturbation vector $\xi_j \geq \phi\frac{c\hat{\eta}}{\lambda_j s_j}$ we obtain an initial lower bound on $\lambda_j(t)s_j(t)$

$$\begin{aligned}
\lambda_j(t)s_j(t) &= (\lambda_j + t\Delta\lambda_j)(s_j + t\Delta s_j) \\
&= \lambda_j s_j(1 - t) + t\xi_j\lambda_j s_j + t^2\Delta\lambda_j\Delta s_j \\
&\geq \psi\hat{\eta}(1 - t) + t\phi c\hat{\eta} - t^2\theta M\hat{\eta}
\end{aligned}$$

Next, we derive an upper bound on $\hat{\eta}(t)$, using the relation (3.77), centrality equation (A.10) and Lemma 11

$$\begin{aligned}
\hat{\eta}(t) &= \lambda(t)^T s(t)/M \\
&= \left(\lambda^T s + t(c - 1)\lambda^T s + t^2\Delta\lambda^T\Delta s\right)/M \\
&\leq \left(\lambda^T s + t(c - 1)\lambda^T s + t^2\theta\lambda^T s\right)/M \\
&= \hat{\eta}\left((1 - t) + tc + t^2\theta\right)
\end{aligned} \tag{A.15}$$

Using the upper bound on $\hat{\eta}(t)$ and lower bound on $\lambda_j(t)s_j(t)$ we can write a sufficient condition for $\lambda_j(t)s_j(t) \geq \psi\hat{\eta}(t)$,

$$\begin{aligned}
\psi\hat{\eta}(1 - t) + t\phi c\hat{\eta} - t^2\theta M\hat{\eta} &\geq \psi\hat{\eta}(1 - t) + t\psi c\hat{\eta} + \psi t^2\theta\hat{\eta} \\
(\phi - \psi)c\hat{\eta} &\geq t\left((\psi + M)\theta\right)\hat{\eta}
\end{aligned}$$

which is identical to condition (A.12).

Next, we tackle the upper bound constraint on $\lambda_j(t)s_j(t) \leq \frac{1}{\psi}\hat{\eta}(t)$ which is equivalent to

$$\begin{aligned}
\lambda_j s_j(1 - t) + t\xi_j\lambda_j s_j + t^2\Delta\lambda_j\Delta s_j \leq \\
\frac{1}{\psi}\left(\hat{\eta}(1 - t) + tc\hat{\eta} + t^2\Delta\lambda^T\Delta s/M\right)
\end{aligned}$$

Seeing that $\lambda_j s_j \leq \frac{1}{\psi}\hat{\eta}$, and $\xi_j \leq \frac{1}{\phi}\frac{c\hat{\eta}}{\lambda_i s_i}$ a sufficient condition is

$$\frac{1}{\psi}\hat{\eta}(1-t) + t\frac{1}{\phi}c\hat{\eta} + t^2\Delta\lambda_j\Delta s_j \leq \frac{1}{\psi}\hat{\eta}(1-t) + t\frac{1}{\psi}c\hat{\eta} + t^2\frac{1}{\psi}\Delta\lambda^T\Delta s/M$$

$$t\frac{1}{\phi}c\hat{\eta} + t^2\Delta\lambda_j\Delta s_j \leq t\frac{1}{\psi}c\hat{\eta} + t^2\frac{1}{\psi}\Delta\lambda^T\Delta s/M$$

$$t^2\left(\Delta\lambda_j\Delta s_j - \frac{1}{\psi}\Delta\lambda^T\Delta s/n\right) \leq tc\hat{\eta}(\frac{1}{\psi} - \frac{1}{\phi})\,.$$

Since $\Delta\lambda_j\Delta s_j \leq \sqrt{2}\theta\lambda_j s_j \leq \sqrt{2}\theta\frac{1}{\psi}\hat{\eta}$ and $\Delta\lambda^T\Delta s \geq 0$ another sufficient condition is

$$t\sqrt{2}\theta\frac{1}{\psi}\hat{\eta} \leq c\hat{\eta}(\frac{1}{\psi} - \frac{1}{\phi})$$
$$t \leq \frac{\psi c(\frac{1}{\psi} - \frac{1}{\phi})}{\sqrt{2}\theta} = \frac{c(\phi-\psi)}{\sqrt{2}\theta\phi}\,.$$

which is equivalent to (A.13).

For a LP and QP, the KKT system is composed of linear equations, except for the complementarity equation. Therefore the Newton step with $r_{\text{dual}} = 0, r_{\text{primal}} = 0$, will preserve the feasibility for any step-size, from which we conclude that a step-size $t \in (0,1]$ satisfying (A.12), (A.13) will keep the iterates within $\mathcal{N}_s$. □

Lemma 13. *For a given point $(x,y,s,\lambda) \in \mathcal{N}_s$, parameter $\alpha \in (0,1)$ and the Newton direction resulting from* (3.76)*, if the step-size t, and the mean reduction parameter c, satisfy the conditions;*

$$1 - c > \alpha, \quad 0 < t \leq \frac{1-c-\alpha}{\theta} \tag{A.16}$$

then the new duality gap η(t) satisfies the reduction requirement

$$\eta(t) \leq (1-\alpha t)\,\eta\,. \tag{A.17}$$

Proof. From (A.15), we have that

$$\eta(t) \leq (1-\gamma(t))\,\eta$$

with $\gamma(t) = t(1-c) - t^2\theta$. Therefore, a sufficient condition for satisfying (A.17) is

$$\gamma(t) = t(1-c) - t^2\theta \geq \alpha t$$
$$\Longrightarrow t^2\theta \leq t(1-c-\alpha) \Longrightarrow\ 0 < t \leq \frac{1-c-\alpha}{\theta}$$

which is equivalent to (A.16). □

In order to obtain the convergence and complexity result, first we will show that there exists a $t \sim \mathcal{O}(\frac{1}{M})$ that satisfies (A.13), (A.12) and (A.16). Then, application of Thorem 3.2 from [Wri97] will provide the

desired result.

