DECENTRALIZED CONTROL OF CONSTRAINED LINEAR SYSTEMS VIA CONVEX OPTIMIZATION METHODS

A Dissertation

Presented to the Faculty of the Graduate School

of Cornell University

in Partial Fulfillment of the Requirements for the Degree of

Doctor of Philosophy

by

Weixuan Lin

May 2020

DECENTRALIZED CONTROL OF CONSTRAINED LINEAR SYSTEMS VIA CONVEX OPTIMIZATION METHODS

Weixuan Lin, Ph.D.

Cornell University 2020

Decentralized control problems naturally arise in the control of large-scale networked systems. Such systems are regulated by a collection of local controllers in a *decentralized manner*, in the sense that each local controller is required to specify its control input based on its locally accessible sensor measurements. In this dissertation, we consider the decentralized control of discrete-time, linear systems subject to exogenous disturbances and polyhedral constraints on the state and input trajectories. The underlying system is composed of a finite collection of dynamically coupled subsystems, each of which is assumed to have a dedicated local controller. The decentralization of information is expressed according to sparsity constraints on the sensor measurements that each local controller has access to. In its most general form, the decentralized control problem amounts to an infinite-dimensional nonconvex program that is, in general, computationally intractable. The primary difficulty of the decentralized control problem stems from the potential informational coupling between the controllers. Specifically, in problems with *nonclassical* information structures, the actions taken by one controller can affect the information acquired by other controllers acting on the system. This gives rise to an incentive for controllers to communicate with each other via the actions that they undertake—the so-called *signaling incentive*. To complicate matters further, there may be hard constraints coupling the actions and local states being regulated by different con-

trollers that must be jointly enforced with limited communication between the local controllers. In this dissertation, we abandon the search for the optimal decentralized control policy, and resort to approximation methods that enable the tractable calculation of feasible decentralized control policies.

We first provide methods for the tractable calculation of decentralized control policies that are affinely parameterized in their measurement history. For problems with *partially nested* information structures, we show that the optimization over such a policy space admits an equivalent reformulation as a semi-infinite convex program. The optimal solution to these semi-inifinite programs can be calculated through the solution of a finite-dimensional conic program. For problems with nonclassical information structures, however, the optimization over such a policy space amounts to a semi-infinite nonconvex program. With the objective of alleviating the nonconvexity in such problems, we propose an approach to decentralized control design in which the information-coupling states are effectively treated as disturbances whose trajectories are constrained to take values in ellipsoidal "contract" sets whose location, scale, and orientation are jointly optimized with the affine decentralized control policy being used to control the system. The resulting problem is a semidefinite program, whose feasible solutions are guaranteed to be feasible for the original decentralized control design problem.

Decentralized control policies that are computed according to such convex optimization methods are, in general, suboptimal. We, therefore, provide a method of bounding the suboptimality of feasible decentralized control policies through an information-based convex relaxation. Specifically, we characterize an expansion of the given information structure, which maximizes the optimal value of the decentralized control design problem associated with the expanded information structure, while guaranteeing that the expanded information structure be partially nested. The resulting decentralized control design problem admits an equivalent reformulation as an infinite-dimensional convex program. We construct a further constraint relaxation of this problem via its partial dualization and a restriction to affine dual control policies, which yields a finite-dimensional conic program whose optimal value is a provable lower bound on the minimum cost of the original decentralized control design problem.

Finally, we apply our convex programming approach to control design to the decentralized control of distributed energy resources in radial power distribution systems. We investigate the problem of designing a fully decentralized disturbance-feedback controller that minimizes the expected cost of serving demand, while guaranteeing the satisfaction of individual resource and distribution system voltage constraints. A direct application of our aforementioned control design methods enables both the calculation of affine controllers and the bounding of their suboptimality through the solution of finite-dimensional conic programs. A case study demonstrates that the decentralized affine controller we compute can perform close to optimal.

BIOGRAPHICAL SKETCH

Weixuan Lin received a B.E. degree in Electrical Engineering with high honors and a B.S. degree in Economics from Tsinghua University in Beijing, China in July 2013. He received a Ph.D. degree in Electrical and Computer Engineering with a minor in Operations Research from Cornell University in Ithaca, NY in May 2020.

He is broadly interested in problems of decentralized decision making. His specific research interest includes algorithmic game theory, decentralized control theory, convex optimization, and their applications in the control of modern power systems and the analysis of market power in power markets. He is a recipient of the Jacobs Fellowship and the Hewlett Packard Fellowship.

To my parents, Guohua Lin and Meina Jin, and Ting.

ACKNOWLEDGEMENTS

First and foremost, I would like to thank my Ph.D. advisor, Professor Eilyan Bitar, for his continued guidance and support throughout the course of my Ph.D.. His deep technical knowledge, elegance and concision in communication, and insights in the thought process has been invaluable in shaping me as a mature researcher. In particular, I am grateful for the enormous amount of time he has investigated in our technical discussions. All the excitements we had upon the discovery of new insights to profound open questions are still fresh in my mind until this day. The experience of collaborating with Professor Bitar has been invaluable, and will have a long-lasting impact on my career. I owe much to him and will always appreciate this Ph.D. experience.

I would also like to thank Professor Adam Wierman with whom I had the opportunity to collaborate with during and after my visit to the Simons Institute for the Theory of Computing. I thoroughly enjoyed our exploration into the area of networked Cournot competition, and learned precious lessons from him on how to communicate complicated ideas in simple and intuitive languages. I am also thankful to Professor Éva Tardos whose course provided me a systematic introduction of algorithmic game theory. Those lectures spurred a collection of interesting papers that use game theoretic solution concepts to identify the sources of market power in various marketplaces. I am also grateful to Professor Shane Henderson and Professor Kevin Tang for serving on my thesis committee and providing helpful feedback to my work.

I could not be more fortunate to have known Raphael Louca, Kia Khezeli, Daniel Munoz Alvarez, Polina Alexeenko, Kursat Mestav, Subhonmesh Bose, and Xiaoshuang Chen both as my colleagues and my close friends. I will always remember the constructive feedback and kind encouragements they provide during the practice talks I gave throughout the years. And working in Rhodes 365 has been a pleasant experience with their companion. I have also made many life-long friends during my days in Ithaca and the last year of my Ph.D. in San Francisco. The interaction with them has made my Ph.D. journey a much more enjoyable experience. Special thanks go to Xiaobo, Rujin, Subrina, Weijia, Qian, Cong, Kecheng, Yuting, Zhe, Jaime, Ibrahim, Nirmal, Kolbeinn, Sattar, Deepan, Sevi, Takula, Tunchanok, and my feline friend Daoji.

All of my achievements to this day would not have been possible without the support from my parents. They have encouraged me to make my decisions independently since my childhood, and have always backed me up on all the decisions I make—both reasonable and ridiculous ones. I would never have the confidence and the ambition to explore all the potentially risky opportunities I have taken without the mindset that I always have a safe haven at home.

Finally, I would like to thank Ting for her unconditional love and emotional support throughout my Ph.D.. In particular, I am grateful for the effort she has made to make sure that I stay positive and keep a forward-looking mindset. I have also been deeply inspired by the self-motivation she displays in the pursuit of her personal and professional goals. She is the person that accepts me for who I am in all aspects—both good and bad ones. The completion of my Ph.D. would not be possible without her.

TABLE OF CONTENTS

LIST OF TABLES

LIST OF FIGURES

CHAPTER 1

INTRODUCTION

1.1 The Decentralized Control Design Problem

Classical optimal control theory is founded on the assumption of *centralized information structures*. That is to say, the system is assumed to be controlled by a single controller that has access to all sensor measurements. Many real-world large-scale systems, however, may fail to satisfy this assumption. Typical examples include the power system, signalized transportation networks, vehicle platoons, supply chains, and digital communication networks. In such systems, the sharing of sensor measurements is limited, due to the geographical separation between different system components, the limited communication capability between the local controllers, and the cost of storage and computation for the local controllers. As a result, the control of such systems is required to be performed in a decentralized fashion, in the sense that each local controller needs to specify its control input using its locally accessible sensor measurements.

The control actions taken by the local controllers are expected to jointly achieve a performance objective of the global system without violating global system constraints. In order to achieve an appropriate coordination between the local controllers, the local control policies need to be jointly designed through the solution of a so-called *decentralized control design problem*. The study of decentralized control design problems dates back to the seminal work by Radner in the early 1960s on team decision problems [81], yet the optimal solutions to most decentralized control design problems have remained unknown to this day. In this dissertation, we consider a family of decentralized control design problems for discrete-time linear systems that operate over a finite time horizon. In such problems, the objective is to specify a decentralized control policy that minimizes the expected value of a convex cost function while guaranteeing the robust satisfaction of constraints on the state and input trajectories. In general, such decentralized control design problems amount to infinite-dimensional non-convex programs that are known to be computationally intractable [89, 101, 109]. In what follows, we provide an overview of the main difficulties in the decentralized control design problems considered in this dissertation.

1.1.1 Nonclassical Information Structures

The primary difficulty in decentralized control problems lies in the potential informational coupling between the controllers. Specifically, in problems with *nonclassical* information structures, the actions taken by one controller can affect the information acquired by other controllers. In these problems, there is an incentive for controllers to communicate with each other through its control actions—the so-called *signaling incentive* in decentralized control problems. As a result of such informational coupling, the optimal decentralized control policy needs to achieve an appropriate three-way tradeoff between *exploration*, *exploitation*, and *communication*, and is, in general, intractable to compute. A typical example to this assertion is the celebrated Witsenhausen's counterexample [109], which shows that even a simple two-stage, two-controller decentralized control design problem would be computationally intractable due to the information coupling between controllers. Specifically, it was shown in [109] that a two-point quantization policy strictly outperforms the optimal affine control policy when the signal-to-noise ratio (SNR) of communicating through the plant is large and the cost of communication is low. Such a result clearly exemplifies the incentive to signal in decentralized control design problems. To this day, the optimal solution to Witsenhausen's counterexample remains unknown [93], indicating that there is little hope in calculating the optimal control policy for more complicated decentralized control problems.

1.1.2 Satisfaction of Hard Constraints

The difficulty in solving problems of decentralized control is further exacerbated by the requirement of enforcing hard constraints on the state and input of the global system with limited communication between controllers. Specifically, there are two major sources of difficulty in enforcing hard constraints in problems of decentralized control. First, the enforcement of such constraints requires an implicit coordination between the control inputs of different controllers. That is, each controller's input is required to be constraint admissible given the "worst-case" realization of other controllers' inputs. Second, the informational coupling between different controllers further complicates the specification of constraint-admissible decentralized control policies. Specifically, a change in the local control policy of one controller might affect the information acquired by other controllers and subsequently lead to a change in their control inputs. Even under the restriction to decentralized affine control policies, enforcing polyheral constraints on the state and input trajectories amounts to the enforcement of a multi-linear inequality constraint on the control gains—a constraint that is, in general, computationally intractable to enforce.

The aforementioned discussions on the difficulty in decentralized control design problems suggests that there might be little hope in obtaining the globally optimal control policy for the decentralized control of constrained dynamical systems. In this dissertation, we relax the requirement that the decentralized control policy be optimal with respect to the broad family of all causal decentralized control policies, and instead search for suboptimal decentralized control policies that can be efficiently computed via convex optimization methods.

1.2 Summary of Contributions and Organization

In this dissertation, we develop tractable inner and outer approximations of constrained decentralized control design problems using convex optimization methods. The primary contributions of this dissertation are two-fold. First, we derive inner approximations of constrained decentralized control design problems that can be tractably computed via the solution of finite-dimensional convex programs. Second, as the decentralized control policy we compute is, in general, suboptimal, we bound the suboptimality of feasible decentralized control policies via the calculation of a lower bound on the optimal value of a decentralized control problem using convex optimization methods. In what follows, we summarize the contents and contributions in each chapter of the dissertation.

In **Chapter 2**, we review the sources of difficulty in constrained decentralized control design problems via an investigation of Witsenhausen's counterexample and its constrained variants. We first provide a review of the classical Witsenhausen's counterexample to unravel the difficulty in decentralized control problems that stems from the incentive to signal between controllers. In particular, our result shows that affine control policies are close to optimal if the incentive to signal—as measured according to the signal-to-noise ratio (SNR) and the cost of "communicating" through the actions—is small, but might be far from optimal if the incentive to signal is large. Additionally, we investigate a constrained variant of Witsenhausen's counterexample, and describe the difficulty in decentralized control associated that arises from the hard constraints on the state and input. We show that the optimal linear control policy has an "assume-guarantee" structure that reflects the implicit coordination between the controllers. Namely, each controller *assumes* that the coupling states and inputs from other controllers behave as disturbances that take value in a given "contract" set, and constrain its control policy in a manner that *guarantees* the consistency between the assumed and actual behaviors of the coupling states and inputs. We illustrate how such a structure enables the tractable calculation of feasible control policies for the constrained variant of Witsenhausen's counterexample, and discuss the application of this technique to more general decentralized control design problems.

In **Chapter 3**, we define the constrained decentralized output-feedback control design problem that we consider, and discuss how the convexity of the problem depends on the underlying information structure. We first construct an equivalent reformulation of the decentralized control design problem using the classical Youla parameterization. We show that the Youla parameterization yields an equivalent reformulation of the decentralized control design problem as a convex program if and only if the information structure is *partially nested*—that is, an information structure in which each controller's local information cannot be affected by any control input. Under the restriction to partially nested information structures, we show that the calculation of the optimal affine control policy amounts to a robust convex program and admits an equivalent reformulation as a finite-dimensional conic program.

In **Chapter 4**, we investigate the design of decentralized control policies that are affine in the state history for problems with arbitrary (possibly nonclassical) information structures. In order to alleviate the nonconvexity arising from the informational coupling between subsystems, we treat the so-called information-coupling states as disturbances whose trajectories are "assumed" to take values in a contract set. To ensure the satisfaction of this assumption, we impose a contractual constraint on the control policy that "guarantees" that the information-coupling states that it induces belong to said contract set. Naturally, this yields an inner approximation of the original decentralized control design problem, where the conservatism of the resulting approximation depends critically on the specification of the contract set. To mitigate this potential conservatism, we formulate a semi-infinite program to co-optimize the decentralized control policy with the location, scale, and orientation of the contract set. We establish a condition on the set of allowable contracts that facilitates the joint optimization the contract set and control policy via the solution of a semidefinite program.

The decentralized control policies we derive in Chapters 3 and 4 are, in general, suboptimal. In **Chapter 5**, we derive a computationally tractable lower bound on the minimum cost of the decentralized control design problem to evaluation the suboptimality of such policies. Our derivation of the lower bound consists of two relaxation steps, which together yield a finite-dimensional convex programming relaxation of the original problem. The first step entails an information relaxation, which eliminates the signaling incentive between controllers by expanding the set of measurements that each controller has access to. Specifically, we characterize an expansion of the given information structure, which ensures its partial nestedness, while maximizing the optimal value of the resulting decentralized control problem under the expanded information structure. The relaxed decentralized control problem admits an equivalent reformulation as an infinite-dimensional convex program via the classical Youla parameterization. Although convex, the resulting optimization problem remains computationally intractable due to its infinite-dimensionality. In the second relaxation step, we obtain a finite-dimensional relaxation of this problem through its partial dualization, and restriction to affine dual control policies. The resulting problem is a finite-dimensional conic program, whose optimal value is guaranteed to be a lower bound on the minimum cost of the original decentralized control design problem.

In **Chapter 6**, we describe an application of the techniques developed in this dissertation to the decentralized control of distributed energy resources in power distribution systems. The problem we consider amounts to the design of a fully decentralized disturbance-feedback controller that minimizes the expected value of a convex quadratic cost function, subject to robust convex quadratic constraints on the system state and input. The optimal control policy for such problems is, in general, intractable to compute. We apply the techniques we developed in Chapter 3 to derive a tractable inner approximation of the decentralized control design problem. This enables the efficient computation of an affine control policy via the solution of a finite-dimensional conic program. As affine control policies are, in general, suboptimal for the family of systems considered, we apply our results in Chapter 5 to bound their suboptimality via the solution of another finite-dimensional conic program. We verify that the decentralized controller we derive are close to optimal for the problem instance considered in the case study.

In **Chapter 7**, we close the dissertation with a summary of main contributions and a discussion on directions for future work.

1.3 Notation

Let $\mathbf{R}$ and $\mathbf{R}_+$ denote the set of real numbers and non-negative real numbers, respectively. Denote the transpose of a vector $x \in \mathbf{R}^n$ by x^T. For any pair of vectors $x = (x_1, .., x_n) \in \mathbf{R}^n$ and $y = (y_1, .., y_m) \in \mathbf{R}^m$, we define their concatenation as $(x, y) = (x_1, .., x_n, y_1, .., y_m) \in \mathbf{R}^{n+m}$. For a vector $x = (x_1, \ldots, x_N)$ that is concatenated from N subvectors $x_1, \ldots, x_N$ and an index set $J \subseteq \{1, \ldots, N\}$, we denote by x_J the subvector of x that is concatenated from the subvectors x_j for $j \in J$. The subvectors in x_J are ordered in ascending order of their indices. For example, if $J = \{1, 3\}$, then $x_J = (x_1, x_3)$. Given a process $\{x(t)\}$ indexed by $t = 0, \ldots, T-1$, we denote by $x^t = (x(0), x(1), \ldots, x(t))$ its history up until and including time t, and by $x^{s:t} = (x(s), x(s+1), \ldots, x(t))$ its history from time s until time t.

We consider block matrices throughout the dissertation. Given a block matrix A whose dimension will be clear from the context, we denote by $[A]_{ij}$ its $(i,j)^{\text{th}}$ block. We denote the trace of a square matrix A by $\operatorname{Tr}(A)$. We denote by $\mathcal{K}$ a proper cone (i.e., convex, closed, and pointed with an nonempty interior). Let $\mathcal{K}^*$ denote its dual cone. We write $x \succeq_{\mathcal{K}} y$ to indicate that $x - y \in \mathcal{K}$. Given a matrix A, we let $A \succeq_{\mathcal{K}} 0$ denote its columnwise inclusion in $\mathcal{K}$. For a set $\mathcal{S} \subseteq \mathbf{R}^n$ and a matrix $A \in \mathbf{R}^{m \times n}$, the image of the set $\mathcal{S}$ under the linear map A

is given by $A\mathcal{S} = \{Ax | x \in \mathcal{S}\}$. For two arbitrary sets $\mathcal{S}, \mathcal{T} \subseteq \mathbf{R}^n$, we denote their Minkowski sum by $\mathcal{S} \oplus \mathcal{T} := \{x + y | x \in \mathcal{S}, y \in \mathcal{T}\}$.

CHAPTER 2

WITSENHAUSEN'S COUNTEREXAMPLE AND ITS VARIANTS

Witsenhausen's counterexample [109] is a two-stage, two-controller decentralized control problem that reveals the difficulty in decentralized control that stems from the incentive to signal between controllers. In this chapter, we restrict ourselves to Witsenhausen's counterexample and its constrained variants. Our investigation of these toy problems reveals the main challenges in solving constrained decentralized control design problems.

The remainder of this chapter is organized as follows. In Section 2.1, we illustrate how *the incentive to signal* between controllers leads to the intractability in solving Witsenhausen's counterexample. We first show that calculation of the optimal affine control policy of Witsenhausen's counterexample requires the solution of a nonconvex program. As affine policies are, in general, suboptimal, we reproduce the information-theoretic lower bounds in [52] on the optimal value of Witsenhausen's counterexample, which provides an upper bound on the suboptimality of affine policies. The resulting bound quantifies the difficulty in solving decentralized control design problems that stems from the incentive to signal—as is measured by the signal-to-noise ratio (SNR) and the cost of communicating through the plant. Namely, we show that affine policies are close to optimal if the SNR of communicating through the plant is small or the cost of communication is high, but might be far from optimal if neither of the aforementioned conditions are satisfied.

In addition to the incentive to signal, another important source of difficulty in constrained decentralized control design problems is the requirement of enforcing hard constraints on the system state and input with limited communication between controllers. In Section 2.2, we investigate a constrained variant of Witsenhausen's counterexample and describe the additional difficulty in solving this problem that stems from the hard constraints on both controllers' inputs. Specifically, the constraint variant of Witsenhausen's counterexample we investigate entails a hard constraint on the inputs from the two controllers that needs to be enforced without explicit communication. We show that the optimal linear control policy has an "assume-guarantee" structure that reflects the implicit coordination between the two controllers. Namely, under the optimal linear control policy, the second controller essentially treats the state that is affected by the first controller's input as a disturbance that is *assumed* to take value in a given "contract set". The control policy of the first controller, on the other hand, is constrained in such a way that *guarantees* the inclusion of the state it affects in the "contract set". Such a mechanism of attaining the implicit coordination between controllers will enable the tractable inner approximation of a general family of decentralized control design problems. We provide an illustration of this approximation technique to the constrained variant of Witsenhausen's counterexample.

2.1 Witsenhausen's Counterexample

In the centralized control of linear systems with quadratic cost and additive Gaussian disturbances (i.e., the LQG problem), it is well-known in the literature that the optimal control policy is affine in the history of measured outputs [92]. For decentralized LQG problems, however, Witsenhausen provided a counterexample—later known as Witsenhausen's counterexample [109]—to the conjecture that affine policies are also optimal to decentralized control problems.

Specifically, Witsenhausen's counterexample is a two-stage, two-controller decentralized LQG problem, in which a two-point quantization policy that communicates the first controller's local information to the second controller outperforms the best affine control policy. This clearly reveals the potential reduction in cost that might be attained via the communication of a controller's local information through its control action—the so-called *signaling incentive*. In what follows, we leverage on Witsenhausen's counterexample to describe the difficulty in decentralized control design problem that stems from the signaling incentive between controllers.

2.1.1 Problem Statement and the Signaling Incentive

Consider a linear system described according to

$$x_1 = x_0 + u_0,$$
$$x_2 = x_1 - u_1,$$

where all state and input variables are assumed to be scalars. The initial condition x_0 is assumed to be a zero mean Gaussian random variable with variance σ^2, where the scalar parameter $\sigma > 0$. The outputs that controller-0 and controller-1 have access to are given by

$$y_0 = x_0,$$
$$y_1 = x_1 + z,$$

where the measurement noise z is assumed to be a zero mean Gaussian random variable with variance 1 that is independent of the initial condition x_0. The control input from each controller is specified as a function of its accessible output.

That is, they are specified according to

$$u_i = \gamma_i(y_i), \quad i = 0, 1,$$

where $\gamma_i(\cdot)$ is a measurable univariate scalar function for $i = 0, 1$. Witsenhausen's counterexample is given by the following decentralized control design problem:

$$\begin{aligned} \text{minimize} \quad & \mathbf{E}[\lambda^2 u_0^2 + x_2^2] \\ \text{subject to} \quad & u_i = \gamma_i(y_i), \quad i = 0, 1 \\ & x_1 = x_0 + u_0, \\ & x_2 = x_1 - u_1, \\ & y_0 = x_0, \\ & y_1 = x_1 + z, \end{aligned} \tag{2.1}$$

where the cost parameter $\lambda > 0$. The information structure in Witsenhausen's counterexample is *nonclassical*, in the sense that controller-1's accessible output is an affine function of controller-0's control input—which controller-1 does not have access to. In Figure 2.1, we provide a cartoon illustration of Witsenhausen's counterexample that is taken from [88]. Specifically, one can think of Witsenhausen's counterexample as a problem of regulating the system initial state x_0 to be close to zero in two steps with minimum cost using a "weak" controller and a "blurry" controller. Here, controller-0 is considered to be "weak", as it is subject to a quadratic cost on its control inputs; and controller-1 is considered to be "blurry", as it only has access to a noisy measurement of the state x_1.

The optimal control input of controller-1 is the conditional expectation of the state x_1 given its noisy measurement $x_1 + z$. Consequently, solving problem (2.1) boils down to specifying the control policy from controller-0 that minimizes

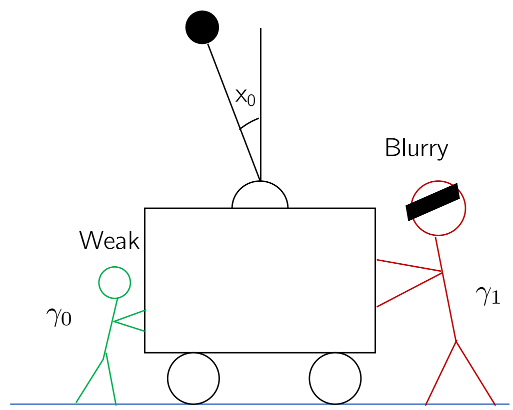

Figure 2.1: A cartoon illustration of Witsenhausen's counterexample (taken from [88]) as the problem of controlling the position of a hypothetical inverted pendulum over two discrete time periods. The controller that acts first is a "weak" controller, who has perfect observation of the system state and a high control cost. The controller that acts second is a "blurry" controller, who has a noisy measurement of the system state and a control cost of zero. The objective is to design control policies for both controllers that regulates the system state close to zero in two steps with a low cost.

the sum of its control cost and the mean squared error (MSE) in estimating the state x_1. In Witsenhausen's seminal paper [109], it was shown that a nonlinear control policy, in which controller-0 transmits the sign of the random variable x_0 to controller-1 using a two-point quantization policy, strictly outperforms the optimal affine control policy in the high "signal-to-noise ratio" regime of the problem parameters. Such a result clearly reveals the incentive for controllers to implicitly communicate to each other through their control inputs—the so-called *incentive to signal*.

Research on Witsenhausen's counterexample has since provided profound insights on the difficulty in decentralized control design problems that arises from the incentive to signal. In spite of its deceptive simplicity, Witsenhausen's counterexample has remained unsolved for 52 years [93]. The hardness for solving Witsenhausen's counterexample is not a result of the inadequacy of available mathematical tools, but in stead, a consequence of its inherent complexity. Specifically, it was shown in [101] that a discrete version of Witsenhausen's counterexample is NP-complete. To this day, the only control policy that is known to be within a constant factor of the optimum of the problem is a multi-point quantization policy [52, 53]. There has been a conjecture that the optimal control policy for Witsenhausen's counterexample is a slopey quantization policy [6, 52, 68], but no formal proof has been obtained for this conjecture.

The computational intractability of Witsenhausen's counterexample suggests that there is little hope in calculating the optimal control policy for more general decentralized control problems that involve the control of multiple subsystems over multiple time periods. In the following subsection, we further strengthen this argument by showing that the problem of computing the optimal decentralized affine control policy requires the solution of a nonconvex (multilinear) program.