The set of admissible step-sizes are given by;

$$\{t \mid 0 < t \leq \min\left(\frac{1-c-\alpha}{\theta}, \frac{c(\phi-\psi)}{(\psi+M)\theta}, \frac{c(\phi-\psi)}{\sqrt{2}\theta\phi}\right)\} \tag{A.18}$$

where $\theta := 2^{-3/2}\max\left(\frac{c}{\phi\psi} - 1, 1 - \phi\psi c\right)$.

Choose $\alpha = 0.1$, $\psi = 0.1$ and $\phi = 0.2$ and fix reduction parameter $c = 0.8$ for every step. For these values θ is given by $2^{-3/2}(\frac{0.8}{0.2*0.1} - 1) \simeq 13.8$ It can be shown that the step-size selection $\frac{1}{200M}$ will be always admissible, which provides $t \sim \mathcal{O}(\frac{1}{M})$.

Proof of Theorem 2

Proof. Given the preceding analysis, proof is a straightforward application of Theorem 3.2 from [Wri97]. We have already shown that for a given parameter set, it is possible to find an admissible $t \sim \mathcal{O}(\frac{1}{M})$. Plugging this into the reduction requirement, we get

$$\eta_{k+1} \leq (1 - \frac{\delta}{M})\eta_k, \quad k = 1, 2, \ldots, \tag{A.19}$$

for some constant $\delta > 0$. From here, we can directly apply Thorem 3.2 from [Wri97] and state the result. We provide the proof for completeness. First we take the log of both sides in (A.19)

$$\log \eta_{k+1} \leq \log(1 - \frac{\delta}{M}) + \log \eta_k \ .$$

Writing the right hand side recursively and using (3.79) we get

$$\log \eta_{k+1} \leq k \log(1 - \frac{\delta}{M}) + \log \eta_0 \leq k \log(1 - \frac{\delta}{M}) + \kappa \log(\frac{1}{\epsilon}) \ .$$

The log function has the property that

$$\log(1 + x) \leq x, \quad \forall x > -1 \ .$$

Since $\alpha, t < 1$, this leads to

$$\log \eta_{k+1} \leq k(-\frac{\delta}{M}) + \kappa \log(\frac{1}{\epsilon}) \ .$$

For the convergence criterion $\eta_k < \epsilon$ to be satisfied, it is sufficient to have

$$k(-\frac{\delta}{M}) + \kappa \log(\frac{1}{\epsilon}) \leq \log \epsilon$$

which is satisfied for all k, such that

$$k \geq K = (1 + \kappa)\frac{M}{\delta} \log \frac{1}{\epsilon} \ .$$

□

A.4 Reachability of steady states

In this section, we will show that the reachability Assumption 4 in Chapter 4 holds for linear systems if the admissible equilibrium points are strictly feasible. The analysis can be extended to nonlinear systems by relying on local controllability of admissible equilibrium points and connectivity of the set $\mathcal{E}$, therefore Assumption 4 is 'mild' in general.

Consider a controllable linear system described by

$$x^+ = Ax + Bu$$

for which the state and input are constrained, $x \in \mathbb{X} \subset \mathbb{R}^{n_x}, u \in \mathbb{U} \subset \mathbb{R}^{n_u}$, where $\mathbb{X}, \mathbb{U}$ are convex, compact sets with non-empty interior. Consider the following set that includes strictly feasible steady states

$$\mathcal{E} := \{(x, u) \in \lambda(\mathbb{X} \times \mathbb{U}),\ x = Ax + Bu\} \tag{A.20}$$

for some $\lambda \in [0, 1)$. We further define $\mathcal{E}_x$ and $\mathcal{E}_u$ which are projections of $\mathcal{E}$ onto x and u spaces respectively, The following lemmas provide results related to reachability of strictly feasible steady states.

Lemma 14. *For any $x_s \in \mathcal{E}_x$, any other state in $\mathcal{E}_x$ is reachable in finite steps, i.e. $\forall (y, z) \in \mathcal{E}_x$ there exists $N < \infty$ such that $z \in \mathcal{R}_N(y)$.*

Proof. For any $x_s \in \mathcal{E}_x$, without loss of generality, we can assume $x_s = 0, u_s = 0$. Let $\hat{N} \geq n_x$. Choose $\epsilon_x > 0$ such that $\mathcal{B}(0, \epsilon_x) \subset \mathbb{X}$. For all $\tilde{x} \in \mathcal{B}(0, \epsilon_x)$, there exists $\boldsymbol{u}$ such that $\tilde{x} = \sum_{i=0}^{\hat{N}-1} A^i B u(i)$ and some $\epsilon_u > 0$ such that $u(i) \in \mathcal{B}(0, \epsilon_u)$ for some $\epsilon_u > 0$. Due to linearity, we can scale the state with some positive $\delta \leq 1$, such that $\delta x = \sum_{i=0}^{\hat{N}-1} A^i B \delta u(i)$ and $\mathcal{B}(0, \delta\epsilon_u) \subset \mathbb{U}$. Therefore any $x \in \mathcal{B}(x_s, \delta\epsilon_x)$ is reachable in $\hat{N}$ steps from x_s and therefore $\mathcal{B}(x_s, \delta\epsilon_x) \subseteq \mathcal{R}_{\hat{N}}(x_s)$. Since $\mathcal{E}$ is a compact subset of $\mathbb{X} \times \mathbb{U}$ there exists a ball of minimum size $\mathcal{B}(0, \gamma)$ with some $\gamma > 0$ such that $\mathcal{B}(x_s, \gamma) \subseteq \mathcal{B}(x_s, \delta\epsilon_x) \subseteq \mathcal{R}_{\hat{N}}(x_s) \subseteq \mathbb{X}$ for all $x_s \in \mathcal{E}$. Since the set $\mathcal{E}$ is connected, for any $(y, z) \in \mathcal{E} \times \mathcal{E}$, $z \in \mathcal{R}_N(y)$ with $N > \hat{N}\|y - z\|/\gamma$. Since $\mathcal{E}$ is also bounded, it follows that N can be chosen finite. □

The following lemma states that for constrained linear systems, a strictly feasible equilibrium point becomes reachable within a small enough neighborhood of that point.

Lemma 15. *For any $N \geq n_x$, there exists a $\kappa > 0$ such that for any $x_s \in \mathcal{E}_x$, $\mathcal{B}(x_s, \kappa) \subset \mathcal{C}_N(x_s)$,i.e. any equilibrium point in $\mathcal{E}$ becomes reachable in finite steps from at least κ distance.*

Proof. Similar to Lemma 14, for any $x_s \in \mathcal{E}_x$, let $x_s = 0, u_s = 0$. Choose $\epsilon_x > 0$ such that $\mathcal{B}(0, \epsilon_x) \subset \mathbb{X}$. For all $x_0 \in \mathcal{B}(0, \epsilon_x)$, there exists $\boldsymbol{u}$ such that $0 = \sum_{i=0}^{N-1} A^i B u(i) + A^i x_0$ and some $\epsilon_u > 0$ such that $u(i) \in \mathcal{B}(0, \epsilon_u)$. Following the same argument in Lemma 14, one can scale $\mathcal{B}(0, \epsilon_x)$ by some positive $\delta \leq 1$ such that $u(i) \in \mathcal{B}(0, \delta\epsilon_u) \subset \mathbb{U}$ and there exists a ball of minimum size $\mathcal{B}(0, \kappa)$ with some $\kappa > 0$ such that $\mathcal{B}(x_s, \kappa) \subseteq \mathcal{B}(x_s, \delta\epsilon_x) \subseteq \mathcal{C}_N(x_s) \subseteq \mathbb{X}$ for all $x_s \in \mathcal{E}$. Therefore $\forall x_s \in \mathcal{E}$, any state within κ distance to x_s is steerable to x_s without violating any constraints, providing $\mathcal{B}(x_s, \kappa) \subseteq \mathcal{C}_N(x_s)$. □

Lemma 16. *Assumption 4 in Chapter 4 holds for convexly constrained, controllable linear systems if $\mathcal{E}$ is limited to steady states that are bounded away form the constraints.*

Proof. Given $\bar{\epsilon} > 0$ and some $\bar{x} \in \mathcal{E}_x$, define the set $T(\bar{x}, \bar{\epsilon}) := \{x \in \mathcal{E}_x \,|l(\tilde{x}, \tilde{u}) \leq \max(l(z^*, v^*), l(\bar{x}, \bar{u}) - \bar{\epsilon})\}$. the set T is non-empty by definition since it always includes z^*. Since $\mathcal{E}_x$ is a bounded set, the distance to $T(\bar{x}, \bar{\epsilon})$ from $\bar{x}$ is bounded from above. From Lemma 14 it follows that there exists a minimum horizon length N_1 such that $\tilde{x} \in \mathcal{R}_{N_1}(\bar{x})$, $\forall \bar{x} \in \mathcal{E}$ and $\forall \tilde{x} \in T(\bar{x}, \bar{\epsilon})$. Furthermore, from Lemma 15 it follows that one can choose some N_2 such that $\bar{x} \in \mathcal{R}_{N_2}(x)$, $\forall x \in \mathcal{B}(\bar{x}, \bar{\delta})$. Combining the horizon lengths and reachable sets we have that $\tilde{x} \in \mathcal{R}_N(x)$, $\forall x \in \mathcal{B}(\bar{x}, \bar{\delta})$ for any $N \geq (N_1 + N_2)$. □

Figure A.1illustrates the proof of Lemma 16.

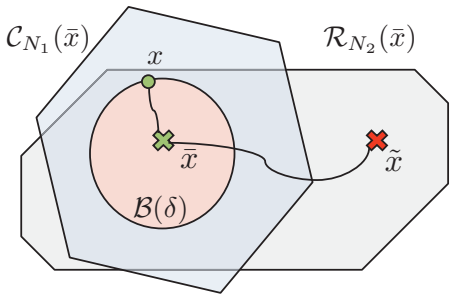

Figure A.1 – Schematic representation of reachability of steady state $\tilde{x}$ from any state $x \in \mathcal{B}(\bar{x}, \delta)$ with a horizon length $N_1 + N_2$.

Bibliography

[AAJ16] A. Alessandretti, A.P. Aguiar, and C.N. Jones. On convergence and performance certification of a continuous-time economic model predictive control scheme with time-varying performance index. *Automatica*, 68:305–313, 2016.

[AAR12] D. Angeli, R. Amrit, and J.B. Rawlings. On Average Performance and Stability of Economic Model Predictive Control. *IEEE Transactions on Automatic Control*, 57(7):1615–1626, July 2012.

[APHW14] M. Annergren, S.K. Pakazad, A. Hansson, and B. Wahlberg. A distributed primal-dual interior-point method for loosely coupled problems using ADMM. *arXiv:1406.2192 [math]*, June 2014. arXiv: 1406.2192.

[Bak08] L. Bakule. Decentralized control: An overview. *Annual Reviews in Control*, 32(1):87–98, April 2008.

[BCC+10] H.Y. Benson, J.J. Cochran, L.A. Cox, P. Keskinocak, J.P. Kharoufeh, and J.C. Smith. Interior-Point Linear Programming Solvers. In *Wiley Encyclopedia of Operations Research and Management Science*. John Wiley & Sons, Inc., 2010.

[Ber08] D. Bertsekas. *Nonlinear Optimization*. Athena Scientific, 2008.

[BGJ17] A. Bitlislioğlu, T.T. Gorecki, and C.N. Jones. Robust Tracking Commitment. *IEEE Transactions on Automatic Control*, 2017.