2.1.2 Optimal Affine Policy

Consider the setting in which the control input from controller-0 is given by

$$u_0 = k_0 x_0 + \overline{u}_0. \tag{2.2}$$

It follows from Eq. (2.2) that the conditional mean $\mathbf{E}[x_1|y_1]$ is affine in y_1. This implies that the optimal affine control policy for controller-1 is given by

$$u_1^* = \mathbf{E}[x_1|y_1] = \overline{u}_0 + \frac{\sigma^2(k_0+1)^2}{1+\sigma^2(k_0+1)^2}(y_1 - \overline{u}_0). \tag{2.3}$$

The terminal state x_2 that results under the control policies in Eqs. (2.2) and (2.3) is given by

$$x_2 = \frac{(k_0+1)x_0 - \sigma^2(k_0+1)^2 z}{\sigma^2(k_0+1)^2+1}.$$

We denote by $J_a(k_0, \overline{u}_0)$ the cost that is incurred by the affine control policies specified in Eqs. (2.2) and (2.3). It is given by

$$J_a(k_0, \overline{u}_0) = \lambda^2(k_0^2\sigma^2 + \overline{u}_0^2) + \frac{\sigma^2(k_0+1)^2}{\sigma^2(k_0+1)^2+1}. \tag{2.4}$$

Eq. (2.4) clearly reveals controller-0's incentive to signal its local information to controller-1, as the cost $J_a(k_0, \overline{u}_0)$ includes one term that is associated with the MSE in estimating the state x_1 given the output y_1.

It is straightforward to show that the optimal open-loop control input from controller-0 is given by $\overline{u}_0^* = 0$. The calculation of the optimal feedback control gain k_0, however, requires the solution of a nonconvex program. We have the following result, which provides a necessary condition that the optimal affine control policy must satisfy.

Lemma 1. The optimal affine control policy for problem (2.1) satisfies

$$u_0 = k_0 y_0 \quad \text{and} \quad u_1 = \frac{\sigma^2(k_0+1)^2}{1+\sigma^2(k_0+1)^2} y_1,$$

where k_0 satisfies

$$-\frac{1}{\lambda} \le k_0 \le \frac{1}{\lambda} \quad \text{and} \quad \lambda^2 k_0 = -\frac{k_0+1}{\left(\sigma^2(k_0+1)^2+1\right)^2}.$$

The proof of Lemma 1 is omitted, as it immediately follows from the stationarity condition $\partial J_a / \partial k_0 = 0$. Lemma 1 shows that even for a decentralized control design problem as simple as Witsenhausen's counterexample, the optimal affine control policy requires the enumeration of all candidate solutions that satisfy the stationarity condition. Moreover, the difficulty in calculating the optimal affine control policy only gets further amplified if we consider the decentralized control of a linear system consisting of multiple local controllers over multiple time periods. For these problems, the calculation of the optimal affine control policy amounts to a multilinear optimization problem and is, in general, computationally intractable.

2.1.3 Lower Bounds and Performance Guarantees

Affine control policies are, in general, suboptimal for Witsenhausen's counterexample. In spite of this, they can perform close to optimal in certain parameter regimes. In this subsection, we bound the suboptimality of the optimal affine control policy by reproducing the information-theoretic lower bound in [52,53]. Of particular relevance to the difficulty in decentralized control that stems from the incentive to signal, this bound shows that affine control policies perform close to optimal when the incentive to signal is small—that is, problem instances in which the signal-to-noise ratio (SNR) of "communicating through the plant" is small or the cost of communication is high.

We begin by reproducing the information-theoretic lower bound on the optimal value of Witsenhausen's counterexample that was initially presented in [52, 53]. Specifically, consider the following constrained variant of Witsenhausen's counterexample, in which we impose a power constraint on the control input from controller-0:

$$\begin{aligned}
\text{minimize} \quad & \mathbf{E}[\lambda^2 u_0^2 + x_2^2] \\
\text{subject to} \quad & u_i = \gamma_i(y_i), \quad i = 0, 1 \\
& \mathbf{E}[u_0^2] = P \\
& x_1 = x_0 + u_0, \\
& x_2 = x_1 - u_1, \\
& y_0 = x_0, \\
& y_1 = x_1 + z,
\end{aligned} \tag{2.5}$$

where the power of controller-0 satisfies $P \geq 0$. Let $J^*(P)$ denote the optimal value for problem (2.5). The derivation of a lower bound on $J^*(P)$ is equivalent to deriving a lower bound on the estimation error $\mathbf{E}[x_2^2]$ under the power constraint $\mathbf{E}[u_0^2] = P$. This follows from the triangular inequality arguments in [52, Lem. 3]. Specifically, we have that

$$\sqrt{\mathbf{E}[(x_0 - u_1)^2]} \leq \sqrt{\mathbf{E}[(x_1 - u_1)^2]} + \sqrt{\mathbf{E}[(x_0 - x_1)^2]} = \sqrt{\mathbf{E}[x_2^2]} + \sqrt{P}.$$

It follows that $\mathbf{E}[x_2^2]$ is lower bounded by

$$\mathbf{E}[x_2^2] \geq \left(\left(\sqrt{\mathbf{E}[(x_0 - u_1)^2]} - \sqrt{P} \right)^+ \right)^2, \tag{2.6}$$

where $(\cdot)^+ := \max\{\cdot, 0\}$. In the following lemma, we provide a lower bound on the estimation error $\mathbf{E}[(x_0 - u_1)^2]$ under the power constraint $\mathbf{E}[u_0^2] = P$. Its proof is identical to the information-theoretic arguments in [7,8,13]. In spite of this, we keep the proof in its full details to make this subsection self-contained.

Lemma 2. Let u_0 and u_1 satisfy all constraints in problem (2.5). It follows that

$$\mathbf{E}[(x_0 - u_1)^2] \geq \frac{\sigma^2}{\left(\sigma + \sqrt{P}\right)^2 + 1}.$$

Proof: It follows from the triangular inequality in [52, Lem. 3] that $\mathbf{E}[x_1^2]$ is upper bounded by

$$\mathbf{E}[x_1^2] \leq \left(\sqrt{\mathbf{E}[x_0^2]} + \sqrt{\mathbf{E}[u_0^2]}\right)^2 = \left(\sigma + \sqrt{P}\right)^2.$$

Additionally, note that x_1 is a function of x_0, and u_1 is a function of $x_1 + z$. It follows from the data processing inequality [110] that

$$I(x_0; u_1) \leq I(x_1; x_1 + z), \tag{2.7}$$

where $I(\cdot;\cdot)$ denotes the mutual information. The mutual information $I(x_1; x_1 + z)$ is upper bounded by the capacity of a Gaussian channel with unit noise variance and power constraint $\mathbf{E}[x_1^2] \leq \left(\sigma + \sqrt{P}\right)^2$. That is to say, we have that

$$I(x_1; x_1 + z) \leq \frac{1}{2} \log \left(1 + \left(\sigma + \sqrt{P}\right)^2\right). \tag{2.8}$$

Additionally, the mutual information $I(x_0, u_1)$ is lower bounded by the rate distortion function (cf. [19]), which is given by

$$I(x_0; u_1) \geq \frac{1}{2} \log \frac{\sigma^2}{\mathbf{E}[(u_1 - x_0)^2]}. \tag{2.9}$$

The combination of (2.7)–(2.9) implies that

$$\frac{1}{2} \log \frac{\sigma^2}{\mathbf{E}[(u_1 - x_0)^2]} \leq \frac{1}{2} \log \left(1 + \left(\sigma + \sqrt{P}\right)^2\right).$$

It follows that the expected distortion $\mathbf{E}[(u_1 - x_0)^2]$ is lower bounded by

$$\mathbf{E}[(u_1 - x_0)^2] \geq \frac{\sigma^2}{\left(\sigma + \sqrt{P}\right)^2 + 1}.$$

This completes the proof. ■

A combination of inequality (2.6) and Lemma 2 implies the following lower bound on the optimal value of problem (2.5):

$$J^*(P) \geq \lambda^2 P + \left(\left(\sqrt{\frac{\sigma^2}{\left(\sigma + \sqrt{P}\right)^2 + 1}} - \sqrt{P}\right)^+\right)^2 \tag{2.10}$$

Let J^* be the optimal value of Witsenhausen's counterexample (2.1). It follows that J^* is lower bounded by

$$J^* \geq \inf_{P \geq 0} \left\{ \lambda^2 P + \left(\left(\sqrt{\frac{\sigma^2}{\left(\sigma + \sqrt{P}\right)^2 + 1}} - \sqrt{P} \right)^+ \right)^2 \right\}. \tag{2.11}$$

We note that the derivation of the lower bound (2.11) relies heavily on the information-theoretic interpretation of Witsenhausen's counterexample, and might not generalize to the decentralized control of arbitrary linear systems over multiple time periods. In Chapter 5, we provide a computationally tractable lower bound for a large family of constrained decentralized control problems that is based on an expansion of each controller's accessible information.

With the lower bound (2.11) in hand, we have the following result, which provides an upper bound on the suboptimality of the optimal affine control policy for Witsenhausen's counterexample.

Proposition 1. Let J_a^* be the cost that is incurred by the optimal affine controller. We have that

$$\frac{J_a^*}{J^*} \leq \left(\frac{\lambda\sqrt{\sigma^2+1}+1}{\left(\lambda\sqrt{\sigma^2+1} - \sqrt{\sigma^2\left(1+1/\lambda\right)^2+1}\right)^+} \right)^2. \tag{2.12}$$

In Figure 2.2a, we plot the suboptimality upper bound in (2.12) as a function of the cost parameter λ for different fixed values of σ. The suboptimality upper bound is decreasing in λ for each fixed value of σ. Moreover, it is straightforward to show that the suboptimality upper bound converges to 1 as $\lambda \to \infty$. This is to be expected, as an increase in the parameter λ leads to an increase in the cost to signal, and subsequently shrinks the possible reduction in cost that controller-0 might achieve via its effort to signal through the plant.

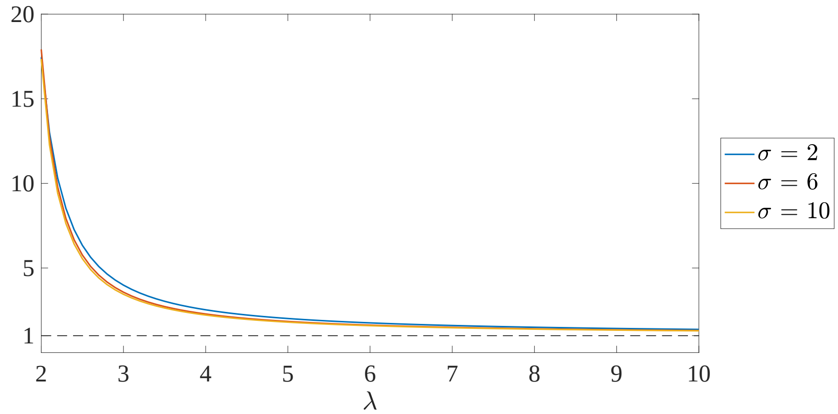

(a) The suboptimality upper bound as a function of the cost parameter λ at different fixed values of the parameter σ.

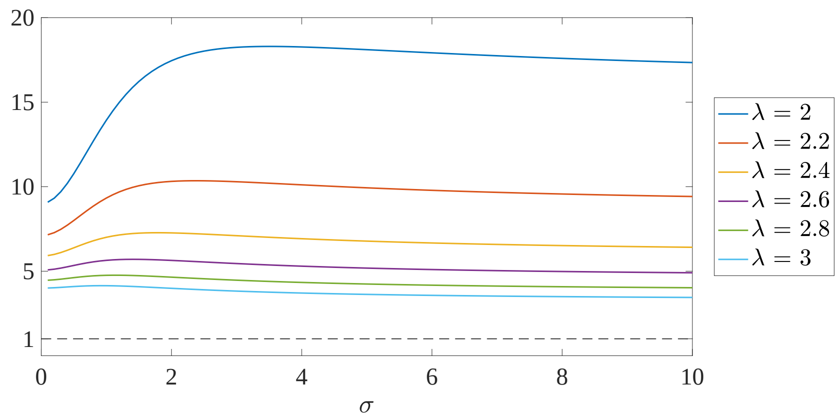

(b) The suboptimality upper bound as a function of the parameter σ at different fixed values of the cost parameter λ.

Figure 2.2: The suboptimality upper bound in (2.12) at different values of the parameters σ and λ.

In order to further examine the dependency of the suboptimality upper bound on the variance of controller-0's accessible signal x_0, we plot the suboptimality upper bound in (2.12) as a function of the parameter σ for different fixed values of the cost parameter λ in Figure 2.2b. The suboptimality upper bound is insensitive to variations in σ when the cost parameter λ is large. However, for $\lambda \leq 2.2$, the minimum value of the suboptimality upper bound for $\sigma \in [0, 10]$ is attained when $\sigma \to 0$. Moreover, one can show that the suboptimality upper bound converges to $1 + 2/(\lambda - 1)^+$ as $\sigma \to 0$. This implies the possibility that the suboptimality of affine control policies are small when the "signal-to-noise ratio" (SNR) of communicating through the plant is small. This is closely related to the classsical result in communication theory that nonlinear modulation policies are optimal only in the high SNR regime [59].

Finally, both of aforementioned observations hint at the possibility that affine control policies be highly suboptimal if the SNR of communication through the plant is high and the cost of communication is low. This is consistent with the results in Witsenhausen's seminal paper [109], in which a two-point quantization policy is shown to outperform the best linear policy asymptotically as the SNR becomes large and control cost becomes small.

We complete this subsection with the proof of Proposition 1.

Proof of Proposition 1: First note that J^* is upper bounded by

$$J^* \leq J_a(0,0) = \frac{\sigma^2}{\sigma^2+1}.$$

Let $P^* \in \operatorname{argmin}_{P\geq 0}\{J^*(P)\}$. We have the following upper bound on P^*:

$$P^* \leq \frac{J^*}{\lambda^2} \leq \frac{\sigma^2}{\lambda^2(\sigma^2+1)}.$$

Inequality (2.10) provides a lower bound on $J^*(P)$ for each $P \geq 0$. In what follows, we provide a further lower bound on $J^*(P)$ under the additional assumption that $P \in [0, \sigma^2/(\lambda^2(\sigma^2+1))]$. We have that

$$\begin{aligned}
J^*(P) &\geq \lambda^2 P + \left(\left(\sqrt{\frac{\sigma^2}{\left(\sigma+\sqrt{P}\right)^2+1}} - \sqrt{P}\right)^+\right)^2 \\
&= \lambda^2 P + \frac{\left(\left(\sigma - \sqrt{P\left(\left(\sigma+\sqrt{P}\right)^2+1\right)}\right)^+\right)^2}{\left(\sigma+\sqrt{P}\right)^2+1} \\
&\geq \lambda^2 P + \frac{\left(\left(\sigma - \sigma\sqrt{\frac{\left(\sigma+\sqrt{P}\right)^2+1}{\lambda^2(\sigma^2+1)}}\right)^+\right)^2}{\left(\sigma+\sqrt{P}\right)^2+1} \\
&\geq \lambda^2 P + \frac{\sigma^2\left(\left(1 - \sqrt{\frac{\left(\sigma\left(1+1/\sqrt{\lambda^2(\sigma^2+1)}\right)\right)^2+1}{\lambda^2(\sigma^2+1)}}\right)^+\right)^2}{\left(\sigma+\sqrt{P}\right)^2+1} \\
&\geq \lambda^2 P + \frac{\sigma^2\left(\left(1 - \sqrt{\frac{(\sigma(1+1/\lambda))^2+1}{\lambda^2(\sigma^2+1)}}\right)^+\right)^2}{\left(\sigma+\sqrt{P}\right)^2+1}.
\end{aligned}$$

where the third and fourth inequalities both follow from the assumption that $P \leq \sigma^2/(\lambda^2(\sigma^2+1))$.

Let $J_a^*(P)$ be the cost associated with the optimal linear control policy that satisfies $\mathbf{E}[u_0^2] = P$. In what follows, we provide an upper bound on $J_a^*(P)$ under the assumption that $0 \leq P \leq \sigma^2/(\lambda^2(\sigma^2+1))$. We have that

$$\begin{aligned}
J_a^*(P) = \min\left\{J_a\left(\sqrt{P/\sigma^2}, 0\right),\ J_a\left(-\sqrt{P/\sigma^2}, 0\right)\right\} &\leq \lambda^2 P + \frac{\left(\sigma+\sqrt{P}\right)^2}{\left(\sigma+\sqrt{P}\right)^2+1} \\
\leq \lambda^2 P + \frac{\sigma^2\left(1+\sqrt{\frac{1}{\lambda^2(\sigma^2+1)}}\right)^2}{\left(\sigma+\sqrt{P}\right)^2+1},
\end{aligned}$$

where the last inequality follows from the assumption that $P \leq \sigma^2/(\lambda^2(\sigma^2+1))$. The combination of the lower bound on $J^*(P)$ and the upper bound on $J_a^*(P)$ implies that

$$\begin{aligned}
\frac{J_a^*}{J^*} &= \frac{\inf_{P\geq 0} J_a^*(P)}{J^*(P^*)} \leq \frac{J_a^*(P^*)}{J^*(P^*)} \\
&\leq \left(\lambda^2 P^* + \frac{\sigma^2\left(1+\sqrt{\frac{1}{\lambda^2(\sigma^2+1)}}\right)^2}{\left(\sigma+\sqrt{P^*}\right)^2+1}\right) \Bigg/ \left(\lambda^2 P^* + \frac{\sigma^2\left(\left(1-\sqrt{\frac{(\sigma(1+1/\lambda))^2+1}{\lambda^2(\sigma^2+1)}}\right)^+\right)^2}{\left(\sigma+\sqrt{P^*}\right)^2+1}\right) \\
&\leq \left(\left(1+\sqrt{\frac{1}{\lambda^2(\sigma^2+1)}}\right) \Bigg/ \left(1-\sqrt{\frac{\sigma^2\left(1+1/\lambda\right)^2+1}{\lambda^2(\sigma^2+1)}}\right)^+\right)^2 \\
&= \left(\frac{\lambda\sqrt{\sigma^2+1}+1}{\left(\lambda\sqrt{\sigma^2+1}-\sqrt{\sigma^2\left(1+1/\lambda\right)^2+1}\right)^+}\right)^2.
\end{aligned}$$

This completes the proof. ■

2.2 Constrained Variants of Witsenhausen's Counterexample

Our analysis of Witsenhausen's counterexample in Section 2.1 reveals the difficulty of decentralized control design problems that arises from the incentive to signal between controllers. To complicate matters further, there may be hard constraints coupling the actions and local states being regulated by different controllers that must be jointly enforced with limited communication between the controllers. In what follows, we introduce a constrained variant of Witsenhausen's counterexample to describe the conceptual idea of attaining an implicit coordination between controllers through the enforcement of an "assume-guarantee contract". In this specific example, the second controller *assumes* that the input from the first controller behave as a disturbance with bounded support, and the first controller's control input is constrained in such a way to *guarantee* the consistency between the assumed and actual behavior of its input. Despite its simplicity, such a conceptual idea of attaining an implicit coordination between controllers will serve as the basis of our convex inner approximation of a large family of constrained decentralized control problems in Chapter 4.

2.2.1 Problem Statement

We consider a variant of Witsenhausen's counterexample, in which we impose a capacity constraint on the sum of the inputs from both controllers. Instead of adopting the Gaussian disturbance model in Witsenhausen's counterexample, we assume that the disturbance has a bounded support, and require that the capacity constraint be satisfied for all realizations of the disturbance. Specifically, consider the following decentralized control design problem[1]:

$$\begin{array}{ll} \text{minimize} & \mathbf{E}[\lambda^2 u_0^2 + x_2^2] \\ \text{subject to} & \left.\begin{array}{l} u_i = \gamma_i(y_i),\ i = 0, 1 \\ x_1 = x_0 - u_0, \\ x_2 = x_1 - u_1, \\ y_0 = x_0, \\ y_1 = x_1 + z, \\ |u_0 + u_1| \leq c \end{array}\right\} \forall (x_0, z) \in W. \end{array} \tag{2.13}$$

[1]Compared to the original Witsenhausen's counterexample, the state x_1 in problem (2.13) is given by $x_1 = x_0 - u_0$ instead. Such a modification to the state equation is made to clarify the interpretation of the constraint $|u_0 + u_1| \leq c$.

Here, we assume that the disturbance (x_0, z) be uniformly distributed on its support $W \subseteq \mathbf{R}^2$, where

$$W := \left\{(x_0, z) \in \mathbf{R}^2 \mid |x_0| \leq \sigma,\ |z| \leq 1\right\}. \tag{2.14}$$

We assume that the capacity $c > 0$.

2.2.2 Optimal Linear Policy

In what follows, we derive the optimal linear control policy for problem (2.13). Specifically, consider decentralized linear controllers of the form

$$u_0 = k_0 y_0, \quad u_1 = k_1 y_1. \tag{2.15}$$

We have the following result, which provides a necessary optimality condition that the optimal linear control policy is required to satisfy.

Lemma 3. Let $(k_0^*, k_1^*) \in \mathbf{R}^2$ be an optimal linear control policy for problem (2.13), and define $q_0^* := k_0^* \sigma$. It follows that k_1^* satisfies

$$k_1^* = \min\left\{\frac{(\sigma - q_0^*)^2}{(\sigma - q_0^*)^2 + 1}, \frac{c - q_0^*}{\sigma - q_0^* + 1}\right\}.$$

Here, q_0^* is a minimizer of the function $J_0(q_0)$ on the interval $[0,\ \min\{c, \sigma\}]$, where the function $J_0(q_0)$ is given by

$$J_0(q_0) = \begin{cases} \lambda^2 q_0^2 + \dfrac{(\sigma - q_0)^2}{(\sigma - q_0)^2 + 1} & \text{if } \dfrac{c - \sigma - 1}{\sigma - q_0 + 1} + \dfrac{1}{(\sigma - q_0)^2 + 1} \geq 0, \\ \lambda^2 q_0^2 + \dfrac{(\sigma - q_0)^2(\sigma - c + 1)^2 + (c - q_0)^2}{(\sigma - q_0 + 1)^2} & \text{otherwise.} \end{cases}$$

The structure of the optimal linear control policy reveals the mechanism through which both controllers coordinate implicitly to guarantee the satisfaction of the capacity constraint that couple their inputs. Specifically, the fact that

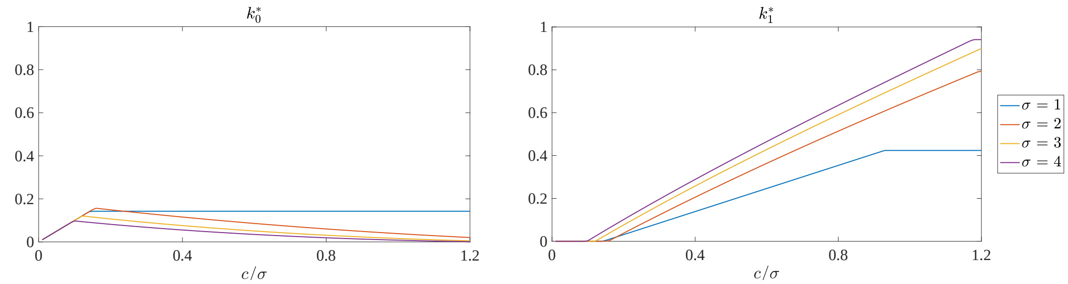

(a) Controller-0's optimal linear policy (b) Controller-1's optimal linear policy

Figure 2.3: The optimal linear control policy for problem (2.13) as a function of the ratio c/σ, evaluated at different fixed values of σ.

$k_1^* \leq (c - q_0^*)/(\sigma - q_0^* + 1)$ implies that the control input u_1 effectively has a max capacity of $c - q_0^*$. That is to say, controller-1 essentially assumes that the input from controller-0 behave as a disturbance taking value in the interval $[-q_0^*, q_0^*]$, and choose its control policy in a way that ensures the satisfaction of the capacity constraint given the "worst-case" realization of controller-0's input. On the other hand, the choice of controller-0's control policy ensures the consistency between the assumed and actual behavior of its control input, as controller-0's control input is guaranteed to take value in the set $[-q_0^*, q_0^*]$ under the policy $u_0 = k_0^* x_0$. In spite of its simplicity, this mechanism of implicit coordinating between controllers will enable the derivation of a family of convex inner approximations to constrained decentralized control problems in Chapter 4.

Additionally, we calculate the optimal linear control policy according to Lemma 3, and investigate the reduction in the controllers' incentive to signal that results from the input capacity constraint. In Figure 2.3, we plot the optimal feedback control gains for both controllers as a function of the ratio c/σ for different values of the signal magnitude σ. As we tighten the input capacity constraint, the input capacity is relocated from controller-1 to controller-0. Once the ratio c/σ gets below a threshold, the optimal linear control policy will not entail any signaling between controllers, as controller-1, the "decoder", is allocated an input capacity of 0. Such a result clearly reflects the reduction in the incentive to signal that results from the tightening of the input capacity constraint.

We close this subsection with the proof of Lemma 3.

Proof of Lemma 3: We first characterize the set of feasible linear control policies for problem (2.13) and provide a closed-form expression for the cost incurred by a linear control policy. Given a linear control policy specified in (2.15), the control inputs u_0 and u_1 are given by

$$u_0 = k_0 x_0$$
$$u_1 = k_1((1-k_0)x_0 + z).$$

It follows that $|u_0 + u_1| \leq c$ for all $(x_0, z) \in W$ if and only if

$$|k_0 + k_1(1-k_0)|\sigma + |k_1| \leq c. \tag{2.16}$$

Additionally, with a slight abuse of notation, let $J_a(k_0, k_1)$ denote the cost incurred by the linear control policy specified in (2.15). It is given by

$$J_a(k_0, k_1) = \lambda^2 k_0^2 \sigma^2 + (1-k_0)^2(1-k_1)^2\sigma^2 + k_1^2$$

The remainder of the proof is divided into two parts. In Part 1, we show that $0 \leq q_0^* \leq \min\{c, \sigma\}$. In Part 2, we establish the condition that the optimal control gains satisfy.

Part 1: Proof of the claim that $0 \leq q_0^* \leq \min\{c, \sigma\}$. In what follows, we prove this claim by showing that $0 \leq k_0^* \leq \min\{c/\sigma, 1\}$.

Part 1.1: We prove by contradiction that $k_0^* \geq 0$. Assume that the linear control policy (k_0, k_1) is feasible for problem (2.13) and satisfies $k_0 < 0$. We show that there exists another feasible linear control policy for problem (2.13) that incurs a strictly lower cost. We consider the following cases.

Case 1: $k_1(1-k_0)+k_0 \le 0$. It is straightforward to verify that the open-loop control policy $u_0 = u_1 = 0$ is feasible and satisfies $J_a(k_0, k_1) > J_a(0, 0)$.