[BGK14] D. Bertsimas, V. Gupta, and N. Kallus. Data-driven robust optimization. *ArXiv e-prints*, December 2014.

[BJ17a] A. Bitlislioğlu and C.N. Jones. On Coordinated Primal-Dual Interior-Point Methods for Multi-Agent Optimization. In *56th IEEE Conference on Decision and Control*, 2017.

[BJ17b] A. Bitlislioğlu and C.N. Jones. On generalized primal-dual interior-point methods with non-uniform complementarity perturbations for quadratic programming. Technical Report 231144, École Polytechnique Fédérale de Lausanne, September 2017.

[BJ18] A. Bitlislioğlu and C.N. Jones. Regularized economic model predictive control with barrier functions. Technical report, École Polytechnique Fédérale de Lausanne, June 2018.

[BLR+10] A Brooks, E. Lu, D. Reicher, C. Spirakis, and B. Weihl. Demand dispatch. *IEEE Power and Energy Magazine*, 8(3):20–29, May 2010.

[BM08] F. Blanchini and S. Miani. *Set-Theoretic Methods in Control*. Springer, New York, 2008.

[BOCT14] M. Balandat, F. Oldewurtel, Mo Chen, and C. Tomlin. Contract design for frequency regulation by aggregations of commercial buildings. In *2014 52nd Annual Allerton Conf. on Communication, Control, and Computing (Allerton)*, pages 38–45, September 2014.

[Boy10] S. Boyd. Distributed Optimization and Statistical Learning via the Alternating Direction Method of Multipliers. *Foundations and Trends in Machine Learning*, 3(1):1–122, 2010.

[BPJ17a] A. Bitlislioğlu, I. Pejcic, and C.N. Jones. Interior point decomposition for multi-agent optimization. In *IFAC 2017 World Congress, Toulouse, France*, July 2017.

[BPJ17b] A. Bitlislioğlu, I. Pejcic, and C.N. Jones. Interior Point Decomposition for Multi-Agent Optimization. In *20th IFAC World Congress*, 2017.

[BT97] D. Bertsekas and J.N. Tsitsiklis. *Parallel and Distributed Computation: Numerical Methods*. Athena Scientific, Belmont, Massachusetts, 1997.

[BTEGN09] A. Ben-Tal, L. El Ghaoui, and A. Nemirovski. *Robust Optimization*. Princeton University Press, 2009.

[BV04] S. Boyd and L. Vandenberghe. *Convex Optimization*. Cambridge University Press, Cambridge, 2004.

[BZ16] S. Bolognani and S. Zampieri. On the Existence and Linear Approximation of the Power Flow Solution in Power Distribution Networks. *IEEE Transactions on Power Systems*, 31(1):163–172, January 2016.

[Car05] C. Cartis. On the convergence of a primal-dual second-order corrector interior point algorithm for linear programming. Technical Report 05/04, Numerical Analysis Group, Computing Laboratory, Oxford University, March 2005.

[CG08] M. Colombo and J. Gondzio. Further development of multiple centrality correctors for interior point methods. *Computational Optimization and Applications*, 41(3):277–305, December 2008.

[CGP09] M.C. Campi, S. Garatti, and M. Prandini. The scenario approach for systems and control design. *Annual Reviews in Control*, 33(2):149–157, December 2009.

[CM11] M. Chen and S. Mehrotra. Self-concordance and Decomposition-based Interior Point Methods for the Two-stage Stochastic Convex Optimization Problem. *SIAM Journal on Optimization*, 21(4):1667–1687, October 2011.

[CS11] E. Camponogara and H.F. Scherer. Distributed Optimization for Model Predictive Control of Linear Dynamic Networks With Control-Input and Output Constraints. *IEEE Transactions on Automation Science and Engineering*, 8(1):233–242, January 2011.

[DAR11] M. Diehl, R. Amrit, and J.B. Rawlings. A Lyapunov Function for Economic Optimizing Model Predictive Control. *IEEE Transactions on Automatic Control*, 56(3):703–707, March 2011.

[DCB13a] S. Di Cairano and F. Borrelli. Constrained tracking with guaranteed error bounds. In *2013 IEEE 52nd Annual Conf. on Decision and Control (CDC)*, pages 3800–3805, December 2013.

[DCB13b] A. Domahidi, E. Chu, and S. Boyd. ECOS: An SOCP solver for embedded systems. In *2013 European Control Conference (ECC)*, pages 3071–3076, July 2013.

[DFH09] M. Diehl, H. J. Ferreau, and N. Haverbeke. Efficient Numerical Methods for Nonlinear MPC and Moving Horizon Estimation. In L. Magni, D. Martino Raimondo, and F. Allgöwer, editors, *Nonlinear Model Predictive Control*, number 384 in Lecture Notes in Control and Information Sciences, pages 391–417. Springer Berlin Heidelberg, 2009.

[DHSS12] M. Diehl, B. Houska, O. Stein, and P. Steuermann. A lifting method for generalized semi-infinite programs based on lower level Wolfe duality. *Computational Optimization and Applications*, 54(1):189–210, June 2012.

[DKDS11] M.D. Doan, T. Keviczky, and B. De Schutter. A dual decomposition-based optimization method with guaranteed primal feasibility for hierarchical MPC problems. *IFAC Proceedings Volumes*, 44(1):392–397, January 2011.

[eur10] Communication from the commission to the european parliament, the council, the european economic and social committee and the committee of the regions energy 2020 a strategy for competitive, sustainable and secure energy, 2010.

[FBGS14] M. Farina, G. Betti, L. Giulioni, and R. Scattolini. An Approach to Distributed Predictive Control for Tracking #x2013;Theory and Applications. *IEEE Transactions on Control Systems Technology*, 22(4):1558–1566, July 2014.

[FE15] C. Feller and C. Ebenbauer. Weight recentered barrier functions and smooth polytopic terminal set formulations for linear model predictive control. In *2015 American Control Conference (ACC)*, pages 1647–1652, July 2015.