Case 2: $k_1(1-k_0)+k_0 > 0$ and $k_1 \in (1, \infty) \cup (-\infty, 0)$. For the case in which $k_1 > 1$, it follows from inequality (2.16) that the linear policy $(k_0, 1)$ is feasible. Additionally, we have that

$$J_a(k_0, k_1) = \lambda^2 k_0^2 \sigma^2 + (1-k_0)^2(1-k_1)^2\sigma^2 + k_1^2 > \lambda^2 k_0^2 \sigma^2 + 1 = J_a(k_0, 1),$$

which implies that the linear policy $(k_0, 1)$ strictly outperforms the policy (k_0, k_1). Similarly, one can show that if $k_1 < 0$, then the linear policy $(k_0, 0)$ is also feasible and satisfies $J_a(k_0, 0) < J_a(k_0, k_1)$.

Case 3: $k_1(1-k_0)+k_0 > 0$ and $k_1 \in [0, 1]$. Define the control gain $k_1' := k_1(1-k_0)+k_0$. It follows that

$$k_1' - k_1 = (1-k_1)k_0 \le 0.$$

Additionally, our assumption that $k_1(1-k_0)+k_0 > 0$ implies that $k_1' > 0$. We have the following lower bound on $J_a(k_0, k_1)$:

$$\begin{aligned} J_a(k_0, k_1) &= \lambda^2 k_0^2 \sigma^2 + (1-k_0-k_1(1-k_0))^2\sigma^2 + k_1^2 \\ &\underset{(c)}{=} \lambda^2 k_0^2 \sigma^2 + (1-k_1')^2\sigma^2 + k_1^2 \underset{(d)}{>} (1-k_1')^2\sigma^2 = J_a(0, k_1'), \end{aligned}$$

where (c) follows from the definition of k_1', and (d) follows from the assumption that $k_0 < 0$ and the fact that $0 < k_1' \le k_1$. Additionally, the feasibility of the policy (k_0, k_1) implies that

$$c \ge |k_0 + k_1(1-k_0)|\sigma + |k_1| \ge |k_1'|\sigma + |k_1'|,$$

where the second inequality follows from the fact that $0 < k_1' \leq k_1$. This verifies that the policy $(0, k_1')$ is feasible, which completes our proof of the claim that $k_0^* \geq 0$.

Part 1.2: We prove by contradiction that $k_0^* \leq \min\{c/\sigma, 1\}$. Specifically, let (k_0, k_1) be a feasible linear control policy for problem (2.13), where $k_0 > \min\{c/\sigma, 1\}$. We show that there exists another feasible linear control policy for problem (2.13) whose cost is strictly lower than $J_a(k_0, k_1)$. Consider the following cases.

Case 1: $c > \sigma$. In this case, it is straightforward to verify that $J_a(k_0, k_1) > J_a(1, 0)$, and that the linear control policy $(1, 0)$ is feasible.

Case 2: $c \leq \sigma$. It is straightforward to verify that $k_1 \neq 0$. Set $k_0' := k_0 + k_1(1 - k_0)$. The feasibility of the policy (k_0, k_1) implies that

$$c \geq |k_0 + k_1(1 - k_0)|\sigma + |k_1| > |k_0 + k_1(1 - k_0)|\sigma = |k_0'|\sigma$$

where the second inequality follows from the fact that $k_1 \neq 0$. This implies that the linear control policy $(k_0', 0)$ is feasible. Additionally, we have that

$$\begin{aligned} J_a(k_0, k_1) &= \lambda^2 k_0^2 \sigma^2 + (1 - k_0 - k_1(1 - k_0))^2 \sigma^2 + k_1^2 \\ &= \lambda^2 k_0^2 \sigma^2 + (1 - k_0')^2 \sigma^2 + k_1^2 > \lambda^2 {k_0'}^2 \sigma^2 + (1 - k_0')^2 \sigma^2, \end{aligned}$$

where the last inequality follows from the fact that $k_0 > c/\sigma \geq |k_0'|$. Consequently, we have that $J_a(k_0, k_1) > J_a(k_0', 0)$, where the linear control policy $(k_0', 0)$. This finishes the proof that $k_0^* \leq \min\{c/\sigma, 1\}$.

Part 2: Specification of optimal control gains. We first specify the optimal control gain for controller-1 given a fixed feedback control gain of controller-0. Assume that the feedback control gain of controller-0 satisfies $0 \leq k_0 \leq$

$\min\{c/\sigma, 1\}$. We fix k_0, and optimize over the feedback control gain k_1 of controller-1. This amounts to the following semi-infinite program

$$\begin{aligned} &\text{minimize} && \mathbf{E}[(x_1 - u_1)^2] \\ &\text{subject to} && \left.\begin{aligned} &u_1 = k_1(x_1 + z), \\ &x_1 = (1 - k_0)x_0, \\ &|u_1 + k_0 x_0| \leq c \end{aligned}\right\} \forall (x_0, z) \in W, \end{aligned} \tag{2.17}$$

where the decision variable is the feedback control gain k_1. Problem (2.17) is a convex program. Its optimal solution is given by

$$k_1^* = \min\left\{\frac{(\sigma - q_0)^2}{(\sigma - q_0)^2 + 1}, \frac{c - q_0}{\sigma - q_0 + 1}\right\},$$

where $q_0 := k_0\sigma$. We denote the optimal value of problem (2.17) by $J_1^*(q_0)$. It is given by

$$J_1^*(q_0) = \begin{cases} \dfrac{(\sigma - q_0)^2}{(\sigma - q_0)^2 + 1} & \text{if } \dfrac{c - \sigma - 1}{\sigma - q_0 + 1} + \dfrac{1}{(\sigma - q_0)^2 + 1} \geq 0, \\ \dfrac{(\sigma - q_0)^2(\sigma - c + 1)^2 + (c - q_0)^2}{(\sigma - q_0 + 1)^2} & \text{otherwise.} \end{cases}$$

It follows that the optimal q_0 is the minimizer of the function $\lambda^2 q_0^2 + J_1^*(q_0)$ over the interval $[0, \min\{c, \sigma\}]$. This completes the proof of Lemma 3. ■

2.2.3 Control Design via Assume-Guarantee Contracts

In the previous subsection, our structural characterization of the optimal linear control policy of problem (2.13) reveals a mechanism of implicit coordination between controllers. Such a mechanism amounts to the introduction of an *assume-guarantee contract* between controllers [4, 67]. Namely, each controller *assumes* that the coupling states and inputs of other controllers behave as disturbances with bounded support, and the decentralized control policy is constrained in a manner that *guarantees* the consistency between the assumed and actual behavior of the coupling states and inputs. In spite of its simplicity, this mechanism will enable a convex inner approximation for a large family of constrained decentralized control design problems. In this subsection, we restrict our attention to problem (2.13), and leverage on the aforementioned mechanism of implicit coordination to derive a family of convex inner approximations.

The specific derivation of the convex inner approximation entails two steps of reasoning. First, controller-1 assumes that the input u_0 and the state x_1 behave as disturbances that are uniformly distributed on their support $[-\alpha, \alpha] \times [-\beta, \beta]$. Second, in order to guarantee the satisfaction of such a modeling assumption, we impose an additional constraint on the controller-0's input, which requires that it induces control input u_0 and state x_1 that are guaranteed to belong to the rectangle $[-\alpha, \alpha] \times [-\beta, \beta]$. This gives rise to the following decentralized control design problem.

$$\begin{array}{ll} \text{minimize} & \mathbf{E}[\lambda^2 u_0^2 + x_2^2] \\ \text{subject to} & \left.\begin{array}{l} u_0 = k_0 x_0 \\ u_1 = k_1(\widetilde{x}_1 + z) \\ x_1 = x_0 - u_0, \\ x_2 = \widetilde{x}_1 - u_1, \\ |u_0| \leq \alpha \\ |x_1| \leq \beta \\ |\widetilde{u}_0 + u_1| \leq c \end{array}\right\} \forall (x_0, z, \widetilde{u}_0, \widetilde{x}_1) \in W \times [-\alpha, \alpha] \times [-\beta, \beta], \end{array} \tag{2.18}$$

where the decision variables are the feedback control gains k_0, k_1. For each pair of parameters $(\alpha, \beta) \in \mathbf{R}_+^2$, problem (2.18) is a convex program. Additionally, it is straightforward to show that problem (2.18) is an inner approximation to problem (2.13). We have the following lemma.

Lemma 4. Let $\alpha, \beta \in \mathbf{R}_+$ be a fixed parameter. It follows that the decentralized control policy

$$u_i = k_i y_i, \quad i = 0, 1$$

is feasible for problem (2.13) if (k_0, k_1) is a feasible solution to problem (2.18).

Proof: It suffices to show that $|u_0 + k_1(x_1 + z)| \leq c$ for all $(x_0, z) \in W$. The feasibility of (k_0, k_1) for problem (2.18) implies that

$$|\widetilde{u}_0 + k_1(\widetilde{x}_1 + z)| \leq c \qquad \forall (z, \widetilde{u}_0, \widetilde{x}_1) \in [-1, 1] \times [-\alpha, \alpha] \times [-\beta, \beta], \tag{2.19}$$

$$|u_0| \leq \alpha \qquad \forall x_0 \in [-\sigma, \sigma], \tag{2.20}$$

$$|x_1| \leq \beta \qquad \forall x_0 \in [-\sigma, \sigma]. \tag{2.21}$$

The desired result follows, as (2.20) and (2.21) imply that $(z, u_0, x_1) \in [-1, 1] \times [-\alpha, \alpha] \times [-\beta, \beta]$ for all $(x_0, z) \in W$. In combination with (2.19), this implies that $|u_0 + k_1(x_1 + z)| \leq c$ for all $(x_0, z) \in W$. ■

Problem (2.18) provides an example on how the introduction of *assume-guarantee contracts* enables the tractable calculation of feasible control policies for decentralized control problems with *nonclassical* information structures. In Chapter 4, we formally describe the convex inner approximation of a general family of decentralized control design problems that is built on this conceptual idea.

Additionally, we note that the performance of the controller that we obtain depends heavily on the specification of the contract. In particular, the solution of problem (2.18) is exactly the optimal linear control policy for problem (2.13) if $\alpha = k_0^*\sigma$ and $\beta = (1 - k_0^*)\sigma$, where k_0^* is controller-0's optimal linear control policy in problem (2.13). Consequently, there is a need to co-design the assume-guarantee contract with the underlying decentralized control policy. Such a problem is formally visited in Chapter 4, in which we derive a method of co-optimize the location, scale, and orientation of the contract set with the underlying decentralized control policy via the solution of a semidefinite program.

CHAPTER 3

CONVEXITY OF DECENTRALIZED CONTROL DESIGN PROBLEMS

In this chapter, we define the constrained decentralized output-feedback control design problem that we consider in this dissertation, and discuss how its convexity depends on the underlying information structure. Specifically, we consider the decentralized output-feedback control of a discrete-time, linear system subject to exogenous disturbances and polyhedral constraints on the state and input trajectories. The underlying system is composed of a finite collection of dynamically coupled subsystems, where each subsystem is assumed to have a dedicated local controller. The decentralization of information is expressed according to sparsity constraints on the information that each local controller has access to. We show that the classical Youla parameterization yields an equivalent reformulation of the decentralized control design problem as a convex program if and only if the information structure is *partially nested*—that is, an information structure in which each controller's local information cannot be affected by control inputs it does not have access to. Given a partially nested information structure, we show that the calculation of the optimal affine control policy amounts to a robust convex program and admits an equivalent reformulation as a finite-dimensional conic program.

3.1 Problem Formulation

3.1.1 System Model

Consider a discrete-time, linear time-varying system consisting of N dynamically coupled subsystems whose dynamics are described by

$$x_i(t+1) = \sum_{j=1}^{N} \Big(A_{ij}(t)x_j(t) + B_{ij}(t)u_j(t) \Big) + G_i(t)w(t), \tag{3.1}$$

for $i = 1, \ldots, N$. The system operate for finite time indexed by $t = 0, \ldots, T-1$, and the initial condition is assumed fixed and known. We associate with each subsystem i a *local state* $x_i(t) \in \mathbf{R}^{n_x^i}$ and *local input* $u_i(t) \in \mathbf{R}^{n_u^i}$. And we denote by $w(t) \in \mathbf{R}^{n_w}$ the stochastic *system disturbance.* We denote by $y_i(t) \in \mathbf{R}^{n_y^i}$ the *local measured output* of subsystem i at time t. It is given by

$$y_i(t) = \sum_{j=1}^{N} C_{ij}(t)x_j(t) + H_i(t)w(t), \tag{3.2}$$

for $i = 1, \ldots, N$. All system matrices are assumed to be real and of compatible dimension. In the sequel, it will be convenient to work with a more compact representation of the system Eqs. (3.1) and (3.2) given by

$$\begin{aligned} x(t+1) &= A(t)x(t) + B(t)u(t) + G(t)w(t) \\ y(t) &= C(t)x(t) + H(t)w(t). \end{aligned}$$

Here, we denote by $x(t) := (x_1(t), .., x_N(t)) \in \mathbf{R}^{n_x}$, $u(t) := (u_1(t), .., u_N(t)) \in \mathbf{R}^{n_u}$, and $y(t) := (y_1(t), .., y_N(t)) \in \mathbf{R}^{n_y}$ the full system state, input, and output at time t, respectively. Their corresponding dimensions are given by $n_x := \sum_{i=1}^{N} n_x^i$, $n_u := \sum_{i=1}^{N} n_u^i$, and $n_y := \sum_{i=1}^{N} n_y^i$. We will occasionally refer to the tuple

$$\Theta := \{A(t), B(t), G(t), C(t), H(t)\}_{t=0}^{T-1}$$

as the system parameter when making reference to the underlying system described by Eqs. (3.1) and (3.2). Additionally, we denote by $x, u, w,$ and y the trajectories of the full system state, input, disturbance, and output, respectively. They are specified according to

$$
\begin{aligned}
x &:= (x(0), \ldots, x(T)) \in \mathbf{R}^{N_x}, & N_x &:= n_x(T+1), \\
u &:= (u(0), \ldots, u(T-1)) \in \mathbf{R}^{N_u}, & N_u &:= n_u T, \\
w &:= (1, w(0), \ldots, w(T-1)) \in \mathbf{R}^{N_w}, & N_w &:= 1 + n_w T \\
y &:= (1, y(0), \ldots, y(T-1)) \in \mathbf{R}^{N_y}, & N_y &:= 1 + n_y T.
\end{aligned}
$$

Notice that in our specification of the both the disturbance and output trajectories, w and y, we have extended each trajectory to include a constant scalar as its initial component. This notational convention will prove useful in simplifying the specification of affine control policies in the sequel. The system trajectories are related according to

$$x = Bu + Gw \quad \text{and} \quad y = Cx + Hw.$$

Here, the block matrices (B, G, C, H) are given by:

$$B := \begin{bmatrix} 0 & & & & \\ A_1^1 B(0) & 0 & & & \\ A_1^2 B(0) & A_2^2 B(1) & 0 & & \\ \vdots & & & \ddots & \\ \vdots & & & & 0 \\ A_1^T B(0) & A_2^T B(1) & \cdots & \cdots & A_T^T B(T-1) \end{bmatrix}$$

$$G := \begin{bmatrix} A_0^0 x(0) & & & & \\ A_0^1 x(0) & A_1^1 G(0) & & & \\ A_0^2 x(0) & A_1^2 G(0) & A_2^2 G(1) & & \\ \vdots & \vdots & & \ddots & \\ A_0^T x(0) & A_1^T G(0) & A_2^T G(1) & \cdots & A_T^T G(T-1) \end{bmatrix}$$

$$C := \begin{bmatrix} 0 & & & \\ C(0) & 0 & & \\ & \ddots & \ddots & \\ & & C(T-1) & 0 \end{bmatrix} \qquad H := \text{diag}(1, H(0), \ldots, H(T-1)),$$

where $A_s^t := \prod_{r=s}^{t-1} A(r)$ for $s < t$, and $A_t^t = I$.

We close this subsection by stating a structural assumption on the system dynamics. Assumption 1, which is assumed to hold throughout the dissertation, ensures that each subsystem's local control input can causally affect its local measured output.

Assumption 1. For each subsystem $i \in \mathcal{V}$, there exist time periods $0 \leq s < t \leq T-1$ such that the matrix $\left[C(t) A_{s+1}^t B(s)\right]_{ii}$ is nonzero.

Here, the matrix $\left[C(t)A_{s+1}^{t}B(s)\right]_{ii}$ refers to the $(i,i)^{\text{th}}$ block of the $N \times N$ block matrix $C(t)A_{s+1}^{t}B(s)$.

3.1.2 Disturbance Model

We model the disturbance trajectory w as a random vector defined according to the probability space $(\mathbf{R}^{N_w}, \mathcal{B}(\mathbf{R}^{N_w}), \mathbf{P})$. Here, the Borel σ-algebra $\mathcal{B}(\mathbf{R}^{N_w})$ denotes the set of all events that are assigned probability according to the measure $\mathbf{P}$. We denote by $\mathcal{L}_n^2 := \mathcal{L}^2(\mathbf{R}^{N_w}, \mathcal{B}(\mathbf{R}^{N_w}), \mathbf{P}; \mathbf{R}^n)$ the space of all $\mathcal{B}(\mathbf{R}^{N_w})$-measurable, square-integrable random vectors taking values in $\mathbf{R}^n$. Also, we use $\mathbf{E}$ to denote expectation taken with respect to the probability measure $\mathbf{P}$. With a slight abuse of notation, we occasionally use w to denote a realization of the random vector w.

In order to ensure the well-posedness of the problem to follow, we require that the disturbance trajectory satisfy the following conditions. First, we assume that the disturbance trajectory w has support $\mathcal{W}$ that is a nonempty and compact subset of $\mathbf{R}^{N_w}$, representable by

$$\mathcal{W} = \{w \in \mathbf{R}^{N_w} \mid w_1 = 1 \text{ and } L_k w \succeq_{\mathcal{K}} 0,\ k = 1, \ldots, \ell\},$$

where $\mathcal{K} \subseteq \mathbf{R}^{N_w}$ is a proper cone and the matrices $L_k \in \mathbf{R}^{N_w \times N_w}$ are assumed to be given for $k = 1, \ldots, \ell$. In addition to compactness, we require that the linear hull of $\mathcal{W}$ spans $\mathbf{R}^{N_w}$. Such assumption is without loss of generality. If it were not the case, then there exists a linear transformation of the disturbance trajectory w, which maps w to a lower dimensional subspace of $\mathbf{R}^{N_w}$ for which this assumption is satisfied. And, it is straightforward to verify that such assumption ensures that the corresponding second-order moment matrix, defined as

$M := \mathbf{E}\left[ww^\mathsf{T}\right]$, is both invertible and positive definite. The fact that the second-order moment matrix M is finite-valued is a consequence of our assumption that the disturbance have compact support.

3.1.3 System Constraints

In characterizing the set of feasible input trajectories, we require that the input and state trajectories respect the following linear inequality constraints $\mathbf{P}$-almost surely,

$$\left.\begin{aligned} &F_x x + F_u u + F_w w + s = 0 \\ &s \geq 0 \end{aligned}\right\} \; \mathbf{P}\text{-a.s.} \tag{3.3}$$

where $F_x \in \mathbf{R}^{m \times N_x}$, $F_u \in \mathbf{R}^{m \times N_u}$, and $F_w \in \mathbf{R}^{m \times N_w}$. Here, $s \in \mathcal{L}^2_m$ is a slack variable that is required to be non-negative $\mathbf{P}$-almost surely. The almost sure constraints in (3.3) requires that the input and the resulting state trajectory satisfy the corresponding linear constraints for all possible realizations of the disturbance, except for a set of probability equal to zero. Under the disturbance model we consider, this amounts to requiring the satisfaction of an infinite number of linear constraints.

3.1.4 Decentralized Control Design

We consider information structures that are specified via *sparsity constraints* on the information that each controller has access to. More specifically, we describe the pattern according to which information is shared between subsystems with a directed graph $\mathcal{G}_I = (\mathcal{V}, \mathcal{E}_I)$, which we refer to as the *information graph* of the system. Here, the node set $\mathcal{V} = \{1, \ldots, N\}$ assigns a distinct node i to each subsystem i, and the directed edge set $\mathcal{E}_I$ determines the pattern of information sharing between subsystems. More precisely, we let $(i, j) \in \mathcal{E}_I$ if and only if for each time t, subsystem j has access to subsystem i's local output $y_i(t)$. We let $\mathcal{V}_I^-(i)$ denote the in-neighborhood of each subsystem $i \in \mathcal{V}$ according to the information graph $\mathcal{G}_I$. We make the following assumption on the structure of the information graph, which ensures that each subsystem i has access to its local output $y_i(t)$ at each time period t.

Assumption 2. The directed edge set $\mathcal{E}_I$ is assumed to contain the self-loop (i, i) for each $i \in \mathcal{V}$.

We also assume that each subsystem has *perfect recall*, i.e., each subsystem has access to its entire history of past information at any given time. Accordingly, we define the *local information* available to each subsystem i at time t as

$$z_i(t) := \{y_j^t \mid (j, i) \in \mathcal{E}_I\}. \tag{3.4}$$

We restrict the local input to subsystem i to be of the form

$$u_i(t) = \gamma_i(z_i(t), t), \tag{3.5}$$

where $\gamma_i(\cdot, t)$ is a measurable function of the local information $z_i(t)$. We define the *local control policy* for subsystem i as $\gamma_i := (\gamma_i(\cdot, 0), \ldots, \gamma_i(\cdot, T-1))$. We refer to the collection of local control policies $\gamma := (\gamma_1, \ldots, \gamma_N)$ as the *decentralized control policy* and define $\Gamma(\mathcal{G}_I)$ as the set of all decentralized control policies respecting the information structure defined by the information graph $\mathcal{G}_I$.

Throughout the dissertation, we consider the following family of constrained decentralized control design problems:

$$\begin{aligned} &\text{minimize} && \mathbf{E}\left[x^{\mathsf{T}} R_x x + u^{\mathsf{T}} R_u u\right] \\ &\text{subject to} && \gamma \in \Gamma(\mathcal{G}_I),\ s \in \mathcal{L}^2_m \\ & && \left.\begin{aligned} &F_x x + F_u u + F_w w + s = 0 \\ &x = Bu + Gw \\ &y = Cx + Hw \\ &u = \gamma(y) \\ &s \geq 0 \end{aligned}\right\} \mathbf{P}\text{-a.s.} \end{aligned} \tag{3.6}$$

Here, the *cost matrices* $R_x \in \mathbf{R}^{N_x \times N_x}$ and $R_u \in \mathbf{R}^{N_u \times N_u}$ are both assumed to be symmetric and positive semidefinite. The tractability of the decentralized control design problem (3.6) is known to depend critically on its information structure. In particular, if the information structure if *partially nested*, then problem (3.6) can be equivalently reformulated as a convex program via the Youla parameterization that we introduce in Section 3.2.1. If, on the other hand, the information structure is nonclassical (i.e., not partially nested), then problem (3.6) amounts to an infinite-dimensional nonconvex program that is known to be computationally intractable, in general [76, 89, 101]. In this dissertation, we develop tractable methods of computing feasible control policies for decentralized control design problems whose information structures are allowed to be nonclassical. As the decentralized control policies we compute are, in general, suboptimal, we provide a method of tractably computing a bound on their suboptimality via an information-based convex relaxation.

3.2 Youla Parameterization and Convexity

In what follows, we describe how to equivalently reformulate the decentralized control design problem (3.6) as a static team problem [56] through a nonlinear change of variables akin to the Youla parameterization. This reformulation is shown to result in a convex program if and only if the underlying information structure is partially nested.

3.2.1 Nonlinear Youla Parameterization

As stated, the decentralized control design problem (3.6) is nonconvex in the decentralized control policy γ. In what follows, we define the nonlinear Youla parameterization [112] of the decentralized control policy by parameterizing the input process as a causal function of the so-called *purified output process*. This yields an equivalent reformulation of the decentralized control design problem, in which the only source of nonconvexity is the potential nonconvexity in the set of Youla parameters.

We first introduce the concept of output purification as defined in [16–18]. Given an input process $u(t)$ and the corresponding output process $y(t)$, define the sequence of purified outputs $\eta(t)$ according to

$$\begin{aligned}\bar{x}(0) &= 0,\\ \bar{x}(t+1) &= A(t)\bar{x}(t) \;+\; B(t)u(t),\\ \bar{y}(t) &= C(t)\bar{x}(t),\\ \eta(t) &= y(t) \;-\; \bar{y}(t),\end{aligned}$$

for $t = 0, \ldots, T-1$. Similar to the definition of the *local information* in (3.4), we define the *local purified information* available to each subsystem i at time t as

$$\zeta_i(t) := \{\eta_j^t \mid (j,i) \in \mathcal{E}_I\}. \tag{3.7}$$

In addition, we define the trajectory of the purified output according to $\eta := (1, \eta(0), \ldots, \eta(T-1)) \in \mathbf{R}^{N_y}$. It is straightforward to establish the following relation, which reveals the purified output trajectory η to be independent of the input trajectory u. Namely,

$$\eta = Pw,$$

where $P := (CG + H) \in \mathbf{R}^{N_y \times N_w}$.

An important property of the purified output process is that it is a causal and invertible function of the output process. As a result, one can equivalently reparameterize the input trajectory as a causal function of the purified output process via a nonlinear Youla parameterization. Specifically, define the nonlinear Youla parameterization of the decentralized control policy $\gamma \in \Gamma(\mathcal{G}_I)$ as

$$\phi := \gamma \circ (I - CB\gamma)^{-1}. \tag{3.8}$$

Note that the map $I - CB\gamma : \mathbf{R}^{N_y} \to \mathbf{R}^{N_y}$ is guaranteed to be invertible, as the decentralized control policy γ is causal, and the matrix CB is strictly block lower triangular. The Youla parameter ϕ satisfies the following two important properties. First, it is an invertible function of the policy γ over $\Gamma(\mathcal{G}_I)$, where its inverse is given by $\gamma = \phi \circ (I + CB\phi)^{-1}$. Note that the required inverse exists, as it is straightforward to verify that $I + CB\phi = (I - CB\gamma)^{-1}$. Second, given an input trajectory induced by $u = \gamma(y)$, it holds that

$$\phi(\eta) = \gamma(y) \tag{3.9}$$

for every disturbance trajectory $w \in \mathcal{W}$. Note that Eq. (3.9) follows from the fact that the output trajectory y and purified output trajectory η are related according to

$$y = CB\gamma(y) + \eta,$$

which in turn implies that $y = (I - CB\gamma)^{-1}(\eta)$.

Together, these two properties reveal that problem (3.6) can be equivalently reformulated as a static team problem by applying the nonlinear change of variables in (3.8). This yields the following optimization problem:

$$\begin{array}{ll} \text{minimize} & \mathbf{E}\left[x^\mathsf{T} R_x x + u^\mathsf{T} R_u u\right] \\ \text{subject to} & \phi \in \Phi(\mathcal{G}_I),\ s \in \mathcal{L}^2_m \\ & \left.\begin{array}{l} F_x x + F_u u + F_w w + s = 0 \\ x = Bu + Gw \\ \eta = Pw \\ u = \phi(\eta) \\ s \geq 0 \end{array}\right\} \mathbf{P}\text{-a.s.} \end{array} \tag{3.10}$$

Here, the set of admissible Youla parameters is given by

$$\Phi(\mathcal{G}_I) := \{\gamma \circ (I - CB\gamma)^{-1} \mid \gamma \in \Gamma(\mathcal{G}_I)\}.$$

The only potential source of nonconvexity in problem (3.10) is in the set of Youla parameters $\Phi(\mathcal{G}_I)$. In particular, problem (3.10) is a convex program if and only if the set $\Phi(\mathcal{G}_I)$ is convex.

3.2.2 Convexity under Partially Nested Information Structures

In what follows, we show that the static team problem (3.10) is a convex program if and only if the information structure is partially nested. Before proceeding, we provide a formal definition of partially nested information structures using the notion of precedence, as defined by Ho and Chu in [57].

Definition 1 (Precedence)**.** Given the information structure defined by $\mathcal{G}_I$, we say subsystem j is a *precedent* to subsystem i, denoted by $j \prec i$, if there exist times $0 \leq s < t \leq T-1$ and subsystem k satisfying $(k, i) \in \mathcal{E}_I$, such that $\left[C(t)A_{s+1}^{t}B(s)\right]_{kj} \neq 0$.

Essentially, subsystem j is a *precedent* to subsystem i, if the local input to subsystem j can affect the local information available to subsystem i at some point in the future. In particular, it follows from Assumption 1 that j is a precedent to i if $(j, i) \in \mathcal{E}_I$. Equipped with the concept of precedence, we now provide the definition of partially nested information structures.

Definition 2 (Partially Nested Information)**.** The information structure defined by $\mathcal{G}_I$ is said to be *partially nested* with respect to the system Θ, if $j \prec i$ implies that $z_j(t) \subseteq z_i(t)$ for all times $t = 0, \ldots, T-1$.

We denote by PN(Θ) the set of information graphs that are partially nested with respect to the the system Θ. The information structure defined by $\mathcal{G}_I$ is said to be *nonclassical* if $\mathcal{G}_I \notin$ PN(Θ). We note that the above definition of partial nestedness is tailored to the setting in which controllers are subject to sparsity constraints on the measured outputs that each controller can access. A more general definition of partial nestedness can be found in [50, 56, 57], which applies to the setting in which controllers are subject to both delay and sparsity constraints on information sharing. An important consequence of the partial nestedness of the information structure is that it guarantees that the local information $z_i(t)$ and the purified local information $\zeta_i(t)$ contain the same information for each subsystem i and time t. In other words, they generate the same σ-algebra. We have the following lemma.

Lemma 5. Let $\gamma \in \Gamma(\mathcal{G}_I)$ be any decentralized control policy. If $\mathcal{G}_I \in \text{PN}(\Theta)$, then the local information $z_i(t)$ and local purified information $\zeta_i(t)$ are functions of each other for each subystem $i = 1, \ldots, N$ and time $t = 0, \ldots, T-1$.

Proof: We prove (by induction in t) that $z_i(t)$ and $\zeta_i(t)$ are functions of each other for all $i = 1, \ldots, N$ and $t = 0, \ldots, T-1$ if $\mathcal{G}_I \in \text{PN}(\Theta)$.

Base step: For $t = 0$, we have that $y(0) = \eta(0)$, so the above claim is true.

Induction step: Assume that the claim is true for times $0, \ldots, t-1$. We now show that the claim is also true for time t. For the sake of brevity, we only show that $z_i(t)$ is a function of $\zeta_i(t)$ for all $i \in \{1, \ldots, N\}$. The proof that $\zeta_i(t)$ is a function of $z_i(t)$ for all $i \in \{1, \ldots, N\}$ follows from identical arguments. Recall that

$$z_i(t) = \{y_j^t \mid j \in N_{\mathcal{G}_I}^-(i)\}.$$

It suffices to show that $y_k(t)$ is a function of $\zeta_i(t)$ for all $i \in \{1, \ldots, N\}$ and $k \in N_{\mathcal{G}_I}^-(i)$. Using the definition of precedents, one can write $y_k(t)$ as

$$y_k(t) = \left(\sum_{s=0}^{t-1} \sum_{j:\, j \prec i} \left[C(t) A_{s+1}^t B(s) \right]_{kj} u_j(s) \right) + \eta_k(t),$$

for all $i \in \{1, \ldots, N\}$ and $k \in N_{\mathcal{G}_I}^-(i)$. To show that $y_k(t)$ is a function of $\zeta_i(t)$, it suffices to show that $u_j(s)$ is a function of $\zeta_i(s)$ for all $j \prec i$ and $s \leq t-1$. By construction, it holds that $u_j(s)$ is a function of $z_j(s)$. It also holds that $z_j(s) \subseteq z_i(s)$, since $j \prec i$ and $\mathcal{G}_I \in \text{PN}(\Theta)$. It follows that $u_j(s)$ is a function of $z_i(s)$. Moreover, by the induction hypothesis, $z_i(s)$ is a function of $\zeta_i(s)$ for all $s \leq t-1$. It follows that $u_j(s)$ is a function of $\zeta_i(s)$ for all $j \prec i$ and $s \leq t-1$. It follows that $y_k(t)$ is a function of $\zeta_i(t)$ for all $i \in \{1, \dots, N\}$ and $k \in N^-_{\mathcal{G}_I}(i)$, thus completing the induction step of the proof. ■

As a result of Lemma 5, one can equivalent reparameterize each subsystem's control input in the purified local information without loss of optimality if the information structure is partially nested. That is to say, if $u_i(t) = \gamma_i(z_i(t), t)$ for a control policy $\gamma \in \Gamma(\mathcal{G}_I)$, then there exist another decentralized control policy $\phi \in \Gamma(\mathcal{G}_I)$, such that $u_i(t) = \phi_i(\zeta_i(t), t)$ for each subsystem i and time t, and vice versa. In other words, the partial nestedness of the information structure implies that $\Gamma(\mathcal{G}_I) = \Phi(\mathcal{G}_I)$. In the following lemma, we further strengthen this result by showing that the set of Youla parameters $\Phi(\mathcal{G}_I)$ is convex if and only if the information structure is partially nested. We omit the proof of Lemma 6, as it directly follows from existing arguments in [86, Thm. 1] and [69, Cor. 7].

Lemma 6. The following statements are equivalent:

(i) $\Phi(\mathcal{G}_I)$ is a convex set,

(ii) $\Phi(\mathcal{G}_I) = \Gamma(\mathcal{G}_I)$,

(iii) $\mathcal{G}_I \in \text{PN}(\Theta)$.

Lemma 6 implies Ho and Chu's classical result [56, Thm. 1] showing that a dynamic team problem with a partially nested information structure can be equivalently reformulated as a static team problem with the same set of admissible policies. It follows from Lemma 6 that the reformulated decentralized control problem in (3.10) is convex if and only if the underlying information structure is partially nested.[1]

3.3 Affine Control Design under Partially Nested Information Structures

In this section, we restrict our attention to finite-dimensional decentralized control policies that are *affine* in the measured output, and explore the extent to which the partial nesting of information might facilitate the efficient optimization over such restricted class of policies. We demonstrate how powerful techniques for centralized affine control design [16, 54, 90, 104, 107] can be extended to decentralized systems to compute optimal affine decentralized policies via finite-dimensional convex optimization.

[1]We note that the convexity result in Lemma 6 does not depend on the structure of the cost matrices or the probability distribution of system disturbance. There is a related literature, which identifies structural conditions on the system and cost matrices and the probability distribution of system disturbance, under which the communication of private information from any controller's precedent to said controller does not lead to a reduction in cost. Under these conditions, the optimal solution of problem (3.10) can be computed via the solution of a convex program when the information structure is nonclassical. See [5, 113–115] for recent advances.

3.3.1 Decentralized Affine Controllers

Given an information structure $\mathcal{G}_I$, we consider affine control policies of the form

$$u_i(t) = \bar{u}_i(t) + \sum_{s=0}^{t} \sum_{j \in \mathcal{V}_I^-(i)} K_{ij}(t,s) y_j(s), \tag{3.11}$$

for each subsystem $i = 1, \cdots, N$ and time $t = 0, \cdots, T-1$. Here, $\bar{u}_i(t) \in \mathbf{R}^{n_u^i}$ represents the open-loop component of the control and $K_{ij}(t,s) \in \mathbf{R}^{n_u^i \times n_y^j}$ the feedback control gain. One can lift the representation in (3.11) to relate the output trajectory y to the input trajectory u under the linear map

$$u = Ky, \quad \text{where} \quad K \in S(\mathcal{G}_I).$$

Here, we define $S(\mathcal{G}_I)$ to be the subspace of causal (lower block triangular) matrices respecting the information structure defined by $\mathcal{G}_I$. That is, for any $K \in S(\mathcal{G}_I)$, the decentralized control policy defined by $\gamma(y) = Ky$ satisfies $\gamma \in \Gamma(\mathcal{G}_I)$. We have the following decentralized affine control design problem:

$$\begin{aligned}
\text{minimize} \quad & \mathbf{E}\left[x^\mathsf{T} R_x x + u^\mathsf{T} R_u u\right] \\
\text{subject to} \quad & K \in S(\mathcal{G}_I) \\
& \left.\begin{aligned} & F_x x + F_u u + F_w w \leq 0 \\ & x = Bu + Gw \\ & y = Cx + Hw \\ & u = Ky \end{aligned}\right\} \forall\, w \in \mathcal{W}.
\end{aligned} \tag{3.12}$$

The affine control design problem (3.12) is known to be nonconvex in the matrix variable K [17, 54, 90]. However, under the additional assumption of partially nested information structure, problem (3.12) admits an equivalent reformulation as a semi-infinite convex program via the Youla parameterziation we introduced in Section 3.2.1. Specifically, such an equivalent reformulation relies on the following lemma.

Lemma 7. If $\mathcal{G}_I \in \text{PN}(\Theta)$, then both of the following statements are true.

(i) Let $K \in S(\mathcal{G}_I)$ and define $Q = K(I - CBK)^{-1}$. Then $Q \in S(\mathcal{G}_I)$ and $Q\eta = Ky$ for all $w \in \mathcal{W}$.

(ii) Let $Q \in S(\mathcal{G}_I)$ and define $K = Q(I + CBQ)^{-1}$. Then $K \in S(\mathcal{G}_I)$ and $Ky = Q\eta$ for all $w \in \mathcal{W}$.

Proof: We only prove part (ii) of the lemma, as part (i) immediately follows from Lemma 6 and Eq. (3.9). Lemma 6 implies that $\Phi(\mathcal{G}_I) = \Gamma(\mathcal{G}_I)$ if $\mathcal{G}_I \in \text{PN}(\Theta)$. In combination with our assumption that $\mathcal{G}_I \in \text{PN}(\Theta)$, this implies that $S(\mathcal{G}_I) \subseteq \Phi(\mathcal{G}_I)$. Recall that the decentralized control policy γ is related to the Youla parameter ϕ according to

$$\gamma = \phi \circ (I + CB\phi)^{-1}.$$

It follows that $K \in \Gamma(\mathcal{G}_I)$, which further implies that $K \in S(\mathcal{G}_I)$. Moreover, the combination of Eq. (3.9) and the fact that $K \in \Gamma(\mathcal{G}_I)$ implies that $Ky = Q\eta$ for all $w \in \mathcal{W}$. ■

Lemma 7 builds on Lemma 6 to reveal that if the information structure is partially nested, then any decentralized affine output feedback controller $K \in S(\mathcal{G}_I)$ can be transformed to an *equivalent* decentralized affine purified output feedback controller $Q \in S(\mathcal{G}_I)$ through an invertible nonlinear transformation, and vice versa. Consequently, one can apply the change of variable specified in Lemma 7 to equivalently reformulate problem (3.12) as a semi-infinite convex program. We have the following result.

Proposition 2. Let Q^* be an optimal solution to the following optimization problem,

$$\begin{array}{ll} \text{minimize} & \mathbf{E}\left[x^\mathsf{T} R_x x + u^\mathsf{T} R_u u\right] \\ \text{subject to} & Q \in S(\mathcal{G}_I) \\ & \left.\begin{array}{l} F_x x + F_u u + F_w w \leq 0 \\ x = Bu + Gw \\ u = QPw \end{array}\right\} \forall\, w \in \mathcal{W}. \end{array} \tag{3.13}$$

Then $K^* = (I + Q^* CB)^{-1} Q^*$ is an optimal solution to problem (3.12).

Proof: It follows from Lemma 7 that for any $Q \in S(\mathcal{G}_I)$, $K = Q(I + CBQ)^{-1}$ satisfies $K \in S(\mathcal{G}_I)$. Let Q^* be an optimal solution to problem (3.13). It follows from Lemma 7 that $K^* = Q^*(I + CBQ^*)^{-1}$ is the affine output feedback controller that results in the same sequence of control inputs as the affine purified output feedback controller Q^*. It follows that K^* is the optimal affine output feedback controller. ■

3.3.2 Conic Programming Reformulation

In the absence of constraints on the state and input trajectories (i.e., $F_x, F_u, F_w = 0$), problem (3.13) reduces to an unconstrained convex quadratic program. Problem (3.13) is in general, however, a semi-infinite convex quadratic program, as it contains infinitely many linear constraints in Q. Given our assumption that the set $\mathcal{W}$ is described by finitely many conic inequalities, one can use techniques grounded in duality to show that problem (3.13) admits an equivalent reformulation as a finite-dimensional conic optimization problem. The underlying approach relies on arguments analogous to those in [54, 65].

Proposition 3. An optimal solution to problem (3.13) can be obtained by solving the following equivalent finite-dimensional conic optimization problem,

$$\begin{aligned}
\text{minimize} \quad & \operatorname{Tr}\left(P^\top Q^\top R Q P M + 2G^\top R_x B Q P M + G^\top R_x G M\right) \\
\text{subject to} \quad & Q \in S(\mathcal{G}_I),\ Z \in \mathbf{R}^{m\times N_w},\ \ \Lambda_k \in \mathbf{R}^{N_w \times m},\ \ \mu \in \mathbf{R}^m_+ \\
& (F_u + F_x B) Q P + F_x G + F_w + Z = 0, \\
& Z = \mu e_1^\top + \sum_{k=1}^{\ell} \Lambda_k^\top L_k, \\
& \Lambda_k \succeq_{\mathcal{K}^*} 0,\ \ k = 1, \ldots, \ell,
\end{aligned} \tag{3.14}$$

where $R = R_u + B^\top R_x B$, and $e_1 = (1, 0, \ldots, 0)$ is a unit vector in $\mathbf{R}^{N_w}$.

The proof of Proposition 3 amounts to directly applying Lemma 22 in Appendix A.2 to equivalently reformulate the robust linear constraints in problem (3.13) as finite-dimensional conic constraints. It is, therefore, omitted for the brevity of exposition. We remark that the conic optimization problem (3.14) can be efficiently solved for a wide range of cones $\mathcal{K}$, including polyhedral and second-order cones.

Finally, we note that Proposition 3 provides a method of tractably calculating decentralized control policies that are *affine* in the purified output via the solution of a finite dimensional conic program. In principle, one can also generalize the techniques in [21, 51] for the calculation of centralized control policies that are *polynomial* or *piecewise affine* in the purified output to decentralized control design problems with partially nested information structures. In this dissertation, however, we omit a detailed treatment of such generalizations, as the focus of this dissertation is the approximation of decentralized control design problems with nonclassical information structures.

CHAPTER 4

DECENTRALIZED CONTROL DESIGN VIA ASSUME-GUARANTEE CONTRACTS

In this chapter, we investigate the design of decentralized control policies that are affinely parameterized in the state history for decentralized control design problems with arbitrary (possibly nonclassical) information structures. In order to alleviate the nonconvexity arising from the informational coupling between subsystems, we propose an approach to decentralized control design in which the information-coupling states are effectively treated as disturbances whose trajectories are constrained to take values in ellipsoidal "contract" sets . To ensure the satisfaction of this assumption, we impose a contractual constraint on the control policy that "guarantees" that the information-coupling states that it induces belong to said contract set. Naturally, this approach yields an inner approximation of the original decentralized control design problem, where the conservatism of the resulting approximation depends on the specification of the contract set. To limit the extent of the suboptimality that may result, we formulate a semi-infinite program to co-optimize the decentralized control policy with the location, scale, and orientation of the contract set. We establish a structural condition on the space of allowable contracts that facilitates the joint optimization of the control policy and the contract set via semidefinite programming.

4.1 Introduction

We consider a special case of problem (3.6), in which the measured output at each subsystem is given by a perfect measurement of the local state. In systems with nonclassical information structures, the derivation of the optimal affine control policy is, in general, computationally intractable due to the signaling incentive between controller. We, therefore, abandon the search for the optimal decentralized affine control policy. Instead, Our objective is to derive suboptimal decentralized control policies that can be efficiently computed via convex programming methods.

Related Literature: There is a large stream of literature that leverages on techniques derived from centralized model predictive control (MPC) to facilitate the design of decentralized controllers for constrained dynamical systems [15, 26, 40, 41, 44, 55, 60, 64, 75, 82–84, 97–99, 105]. Compared to their centralized counterparts, the additional challenge in decentralized MPC lies in the need to implicitly coordinate the control inputs from all subsystems (without communication) to achieve good performance and guarantee constraint satisfaction for the *full system*. Many earlier works on decentralized MPC achieve such a coordination via the sharing of point forecasts of each subsystem's states and control inputs in the future [15, 26, 60, 64, 105]. The typical approach to the calculation of control inputs in these papers is to first decompose the decentralized control problem into a collection of decoupled local control problems, in which the coupling states and inputs associated with each subsystem's "neighbors" are assumed to be equal to their point forecasts. Each of the resulting local control problems can be subsequently solved using centralized MPC methods. The main drawback of this approach, however, is that it cannot guarantee the satisfaction of input constraints that couple across subsystems and state constraints, as the point forecast of other subsystems' states and inputs will, in general, differ from their true values.

In order to deal with such drawbacks, more recent papers on decentralized

MPC have adopted a "tube-based" decentralized MPC approach [40, 41, 44, 55, 75, 82–84, 97–99]. Specifically, in constructing the aforementioned decomposition of the global decentralized control problem, each of the local controllers, instead, treats the coupling states and inputs as independent exogenous disturbances, which are assumed to take values in the given *state and input constraint sets*. Given the resulting collection of decoupled local control problems, centralized tube-MPC methods can be applied to compute local control policies that are guaranteed to be feasible for each sub-problem. Although decentralized control policies calculated according to such decomposition methods are guaranteed to be feasible for the full problem, they may result in behaviors that are overly conservative in terms of the cost they incur for a number of reasons. *First*, the treatment of the coupling states and inputs as independent disturbances ignores the potential dynamical coupling between these variables. *Second*, the over approximation of the coupling state and input trajectory sets by their corresponding state and input constraint sets will likely be very loose for many problem instances. More importantly, the over approximation of the coupling state and input trajectory sets in this manner ignores the fact that these sets depend on the control policy being used to regulate the system, and, therefore, neglects the possibility of co-optimizing their specification with the control policy.

Summary of results: We provide a computationally tractable method to calculate control policies that are guaranteed to be feasible for constrained decentralized control problems with nonclassical information structures. Loosely speaking, the proposed approach seeks to neutralize the nonconvexity arising from the informational coupling between subsystems by treating the information-coupling states as disturbances whose trajectories are "assumed" to take values in a certain "contract" set. To ensure the satisfaction of this assumption, we impose a contractual constraint on the control policy that "guarantees" that the information-coupling states that it induces indeed belong to said contract set. Naturally, this approach yields an inner approximation of the original decentralized control design problem, where the conservatism of the resulting approximation depends on the specification of the contract set. To limit the extent of the suboptimality that may result, we formulate a semi-infinite program to co-optimize the decentralized control policy with the location, scale, and orientation of an ellipsoidal contract set. We establish a condition on the set of allowable contracts that facilitates the joint optimization of the control policy and the contract set via semidefinite programming.

We note that there are several related papers appearing in the literature that investigate a similar approach to decentralized control design via the co-optimization of control policies and contract sets [36, 100]. Importantly, the techniques developed in these papers only permit the scaling and translation of a base contract set when co-optimizing it with the control policy. To the best of our knowledge, the method proposed in this paper provides the first systematic approach to co-optimize the control policy with the location, scale, and *orientation* of the contract set, expanding substantially the family of contracts that can be efficiently optimized.

Organization: The remainder of this chapter is organized as follows. In Section 4.2, we specify the formulation of the decentralized state-feedback control design problem considered in this chapter.[1] In Section 4.3, we provide a decomposition of the local information available to each subsystem into a partially

[1]While the state-feedback control design problem is a special case of the decentralized output-feedback control design problem (3.6), stating a problem formulation that is tailored to the results in this chapter will substantially improve the clarify of our exposition.

nested subset and its complement, which enables the isolation of the subset of each subsystem's accessible state measurements that give rise to the informational coupling between subsystems. In Section 4.4, we construct a convex inner approximation of the decentralized control design problem where the information coupling states are "assumed" to behave as disturbances taking value in a "contract" set and the control policy is constrained in such a way to induce information coupling states that are "guaranteed" to take value in the aforementioned contract set. In Section 4.5, we describe a method of co-optimizing the decentralized control policy with the location, scale, and orientation of the contract set via the solution of a semidefinite program. Section 4.6 concludes the chapter with a discussion on future work.

Finally, we note that some of the notation introduced in Chapter 3 might be reloaded in this chapter. The notation convention we introduced in Chapter 3 will be re-inherited in Chapter 5.

4.2 Problem Formulation

4.2.1 System Model

Consider the following variant of the discrete-time, linear time varying system described in Eq. (3.1), where the state of subsystem i evolves according to

$$x_i(t+1) = \sum_{j=1}^{N} \left(A_{ij}(t)x_j(t) + B_{ij}(t)u_j(t)\right) + w_i(t), \tag{4.1}$$

for $i = 1, \ldots, N$. For each subsystem i, we denote its *local state*, *local input*, and *local disturbance* at time t by $x_i(t) \in \mathbf{R}^{n_x^i}$, $u_i(t) \in \mathbf{R}^{n_u^i}$, $w_i(t) \in \mathbf{R}^{n_x^i}$, respectively.

The system is assumed to evolve over a finite time horizon T, and the initial condition is assumed to be a random vector with known probability distribution. In the sequel, we will work with a more compact representation of the full system dynamics given by

$$x(t+1) = A(t)x(t) + B(t)u(t) + w(t).$$

Here, we denote by $x(t) := (x_1(t), .., x_N(t)) \in \mathbf{R}^{n_x}$, $u(t) := (u_1(t), .., u_N(t)) \in \mathbf{R}^{n_u}$, and $w(t) := (w_1(t), .., w_N(t)) \in \mathbf{R}^{n_x}$ the full system state, input, and disturbance at time t. The dimensions of the system state and input are given by $n_x := \sum_{i=1}^{N} n_x^i$ and $n_u := \sum_{i=1}^{N} n_u^i$, respectively.

We define the system state, input, and disturbance trajectories according to

$$x := (x(0), \ldots, x(T)) \in \mathbf{R}^{N_x}, \quad N_x := n_x(T+1), \tag{4.2}$$

$$u := (u(0), \ldots, u(T-1)) \in \mathbf{R}^{N_u}, \quad N_u := n_u T, \tag{4.3}$$

$$w := (w(-1), w(0), \ldots, w(T-1)) \in \mathbf{R}^{N_x}, \tag{4.4}$$

respectively, where the initial component $w(-1)$ of the system disturbance trajectory is given by $w(-1) = x(0)$. Additionally, we denote by $\mathcal{W}$ the support of the system disturbance trajectory w, which we assume to be a convex and compact subset of $\mathbf{R}^{N_x}$. Finally, the input and disturbance trajectories are related to the state trajectory according to

$$x = Bu + Gw, \tag{4.5}$$

where the matrices B and G are given by

$$B := \begin{bmatrix} 0 & & & & \\ A_1^1 B(0) & 0 & & & \\ A_1^2 B(0) & A_2^2 B(1) & 0 & & \\ \vdots & & \ddots & \ddots & \\ \vdots & & & \ddots & 0 \\ A_1^T B(0) & A_2^T B(1) & \cdots & \cdots & A_T^T B(T-1) \end{bmatrix}, \quad G := \begin{bmatrix} A_0^0 & & & \\ A_0^1 & A_1^1 & & \\ \vdots & & \ddots & \\ A_0^T & A_1^T & \cdots & A_T^T \end{bmatrix}.$$

4.2.2 System Constraints

We consider a general family of polyhedral constraints on the state and input trajectories of the form

$$F_x x + F_u u + F_w w \leq g \qquad \forall w \in \mathcal{W}, \tag{4.6}$$

where $F_x \in \mathbf{R}^{m \times N_x}$, $F_u \in \mathbf{R}^{m \times N_u}$, $F_w \in \mathbf{R}^{m \times N_x}$, $g \in \mathbf{R}^m$ are assumed to be given. Note that such constraints may couple states and inputs across subsystems and time periods.

4.2.3 Decentralized Control Design

The information structure we consider in this chapter is identical to the one considered in problem (3.6). Specifically, let $\mathcal{G}_I = (\mathcal{V}, \mathcal{E}_I)$ be the information graph of the system. Under the assumption that each subsystem's measured output is a perfect measurement of its local state, the local information of each subsystem i at time t as

$$z_i(t) := \{x_j^{0:t} \mid (j,i) \in \mathcal{E}_I\}. \tag{4.7}$$

It follows that the local input to subsystem i at time t is of the form

$$u_i(t) = \gamma_i(z_i(t), t), \tag{4.8}$$

where $\gamma_i(\cdot, t)$ is a measurable function of the local information $z_i(t)$. We define the *local control policy* for subsystem i as $\gamma_i := (\gamma_i(\cdot, 0), \ldots, \gamma_i(\cdot, T-1))$. We refer to the collection of local control policies $\gamma := (\gamma_1, \ldots, \gamma_N)$ as the *decentralized control policy*, which relates the state trajectory x to the input trajectory u according to $u = \gamma(x)$. Finally, we let Γ denote the set of all decentralized control policies respecting the information constraints encoded in Eq. (4.8).

We consider the following family of constrained decentralized control design problems:

$$\begin{array}{ll} \text{minimize} & \mathbf{E}\left[x^\mathsf{T} R_x x + u^\mathsf{T} R_u u\right] \\ \text{subject to} & \gamma \in \Gamma \\ & \left.\begin{array}{l} u = \gamma(x) \\ x = Bu + Gw \\ F_x x + F_u u + F_w w \leq g \end{array}\right\} \forall w \in \mathcal{W}. \end{array} \tag{4.9}$$

Here, the cost matrices $R_x \in \mathbf{R}^{N_x \times N_x}$ and $R_u \in \mathbf{R}^{N_u \times N_u}$ are both assumed to be symmetric and positive semidefinite. In this chapter, out objective is to develop a tractable method to compute feasible control policies for decentralized control design problems with arbitrary (and possibly nonclassical) information structures.