[FE16a] C. Feller and C. Ebenbauer. Relaxed Logarithmic Barrier Function Based Model Predictive Control of Linear Systems. *IEEE Transactions on Automatic Control*, PP(99):1–1, 2016.

[FE16b] C. Feller and C. Ebenbauer. A stabilizing iteration scheme for model predictive control based on relaxed barrier functions. *arXiv:1603.04605 [math]*, March 2016. arXiv: 1603.04605.

[FGM18] T. Faulwasser., L. Grüne., and M. A. Müller. Economic nonlinear model predictive control. *Foundations and Trends® in Systems and Control*, 5(1):1–98, 2018.

[FGQ+18] L. Fabietti, T.T. Gorecki, F.A. Qureshi, A. Bitlislioğlu, I. Lymperopoulos, and C.N. Jones. Experimental Implementation of Frequency Regulation Services Using Commercial Buildings. *IEEE Transactions on Smart Grid*, 9(3):1657–1666, May 2018.

[FM13] P. Falugi and D.Q. Mayne. Model predictive control for tracking random references. In *European Control Conf. (ECC), 2013 European*, pages 518–523, July 2013.

[Fra77] B. Francis. The Linear Multivariable Regulator Problem. *SIAM Journal on Control and Optimization*, 15(3):486–505, May 1977.

[FRLC10] A. Ferramosca, J.B. Rawlings, D. Limón, and E.F. Camacho. Economic MPC for a changing economic criterion. In *Decision and Control (CDC), 2010 49th IEEE Conference on*, pages 6131–6136. IEEE, 2010.

[FT13] L. Fagiano and A.R. Teel. Generalized terminal state constraint for model predictive control. *Automatica*, 49(9):2622–2631, September 2013.

[GBSJ15] T.T. Gorecki, A. Bitlislioğlu, G. Stathopoulos, and C.N. Jones. Guaranteeing input tracking for constrained systems: theory and application to demand response. In *American Control Conference (ACC), 2015*, pages 232–237. IEEE, 2015.

[GKM06] P.J. Goulart, E.C. Kerrigan, and J.M. Maciejowski. Optimization over state feedback policies for robust control with constraints. *Automatica*, 42(4):523–533, April 2006.

[Gon96] J. Gondzio. Multiple centrality corrections in a primal-dual method for linear programming. *Computational Optimization and Applications*, 6(2):137–156, September 1996.

[Gon12] J. Gondzio. Interior point methods 25 years later. *European Journal of Operational Research*, 218(3):587–601, May 2012.

[Gon13] J. Gondzio. Convergence Analysis of an Inexact Feasible Interior Point Method for Convex Quadratic Programming. *SIAM Journal on Optimization*, 23(3):1510–1527, January 2013.

[Gor17] T.T. Gorecki. *Predictive Control methods for Building Control and Demand Response.* PhD thesis, EPFL, Lausanne, 2017.

[GQJ15a] T. T. Gorecki, F.A. Qureshi, and C.N. Jones. Openbuild : An integrated simulation environment for building control. In *Control Applications (CCA), 2015 IEEE Conference on*, pages 1522–1527, 2015.

[GQJ15b] T.T. Gorecki, F.A. Qureshi, and C.N. Jones. OpenBuild : An integrated simulation environment for building control. In *2015 IEEE Conference on Control Applications (CCA)*, pages 1522–1527, September 2015.

[Gr3] L. Grüne. Economic receding horizon control without terminal constraints. *Automatica*, 49(3):725–734, 2013.

[GVRSS08] F. Guerra Vázquez, J.-J. Rückmann, O. Stein, and G. Still. Generalized semi-infinite programming: A tutorial. *Journal of Computational and Applied Mathematics*, 217(2):394–419, August 2008.

[GWK14] A. Georghiou, W. Wiesemann, and D. Kuhn. Generalized decision rule approximations for stochastic programming via liftings. *Mathematical Programming*, pages 1–38, May 2014.

[GYdH15] B.L. Gorissen, İ. Yanikoğlu, and D. den Hertog. A practical guide to robust optimization. *Omega*, 53:124–137, June 2015.

[HB16] S.M. Harwood and P.I. Barton. Lower level duality and the global solution of generalized semi-infinite programs. *Optimization*, 65(6):1129–1149, June 2016.

[HSPV14] H. Hao, B.M. Sanandaji, K. Poolla, and T.L. Vincent. Aggregate Flexibility of Thermostatically Controlled Loads. *IEEE Transactions on Power Systems*, Early Access Online, 2014.

[JK12] D. Jahn and S. Korolczuk. German exceptionalism: the end of nuclear energy in Germany! *Environmental Politics*, 21(1):159–164, February 2012.

[Jon16] C.N. Jones. *Lectures on Model Predictive Control.* EFPL, 2016.

[KG14] E. Klintberg and S. Gros. A primal-dual Newton method for distributed Quadratic Programming. In *53rd IEEE Conference on Decision and Control*, pages 5843–5848, December 2014.

[KMA16] P.N. Köhler, M.A. Müller, and F. Allgöwer. A distributed economic MPC scheme for coordination of self-interested systems. In *2016 American Control Conference (ACC)*, pages 889–894, July 2016.

[LAAC08] D. Limon, I. Alvarado, T. Alamo, and E.F. Camacho. MPC for tracking piecewise constant references for constrained linear systems. *Automatica*, 44(9):2382–2387, September 2008.

[LAAC10] D. Limon, I. Alvarado, T. Alamo, and E.F. Camacho. Robust tube-based MPC for tracking of constrained linear systems with additive disturbances. *Journal of Process Control*, 20(3):248–260, March 2010.

[LQN+15] I. Lymperopoulos, F. A. Qureshi, T. Nghiem, A. A. Khatir, and C. N. Jones. Providing ancillary service with commercial buildings: the swiss perspective. *IFAC-PapersOnLine*, 48(8):6 – 13, 2015. 9th (IFAC) Symposium on Advanced Control of Chemical Processes (ADCHEM), June 2015.

[LS15] A. Lorca and X.A. Sun. Adaptive Robust Optimization With Dynamic Uncertainty Sets for Multi-Period Economic Dispatch Under Significant Wind. *IEEE Transactions on Power Systems*, 30(4):1702–1713, July 2015.

[MA14] M.A. Müller and F. Allgöwer. Distributed economic MPC: a framework for cooperative control problems. *IFAC Proceedings Volumes*, 47(3):1029–1034, January 2014.

[MAA13] M.A. Müller, D. Angeli, and F. Allgöwer. Economic model predictive control with self-tuning terminal cost. *European Journal of Control*, 19(5):408–416, September 2013.

[MAA14] M.A. Müller, D. Angeli, and F. Allgöwer. On the performance of economic model predictive control with self-tuning terminal cost. *Journal of Process Control*, 24(8):1179–1186, August 2014.

[MB12] J. Mattingley and S. Boyd. CVXGEN: a code generator for embedded convex optimization. *Optimization and Engineering*, 13(1):1–27, March 2012.

[Meh92] S. Mehrotra. On the Implementation of a Primal-Dual Interior Point Method. *SIAM Journal on Optimization*, 2(4):575–601, November 1992.

[MGL14] K. Margellos, P. Goulart, and J. Lygeros. On the Road Between Robust Optimization and the Scenario Approach for Chance Constrained Optimization Problems. *IEEE Transactions on Automatic Control*, 59(8):2258–2263, August 2014.

[ML05] S. Mehrotra and Z. Li. Convergence Conditions and Krylov Subspace–Based Corrections for Primal-Dual Interior-Point Method. *SIAM Journal on Optimization*, 15(3):635–653, January 2005.

[MM10] U. Maeder and M. Morari. Offset-free reference tracking with model predictive control. *Automatica*, 46(9):1469–1476, September 2010.

[MO09a] S. Mehrotra and M. Ozevin. On the Implementation of Interior Point Decomposition Algorithms for Two-Stage Stochastic Conic Programs. *SIAM Journal on Optimization*, 19(4):1846–1880, January 2009.

[MO09b] S. Mehrotra and M.G. Ozevin. Decomposition Based Interior Point Methods for Two-Stage Stochastic Convex Quadratic Programs with Recourse. *Operations Research*, 57(4):964–974, May 2009.

[MRRS00] D.Q. Mayne, J.B. Rawlings, C.V. Rao, and P.O.M. Scokaert. Constrained model predictive control: Stability and optimality. *Automatica*, 36(6):789 – 814, 2000.

[MRSVC14] M. Maasoumy, C. Rosenberg, A. Sangiovanni-Vincentelli, and D.S. Callaway. Model predictive control approach to online computation of demand-side flexibility of commercial buildings HVAC systems for Supply Following. In *American Control Conf. (ACC), 2014*, pages 1082–1089, June 2014.

[MSR05] D.Q. Mayne, M.M. Seron, and S.V. Raković. Robust model predictive control of constrained linear systems with bounded disturbances. *Automatica*, 41(2):219–224, February 2005.

[MT14] A. Mitsos and A. Tsoukalas. Global optimization of generalized semi-infinite programs via restriction of the right hand side. *Journal of Global Optimization*, 61(1):1–17, January 2014.

[MWM+14] D. Manz, R. Walling, N. Miller, B. LaRose, R. D'Aquila, and B. Daryanian. The grid of the future: Ten trends that will shape the grid over the next decade. *IEEE Power and Energy Magazine*, 12(3):26–36, May 2014.

[NBG+16] T.X. Nghiem, A. Bitlislioğlu, T.T. Gorecki, F.A. Qureshi, and C.N. Jones. OpenBuildNet framework for distributed co-simulation of smart energy systems. In *2016 14th International Conference on Control, Automation, Robotics and Vision (ICARCV)*, pages 1–6, November 2016.

[NM14] R.R. Negenborn and J.M. Maestre. Distributed Model Predictive Control: An Overview and Roadmap of Future Research Opportunities. *IEEE Control Systems*, 34(4):87–97, August 2014.

[NN94] Y. Nesterov and A. Nemirovskii. *Interior-point polynomial algorithms in convex programming*. SIAM, Philedelphia, 1994.

[NND11] I. Necoara, V. Nedelcu, and I. Dumitrache. Parallel and distributed optimization methods for estimation and control in networks. *Journal of Process Control*, 21(5):756–766, June 2011.

[NNPPV14] A. Nayyar, M. Negrete-Pincetic, K. Poolla, and P. Varaiya. Duration-differentiated energy services with a continuum of loads. In *53rd IEEE Conference on Decision and Control*, pages 1714–1719, December 2014.

[NÖ09] A. Nedic and A. Özdağlar. Distributed Subgradient Methods for Multi-Agent Optimization. *IEEE Transactions on Automatic Control*, 54(1):48–61, January 2009.

[NOP10] A. Nedic, A. Özdağlar, and P.A. Parrilo. Constrained Consensus and Optimization in Multi-Agent Networks. *IEEE Transactions on Automatic Control*, 55(4):922–938, April 2010.

[NS09] I. Necoara and J.A.K. Suykens. Interior-Point Lagrangian Decomposition Method for Separable Convex Optimization. *Journal of Optimization Theory and Applications*, 143(3):567, May 2009.

[NW00] J. Nocedal and S.J. Wright. *Numerical Optimization*. Springer, New York, 2000.

[OBB+13] F. Oldewurtel, T. Borsche, M. Bucher, P. Fortenbacher, M.G.V.T. Haring, J.L. Mathieu, O. Megel, E. Vrettos, and G. Andersson. A framework for and assessment of demand response and energy storage in power systems. In *Bulk Power System Dynamics and Control - IX Optimization, Security and Control of the Emerging Power Grid (IREP), 2013 IREP Symposium*, pages 1–24, 2013.