4.3 Information Decomposition

The primary difficulty in solving decentralized control design problems stems from the informational coupling that emerges when a subsystem's local information is affected by prior control actions that it cannot access or reconstruct. With the aim of isolating the effects of these actions on the information available to each subsystem, we propose an information decomposition that partitions the local information available to each subsystem into a partially nested subset (i.e., an information subset that is unaffected by control actions previously applied to the system) and its complement. This information decomposition enables an equivalent reformulation of the decentralized control design problem where the control policy is expressed as an explicit function of the system disturbance and the so called *information-coupling* states. This reformulation will serve as the foundation for the contract-based approach to decentralized control design proposed in Section 4.4.

4.3.1 Decomposition of Local Information

We decompose the local information available to each subsystem according to a partition of its in-neighbors as defined by the given information graph $\mathcal{G}_I$. More specifically, for each subsystem $i \in \mathcal{V}$, we let

$$\mathcal{N}(i) \subseteq \mathcal{V}_I^-(i)$$

denote the set of in-neighboring subsystems such that the information conveyed by their local state measurements is unaffected by the prior control actions of any subsystem. This requirement is satisfied if the local information of subsystem i is such that it permits the reconstruction of all states and control actions directly affecting the local states of all subsystems belonging to $\mathcal{N}(i)$. We denote the complement of this set by $\mathcal{C}(i) := \mathcal{V}_I^-(i) \setminus \mathcal{N}(i)$ for each subsystem $i \in \mathcal{V}$.

With the goal of providing an explicit characterization of these sets, we first provide a characterization of the physical coupling between different subsystems as reflected by the block sparsity patterns of the system matrices A and B. We describe this coupling in terms of a pair of directed graphs, $\mathcal{G}_A := (\mathcal{V}, \mathcal{E}_A)$ and $\mathcal{G}_B := (\mathcal{V}, \mathcal{E}_B)$, whose edge sets are defined according to

$$\mathcal{E}_A := \{(j, i) \in \mathcal{V} \times \mathcal{V} \mid \exists t = 0, \dots, T-1 \text{ s.t. } A_{ij}(t) \neq 0\},$$
$$\mathcal{E}_B := \{(j, i) \in \mathcal{V} \times \mathcal{V} \mid \exists t = 0, \dots, T-1 \text{ s.t. } B_{ij}(t) \neq 0\}.$$

We let $\mathcal{V}_A^-(i)$ and $\mathcal{V}_B^-(i)$ denote the in-neighborhoods associated with each node $i \in \mathcal{V}$ in $\mathcal{G}_A$ and $\mathcal{G}_B$, respectively.

Building on these representations, we have the following definition that formalizes the class of information decompositions considered in this paper. For each subsystem $i \in \mathcal{V}$, define the set

$$\mathcal{N}(i) := \{j \in \mathcal{V}_I^-(i) \mid \text{(4.10), (4.11) are satisfied}\},$$

where the above conditions are given by

$$\mathcal{V}_A^-(j) \subseteq \mathcal{V}_I^-(i) \tag{4.10}$$

and

$$\bigcup_{k \in \mathcal{V}_B^-(j)} \mathcal{V}_I^-(k) \subseteq \mathcal{V}_I^-(i). \tag{4.11}$$

Condition (4.10) requires that subsystem i has access to all states that directly affect subsystem j's state through the system dynamics. Condition (4.11) requires that subsystem i has access to the local information of each subsystem whose control actions directly affect subsystem j's state. This ensures that subsystem i is able to reconstruct all control actions that directly affect subsystem j's state. Collectively, conditions (4.10) and (4.11) can be interpreted as a requirement on the *local nesting of information*, in the sense that if $j \in \mathcal{N}(i)$, then subsystem i is assumed to have access to all states and control actions that directly affect subsystem j's state through the state equation. As a result, subsystem i can explicitly reconstruct the local disturbance $w_j(t-1)$ acting on any subsystem $j \in \mathcal{N}(i)$ based only on its local information $z_i(t)$ as follows:

$$\begin{aligned} w_j(t-1) = x_j(t) &- \sum_{k \in \mathcal{V}_A^-(j)} A_{jk}(t-1) x_k(t-1) \\ &- \sum_{k \in \mathcal{V}_B^-(j)} B_{jk}(t-1) u_k(t-1). \end{aligned}$$

The local states of subsystems not belonging to $\mathcal{N}(i)$, on the other hand, may contain information that can be influenced by prior control actions. We refer to these states as the *information-coupling states* associated with subsystem i at stage t, denoting them by $x_{\mathcal{C}(i)}(t)$ where

$$\mathcal{C}(i) := \mathcal{V}_I^-(i) \setminus \mathcal{N}(i).$$

The collection of information-coupling states across all subsystems are denoted by the $x_{\mathcal{C}}(t) \in \mathbf{R}^{n_x^{\mathcal{C}}}$, where

$$\mathcal{C} := \bigcup_{i \in \mathcal{V}} \mathcal{C}(i). \tag{4.12}$$

The trajectory of information-coupling states is denoted by

$$x_{\mathcal{C}} := (x_{\mathcal{C}}(0), \ldots, x_{\mathcal{C}}(T)) \in \mathbf{R}^{N_x^{\mathcal{C}}},$$

where $N_x^{\mathcal{C}} := n_x^{\mathcal{C}}(T+1)$. Finally, it will be notationally convenient to express the mapping from the state trajectory x to its subvector $x_{\mathcal{C}}$ in terms of the projection operator $\Pi_{\mathcal{C}} : \mathbf{R}^{N_x} \rightarrow \mathbf{R}^{N_x^{\mathcal{C}}}$, where $x_{\mathcal{C}} = \Pi_{\mathcal{C}} x$.

Remark 1 (Partially Nested Information)**.** It can be shown that the given information structure is *partially nested* if and only if the set of information coupling states is empty, i.e., $\mathcal{C} = \emptyset$. It is well known that such information structures permit the equivalent reformulation of problem (4.9) as a convex optimization problem in the space of disturbance-feedback control policies.

4.3.2 Control Input Reparameterization

The proposed information decomposition suggests a natural reparameterization of the control policy in terms of the following equivalent information set.

Lemma 8 (Equivalence of Information)**.** Define the information set $\zeta_i(t)$ according to

$$\zeta_i(t) := \{x_j^{0:t} | j \in \mathcal{C}(i)\} \cup \{w_j^{-1:t-1} | j \in \mathcal{N}(i)\}.$$

The sets $z_i(t)$ and $\zeta_i(t)$ are functions of each other for each subsystem i and time t.

The proof of Lemma 8 is omitted, as it mirrors that of Lemma 5. Lemma 8 implies the following equivalent reformulation of the local control input:

$$u_i(t) = \phi_i(\zeta_i(t), t), \tag{4.13}$$

where $\phi_i(\cdot, t)$ is a measurable function of its arguments. We let $\phi_i := (\phi_i(\cdot, 0), \ldots, \phi_i(\cdot, T-1))$ and $\phi := (\phi_1, \ldots, \phi_N)$ denote the reparameterized control policy associated with each subsystem $i \in \mathcal{V}$ and the full system, respectively. With a slight abuse of notation, we express the input trajectory induced by the reparameterized control policy ϕ as

$$u = \phi(w, x_{\mathcal{C}}).$$

Finally, we denote by Φ the set of reparameterized decentralized control policies that respect the information constraints encoded in Eq. (4.13).

The reparameterization of the control input according to Eq. (4.13) results in the following equivalent reformulation of the original decentralized control problem (4.9):

$$\begin{array}{ll} \text{minimize} & \mathbf{E}\left[x^\mathsf{T} R_x x + u^\mathsf{T} R_u u\right] \\ \text{subject to} & \phi \in \Phi \\ & \left.\begin{array}{l} u = \phi(w, x) \\ x = Bu + Gw \\ F_x x + F_u u + F_w w \leq g \end{array}\right\} \forall w \in \mathcal{W}. \end{array} \tag{4.14}$$

Clearly, problem (4.14) remains nonconvex, in general, if the set of information-coupling subsystems is nonempty, i.e., $\mathcal{C} \neq \emptyset$. In Section 4.4, we construct a convex inner approximation to problem (4.14) where the information-coupling states are assumed to behave as disturbances with bounded support, and the control policy is constrained in a manner that ensures consistency between the assumed and actual behaviors of the information-coupling states.

4.4 Decentralized Control Design via Contracts

In this section, we construct a convex inner approximation of the decentralized control design problem (4.14) via the introduction of an assume-guarantee contractual constraint on the information-coupling states $x_{\mathcal{C}}$. We do so by introducing a surrogate information structure in which the information-coupling states are modeled as fictitious disturbances that are "assumed" to take values in a "contract" set. To "guarantee" the satisfaction of this assumption, we impose a contractual constraint on the control policy requiring that the actual information-coupling states induced by the control policy belong to the contract set. Given a fixed contract set, the resulting problem is a convex disturbance-feedback control design problem, whose feasible policies are guaranteed to be feasible for original problem (4.14).

4.4.1 Surrogate Information

We associate a *fictitious disturbance* $\xi_i(t) \in \mathbf{R}^{n_x^i}$ with each subsystem $i \in \mathcal{V}$ and time $t = 0, \ldots, T$. We let $\xi \in \mathbf{R}^{N_x}$ denote the corresponding fictitious disturbance trajectory induced by these individual elements, which we model as a random vector whose support $\Xi \subset \mathbf{R}^{N_x}$ is assumed to be a convex and compact set. We also assume that the fictitious disturbance trajectory ξ is independent of the system disturbance trajectory w.

Letting the collection of fictitious disturbances serve as a surrogates for the information-coupling states, we define the *surrogate local information* for subsystem i as

$$\widetilde{\zeta}_i(t) := \{\xi_j^{0:t} | j \in \mathcal{C}(i)\} \cup \{w_j^{-1:t-1} | j \in \mathcal{N}(i)\}.$$

Given a decentralized control policy $\phi \in \Phi$, the surrogate local information induces a *surrogate control input* for each subsystem i defined according to

$$\widetilde{u}_i(t) := \phi_i(\widetilde{\zeta}_i(t), t).$$

Additionally, the *surrogate input trajectory* induced by the surrogate information structure is given by

$$\widetilde{u} := \phi(w, \xi_{\mathcal{C}}),$$

where $\xi_{\mathcal{C}} := \Pi_{\mathcal{C}} \xi$.

4.4.2 Surrogate Dynamics

The treatment of the information coupling states as fictitious disturbances induces a *surrogate system state* that evolves according to the following surrogate state equation:

$$\widetilde{x}_i(t+1) = \sum_{j \in \mathcal{V} \setminus \mathcal{C}(i)} A_{ij}(t) \widetilde{x}_j(t) + \sum_{j \in \mathcal{C}(i)} A_{ij}(t) \xi_j(t) + \sum_{j=1}^{N} B_{ij}(t) \widetilde{u}_j(t) + w_i(t), \quad (4.15)$$

where $\widetilde{x}_i(t)$ denotes the surrogate state of subsystem i at time t. We require that the initial condition of the surrogate system equal that of the true system states agree, i.e., $\widetilde{x}_i(0) = x_i(0)$ for each subsystem i. Moving forward, it will be convenient to express the surrogate state dynamics in terms of trajectories as follows:

$$\widetilde{x} = \widetilde{B} \widetilde{u} + \widetilde{G} w + \widetilde{H} \xi_{\mathcal{C}}. \quad (4.16)$$

Here, the matrices $\widetilde{B}$, $\widetilde{G}$, and $\widetilde{H}$ are defined according to

$$\widetilde{B} := \begin{bmatrix} 0 & & & & \\ \widetilde{A}_1^1 B(0) & 0 & & & \\ \widetilde{A}_1^2 B(0) & \widetilde{A}_2^2 B(1) & 0 & & \\ \vdots & & \ddots & \ddots & \\ \vdots & & & \ddots & 0 \\ \widetilde{A}_1^T B(0) & \widetilde{A}_2^T B(1) & \cdots & \cdots & \widetilde{A}_T^T B(T-1) \end{bmatrix},$$

$$\widetilde{G} := \begin{bmatrix} \widetilde{A}_0^0 & & & \\ \widetilde{A}_0^1 & \widetilde{A}_1^1 & & \\ \vdots & & \ddots & \\ \widetilde{A}_0^T & \widetilde{A}_1^T & \cdots & \widetilde{A}_T^T \end{bmatrix},$$

$$\widetilde{H} := \begin{bmatrix} 0 & & & & & \\ \widetilde{A}_1^1 \widetilde{H}(0) & 0 & & & & \\ \widetilde{A}_1^2 \widetilde{H}(0) & \widetilde{A}_2^2 \widetilde{H}(1) & 0 & & & \\ \vdots & & \ddots & \ddots & & \\ \vdots & & & \ddots & 0 & \\ \widetilde{A}_1^T \widetilde{H}(0) & \widetilde{A}_2^T \widetilde{H}(1) & \cdots & \cdots & \widetilde{A}_T^T \widetilde{H}(T-1) & 0 \end{bmatrix} \Pi_{\mathcal{C}}^{\mathsf{T}},$$

where $\widetilde{A}_s^t := \prod_{r=s}^{t-1} \widetilde{A}(r)$ for $s < t$, and $\widetilde{A}_t^t = I$, and the matrices $\widetilde{A}(t)$ and $\widetilde{H}(t)$ according to

$$\widetilde{A}_{ij}(t) = \begin{cases} A_{ij}(t) & \text{if } j \in \mathcal{V} \setminus \mathcal{C}(i), \\ 0 & \text{otherwise,} \end{cases}$$

$$\widetilde{H}_{ij}(t) = A_{ij}(t) - \widetilde{A}_{ij}(t),$$

for $i, j = 1, \ldots, N$.

We close this subsection with a lemma that establishes conditions for the equivalence between the surrogate and actual state trajectories. We omit the proof, as it directly follows from the definition of the surrogate state equation (4.16).

Lemma 9. Let $u \in \mathbf{R}^{N_u}$ and $w \in \mathbf{R}^{N_x}$. It holds that $x = Bu + Gw$ if and only if $x = \widetilde{B}u + \widetilde{G}w + \widetilde{H}x_{\mathcal{C}}$.

4.4.3 Assume-Guarantee Contracts

Thus far, we have treated the information-coupling states as fictitious disturbances that are assumed to take values in a given set Ξ. Leveraging on concepts grounded in assume-guarantee reasoning [4, 67], we guarantee the satisfaction of this assumption by imposing a contractual constraint on the control policy, which ensure that it induces information-coupling states that belong to Ξ. We formalize the notion of an *assume-guarantee contract* in the following definition.

Definition 3 (Assume-guarantee Contract)**.** A control policy $\phi \in \Phi$ is said to satisfy the *assume-guarantee contract* speficied in terms of the *contract set* $\Xi_{\mathcal{C}} \subseteq \mathbf{R}^{N_x^{\mathcal{C}}}$ if

$$\Pi_{\mathcal{C}}\widetilde{x} \in \Xi_{\mathcal{C}} \quad \forall\, (w, \xi_{\mathcal{C}}) \in \mathcal{W} \times \Xi_{\mathcal{C}},$$

where $\widetilde{x} = \widetilde{B}\phi(w, \xi_{\mathcal{C}}) + \widetilde{G}w + \widetilde{H}\xi_{\mathcal{C}}$.

Here, the set $\Xi_{\mathcal{C}}$ is referred to as a *contract set*, as it specifies the set that the information-coupling states are both assumed and required to belong to. The satisfaction of the assume-guarantee contract guarantees that the *surrogate* information-coupling states $\widetilde{x}_{\mathcal{C}} := \Pi_{\mathcal{C}}\widetilde{x}$ belong to the contract set. In the following lemma, we show that the *actual* information-coupling states that result under the policy $u = \phi(w, x_{\mathcal{C}})$ are guaranteed to belong to the contract set if the assume-guarantee contract is satisfied.

Lemma 10. Let $\phi \in \Phi$ be a control policy that satisfies the assume-guarantee contract specified in terms of the contract set $\Xi_{\mathcal{C}} \subseteq \mathbf{R}^{N_x^{\mathcal{C}}}$. It follows that $\Pi_{\mathcal{C}}x \in \Xi_{\mathcal{C}}$ for all $w \in \mathcal{W}$, where $x = B\phi(w, x_{\mathcal{C}}) + Gw$.

Proof: The proof centers on a fixed-point characterization of the state trajectory x that results under the input trajectory $u = \phi(w, x_{\mathcal{C}})$. Using such a fixed-point characterization, we prove $x_{\mathcal{C}} \in \Xi_{\mathcal{C}}$ by induction in time.

We first require several preliminary definitions. Define the projection operator $\Pi_{\mathcal{C}}^{0:t}$ according to:

$$\Pi_{\mathcal{C}}^{0:t} := \begin{bmatrix} I_{n_x^{\mathcal{C}}(t+1)} & 0_{n_x^{\mathcal{C}}(t+1)\times n_x^{\mathcal{C}}(T-t)} \end{bmatrix} \Pi_{\mathcal{C}}$$

for $t = 0, \ldots, T$. It follows $x_{\mathcal{C}}^{0:t} = \Pi_{\mathcal{C}}^{0:t}x$ for each time $t \in \{0, \ldots, T\}$.

Define the function $f_\phi : \mathbf{R}^{N_x} \times \mathbf{R}^{N_x^{\mathcal{C}}} \to \mathbf{R}^{N_x}$ according to

$$f_\phi(w, \xi_{\mathcal{C}}) := \widetilde{B}\phi(w, \xi_{\mathcal{C}}) + \widetilde{G}w + \widetilde{H}\xi_{\mathcal{C}}. \tag{4.17}$$

In the following lemma, we show that the function f_ϕ is strictly causal in $\xi_{\mathcal{C}}$, and establish a fixed-point property that the function f_ϕ satisfies. This lemma is stated without proof, as it follows directly from the definition of the matrix $\widetilde{H}$ and Lemma 9.

Lemma 11. Let $\phi \in \Phi$, and let $x = B\phi(w, x_{\mathcal{C}}) + Gw$. We have that

(i) For each time $t \in \{0, \ldots, T\}$ and each $w \in \mathbf{R}^{N_x}$,

$$\Pi_{\mathcal{C}}^{0:t-1}(\xi - \xi') = 0 \Longrightarrow \Pi_{\mathcal{C}}^{0:t} f_\phi(w, \xi_{\mathcal{C}}) = \Pi_{\mathcal{C}}^{0:t} f_\phi(w, \xi'_{\mathcal{C}}).$$

(ii) The state trajectory x satisfies the fixed point condition $x = f_\phi(w, x_\mathcal{C})$ for all $w \in \mathbf{R}^{N_x}$.

Fix $w \in \mathcal{W}$, and let $x = B\phi(w, x_\mathcal{C}) + Gw$. It suffices to prove by induction in t that $x_\mathcal{C}^{0:t} \in \Pi_\mathcal{C}^{0:t}\Pi_\mathcal{C}^\mathsf{T}\Xi_\mathcal{C}$ for each time $t \in \{0, \ldots, T\}$, as the satisfaction of this condition for $t = T$ implies that $x_\mathcal{C} \in \Pi_\mathcal{C}\Pi_\mathcal{C}^\mathsf{T}\Xi_\mathcal{C} = \Xi_\mathcal{C}$. Our inductive proof follows.

Base step: Let $t = 0$. Fix $\xi \in \mathbf{R}^{N_x}$ that satisfies $\xi_\mathcal{C} \in \Xi_\mathcal{C}$. It is straightforward to verify that

$$\begin{aligned} x_\mathcal{C}(0) &= \Pi_\mathcal{C}^{0:0}x = \Pi_\mathcal{C}^{0:0}f_\phi(w, x_\mathcal{C}) = \Pi_\mathcal{C}^{0:0}f_\phi(w, \xi_\mathcal{C}) \\ &= \Pi_\mathcal{C}^{0:0}\Pi_\mathcal{C}^\mathsf{T}\Pi_\mathcal{C}f_\phi(w, \xi_\mathcal{C}) \in \Pi_\mathcal{C}^{0:0}\Pi_\mathcal{C}^\mathsf{T}\Xi_\mathcal{C}. \end{aligned}$$

Here, the second equality follows from property (ii) in Lemma 11; the third equality follows from property (i) in Lemma 11; the fourth equality follows from the identity that $\Pi_\mathcal{C}^{0:t}\Pi_\mathcal{C}^\mathsf{T}\Pi_\mathcal{C} = \Pi_\mathcal{C}^{0:t}$ for each time t; and the last inclusion condition follows from the assumption the policy ϕ satisfies the assume-guarantee contract specified by the contract set $\Xi_\mathcal{C}$. It follows that our claim is true for the base step.

Induction step: Assume that $x_\mathcal{C}^{0:t-1} \in \Pi_\mathcal{C}^{0:t-1}\Pi_\mathcal{C}^\mathsf{T}\Xi_\mathcal{C}$. We now show that $x_\mathcal{C}^{0:t} \in \Pi_\mathcal{C}^{0:t}\Pi_\mathcal{C}^\mathsf{T}\Xi_\mathcal{C}$. Fix $\xi \in \mathbf{R}^{N_x}$ that satisfies $\Pi_\mathcal{C}\xi \in \Xi_\mathcal{C}$ and $x_\mathcal{C}^{0:t-1} = \Pi_\mathcal{C}^{0:t-1}\xi$ (which is guaranteed to exist by our induction hypothesis). Using arguments analogous to our proof in the base step, we have that

$$\begin{aligned} x_\mathcal{C}^{0:t} &= \Pi_\mathcal{C}^{0:t}x = \Pi_\mathcal{C}^{0:t}f_\phi(w, x_\mathcal{C}) = \Pi_\mathcal{C}^{0:t}f_\phi(w, \xi_\mathcal{C}) \\ &= \Pi_\mathcal{C}^{0:t}\Pi_\mathcal{C}^\mathsf{T}\Pi_\mathcal{C}f_\phi(w, \xi_\mathcal{C}) \in \Pi_\mathcal{C}^{0:t}\Pi_\mathcal{C}^\mathsf{T}\Xi_\mathcal{C}, \end{aligned}$$

where the third equality follows from a combination of property (i) in Lemma

11 and the fact that $x_{\mathcal{C}}^{0:t-1} = \Pi_{\mathcal{C}}^{0:t-1}\xi$. This completes the induction step of the proof. ■

In the following proposition, we provide an inner approximation of the decentralized control design problem (4.14) via the introduction of an assume-guarnatee contractual constraint.

Proposition 4. Let $\phi \in \Phi$ be a feasible control policy for the following problem:

$$\begin{aligned} &\text{minimize} && \mathbf{E}\left[\widetilde{x}^{\mathsf{T}} R_x \widetilde{x} + \widetilde{u}^{\mathsf{T}} R_u \widetilde{u}\right] \\ &\text{subject to} && \phi \in \Phi \\ & && \left.\begin{aligned} &\widetilde{u} = \phi(w, \xi_{\mathcal{C}}) \\ &\Pi_{\mathcal{C}} \widetilde{x} \in \Xi_{\mathcal{C}} \\ &\widetilde{x} = \widetilde{B}\widetilde{u} + \widetilde{G}w + \widetilde{H}\xi_{\mathcal{C}} \\ &F_x \widetilde{x} + F_u \widetilde{u} + F_w w \leq g \end{aligned}\right\} \forall (w, \xi_{\mathcal{C}}) \in \mathcal{W} \times \Xi_{\mathcal{C}}, \end{aligned} \tag{4.18}$$

It follows that ϕ is also feasible for problem (4.14).

Proof: Let ϕ be a feasible policy for problem (4.18), and $x = B\phi(w, x_{\mathcal{C}}) + Gw$. It follows from Lemma 10 that the state trajectory x satisfies

$$x_{\mathcal{C}} \in \Xi_{\mathcal{C}} \quad \forall w \in \mathcal{W}. \tag{4.19}$$

Additionally, the feasibility of the policy ϕ for problem (4.18) implies that

$$F_x f_\phi(w, \xi_{\mathcal{C}}) + F_u \phi(w, \xi_{\mathcal{C}}) + F_w w \leq g \tag{4.20}$$

for all $(w, \xi_{\mathcal{C}}) \in \mathcal{W} \times \Xi_{\mathcal{C}}$, where the function f_ϕ is defined in Eq. (4.17). In combination with the fact that $x_{\mathcal{C}} \in \Xi_{\mathcal{C}}$ for all $w \in \mathcal{W}$, this implies that

$$F_x f_\phi(w, x_{\mathcal{C}}) + F_u \phi(w, x_{\mathcal{C}}) + F_w w \leq g \quad \forall w \in \mathcal{W}. \tag{4.21}$$

Finally, it follows from Lemma 11 the trajectory x satisfies the fixed-point condition $x = f_\phi(w, x_\mathcal{C})$ for all $w \in \mathcal{W}$. In combination with the satisfaction of the robust inequality constraint (4.21), this implies that

$$F_x x + F_u \phi(w, x_\mathcal{C}) + F_w w \leq g \quad \forall w \in \mathcal{W},$$

which shows that the policy ϕ is feasible for problem (4.14). ■

Problem (4.18) is a convex disturbance feedback control design problem, given a fixed contract set $\Xi_\mathcal{C}$. The choice of the contract set does, however, play an important role in determining the performance of the control policies that it gives rise to via solutions to problem (4.18). In Section 4.5, we develop a systematic approach to enable the joint optimization of the contract set with the control policy via semidefinite programming.

4.5 Policy-Contract Optimization

In this section, we provide a semidefinite programming-based method to co-optimize the design of the decentralized control policy together with the contract set that constrains its design. As part of the proposed approach, we consider a restricted family of control policies that are affinely parameterized in both the disturbance and fictitious disturbance histories. We also parameterize the fictitious disturbance process as a causal affine function of a given (primitive) disturbance process—an approach that is similar in nature to the class of parameterizations that have been recently studied in the context of robust optimization with adjustable uncertainty sets [118]. As one of our primary results in this section, we identify a structural condition on the family of allowable contract sets that permits the inner approximation of the resulting policy-contract optimization problem as a semidefinite program.

4.5.1 Affine Control Policies

We restrict our attention to affine decentralized disturbance-feedback control policies of the form

$$\widetilde{u}_i(t) = u_i^o(t) + \sum_{j\in\mathcal{N}(i)} \sum_{s=-1}^{t-1} Q_{ij}^w(t, s+1) w_j(s) + \sum_{j\in\mathcal{C}(i)} \sum_{s=0}^{t} Q_{ij}^\xi(t, s)\xi_j(s), \qquad (4.22)$$

for $t = 0, \ldots, T-1$ and $i = 1, \ldots, N$. Here, $u_i^o(t)$ denotes the open-loop control input, and the matrices $Q_{ij}^w(t, s+1)$ and $Q_{ij}^\xi(t, s)$ denote the feedback control gains. The affine control policy specified in Eq. (4.22) can be expressed in terms of trajectories as

$$\widetilde{u} = u^o + Q^w w + Q^\xi \xi, \qquad (4.23)$$

where the gain matrices Q^w and Q^ξ are both $T \times (T+1)$ block matrices, whose (t, s)-th block is defined according to

$$[Q^w(t,s)]_{ij} = \begin{cases} Q_{ij}^w(t,s) & \text{if } j \in \mathcal{N}(i),\ t \geq s, \\ 0 & \text{otherwise,} \end{cases} \qquad (4.24)$$

$$[Q^\xi(t,s)]_{ij} = \begin{cases} Q_{ij}^\xi(t,s) & \text{if } j \in \mathcal{C}(i),\ t \geq s, \\ 0 & \text{otherwise.} \end{cases} \qquad (4.25)$$

for $i, j = 1, \ldots, N$. We let $\mathcal{Q}_N$ and $\mathcal{Q}_C$ denote the matrix subspaces respecting the block sparsity patterns specified according to Eqs. (4.24) and (4.25), respectively.

4.5.2 Affine Parameterization of the Fictitious Disturbance

We focus our analysis on fictitious disturbances that are expressed according to affine transformations of a *primitive disturbance*. Such a parameterization yields contract sets that have adjustable location, scale, and orientation. Specifically, we let the random vector v denote the *primitive disturbance trajectory*, which is assumed to be an i.i.d. copy of the system disturbance trajectory w. We parameterize the fictitious disturbance trajectory affinely in the primitive disturbance as

$$\xi := \overline{\xi} + Zv. \tag{4.26}$$

Here, the parameters $\overline{\xi} \in \mathbf{R}^{N_x}$ and $Z \in \mathbf{R}^{N_x \times N_x}$ can be adjusted to control the shape of the resulting contract set $\Xi_{\mathcal{C}}$, which takes the form of

$$\Xi_{\mathcal{C}} = \Pi_{\mathcal{C}} \left(\overline{\xi} \oplus Z\mathcal{W} \right). \tag{4.27}$$

Throughout the paper, we will restrict our attention to transformations (4.26) in which the matrix parameter Z is both lower triangular and invertible. We denote the set of all such matrices by $\mathcal{Z} \subset \mathbf{R}^{N_x \times N_x}$.

The specification of the fictitious disturbance according to Eq. (4.26) induces the following the surrogate control input:

$$\widetilde{u} = u^o + Q^{\xi}\overline{\xi} + Q^w w + Q^{\xi} Z v. \tag{4.28}$$

We eliminate the bilinear terms in Eq. (4.28) through the following the change of variables:

$$\overline{u} := u^o + Q^{\xi}\overline{\xi} \text{ and } Q^v := Q^{\xi} Z. \tag{4.29}$$

This change of variables gives rise to a reparameterization of the surrogate input trajectory as

$$\widetilde{u} = \overline{u} + Q^w w + Q^v v, \tag{4.30}$$

where the matrix $Q^v \in \mathbf{R}^{N_u \times N_x}$ must satisfy the sparsity constraint

$$Q^v Z^{-1} \in \mathcal{Q}_C$$

in order to ensure the satisfaction of the original sparsity constraint that $Q^\xi \in \mathcal{Q}_C$.

The parameterization of the contract set and control policy in this manner permits their co-optimization as follows:

$$\begin{array}{ll} \text{minimize} & \mathbf{E}\left[\widetilde{x}^\mathsf{T} R_x \widetilde{x} + \widetilde{u}^\mathsf{T} R_u \widetilde{u}\right] \\ \text{subject to} & Q^w \in \mathcal{Q}_N, Q^v \in \mathbf{R}^{N_u \times N_x}, Z \in \mathcal{Z} \\ & \overline{u} \in \mathbf{R}^{N_u},\ \overline{\xi} \in \mathbf{R}^{N_x}, \\ & Q^v Z^{-1} \in \mathcal{Q}_C \\ & \left.\begin{array}{l} \xi = \overline{\xi} + Zv \\ \widetilde{u} = \overline{u} + Q^w w + Q^v v \\ \widetilde{x} = \widetilde{B}\widetilde{u} + \widetilde{G}w + \widetilde{H}\xi_{\mathcal{C}} \\ \Pi_{\mathcal{C}}\widetilde{x} \in \Pi_{\mathcal{C}}\left(\overline{\xi} \oplus Z\mathcal{W}\right) \\ F_x\widetilde{x} + F_u\widetilde{u} + F_w w \leq g \end{array}\right\} \forall (w,v) \in \mathcal{W}^2. \end{array} \tag{4.31}$$

where $\mathcal{W}^2 := \mathcal{W} \times \mathcal{W}$. Problem (4.31) is a nonconvex semi-infinite program, where the nonconvexity is due to the sparsity constraint on the matrix $Q^v Z^{-1}$ and the contractual constraint on the affine control policy. In what follows, we provide convex inner approximations of these constraints, which yields an inner approximation of problem (4.31) as a semidefinite program.

4.5.3 Restricting the Contract Set

In what follows, we introduce an additional restriction on the set of allowable matrix parameters Z that guarantees the invariance of the subspace $\mathcal{Q}_C$ under multiplication by such matrices, thereby allowing the equivalent reformulation of the bilinear constraint $Q^v Z^{-1} \in \mathcal{Q}_C$ as $Q^v \in \mathcal{Q}_C$.

Specifically, we require that the matrix Z be of the form

$$Z = \lambda I - Y, \tag{4.32}$$

where $\lambda \geq 1$ is scalar parameter and $Y \in \mathbf{R}^{N_x \times N_x}$ is a $(T+1) \times (T+1)$ strictly block lower triangular matrix of the form

$$Y = \begin{bmatrix} 0 & & & \\ Y(1,0) & 0 & & \\ \vdots & \ddots & \ddots & \\ Y(T,0) & \cdots & Y(T,T-1) & 0 \end{bmatrix}. \tag{4.33}$$

Furthermore, each block of the matrix Y is an $N \times N$ block matrix, whose (i,j)-th block is of dimension $n_x^i \times n_x^j$. We impose an additional restriction on the structure of the matrix Y in the form of sparsity constraints (that reflect the pattern of informational coupling between subsystems) on each of its blocks.

More specifically, we encode the pattern of informational coupling between subsystems according to a directed graph $\mathcal{G}_C := (\mathcal{V}, \mathcal{E}_C)$, whose directed edge set $\mathcal{E}_C$ is defined as

$$\mathcal{E}_C := \{(j,i) \in \mathcal{E}_I \mid j \in \mathcal{C}(i)\}.$$

We let $\mathcal{V}_C^+(i)$ denote the out-neighborhood of a node $i \in \mathcal{V}$ in the graph $\mathcal{G}_C$. Using this graph, we impose a sparsity constraint on each block of the matrix Y

of the form:

$$[Y(t,s)]_{ij} = 0 \quad \text{if } \mathcal{V}_C^+(i) \not\subseteq \mathcal{V}_C^+(j) \tag{4.34}$$

for all $i, j = 1, \ldots, N$, and $t, s = 0, \ldots, T$. We let $\mathcal{Y}(\mathcal{G}_C)$ denote the subspace of all matrices that respect the sparsity constraints implied by Eqs. (4.33) and (4.34).

We have the following result establishing the invariance of the subspace $\mathcal{Q}_C$ under multiplication by matrices $Y \in \mathcal{Y}(\mathcal{G}_C)$.

Lemma 12. If $Q \in \mathcal{Q}_C$ and $Y \in \mathcal{Y}(\mathcal{G}_C)$, then $QY \in \mathcal{Q}_C$.

Proof: Fix $Q \in \mathcal{Q}_C$ and $Y \in \mathcal{Y}(\mathcal{G}_C)$, and set $P = QY$. It follows that P is a $T \times (T+1)$ block matrix, whose (t,s)th block is given by

$$P(t,s) = \begin{cases} \sum_{r=s+1}^{t} Q(t,r)Y(r,s) & \text{if } t > s \\ 0 & \text{otherwise} \end{cases}$$

We complete the proof by showing that

$$[P(t,s)]_{ij} \neq 0 \Longrightarrow j \in \mathcal{C}(i).$$

The condition that $[P(t,s)]_{ij} \neq 0$ implies the existence of a time $r \in \{s+1, \ldots, t\}$ and a subsystem $k \in \mathcal{V}$, such that

$$[Q(t,r)]_{ik} \neq 0, \quad \text{and} \quad [Y(r,s)]_{kj} \neq 0.$$

The fact that $[Q(t,r)]_{ik}$ is nonzero implies that $k \in \mathcal{C}(i)$, as the matrix Q satisfies $Q \in \mathcal{Q}_C$. The fact that $[Y(r,s)]_{kj}$ is nonzero implies that $\mathcal{V}_C^+(k) \subseteq \mathcal{V}_C^+(j)$, as the matrix Y satisfies $Y \in \mathcal{Y}(\mathcal{G}_C)$. It follows from the definition of out-neighborhoods that $i \in \mathcal{V}_C^+(k) \subseteq \mathcal{V}_C^+(j)$, which subsequently implies that $j \in \mathcal{C}(i)$. The desired result follows. ∎

We have the following result as an immediate consequence of Lemma 12.

Lemma 13. Let $Y \in \mathcal{Y}(\mathcal{G}_C)$ and $\lambda \in [1, \infty)$. It follows that

$$\mathcal{Q}_C = \left\{ Q^v(\lambda I - Y)^{-1} \,|\, Q^v \in \mathcal{Q}_C \right\}.$$

Proof: Fix $\lambda \in [1, \infty)$ and $Y \in \mathcal{Y}(\mathcal{G}_C)$. Define the finite-dimensional linear map $F : \mathbf{R}^{N_u \times N_x} \to \mathbf{R}^{N_u \times N_x}$ according to

$$F(Q) := Q(\lambda I - Y)^{-1}.$$

It suffices to show that $F(\mathcal{Q}_C) = \mathcal{Q}_C$. The invertibility of the linear map F, in combination with the fact that the subspace $\mathcal{Q}_C$ is finite-dimensional, implies that $F(\mathcal{Q}_C) = \mathcal{Q}_C$ if and only if $F^{-1}(\mathcal{Q}_C) \subseteq \mathcal{Q}_C$. Note that the inverse of the linear map F is given by

$$F^{-1}(Q) = Q(\lambda I - Y).$$

The desired result follows, as Lemma 12 implies that $QY \in \mathcal{Q}_C$ for each $Q \in \mathcal{Q}_C$ and $Y \in \mathcal{Y}(\mathcal{G}_C)$. ■

It follows from Lemma 13 that the constraint $Q^v Z^{-1} \in \mathcal{Q}_C$ is equivalent to $Q^v \in \mathcal{Q}_C$ if $Z = \lambda I - Y$, where $Y \in \mathcal{Y}(\mathcal{G}_C)$ and $\lambda \geq 1$.

We conclude this subsection with a discussion on the family of correlations in the fictitious disturbances that can be adjusted under the sparsity constraint $Y \in \mathcal{Y}(\mathcal{G}_C)$. First note that the block lower triangular structure of the matrix Y enables one to adjust the inter-temporal correlation in the fictitious disturbance process. Such a property is desirable, as the information coupling states will, in general, couple across time. Additionally, the sparsity constraint in Eq. (4.34) imposes the restriction that one is allowed to adjust the correlation between the fictitious disturbances from two subsystems if and only if the out-neighborhood of one subsystem in the graph $\mathcal{G}_C$ is contained in that of the other. This sparsity constraint is a technical condition that is required to guarantee the invariance of the subspace $\mathcal{Q}_C$ under multiplication with any matrix belonging to the set $\mathcal{Y}(\mathcal{G}_C)$. However, such a requirement might be restrictive, as the aforementioned containment condition on two subsystems' out-neighborhoods might not be satisfied by two subsystems that are dynamically coupled through the system A and B matrices. The question as to how one might address such a drawback remains an important direction for future work.

4.5.4 Semidefinite Programming Approximation

In this section, we require an additional assumption on the support of the disturbance trajectory[2].

Assumption 3. The disturbance trajectory w is a zero-mean random vector whose support is an ellipsoid given by

$$\mathcal{W} := \left\{ z \in \mathbf{R}^{N_x} \;\middle|\; z^\mathsf{T} \Sigma^{-1} z \leq 1 \right\},$$

where the matrix Σ is symmetric and positive definite.

In the remainder of this section, we illustrate how Assumption 3 enables both the equivalent reformulation of the robust linear constraint in problem (4.31) as second order cone constraints and the inner approximation of the contractual constraint in problem (4.31) as linear matrix inequality constraints.

[2]The assumption that the disturbance trajectory w be zero mean and have support centered at the origin is made to simplify the statement of our subsequent results in this section. The relaxation of this assumption only requires additional algebraic manipulations, and is omitted to facilitate exposition.

To lighten notation, we write the surrogate state trajectory $\widetilde{x}$ more compactly as

$$\widetilde{x} = \overline{x} + P^w w + P^v v,$$

where $\overline{x} := \widetilde{B}\overline{u} + \widetilde{H}\overline{\xi}$, $P^w := \widetilde{B}Q^w + \widetilde{G}$, and $P^v := \widetilde{B}Q^v + \widetilde{H}(\lambda I - Y)$.

We first address the robust linear constraints in problem (4.31). The following result provides an equivalent reformulation as second-order cone constraints. Its proof is omitted, as it is an immediate consequence of the identity $\sup_{w \in \mathcal{W}} c^\mathsf{T} w = \|\Sigma^{1/2} c\|_2$ for all $c \in \mathbf{R}^{N_x}$.

Lemma 14. The semi-infinite constraint $F_x \widetilde{x} + F_u \widetilde{u} + F_w w \leq g$ for all $(w, v) \in \mathcal{W}^2$ is satisfied if and only if

$$\begin{aligned} &\left\|\Sigma^{1/2} e_i^\mathsf{T}(F_x P^w + F_u Q^w + F_w)\right\|_2 + \left\|\Sigma^{1/2} e_i^\mathsf{T}(F_x P^v + F_u Q^v)\right\|_2 \\ &\quad \leq e_i^\mathsf{T}(g - F_x \overline{x} - F_u \overline{u}), \quad i = 1, \ldots, m, \end{aligned} \tag{4.35}$$

where e_i is the ith standard basis vector in $\mathbf{R}^m$.

We now address the nonconvexity that stems from the contractual constraint in problem (4.31). First, notice that the contractual constraint is equivalent to the following set containment constraint

$$\Pi_\mathcal{C}\left(\overline{x} \oplus P^w \mathcal{W} \oplus P^v \mathcal{W}\right) \subseteq \Pi_\mathcal{C}\left(\overline{\xi} \oplus Z\mathcal{W}\right). \tag{4.36}$$

The set containment constraint (4.36) amounts to requiring that the Minkowski sum of two ellipsoids be contained within another ellipsoid. It follows from [42][Theorem 4.2] that this class of set containment constraints can be approximated from within by a quadratic matrix inequality. We first state [42][Theorem 4.2] in the following lemma.

Lemma 15. Let Assumption 3 hold, and let $L_1, L_2, L_3 \in \mathbf{R}^{m \times N_x}$, where $m \leq N_x$. We have that $L_1\mathcal{W} \oplus L_2\mathcal{W} \subseteq L_3\mathcal{W}$ if there exist a scalar $\alpha \in [0,1]$ such that

$$\alpha^{-1}L_1\Sigma L_1^\mathsf{T} + (1-\alpha)^{-1}L_2\Sigma L_2^\mathsf{T} \preceq L_3\Sigma L_3^\mathsf{T}. \tag{4.37}$$

The outer approximation of the Minkowski sum in Lemma 15 is *tight*. That is to say, there exists a scalar $\alpha \in [0,1]$ for which the boundary of the ellipsoid $L_3\mathcal{W}$ intersects the boundary of the Minkowski sum $L_1\mathcal{W} \oplus L_2\mathcal{W}$. A direct application of Lemma 15 implies that the set containment constraint (4.36) is satisfied if there exists $\alpha \in [0,1]$ such that

$$\Pi_\mathcal{C}(\overline{\xi} - \overline{x}) = 0 \tag{4.38}$$

$$\Pi_\mathcal{C}\left(\alpha^{-1}P^w\Sigma P^{w\mathsf{T}} + (1-\alpha)^{-1}P^v\Sigma P^{v\mathsf{T}}\right)\Pi_\mathcal{C}^\mathsf{T} \preceq \Pi_\mathcal{C}(\lambda I - Y)\Sigma(\lambda I - Y)^\mathsf{T}\Pi_\mathcal{C}^\mathsf{T}. \tag{4.39}$$

The quadratic matrix inequality (4.39) is still non-convex in the decision variables λ and Y. We have the following result, which leverages on Schur's Lemma to construct a further inner approximation of the matrix inequality (4.39) as a linear matrix inequality.

Lemma 16. Let Assumption 3 hold. The set containment constraint (4.36) is satisfied if there exists a scalar $\beta \in [0, \lambda]$ such that

$$\Pi_\mathcal{C}\left(\overline{x} - \overline{\xi}\right) = 0, \tag{4.40}$$

$$\begin{bmatrix} \Pi_\mathcal{C}\widetilde{\Sigma}\Pi_\mathcal{C}^\mathsf{T} & \Pi_\mathcal{C}P^w & \Pi_\mathcal{C}P^v \\ P^{w\mathsf{T}}\Pi_\mathcal{C}^\mathsf{T} & \beta\Sigma^{-1} & 0 \\ P^{v\mathsf{T}}\Pi_\mathcal{C}^\mathsf{T} & 0 & (\lambda-\beta)\Sigma^{-1} \end{bmatrix} \succeq 0, \tag{4.41}$$

where $\widetilde{\Sigma} = \lambda\Sigma - Y\Sigma - \Sigma Y^\mathsf{T}$.

Proof: It suffices to show that the matrix inequality (4.39) is satisfied if the LMI (4.41) is satisfied. Define $\beta := \alpha\lambda$, and divide both sides of the matrix inequality

(4.39) by λ. It follows that the matrix inequality (4.39) is satisfied if and only if there exists $\beta \in [0, \lambda]$, such that

$$\beta^{-1}\Pi_{\mathcal{C}} P^{w} \Sigma P^{w\mathsf{T}} \Pi_{\mathcal{C}}^{\mathsf{T}} + (\lambda - \beta)^{-1} \Pi_{\mathcal{C}} P^{v} \Sigma P^{v\mathsf{T}} \Pi_{\mathcal{C}}^{\mathsf{T}}$$
$$\preceq \Pi_{\mathcal{C}}(\lambda\Sigma - Y\Sigma - \Sigma Y^{\mathsf{T}} + \lambda^{-1} Y \Sigma Y^{\mathsf{T}}) \Pi_{\mathcal{C}}^{\mathsf{T}}. \tag{4.42}$$

Recall that the matrix Σ is positive definite. Consequently, it follows from Schur's Lemma that the matrix inequality (4.42) is satisfied if and only if

$$\begin{bmatrix} \Pi_{\mathcal{C}}(\widetilde{\Sigma} + \lambda^{-1} Y \Sigma Y^{\mathsf{T}}) \Pi_{\mathcal{C}}^{\mathsf{T}} & \Pi_{\mathcal{C}} P^{w} & \Pi_{\mathcal{C}} P^{v} \\ P^{w\mathsf{T}} \Pi_{\mathcal{C}}^{\mathsf{T}} & \beta \Sigma^{-1} & \\ P^{v\mathsf{T}} \Pi_{\mathcal{C}}^{\mathsf{T}} & & (\lambda - \beta)\Sigma^{-1} \end{bmatrix} \succeq 0, \tag{4.43}$$

where the matrix $\widetilde{\Sigma} = \lambda\Sigma - Y\Sigma - \Sigma Y^{\mathsf{T}}$. The fact that $\lambda^{-1} Y \Sigma Y^{\mathsf{T}} \succeq 0$ implies that the matrix inequality (4.43) is satisfied if the linear matrix inequality (4.41) is satisfied. This completes the proof. ∎

By applying Lemmas 13, 14, and 16, one can approximate the nonconvex semi-infinite program (4.31) from within as the following finite-dimensional semidefinite program.

Proposition 5. Each feasible solution to the following semidefinite program is feasible for problem (4.31):

$$\begin{aligned}
\text{minimize} \quad & \operatorname{Tr}\left(P^{v\mathsf{T}} R_x P^v M + P^{w\mathsf{T}} R_x P^w M\right) + \overline{x}^\mathsf{T} R_x \overline{x} \\
& + \operatorname{Tr}\left(Q^{w\mathsf{T}} R_u Q^w M + Q^{v\mathsf{T}} R_u Q^v M\right) + \overline{u}^\mathsf{T} R_u \overline{u} \\
\text{subject to} \quad & Q^w \in \mathcal{Q}(\mathcal{G}_N),\ Q^v \in \mathcal{Q}(\mathcal{G}_C),\ Y \in \mathcal{Y}(\mathcal{G}), \\
& \overline{u} \in \mathbf{R}^{N_u}, \overline{\xi},\ \overline{x} \in \mathbf{R}^{N_x}, \lambda,\ \beta \in \mathbf{R}_+,\ P^w,\ P^v \in \mathbf{R}^{N_x \times N_x}, \\
& \lambda \geq \max\{1, \beta\}, \\
& \overline{x} = \widetilde{B}\overline{u} + \widetilde{H}\overline{\xi} \\
& P^w = \widetilde{B}Q^w + \widetilde{G} \\
& P^v = \widetilde{B}Q^v + \widetilde{H}(\lambda I - Y) \\
& \text{(4.35)},\ \text{(4.40)}, \text{(4.41)},
\end{aligned} \tag{4.44}$$

where $M := \mathbf{E}[ww^\mathsf{T}]$ is the second moment matrix of the disturbance trajectory w.

The decision variables for problem (4.44) are the matrices Q^w, Q^v, Y, P^w, P^v, the vectors $\overline{u}$, $\overline{v}$, $\overline{x}$, and the scalars λ and β. Problem (4.44) is a convex inner approximation of the reformulated decentralized control design problem (4.14), in the sense that each feasible solution of problem (4.44) can be mapped to a feasible affine control policy for problem (4.14) via the change of variables specified in (4.29).

4.6 Conclusion

We provide a method to compute feasible control policies for constrained decentralized control design problems via the introduction of assume-guarantee contracts. At the heart of this approximation is the treatment of information-coupling states as a fictitious disturbances that are "assumed" to take values in a contract set. We "guarantee" the inclusion of the information-coupling states in the contract set by imposing an assume-guarantee contractual constraint on the control policy. The introduction of such assume-guarantee contracts gives rise to an inner approximation of the decentralized control design problem, whose quality depends on the specification of the contract set. We provide a method of co-optimizing the decentralized control policy with the location, scale, and orientation of the contract set via semidefinite programming.

We conclude this chapter with a discussion on several interesting directions for future work. *First*, one potential drawback of the proposed technique is its explicit reliance on the assumption that the system under consideration operates over a finite time horizon. It would be of interest to explore the possibility of guaranteeing the *recursive feasibility* of the controller we design by imposing additional constraints on the terminal state, in a similar spirit to several existing work on tube-based decentralized model predictive control (MPC) [75, 97, 100]. *Second*, our semidefinite programming inner approximation to the policy-contract optimization problem relies explicitly on the assumption that the support of the disturbance trajectory is an ellipsoid having non-empty interior. In the future, it would be of interest to accommodate other family of disturbance trajectories whose supports are convex polyhedra or product of ellipsoids. *Finally*, the techniques developed in this paper are tailored to the setting in which each subsystem's state is perfectly observed by some other subsystems. It would be of interest to generalize our proposed technique to the decentralized control of partially observed linear systems.

CHAPTER 5

CONVEX INFORMATION RELAXATION AND PERFORMANCE BOUNDS

The decentralized control policies we derive in Chapters 3 and 4 are, in general, suboptimal. In this chapter, we provide a method of bounding the suboptimality of such policies via the derivation of a computationally tractable lower bound on the optimal value of the original decentralized control design problem. Specifically, given a decentralized control problem with nonclassical information, we characterize an expansion of the given information structure, which ensures its partial nestedness, while maximizing the optimal value of the resulting decentralized control problem under the expanded information structure. The resulting decentralized control problem is cast as an infinite-dimensional convex program, which is further relaxed via a partial dualization and restriction to affine dual control policies. The resulting problem is a finite-dimensional conic program whose optimal value is a provable lower bound on the minimum cost of the original constrained decentralized control problem.

5.1 Introduction

We provide a method of tractably bounding the suboptimality of a feasible decentralized control policy. Our approach is to derive a tractably computable lower bound on the optimal value of the original decentralized control design problem. In order to state our result in its full generality, we resort to the general formulation of decentralized output-feedback control design problems specified in problem (3.6). Specifically, we restate the decentralized control design problem we consider as follows:

$$
\begin{aligned}
&\text{minimize} && \mathbf{E}\left[x'R_x x + u'R_u u\right] \\
&\text{subject to} && \gamma \in \Gamma(\mathcal{G}_I),\ s \in \mathcal{L}_m^2 \\
& && \left.\begin{aligned}
&F_x x + F_u u + F_w w + s = 0 \\
&x = Bu + Gw \\
&y = Cx + Hw \\
&u = \gamma(y) \\
&s \geq 0
\end{aligned}\right\} \ \mathbf{P}\text{-a.s.}.
\end{aligned}
$$

We denote the *optimal value* of problem (3.6) by $J^*(\mathcal{G}_I)$, where we have made explicit the dependence of the optimal value of problem (3.6) on the underlying information graph $\mathcal{G}_I$. Our objective in this chapter is to provide a non-trivial and tractably computable lower bound on $J^*(\mathcal{G}_I)$.

Related Work: As the tractable computation of optimal policies for the majority of decentralized control problems with nonclassical information structures remains out of reach [76], there is a practical need to quantify the suboptimality of feasible policies via the derivation of lower bounds on the optimal values of such problems. Focusing on Witsenhausen's counterexample [109] and its variants, there are several results in the literature, which establish lower bounds using information-theoretic techniques (e.g., using the data processing inequality) [7,52,66], and linear programming-based relaxations [61]. However, looking beyond Witsenhausen's counterexample, it is unclear as to how one might extend these techniques to establish computationally tractable lower bounds for the more general family of decentralized control problems considered in this note. More closely related to the approach adopted in this note, there is another stream of literature that investigates the derivation of computationally tractable lower bounds via information relaxations that increase the amount of information to which each controller has access to ensure the partial nestedness [5, 29, 77, 91, 113] or quadratic invariance [87] of the expanded information structure.