[Oht04] T. Ohtsuka. A continuation/GMRES method for fast computation of nonlinear receding horizon control. *Automatica*, 40(4):563–574, April 2004.

[PHA15] S.K. Pakazad, A. Hansson, and M.S. Andersen. Distributed Primal-dual Interior-point Methods for Solving Loosely Coupled Problems Using Message Passing. *arXiv:1502.06384 [math]*, February 2015. arXiv: 1502.06384.

[PHAN17] S.K. Pakazad, A. Hansson, M.S. Andersen, and I. Nielsen. Distributed primal–dual interior-point methods for solving tree-structured coupled convex problems using message-passing. *Optimization Methods and Software*, 32(3):401–435, May 2017.

[PW00] F.A. Potra and S.J. Wright. Interior-point methods. *Journal of Computational and Applied Mathematics*, 124(1–2):281–302, December 2000.

[Qur17] F.A. Qureshi. *Predictive Control of Buildings for Demand Response and Ancillary Services Provision.* PhD thesis, EPFL, Lausanne, 2017.

[RAB12] J.B. Rawlings, D. Angeli, and C.N. Bates. Fundamentals of economic model predictive control. In *2012 IEEE 51st IEEE Conference on Decision and Control (CDC)*, pages 3851–3861, December 2012.

[Rak12] S. Raković. Invention of Prediction Structures and Categorization of Robust MPC Syntheses. In M. Lazar, editor, *Proc. 4th IFAC Nonlinear Model Predictive Control Conf.*, pages 245–273, August 2012.

[RFFT13] S. Riverso, M. Farina, and G. Ferrari-Trecate. Plug-and-Play Decentralized Model Predictive Control for Linear Systems. *IEEE Transactions on Automatic Control*, 58(10):2608–2614, October 2013.

[RFT12] S. Riverso and G. Ferrari-Trecate. Hycon2 Benchmark: Power Network System. *arXiv:1207.2000 [cs]*, July 2012. arXiv: 1207.2000.

[RHF+08] G. Resch, A. Held, T. Faber, C. Panzer, F. Toro, and R. Haas. Potentials and prospects for renewable energies at global scale. *Energy Policy*, 36(11):4048–4056, November 2008.

[RKC+12] S.V. Raković, B. Kouvaritakis, M. Cannon, C. Panos, and R. Findeisen. Parameterized Tube Model Predictive Control. *IEEE Transactions on Automatic Control*, 57(11):2746–2761, November 2012.

[RKFC12] S. V. Raković, B. Kouvaritakis, R. Findeisen, and M. Cannon. Homothetic tube model predictive control. *Automatica*, 48(8):1631–1638, August 2012.

[RKMK07] S.V. Raković, E.C. Kerrigan, D.Q. Mayne, and K.I. Kouramas. Optimized robust control invariance for linear discrete-time systems: Theoretical foundations. *Automatica*, 43(5):831–841, May 2007.

[RKML06] S.V. Raković, E.C. Kerrigan, D.Q. Mayne, and J. Lygeros. Reachability analysis of discrete-time systems with disturbances. *IEEE Transactions on Automatic Control*, 51(4):546–561, April 2006.

[RKTR07] Y.G. Rebours, D.S. Kirschen, M. Trotignon, and S. Rossignol. A Survey of Frequency and Voltage Control Ancillary Services mdash;Part I: Technical Features. *IEEE Transactions on Power Systems*, 22(1):350–357, February 2007.

[RW09] R.T. Rockafellar and R.J-B. Wets. *Variational Analysis*. 2009.

[SDR09] A. Shapiro, D. Dentcheva, and A.P. Ruszczyński. *Lectures on Stochastic Programming: Modeling and Theory*. Society for Industrial and Applied Mathematics, Philadelphia, 2009.

[SMSS09] A.T. Sarić, F.H. Murphy, A.L. Soyster, and A.M. Stanković. Two-stage stochastic programming model for market clearing with contingencies. *IEEE Transactions on Power Systems*, 24(3):1266–1278, 2009.

[SPT07] M. Salahi, J. Peng, and T. Terlaky. On Mehrotra-Type Predictor-Corrector Algorithms. *SIAM Journal on Optimization*, 18(4):1377–1397, December 2007.

[sta16] OpenEI Datasets, 2016. http://en.openei.org/datasets, accessed on 2016-05-18.

[Sti04] G. Still. Solving generalized semi-infinite programs by reduction to simpler problems. *Optimization*, 53(1):19–38, February 2004.

[Swi15a] Swissgrid. Overview of ancillary service products, 2015. https://www.swissgrid.ch/dam/swissgrid/experts/ancillary_services/Dokumente/D100412_AS-concept_V1R0_en.pdf, accessed on 2015-04-15.

[Swi15b] Swissgrid. Test for secondary control capability, 2015. https://www.swissgrid.ch/dam/swissgrid/experts/ancillary_services/prequalification/D130422_Test-for-secondary-control-capability_V2R1_EN.pdf, accessed on 2015-04-15.

[UR] US Department of Energy: Office of Energy Efficiency and Renewable Energy. Commercial Reference Buildings. http://energy.gov/eere/buildings/commercial-reference-buildings, accessed on 2014-01-02.

[VOA16] E. Vrettos, F. Oldewurtel, and G. Andersson. Robust Energy-Constrained Frequency Reserves From Aggregations of Commercial Buildings. *IEEE Transactions on Power Systems*, PP(99):1–14, 2016.

[VOAZ14] E. Vrettos, F. Oldewurtel, G. Andersson, and F. Zhu. Robust Provision of Frequency Reserves by Office Building Aggregations. In *Proceedings of the 19th IFAC World Congress*, Cape Town, South Africa, 2014.

[WB10] Y. Wang and S. Boyd. Fast Model Predictive Control Using Online Optimization. *IEEE Transactions on Control Systems Technology*, 18(2):267–278, March 2010.