Summary of Results: In this chapter, we develop a tractable approach to the computation of tight lower bounds on the minimum cost of constrained decentralized control problems with nonclassical information structures. The proposed approach is predicated on two relaxation steps, which together yield a finite-dimensional convex programming relaxation of the original problem. The first step entails an information relaxation, which eliminates the signaling incentive between controllers by expanding the set of measurements that each controller has access to. Specifically, we characterize an expansion of the given information structure, which ensures its partial nestedness, while maximizing the optimal value of the resulting decentralized control problem under the expanded information structure. The relaxation is also shown to be *tight*, in the sense that the lower bound induced by the relaxation is achieved for several families of decentralized control problems with nonclassical information. The relaxed decentralized control problem is then recast as an equivalent convex, infinite-dimensional program via the classical Youla parameterization. Although convex, the resulting optimization problem remains computationally intractable due to its infinite-dimensionality. As part of the second relaxation step, we obtain a finite-dimensional relaxation of this problem through its partial dualization, and restriction to affine dual control policies. The resulting problem is a finite-dimensional conic program, whose optimal value is guaranteed to be a lower bound on the minimum cost of the original decentralized control design problem. To the best of our knowledge, such result is the first to offer an efficiently computable (and nontrivial) lower bound on the optimal cost of a decentralized control design problem with multiple subsystems, multiple time periods, and polyhedral constraints on state and input. If the gap between the cost incurred by an admissible policy and the proposed lower bound is small, then one may conclude that said policy is near-optimal.

5.2 A Convex Information Relaxation

In what follows, we provide a method of convexifying decentralized control design problems with nonclassical information structures via information-based relaxations. Specifically, we characterize an expansion of the given information graph that guarantees the partial nestedness of the relaxed information structure, while maximizing the optimal value of the relaxed problem. It is given by the optimal solution to:

$$\underset{\mathcal{G} \supseteq \mathcal{G}_I}{\text{maximize}} \quad J^*(\mathcal{G}) \quad \text{subject to} \quad \mathcal{G} \in \text{PN}(\Theta). \tag{5.1}$$

Here, the main difficulty in solving problem (5.1) is the requirement that any feasible solution both induce a partially nested information structure and be a supergraph of $\mathcal{G}_I$. We require a few definitions before stating the solution to problem (5.1).

Definition 4 (Precedence Graph)**.** We define the *precedence graph* associated with the system Θ and the information graph $\mathcal{G}_I$ as the directed graph $\mathcal{G}_P(\Theta, \mathcal{G}_I) = (\mathcal{V}, \mathcal{E}_P(\Theta, \mathcal{G}_I))$ whose directed edge set is defined as

$$\mathcal{E}_P(\Theta, \mathcal{G}_I) := \{(i, j) \mid i, j \in \mathcal{V},\ i \prec j \text{ with respect to } (\Theta, \mathcal{G}_I)\}.$$

Essentially, the precedence graph provides a directed graphical representation of the precedence relations between all subsystems, as specified in Definition 1.

Definition 5 (Transitive Closure)**.** The *transitive closure* of a directed graph $\mathcal{G} = (\mathcal{V}, \mathcal{E})$ is defined as the directed graph $\overline{\mathcal{G}} = (\mathcal{V}, \overline{\mathcal{E}})$, where $(i, j) \in \overline{\mathcal{E}}$ if and only if there exists a directed path in $\mathcal{G}$ from node i to node j.

The transitive closure of a directed graph can be efficiently computed using Warshall's algorithm [108]. Equipped with these definitions, we state the following result, which provides a 'closed-form' solution to problem (5.1).

Theorem 1 (Information Relaxation)**.** An optimal solution to (5.1) is given by $\overline{\mathcal{G}_P(\Theta, \mathcal{G}_I)}$, the *transitive closure of the precedence graph.*

Theorem 1 implies the following lower bound on the optimal value of the original decentralized control problem (3.6):

$$J^*\left(\overline{\mathcal{G}_P(\Theta, \mathcal{G}_I)}\right) \leq J^*(\mathcal{G}_I). \tag{5.2}$$

Moreover, this lower bound can be computed via the solution of the convex infinite-dimensional optimization problem (3.10) that we specified in Section 3.2.1. In Theorem 2, we provide a finite-dimensional relaxation of problem (3.10) to enable the tractable approximation of the corresponding lower bound.

It is also worth noting that the transitive closure of the precedence graph induces an information structure under which each subsystem is guaranteed to have access to the information of those subsystems whose control input can directly or indirectly affect its information. This implies that the information relaxation $\overline{\mathcal{G}_P(\Theta, \mathcal{G}_I)}$ yields a partially nested information structure—a result that was originally shown in [29]. Theorem 1 improves upon this result by establishing the optimality of such a relaxation, in the sense that it is shown to yield the best lower bound among all partially nested information relaxations.

Remark 2 (Tightness of the Relaxation)**.** We also note that the information relaxation in Theorem 1 is *tight*. That is, $J^*\big(\overline{\mathcal{G}_P(\Theta, \mathcal{G}_I)}\big) = J^*(\mathcal{G}_I)$ for certain families of nonclassical control problems. In particular, it is known that signaling is *performance irrelevant* if the partially nested information relaxation only introduces additional information that is superfluous in terms of cost reduction—i.e., the additional information does not contribute to an improvement in performance. For such problems, one can establish the existence of an optimal policy under the partially nested information relaxation that also respects the original (nonclassical) information structure—implying the tightness of the relaxation. We refer the reader to [77], [113], [114, Sec. 3.5] for a rigorous explication of such claims. It can also be shown that the lower bound (5.2) is achieved by nonclassical LQG control problems that satisfy the so-called *substitutability condition*. See [5, Sec. 3] for a formal proof of this claim.

In Lemma 17, we present an alternative characterization of partially nested information structures that will prove useful in the proof of Theorem 1.

Lemma 17. $\mathcal{G} \in \mathrm{PN}(\Theta)$ if and only if $\mathcal{G} = \overline{\mathcal{G}_P(\Theta, \mathcal{G})}$.

The graph theoretic fixed-point condition in Lemma 17 implies that an information structure is partially nested if and only if the given information graph is equal to the transitive closure of the precedence graph that it induces. We also note that Lemma 17 is closely related to the graph theoretic necessary and sufficient condition for quadratic invariance presented in [95], which requires that the information graph be equal to its transitive closure, and be a supergraph of the transitive closure of the so-called plant graph.

Proof of Lemma 17: The proof of the "if" direction is straigthforward, and is omitted for brevity. We prove the "only if" direction of the statement. Let $\mathcal{G} = (\mathcal{V}, \mathcal{E})$. Assume that $\mathcal{G} \in \text{PN}(\Theta)$. It follows from Assumption 1 that $j \prec i$ if $(j, i) \in \mathcal{E}$. This implies that $\mathcal{G} \subseteq \mathcal{G}_P(\Theta, \mathcal{G})$, which in turn implies that $\mathcal{G} \subseteq \overline{\mathcal{G}_P(\Theta, \mathcal{G})}$.

To finish the proof, we will show that $\mathcal{G} \supseteq \overline{\mathcal{G}_P(\Theta, \mathcal{G})}$. This amounts to showing that $(j, i) \in \overline{\mathcal{E}_P(\Theta, \mathcal{G})}$ implies that $(j, i) \in \mathcal{E}$. Note that $(j, i) \in \overline{\mathcal{E}_P(\Theta, \mathcal{G})}$ implies that j is path connected to i in the corresponding precedence graph $\mathcal{G}_P(\Theta, \mathcal{G})$. That is, there exist $m \geq 1$ distinct vertices $v_1, \ldots, v_m \in \mathcal{V}$ that satisfy $j = v_1 \prec v_2 \prec \cdots \prec v_m = i$. Since $\mathcal{G} \in \text{PN}(\Theta)$, it also holds that $z_{v_1}(t) \subseteq z_{v_2}(t) \subseteq \cdots \subseteq z_{v_m}(t)$ for each time t. In particular, it holds that $z_j(t) \subseteq z_i(t)$ for each time t. This nesting of information, in combination with Assumption 2, implies that $(j, i) \in \mathcal{E}$. It follows that $\mathcal{G} \supseteq \overline{\mathcal{G}_P(\Theta, \mathcal{G})}$, which completes the proof. ∎

We have the following Corollary to Lemma 17 showing that any graph, which is feasible for problem (5.1), must also be a supergraph of the transitive closure of the precedence graph. In other words, this result precludes the existence of feasible information graph relaxations that do not contain $\overline{\mathcal{G}_P(\Theta, \mathcal{G}_I)}$ as a subgraph.

Corollary 1. If $\mathcal{G} \in \text{PN}(\Theta)$ and $\mathcal{G} \supseteq \mathcal{G}_I$, then $\mathcal{G} \supseteq \overline{\mathcal{G}_P(\Theta, \mathcal{G}_I)}$.

Proof of Corollary 1: Lemma 17 implies that $\mathcal{G} = \overline{\mathcal{G}_P(\Theta, \mathcal{G})}$. The result follows, as $\mathcal{G} \supseteq \mathcal{G}_I$ implies that $\overline{\mathcal{G}_P(\Theta, \mathcal{G})} \supseteq \overline{\mathcal{G}_P(\Theta, \mathcal{G}_I)}$. ∎

Proof of Theorem 1: Corollary 1 implies that $J^*\big(\overline{\mathcal{G}_P(\Theta, \mathcal{G}_I)}\big) \geq J^*(\mathcal{G})$ for every graph $\mathcal{G}$ that is feasible for problem (5.1). Hence, to prove the result, it suffices to show that the graph $\overline{\mathcal{G}_P(\Theta, \mathcal{G}_I)}$ is also feasible for problem (5.1). We previously showed in the proof of Lemma 17 that $\overline{\mathcal{G}_P(\Theta, \mathcal{G}_I)} \supseteq \mathcal{G}_I$. We complete the proof by showing that $\overline{\mathcal{G}_P(\Theta, \mathcal{G}_I)} \in \text{PN}(\Theta)$. It is not difficult to show that

$$\overline{\mathcal{G}_P(\Theta, \mathcal{G}_I)} = \overline{\mathcal{G}_P\left(\Theta, \overline{\mathcal{G}_P(\Theta, \mathcal{G}_I)}\right)}. \tag{5.3}$$

This follows from the observation that each precedence relation $i \prec j$ induced by $\overline{\mathcal{G}_P(\Theta, \mathcal{G}_I)}$ necessarily corresponds to an edge $(i, j) \in \overline{\mathcal{E}_P(\Theta, \mathcal{G}_I)}$. It follows from (5.3) and Lemma 17 that $\overline{\mathcal{G}_P(\Theta, \mathcal{G}_I)} \in \text{PN}(\Theta)$, which completes the proof. ■

5.3 A Dual Approach to Constraint Relaxation

The information relaxation developed in Section 5.2 provides a convex programming relaxation of the original decentralized control design problem (3.6). Despite its convexity, the resulting optimization problem remains computationally intractable due to its infinite-dimensionality. In what follows, we employ a general technique from robust optimization [51, 54, 65] to obtain a finite-dimensional relaxation of this problem through its partial dualization, and restriction to affine dual control policies. The resulting problem is a finite-dimensional conic program, whose optimal value is guaranteed to be a lower bound on the minimum cost of the original decentralized control design problem (3.6).

For the remainder of this section, we assume that the given information structure is partially nested, i.e., $\mathcal{G}_I \in \text{PN}(\Theta)$.

5.3.1 Restriction to Affine Dual Control Policies

The derivation of our lower bound centers on a partial Lagrangian relaxation of problem (3.6). We do so by introducing a *dual control policy* $v \in \mathcal{L}^2_m$, and dualizing the linear equality constraints on the state and input trajectories. This gives rise to the following min-max problem, which is equivalent to problem (3.6):

$$\begin{array}{ll} \text{minimize} & \sup_{v \in \mathcal{L}^2_m} \mathbf{E}\Big[x^\mathsf{T} R_x x + u^\mathsf{T} R_u u \\ & \qquad + \ v^\mathsf{T}(F_x x + F_u u + F_w w + s)\Big] \\ \text{subject to} & \gamma \in \Gamma(\mathcal{G}_I), \ \ s \in \mathcal{L}^2_m \\ & \left.\begin{array}{l} x = Bu + Gw \\ \eta = Pw \\ u = \gamma(\eta) \\ s \geq 0 \end{array}\right\} \ \mathbf{P}\text{-a.s.} \end{array} \tag{5.4}$$

In presenting the equivalent min-max reformulation of problem (3.6), we have used the fact that problem (3.6) is equivalent to problem (3.10); and Lemma 6, which implies that $\Phi(\mathcal{G}_I) = \Gamma(\mathcal{G}_I)$ if $\mathcal{G}_I \in \text{PN}(\Theta)$.

In order to obtain a tractable relaxation of problem (5.4), we restrict ourselves to dual control policies that are affine in the disturbance trajectory, i.e., $v = Vw$ for some $V \in \mathbf{R}^{m \times N_w}$. With this restriction, it is possible to derive a closed-form solution for the inner maximization in problem (5.4). This yields another minimization problem, whose optimal value stands as a lower bound on that of problem (5.4). We have the following result, which clarifies this claim.

Proposition 6. The optimal value of the following problem is a lower bound on the optimal value of problem (5.4):

$$
\begin{aligned}
\text{minimize} \quad & \sup_{V \in \mathbf{R}^{m \times N_w}} \mathbf{E}\Big[x^\mathsf{T} R_x x + u^\mathsf{T} R_u u \\
& \qquad\qquad + w^\mathsf{T} V^\mathsf{T} (F_x x + F_u u + F_w w + s)\Big] \\
\text{subject to} \quad & \gamma \in \Gamma(\mathcal{G}_I), \quad s \in \mathcal{L}_m^2 \\
& \left.\begin{aligned} & x = Bu + Gw \\ & \eta = Pw \\ & u = \gamma(\eta) \\ & s \geq 0 \end{aligned}\right\} \mathbf{P}\text{-a.s.}
\end{aligned}
\tag{5.5}
$$

Moreover, the optimal value of problem (5.5) equals that of the following optimization problem:

$$
\begin{aligned}
\text{minimize} \quad & \mathbf{E}\left[x^\mathsf{T} R_x x + u^\mathsf{T} R_u u\right] \\
\text{subject to} \quad & \gamma \in \Gamma(\mathcal{G}_I), \quad s \in \mathcal{L}_m^2 \\
& \mathbf{E}\left[(F_x x + F_u u + F_w w + s) w^\mathsf{T}\right] = 0 \\
& \left.\begin{aligned} & x = Bu + Gw \\ & \eta = Pw \\ & u = \gamma(\eta) \\ & s \geq 0 \end{aligned}\right\} \mathbf{P}\text{-a.s.}
\end{aligned}
\tag{5.6}
$$

Proof of Proposition 6: The fact that the optimal value of problem (5.5) lower bounds that of (5.4) is straightforward, since any dual affine control policy $v = Vw$ is feasible for the inner maximization problem in (5.4). To see that the optimal values of problem (5.5) and (5.6) are equal, we note that

$$\begin{aligned}
&\sup_{V\in\mathbf{R}^{m\times N_w}} \mathbf{E}\Big[w^\mathsf{T} V^\mathsf{T}(F_x x + F_u u + F_w w + s)\Big] \\
&\qquad = \sup_{V\in\mathbf{R}^{m\times N_w}} \mathbf{E}\Big[\mathrm{Tr}\left(V^\mathsf{T}(F_x x + F_u u + F_w w + s)w^\mathsf{T}\right)\Big] \\
&\qquad = \begin{cases} 0, & \text{if } \mathbf{E}\left[(F_x x + F_u u + F_w w + s)w^\mathsf{T}\right] = 0, \\ +\infty, & \text{otherwise}. \end{cases}
\end{aligned}$$

■

5.3.2 Relaxation to a Finite-dimensional Conic Program

Problem (5.6) appears to be intractable, as it entails the optimization over an infinite-dimensional function space. In what follows, we show that it admits a relaxation in the form of a second-order conic program under the additional assumption that the disturbance trajectory has an elliptically contoured distribution.

Assumption 4 (Elliptically Contoured Disturbance)**.** The disturbance trajectory w is assumed to have an elliptically contoured distribution.

In Appendix B, we provide a formal definition of elliptically contoured distributions and discuss its properties that are useful in optimal control problems. Note that Assumption 4 guarantees that the support of the disturbance trajectory w admits a representation of the form

$$\mathcal{W} = \{w \in \mathbf{R}^{N_w} \mid w_1 = 1 \text{ and } Lw \succeq_{\mathcal{K}_2} 0\},$$

where $L \in \mathbf{R}^{N_w\times N_w}$, and $\mathcal{K}_2$ denotes a second order cone of compatible dimension.

The family of elliptically contoured distributions is broad. It includes the multivariate Gaussian distribution, multivariate t-distribution, their truncated versions, and uniform distributions on ellipsoids. If w has an elliptically contoured distribution, then the conditional expectation of w given a subvector of w is affine in this subvector. And any linear transformation of w also follows an elliptically contoured distribution [25]. Such properties will play an integral role in the derivation of our relaxation of problem (5.6) as a second-order conic program.

Additionally, we require a formal definition[1] of the subspace of causal affine controllers respecting the information structure defined by $\mathcal{G}_I$.

Definition 6. Define $S(\mathcal{G}_I) \subseteq \mathbf{R}^{N_u \times N_y}$ to be the linear subspace of all causal affine controllers respecting the information structure defined by $\mathcal{G}_I$.

In other words, for all $K \in S(\mathcal{G}_I)$, the decentralized control policy defined by $\gamma(y) := Ky$ satisfies $\gamma \in \Gamma(\mathcal{G}_I)$. Equipped with this definition, we state the following result, which provides a finite-dimensional relaxation of problem (5.6) as a conic program. We note that the proposed conic relaxation is largely inspired by the duality-based relaxation methods originally developed in the context of centralized control design problems [54,65]. We provide a proof of Theorem 2 in the next subsection, which extends these techniques to accommodate the added complexity of decentralized information constraints on the controller.

[1]Note that the subspace $S(\mathcal{G}_I)$ has been previously defined in Section 3.3.1. We redefine this subspace in this chapter in order to guarantee that this chapter is self-contained.

Theorem 2. Let Assumption 4 hold. If $\mathcal{G}_I \in \text{PN}(\Theta)$, then the optimal value of the following problem is a lower bound on the optimal value of problem (3.6):

$$\begin{aligned} \text{minimize} \quad & \text{Tr}\left(P^\mathsf{T} Q^\mathsf{T} R Q P M + 2G^\mathsf{T} R_x B Q P M + G^\mathsf{T} R_x G M\right) \\ \text{subject to} \quad & Q \in S(\mathcal{G}_I), \quad Z \in \mathbf{R}^{m \times N_w} \\ & (F_u + F_x B) Q P + F_x G + F_w + Z = 0, \\ & L M Z^\mathsf{T} \succeq_{\mathcal{K}_2} 0, \\ & e_1^\mathsf{T} M Z^\mathsf{T} \geq 0, \end{aligned} \tag{5.7}$$

where $R = R_u + B^\mathsf{T} R_x B$, and $e_1 = (1, 0, \ldots, 0)$ is a unit vector in $\mathbf{R}^{N_w}$.

Let $J^d(\mathcal{G}_I)$ denote the optimal value of the finite-dimensional conic program (5.7). Theorem 2 states that $J^d(\mathcal{G}_I) \leq J^*(\mathcal{G}_I)$ if $\mathcal{G}_I \in \text{PN}(\Theta)$. The following result—an immediate corollary to Theorems 1 and 2—provides a computationally tractable lower bound for problems with nonclassical information structures.

Corollary 2. Let $J^*(\mathcal{G}_I)$ denote the optimal value of the decentralized control design problem (3.6). It follows that

$$J^d(\overline{\mathcal{G}_P(\Theta, \mathcal{G}_I)}) \leq J^*(\mathcal{G}_I),$$

where $\overline{\mathcal{G}_P(\Theta, \mathcal{G}_I)}$ refers to the transitive closure of the precedence graph associated with problem (3.6).

5.3.3 Proof of Theorem 2

The crux of the proof centers on the introduction of new finite-dimensional decision variables, which enable the removal of the infinite-dimensional decision variables in problem (5.6). Consider the following result.

Lemma 18. Let Assumption 4 hold. For each $s \in \mathcal{L}^2_m$, there exists a matrix $Z \in \mathbf{R}^{m \times N_w}$ that satisfies

$$ZM = \mathbf{E}\big[sw^\mathsf{T}\big]. \tag{5.8}$$

For each $\gamma \in \Gamma(\mathcal{G}_I)$, there exists a matrix $Q \in S(\mathcal{G}_I)$ that satisfies

$$QPM = \mathbf{E}\big[uw^\mathsf{T}\big], \tag{5.9}$$

where $u = \gamma(\eta)$.

Proof: This proof extends arguments originally developed in [54, Lem 4.4] to accommodate the more general setting considered in this note, where the affine controller Q is subject to a decentralized information constraint.

Proof of the first part: Fix $s \in \mathcal{L}^2_m$. The matrix M is invertible, since it is assumed to be positive definite. Setting $Z = \mathbf{E}[sw^\mathsf{T}]M^{-1}$ yields the desired result in (5.8).

Proof of the second part: We first introduce the notion of a truncation operator. Given a nonempty set of indices $J \subseteq \{1, \ldots, N_y\}$ we define the *truncation operator* $\Pi_J : \mathbf{R}^{N_y} \rightarrow \mathbf{R}^{|J|}$ as the mapping from a vector x to its subvector x_J, i.e., $x_J = \Pi_J x$.

Now, fix $\gamma \in \Gamma(\mathcal{G}_I)$, and let $u = \gamma(\eta)$. The following Lemma will prove useful in establishing the existence of a matrix $Q \in S(\mathcal{G}_I)$ satisfying Eq. (5.9).

Lemma 19. Let Assumption 4 hold. Let $z \in \mathcal{L}^2_1$ be random variable that is a (possibly nonlinear) function of the random vector $\eta_J = \Pi_J \eta$, where $\{1\} \subseteq J \subseteq \{1, \ldots, N_y\}$ is a given index set. Then, there exists another random variable

$\widetilde{z} \in \mathcal{L}_1^2$, which is an *affine*[2] function of η_J, and satisfies $\mathbf{E}\big[\widetilde{z}\eta^\mathsf{T}\big] = \mathbf{E}\big[z\eta^\mathsf{T}\big]$.

Proof: Define the vector $r \in \mathbf{R}^{|J|}$ according to

$$r^\mathsf{T} := \mathbf{E}\big[z\eta_J^\mathsf{T}\big]\big(P_J M P_J^\mathsf{T}\big)^\dagger, \tag{5.10}$$

where $P_J := \Pi_J P$, and $(\cdot)^\dagger$ denotes the Moore-Penrose pseudoinverse of a matrix. We first show that the vector r satisfies

$$\mathbf{E}\big[z\eta_J^\mathsf{T}\big] = r^\mathsf{T} P_J M P_J^\mathsf{T}. \tag{5.11}$$

Define the matrix $\Psi := P_J M^{1/2}$, where $M^{1/2}$ is the unique square root of the symmetric positive definite matrix M. Note that the matrix $M^{1/2}$ is symmetric and positive definite (and hence invertible). It holds that

$$\begin{aligned}
r^\mathsf{T} P_J M P_J^\mathsf{T} &= \mathbf{E}\big[z\eta_J^\mathsf{T}\big]\big(P_J M P_J^\mathsf{T}\big)^\dagger P_J M P_J^\mathsf{T} \\
&= \mathbf{E}\big[zw^\mathsf{T}\big] P_J^\mathsf{T}\big(P_J M P_J^\mathsf{T}\big)^\dagger P_J M P_J^\mathsf{T} \\
&= \mathbf{E}\big[zw^\mathsf{T}\big] M^{-1/2} M^{1/2} P_J^\mathsf{T}\big(P_J M P_J^\mathsf{T}\big)^\dagger P_J M P_J^\mathsf{T} \\
&= \mathbf{E}\big[zw^\mathsf{T}\big] M^{-1/2} \Psi^\mathsf{T}\big(\Psi\Psi^\mathsf{T}\big)^\dagger \Psi\Psi^\mathsf{T} \\
&= \mathbf{E}\big[zw^\mathsf{T}\big] M^{-1/2} \Psi^\dagger \Psi\Psi^\mathsf{T} \\
&= \mathbf{E}\big[zw^\mathsf{T}\big] M^{-1/2} \Psi^\mathsf{T} = \mathbf{E}\big[zw^\mathsf{T}\big] P_J^\mathsf{T} = \mathbf{E}\big[z\eta_J^\mathsf{T}\big]
\end{aligned}$$

The second and the last equalities both follow from the fact that $\eta_J = \Pi_J P w = P_J w$. The fourth equality follows from the definition of the matrix Ψ and the symmetry of the matrix $M^{1/2}$. The fifth equality follows from the fact [12, Prop. 3.2] that $\Psi^\mathsf{T}\big(\Psi\Psi^\mathsf{T}\big)^\dagger = \Psi^\dagger$. The sixth equality follows from the fact [12, Prop. 3.1] that $\Psi^\dagger\Psi\Psi^\mathsf{T} = \Psi^\mathsf{T}$. It follows that the vector r satisfies Eq. (5.11).

[2]As a matter of notational convenience, we have required that $1 \in J$. This enables one to represent affine functions of η_J as linear functions of η_J, since $\eta_1 = 1$ by construction.