[WGMM13] J. Warrington, P. Goulart, S. Mariethoz, and M. Morari. Policy-Based Reserves for Power Systems. *IEEE Transactions on Power Systems*, 28(4):4427–4437, November 2013.

[WH04] A.G. Wills and W.P. Heath. Barrier function based model predictive control. *Automatica*, 40(8):1415–1422, August 2004.

[WO12] E. Wei and A. Ozdaglar. Distributed Alternating Direction Method of Multipliers. In *51st IEEE Conference on Decision and Control (CDC)*, pages 5445–5450, December 2012.

[Wri97] S.J. Wright. *Primal-Dual Interior-Point Methods*. SIAM, Philadelphia, 1997.

[ZdH15] J. Zhen and D. den Hertog. Computing the Maximum Volume Inscribed Ellipsoid of a Polytopic Projection. *CentER Discussion Paper Series*, 2015-004, January 2015.

[Zha01] G. Zhao. A Log-Barrier method with Benders decomposition for solving two-stage stochastic linear programs. *Mathematical Programming*, 90(3):507–536, May 2001.

[Zha05] G. Zhao. A Lagrangian Dual Method with Self-Concordant Barriers for Multi-Stage Stochastic Convex Programming. *Mathematical Programming*, 102(1):1–24, January 2005.

[ZKG+17] X. Zhang, M. Kamgarpour, A. Georghiou, P. Goulart, and J. Lygeros. Robust optimal control with adjustable uncertainty sets. *Automatica*, 75:249–259, January 2017.

[ZKGL14] X. Zhang, M. Kamgarpour, P. Goulart, and J. Lygeros. Selling robustness margins: A framework for optimizing reserve capacities for linear systems. In *2014 IEEE 53rd Annual Conf. on Decision and Control (CDC)*, pages 6419–6424, December 2014.

[ZQJD17] A. Zanelli, R. Quirynen, J. Jerez, and M. Diehl. A Homotopy-based Nonlinear Interior-Point Method for NMPC. *IFAC-PapersOnLine*, 50(1):13188–13193, July 2017.

[ZZ16] L. Zhao and W. Zhang. A geometric approach to virtual battery modeling of thermostatically controlled loads. In *2016 American Control Conference (ACC)*, pages 1452–1457, July 2016.

Altuğ Bitlislioğlu

Curriculum Vitae

Avenue de Riant-Mont 1
Lausanne, Switzerland
+41 78 907 61 55
abitlislioglu@gmail.com, altug.bitlislioglu@epfl.ch
http://people.epfl.ch/altug.bitlislioglu
Birthday : 09.11.1988
Nationality : Turkish

Areas of expertise

Control	Model Predictive Control, Classical Feedback Control
Modeling	Mechatronics, Thermodynamics, Power grids
Identification	Linear & Nonlinear systems, Parameter Identification
Optimization	Convex Optimization, Robust & Distributed Optimization
Applications	Climate control, Robotics, Power Systems

Education

2013–2018 **PhD**, *Electrical Engineering (Control)*, EPFL, Lausanne - Switzerland.

2011–2013 **MSc**, *Robotics, Systems and Control*, ETHZ, Zurich - Switzerland.

2007–2010 **BSc**, *Mechanical Engineering*, Middle East Technical University, Ankara - Turkey.

Phd Thesis

Title *Coordinated Optimization and Control for Smart Grids*

Supervisor Prof. Colin Jones

Description The aim of this thesis is to develop model based control and optimization methods for enabling fleets of coordinated building climate control systems to provide supporting services to the electricity grid.

Masters Thesis

Title *Adaptive Predictive Control of the Human Cardiovascular System*

Supervisor Prof. Robert Riener

Description This thesis focused on the combined control of heart rate and blood pressure signals with only physical stimuli. An adaptive model predictive control strategy is implemented and tested on an experimental set-up.

Experience

2013-Present **Research Assistant**, *Labratoire d'Automatique - EPFL*, Lausanne.

BuildNet Project - ERC Grant Agreement n. 307608, Provision of ancillary service to the power grid through distributed control of building networks.

- Development of novel optimization and control methods for power consumption flexibility assessment of building HVAC systems and coordinated control of building fleets for power grid related services
- Support of simulation platform for distributed systems, OpenBuildNet (OBN), focusing on building thermal and HVAC models.

- Design and deployment of smart electrical heaters as part of an experimental platform for demand response research - Laboratoire d'automatique Demand Response (LADR)

2013-Present **Teaching Assistant**, *School of Engineering - EPFL*, Lausanne.

Course assistant

- Model Predictive Control
- Control Systems Laboratory
- Complex Analysis

Thesis supervision

Supervision of Masters Theses (30 ECTS credits) and Semester Theses (10 ECTS credits) for students from mechanical, micro-engineering and energy management master programs.

- Quadcopters (2 Master, 3 Semester Projects)

Construction, modeling and control of autonomous quadcopters.

- Hydrofoil Boats (1 Master, 2 Semester Projects)

State estimation, control and optimization of a hydrofoil boat, as part of the HYRDOS Project.

- Process Control and Optimization (5 Semester Projects)

Modeling, control and optimization of battery management, chemical production, wine fermentation and greenhouse climate control systems.

06.2012-12.2012 **Intern**, *Sulzer Metco Ag.*, Wohlen.

- Development of a simulation platform for benchmarking and designing automation of jet propulsion based thermal spray systems for metal coating.

Autumn 2011 **Teaching Assistant**, *Sensory Motor Systems Lab - ETHZ*, Zurich.

- Supervision of laboratory practice : haptic control of a rehabilitation robot.

Computer skills

Language PYTHON,MATLAB/SIMULINK,C# C,C++,
CAD SOLIDWORKS, CADKEY

Languages

English **Fluent**
German **Intermediate**
French **Intermediate**
Turkish **Native**