Now, define the random variable $\widetilde{z} := r^{\mathsf{T}}\eta_J$. Clearly, $\widetilde{z}$ is an affine function of η_J. We complete the proof by showing that $\widetilde{z}$ satisfies

$$\mathbf{E}\big[\widetilde{z}\eta^{\mathsf{T}}\big] = \mathbf{E}\big[z\eta^{\mathsf{T}}\big]. \tag{5.12}$$

First note that the combination of Assumption 4 and Lemma 25 in Appendix B implies that the random vector $(w, \eta) = (w, Pw)$ has an elliptically contoured distribution. Hence, it follows from Proposition 11 that the conditional expectation of η given η_J is affine in η_J. The assumption that $1 \in J$ guarantees that the first entry of η_J equals 1. Hence, there exists a matrix $Y_J \in \mathbf{R}^{N_y \times |J|}$, such that

$$\mathbf{E}\left[\eta \,|\eta_J\right] = Y_J\eta_J \quad \mathbf{P}\text{-a.s.}. \tag{5.13}$$

It holds that

$$\mathbf{E}\big[z\eta^{\mathsf{T}}\big] = \mathbf{E}\Big[\mathbf{E}\big[z\eta^{\mathsf{T}}\big|\eta_J\big]\Big] = \mathbf{E}\big[z\eta_J^{\mathsf{T}}\big]Y_J^{\mathsf{T}} = r^{\mathsf{T}}P_JMP_J^{\mathsf{T}}Y_J^{\mathsf{T}}$$

Here, the first equality follows from the law of iterated expectations. The second equality follows from a combination of Eq. (5.13) and the assumption that z is a function of η_J. The third equality follows from Eq. (5.11). It also holds that

$$r^{\mathsf{T}}P_JMP_J^{\mathsf{T}}Y_J^{\mathsf{T}} = r^{\mathsf{T}}\mathbf{E}\big[\eta_J\eta_J^{\mathsf{T}}\big]Y_J^{\mathsf{T}} = r^{\mathsf{T}}\mathbf{E}\Big[\mathbf{E}\big[\eta_J\eta^{\mathsf{T}}\big|\eta_J\big]\Big] = r^{\mathsf{T}}\mathbf{E}\big[\eta_J\eta^{\mathsf{T}}\big] = \mathbf{E}\big[\widetilde{z}\eta^{\mathsf{T}}\big],$$

which completes the proof. ∎

Stated in other words, Lemma 19 asserts the existence of a vector $q \in \mathbf{R}^{N_y}$ that satisfies

$$\mathbf{E}\big[q^{\mathsf{T}}\eta\eta^{\mathsf{T}}\big] = \mathbf{E}\big[z\eta^{\mathsf{T}}\big], \tag{5.14}$$

where the vector q respects the sparsity pattern encoded by the index set J, i.e., $q = \Pi_J^{\mathsf{T}}\Pi_J q$. It follows that one can apply Lemma 19 to each row of the matrix $\mathbf{E}[u\eta^{\mathsf{T}}]$ to establish the existence of a matrix $Q \in S(\mathcal{G}_I)$ that satisfies

$$\mathbf{E}\big[Q\eta\eta^{\mathsf{T}}\big] = \mathbf{E}\big[u\eta^{\mathsf{T}}\big]. \tag{5.15}$$

Consider a matrix $Q \in S(\mathcal{G}_I)$ that satisfies Eq. (5.15). We complete the proof by showing that this matrix also satisfies Eq. (5.9). First recall that we have argued in the proof of Lemma 19 that (w, η) admits an elliptically contoured distribution. Hence, it follows from Proposition 11 that the conditional expectation of w given η is an affine function of η. The definition of the matrix P ensures that $\eta_1 = 1$. Hence, the conditional expectation can be expressed as

$$\mathbf{E}[w|\eta] = Y\eta \quad \mathbf{P}\text{-a.s.} \tag{5.16}$$

for some matrix $Y \in \mathbf{R}^{N_w \times N_y}$. It holds that

$$\mathbf{E}\big[uw^\mathsf{T}\big] = \mathbf{E}\Big[\mathbf{E}\big[uw^\mathsf{T}|\eta\big]\Big] = \mathbf{E}\big[u\eta^\mathsf{T}\big]Y^\mathsf{T} = Q\mathbf{E}\big[\eta\eta^\mathsf{T}\big]Y^\mathsf{T}.$$

Here, the first equality follows from the law of iterated expectations; the second equality follows from the fact that $u = \gamma(\eta)$ and a direct application of Eq. (5.16); and the third equality follows from Lemma 19. It also holds that

$$\mathbf{E}\big[\eta\eta^\mathsf{T}\big]Y^\mathsf{T} = \mathbf{E}\Big[\eta\mathbf{E}\big[w^\mathsf{T}|\eta\big]\Big] = \mathbf{E}\Big[\mathbf{E}\big[\eta w^\mathsf{T}|\eta\big]\Big] = \mathbf{E}\big[\eta w^\mathsf{T}\big] = PM,$$

which completes the proof. ■

With Lemma 18 in hand, we obtain an equivalent reformulation of problem (5.6) as the following optimization problem—via the introduction of the finite-dimensional decision variables Z and Q through the constraints (5.8) and (5.9), respectively.

$$
\begin{aligned}
\text{minimize} \quad & \mathbf{E}\left[u^{\mathsf{T}} R u\right] + \operatorname{Tr}\left(2G^{\mathsf{T}} R_x B Q P M + G^{\mathsf{T}} R_x G M\right) \\
\text{subject to} \quad & \gamma \in \Gamma(\mathcal{G}_I), \quad s \in \mathcal{L}_m^2, \quad Q \in S(\mathcal{G}_I), \quad Z \in \mathbf{R}^{m \times N_w} \\
& (F_u + F_x B) Q P M + F_x G M + F_w M + Z M = 0 \\
& QPM = \mathbf{E}\big[u w^{\mathsf{T}}\big] \\
& ZM = \mathbf{E}\big[s w^{\mathsf{T}}\big] \\
& \left.\begin{aligned} & \eta = P w \\ & u = \gamma(\eta) \\ & s \geq 0 \end{aligned}\right\} \mathbf{P}\text{-a.s.}
\end{aligned}
\tag{5.17}
$$

where $R = R_u + B^{\mathsf{T}} R_x B$.

We now introduce two technical Lemmas that permit us to construct a finite-dimensional relaxation of problem (5.17).

Lemma 20. Fix the matrix $Q \in S(\mathcal{G}_I)$. It follows that $\gamma(\eta) = Q\eta$ is an optimal solution to the following optimization problem:

$$
\begin{aligned}
\text{minimize} \quad & \mathbf{E}\left[u^{\mathsf{T}} R u\right] \\
\text{subject to} \quad & \gamma \in \Gamma(\mathcal{G}_I) \\
& QPM = \mathbf{E}\big[u w^{\mathsf{T}}\big] \\
& \left.\begin{aligned} & \eta = P w \\ & u = \gamma(\eta) \end{aligned}\right\} \mathbf{P}\text{-a.s.}
\end{aligned}
$$

We omit the proof of Lemma 20, as it is an immediate corollary of [54, Lem. 4.5]. A direct application of Lemma 20 yields the following equivalent reformulation of problem (5.17) as:

$$
\begin{aligned}
\text{minimize} \quad & \text{Tr}\left(P^\mathsf{T} Q^\mathsf{T} R Q P M + 2G^\mathsf{T} R_x B Q P M + G^\mathsf{T} R_x G M\right) \\
\text{subject to} \quad & s \in \mathcal{L}^2_m, \;\; Q \in S(\mathcal{G}_I), \;\; Z \in \mathbf{R}^{m \times N_w} \\
& (F_u + F_x B) Q P + F_x G + F_w + Z = 0 \\
& ZM = \mathbf{E}\left[s w^\mathsf{T}\right] \\
& s \geq 0 \quad \mathbf{P}\text{-a.s.}
\end{aligned}
\tag{5.18}
$$

Note that, in reformulating problem (5.17), we have eliminated the second-order moment matrix M from the equality constraint $(F_u + F_x B)QPM + F_x GM + F_w M + ZM = 0$, as M is assumed to be positive definite, and, therefore, invertible.

Lemma 21 provides a conic relaxation of the constraints in problem (5.18) involving the infinite-dimensional decision variable $s \in \mathcal{L}^2_m$.

Lemma 21. If $s \in \mathcal{L}^2_m$ and $Z \in \mathbf{R}^{m \times N_w}$ satisfy $ZM = \mathbf{E}[sw^\mathsf{T}]$ and $s \geq 0$ $\mathbf{P}$-a.s., then $LMZ^\mathsf{T} \succeq_{\mathcal{K}_2} 0$ and $e_1^\mathsf{T} M Z^\mathsf{T} \geq 0$.

Proof: It follows from the symmetry of the matrix M that $MZ^\mathsf{T} = (ZM)^\mathsf{T} = \mathbf{E}[ws^\mathsf{T}]$. It, therefore, holds that

$$e_1^\mathsf{T} M Z^\mathsf{T} = e_1^\mathsf{T} \mathbf{E}[ws^\mathsf{T}] = \mathbf{E}[e_1^\mathsf{T} w s^\mathsf{T}] = \mathbf{E}[s^\mathsf{T}] \geq 0.$$

The last equality follows from the fact that $e_1^\mathsf{T} w = 1$ $\mathbf{P}$-almost surely. To show that $LMZ^\mathsf{T} \succeq_{\mathcal{K}_2} 0$, it suffices to show columnwise inclusion in the second-order cone, i.e.,

$$L\mathbf{E}[s_i w] \succeq_{\mathcal{K}_2} 0, \quad \text{for } \; i = 1, \ldots, m,$$

where $s_i \in \mathcal{L}^2_1$ is the i^{th} element of the random vector s. By definition, we have that $Lw \succeq_{\mathcal{K}_2} 0$ for all $w \in \mathcal{W}$. Also, since $s_i \geq 0$ almost surely, we have that

$L(s_i w) \succeq_{\mathcal{K}_2} 0$ almost surely. It follows from the convexity of the second-order cone that $L\mathbf{E}[s_i w] \succeq_{\mathcal{K}_2} 0$. ■

We complete the proof with the following string of inequalities and equalities relating the optimal values of the various optimization problems formulated thus far.

$$\text{(5.7)} \underset{(a)}{\leq} \text{(5.18)} \underset{(b)}{=} \text{(5.17)} \underset{(c)}{=} \text{(5.6)} \underset{(d)}{\leq} \text{(5.4)} \underset{(e)}{=} \text{(3.10)} = \text{(3.6)}$$

Inequality (a) follows from Lemma 21, which implies that problem (5.7) is a relaxation of problem (5.18). Equality (b) follows from Lemma 20. Equality (c) follows from Lemma 18. Inequality (d) follows from Proposition 6. Finally, Equality (e) follows from Lemma 6, as the assumption of a partially nested information structure implies equivalence between the optimal values of problems (5.4) and (3.10). The equivalence between (3.10) and (3.6) is argued in Section 3.2.1.

CHAPTER 6

DECENTRALIZED CONTROL OF DISTRIBUTED ENERGY RESOURCES

In this chapter, we describe an application of the techniques developed in this dissertation to the decentralized control of distributed energy resources in power distribution systems. We consider the decentralized control of radial distribution systems with controllable photovoltaic inverters and energy storage resources. For such systems, we investigate the problem of designing fully decentralized controllers that minimize the expected cost of balancing demand, while guaranteeing the satisfaction of individual resource and distribution system voltage constraints. Employing a linear approximation of the branch flow model, we formulate this problem as the design of a decentralized disturbance-feedback controller that minimizes the expected value of a convex quadratic cost function, subject to robust convex quadratic constraints on the system state and input. The optimal control policy for such problems is, in general, intractable to compute. We apply the techniques we developed in Chapter 3 to derive a tractable inner approximation of the decentralized control design problem, which enables the efficient computation of an affine control policy via the solution of a finite-dimensional conic program. As affine control policies are, in general, suboptimal for the family of systems considered, we apply our results in Chapter 5 to bound their suboptimality via the solution of another finite-dimensional conic program. A case study of a 12 kV radial distribution system demonstrates that decentralized affine controllers can perform close to optimal.

6.1 Introduction

The increasing penetration of distributed and renewable energy resources introduces challenges to the operation of power distribution systems, including rapid fluctuations in bus voltage magnitudes, reverse power flows at distribution substations, and deteriorated power quality due to the intermittency of supply from renewables. These challenges are exasperated by the fact that traditional techniques for distribution system management, including the deployment of on-load tap changing (OLTC) transformers and shunt capacitors, cannot effectively deal with the rapid variation in the power supplied from renewable resources [28]. In this chapter, we illustrate how the decentralized control design techniques developed in this dissertation might be applied to address such challenges in the operation of power distribution systems. Specifically, our objective is to develop a systematic approach to the design of decentralized feedback controllers for distribution networks with a high penetration of distributed solar and energy storage resources, which minimizes the expected cost of meeting demand over a finite horizon, while respecting network and resource constraints.

Related Work: Although current industry standards require that photovoltaic (PV) inverters operate at a unity power factor [1], the latent reactive power capacity of PV inverters can be utilized to regulate voltage profiles [45,47,70,85,102,119], and reduce active power losses [14,22,32–35,46,62,94,116] in distribution networks. A large swath of the literature on the reactive power management of PV inverters prescribes the solution of an optimal power flow

(OPF) problem to determine the reactive power injections of PV inverters in real time [22, 32–34, 45–47, 70, 85, 94, 102, 116, 119]. The resulting OPF problem must be repeatedly solved over fast time scales (e.g., every minute) to accommodate the rapid fluctuations in the active power supplied from the PV resources. In the presence of a large number of PV resources, the sheer size of the resulting OPF problem that needs to be solved, and the communication requirements it entails, gives rise to the need for distributed optimization methods [22, 33, 34, 47, 70, 85, 94, 116, 119]. In particular, there has emerged a recent stream of literature developing distributed optimization methods, which enable the real-time control of reactive power injections of PV inverters using only local measurements of bus voltage magnitudes [47,70,85,119]. Under the assumption that the underlying OPF problem being solved is time-invariant, such methods are guaranteed to asymptotically converge to the globally optimal reactive power injection profile. There is, however, no guarantee on the performance or constraint-satisfaction of these methods in finite time. The aforementioned methods can be interpreted as being fully decentralized, in that they do not require the explicit exchange of information between local controllers. Instead, the local controllers can be interpreted as communicating implicitly through the distribution network, which couples them physically. There exists another class of distributed optimization methods, which rely on the explicit exchange of information between neighboring controllers through a digital communication network [22, 33, 34, 94, 116]. Additionally, there exists a related stream of literature, which aims to explicitly treat uncertainty in renewable supply and demand by leveraging on methods grounded in stochastic optimization [14, 35, 62].

In addition to the reactive power control of PV inverters, one can imagine a future power system in which a broader class of distributed energy resources with storage capability (e.g., electric vehicles, standalone battery packs) are actively controlled to mitigate voltage fluctuations and distribution system losses [27, 28, 31, 37, 49, 74, 78, 96, 106]. As the set of feasible power injections to and withdrawals from energy storage systems are naturally coupled across time, the problem of managing their operation amounts to a multi-period, constrained stochastic control problem [27, 28, 31, 37, 49, 74, 106]. In the presence of network constraints and uncertainty in demand and renewable supply, the calculation of the optimal control policy is, in general, computationally intractable. These computational difficulties in control design are underscored by a recent report from the U.S. Department of Energy pointing to an apparent lack of effective control methods capable of "integrating [PV] inverter controls with control of other DERs or the management of uncertainty from intermittent generation" [103][p. 31]. The development of computational methods to enable the tractable calculation of feasible control policies with computable bounds on their suboptimality is therefore desired, and stands as the primary subject of this chapter.

Contribution: The setting we consider entails the decentralized control of distributed energy resources spread throughout a radial distribution network, subject to uncertainty in demand and renewable supply. The power flow equations over the radial network are described according to a linearized branch flow model. Our objective is to minimize the expected amount of active power supplied at the substation required to meet demand, while guaranteeing the satisfaction of network and individual resource constraints. For the setting considered, this is technically equivalent to minimizing the expected active power losses plus the terminal energy stored in the distribution network. The determination of an optimal decentralized control policy for such problems is, in general, computationally intractable, due to the presence of stochastic disturbances and hard constraints on the system state and and input. Our primary contributions are two-fold. First, we develop a convex programming approach to the design of decentralized, affine disturbance-feedback controllers. Second, as such control policies are, in general, suboptimal, we provide a technique to bound their suboptimality through the solution of another convex program. We verify that the decentralized affine policies we derive are close to optimal for the problem instance considered in our case study.

Organization: The remainder of this chapter is organized as follows. Section 6.2 describes the architecture of the controllers that might be used in the control of distributed energy resources in a large-scale power distribution network. Section 6.3 describes our models of the distribution network and the distributed energy resources. Section 6.4 formally states the decentralized control design problem. Section 6.5 describes an approach to the computation of decentralized affine control policies via a finite-dimensional conic program. Section 6.6 describes a method to implement the decentralized affine control policy over a time-scale that is more fine-grained that the one used in the control design. Section 6.7 describes an approach to the tractable calculation of guaranteed bounds on the suboptimality incurred by these affine control policies via another finite-dimensional conic program. Section 6.8 demonstrates the proposed techniques with a numerical study of a 12 kV radial distribution network. Section 6.9 concludes this chapter.

6.2 The Controller Architecture

In this section, we describe a two-layer control architecture that fulfills several (possibly conflicting) requirements on the control of distributed energy resources in power distribution systems. We start with a summary of such requirements. First, the control actions need to be performed in real-time, as active power supply from distributed energy resources and load might change at a fast time-scale (seconds to minutes). Second, the controller is required to have provable guarantees on constraint satisfaction, in the sense that the controller we design should induce control inputs that are guaranteed to respect individual resource capacity and network voltage magnitude constraints. Finally, the controller is expected to have close-to-optimal performance. That is to say, the cost that is incurred by the controller we design should be close to the minimum cost that can be attained by a feasible controller.

In Figure 6.1, we describe a two-layer controller architecture that fulfills the aforementioned requirements. We assume that the model parameters of the power distribution network under consideration is known exactly. Given the model parameters, the task of the *optimization* layer is to solve a decentralized control design problem, which amounts to computing a decentralized control policy that minimizes the expected amount of active power that is supplied at the substation, while guaranteeing the satisfaction of individual resource capacity and network voltage magnitude constraints. On the other hand, the task of the *direct control* layer is to deploy and implement the decentralized control policy that is computed at the optimization layer over a fast time-scale (seconds to minutes). Specifically, the direct control layer directly interacts with the distributed energy resources that are located at the "grid-edge", and specifies their

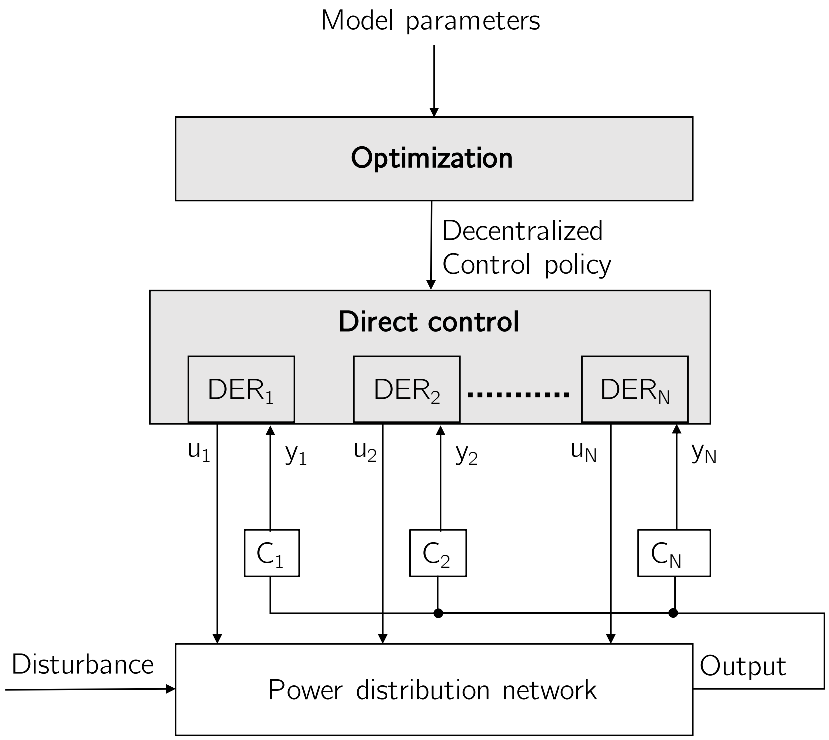

Figure 6.1: A two-layer architecture for the control of distributed energy resources in power distribution systems.

active and reactive power injections in real-time based on the local sensor measurements.

We note that one potential limitation of this two-layer controller architecture is the assumption that the system model parameters are known exactly. Such an assumption is not necessarily satisfied in the control of power distribution networks. For example, the topology of many practical power distribution networks might be unknown [38]. Even if the network topology is known, the impedance of the distribution lines might be unknown or inexact. The question as to how one might attain a proper tradeoff between learning the system parameters (exploration) and controlling the system based on current information (exploitation) is left as an important direction for future work—see Chapter 7 for some initial thoughts.

6.3 Network and Resource Models

6.3.1 Branch Flow Model

Consider a radial distribution network whose topology is described by a *rooted tree* $\mathcal{G} = (\mathcal{V}, \mathcal{E})$, where $\mathcal{V} = \{0, 1, .., n\}$ denotes its set of (nodes) buses, and $\mathcal{E}$ its set of (directed edges) distribution lines. In particular, bus 0 is defined as the root of the network, and represents the substation that connects to the external power system. Each directed distribution line admits the natural orientation, i.e., away from the root. For each distribution line $(i, j) \in \mathcal{E}$, we denote by $r_{ij} + \mathbf{i}x_{ij}$ its *impedance*. In addition, define I_{ij} as the *complex current* flowing from bus i to j, and $p_{ij} + \mathbf{i}q_{ij}$ as the *complex power* flowing from bus i to j. For each bus $i \in \mathcal{V}$, let V_i denote its *complex voltage*, and $p_i + \mathbf{i}q_i$ the *complex power injection* at this bus. We assume that the complex voltage V_0 at the substation is fixed and known.

We employ the *branch flow model* proposed in [10, 11] to describe the steady-state, single-phase AC power flow equations associated with this radial distribution network. In particular, for each bus $j = 1, \ldots, n$, and its unique *parent* $i \in \mathcal{V}$, we have

$$-p_j = p_{ij} - r_{ij}\ell_{ij} - \sum_{k:(j,k)\in\mathcal{E}} p_{jk}, \tag{6.1}$$

$$-q_j = q_{ij} - x_{ij}\ell_{ij} - \sum_{k:(j,k)\in\mathcal{E}} q_{jk}, \tag{6.2}$$

$$v_j^2 = v_i^2 - 2(r_{ij}p_{ij} + x_{ij}q_{ij}) + (r_{ij}^2 + x_{ij}^2)\ell_{ij}, \tag{6.3}$$

$$\ell_{ij} = (p_{ij}^2 + q_{ij}^2)/v_i^2, \tag{6.4}$$

where $\ell_{ij} = |I_{ij}|^2$ and $v_i = |V_i|$. We note that the branch flow model is well defined only for radial distribution networks, as we require that each bus j (excluding the substation) have a unique parent $i \in \mathcal{V}$.

For the remainder of the chapter, we consider a linear approximation of the branch flow model (6.1)-(6.4) based on the Simplified Distflow method developed in [9]. The derivation of this approximation relies on the assumption that $\ell_{ij} = 0$ for all $(i, j) \in \mathcal{E}$, as the active and reactive power losses on distribution lines are considered small relative to the power flows. According to [45, 58], such an approximation tends to introduces a relative model error of 1-5% for practical distribution networks. Under this assumption, Eq. (6.1)-(6.3) can be simplified to

$$-p_j \;=\; p_{ij} - \sum_{k:(j,k)\in\mathcal{E}} p_{jk}, \tag{6.5}$$

$$-q_j \;=\; q_{ij} - \sum_{k:(j,k)\in\mathcal{E}} q_{jk}, \tag{6.6}$$

$$v_j^2 \;=\; v_i^2 - 2(r_{ij}p_{ij} + x_{ij}q_{ij}). \tag{6.7}$$

The linearized branch flow Eq. (6.5)-(6.7) can be written more compactly as

$$v^2 = Rp + Xq + v_0^2\mathbf{1}. \tag{6.8}$$

Here, $v^2 = (v_1^2, .., v_n^2)$, $p = (p_1, .., p_n)$, and $q = (q_1, .., q_n)$ denote the vectors of squared bus voltage magnitudes, real power injections, and reactive power injections, respectively, and $\mathbf{1} = (1, .., 1)$ is a vector of all ones in $\mathbf{R}^n$. The matrices $R, X \in \mathbf{R}^{n\times n}$ are defined according to

$$R_{ij} = 2\sum_{(h,k)\in\mathcal{P}_i\cap\mathcal{P}_j} r_{hk},$$

$$X_{ij} = 2\sum_{(h,k)\in\mathcal{P}_i\cap\mathcal{P}_j} x_{hk},$$

where $\mathcal{P}_i \subset \mathcal{E}$ is defined as the set of edges on the unique path from bus 0 to i.

In the sequel, we will consider the control of the distribution system over discrete time periods indexed by $t = 0, \cdots, T-1$. Each discrete time period t is defined over a time interval of length Δ. We require the vector of bus voltage magnitudes $v(t) = (v_1(t), .., v_n(t)) \in \mathbf{R}^n$ at each time period t to satisfy

$$\underline{v} \leq v(t) \leq \overline{v}, \tag{6.9}$$

where the allowable range of voltage magnitudes is defined by $\underline{v}, \overline{v} \in \mathbf{R}^n$.

6.3.2 Energy Storage Model

We consider a distribution system consisting of n perfectly efficient energy storage devices, where each bus i (excluding the substation) is assumed to have an energy storage capacity of $b_i \in \mathbf{R}$. The dynamic evolution of each energy storage device i is described according to the state equation

$$x_i(t+1) = x_i(t) - \Delta p_i^S(t), \quad t = 0, \ldots, T-1, \tag{6.10}$$

where the state $x_i(t) \in \mathbf{R}$ denotes the amount of energy stored in storage device i just preceding period t, and $p_i^S(t) \in \mathbf{R}$ denotes the active power extracted from device i during period t. For ease of exposition, we assume that the initial condition $x_i(0)$ of each storage device is fixed and known.[1] We impose state and input constraints of the form

$$0 \leq x_i(t) \leq b_i, \quad t = 0, \ldots, T \tag{6.11}$$

$$\underline{p}_i^S \leq p_i^S(t) \leq \overline{p}_i^S, \quad t = 0, \ldots, T-1. \tag{6.12}$$

[1] We emphasize that all results presented in this chapter are easily generalized to the setting in which the initial condition $x_i(0)$ is modeled as a random variable with known distribution. In particular, one can treat the initial condition as an additive disturbance to the state equation at time period $t = 0$. We refer the readers to [72] for a detailed treatment of such systems.

for $i = 1, .., n$. The interval $[\underline{p}_i^S,\ \overline{p}_i^S] \subset \mathbf{R}$ defines the range of allowable inputs for storage device i at each time period t.

6.3.3 Photovoltaic Inverter Model

We assume that, in addition to energy storage capacity, each bus i (excluding the substation) has a photovoltaic (PV) inverter whose reactive power injection can be actively controlled. We denote by $w_i^I(t) \in \mathbf{R}$ the active power injection, and by $q_i^I(t) \in \mathbf{R}$ the reactive power injection from the PV inverter at bus i and time t. Due to the intermittency of solar irradiance, we will model $w_i^I(t)$ as a discrete-time stochastic process, whose precise specification is presented in Section 6.3.5. Additionally, we require that the reactive power injections respect capacity constraints of the form

$$\left|q_i^I(t)\right| \leq \sqrt{{s_i^I}^2 - w_i^I(t)^2}, \quad i = 1, \ldots, n, \tag{6.13}$$

for $t = 0, \ldots, T-1$. Here, $s_i^I \in \mathbf{R}$ denotes the apparent power capacity of PV inverter i. Clearly, it must hold that $w_i^I(t) \leq s_i^I$.

6.3.4 Load Model

Each bus in the distribution network is assumed to have a constant power load, which we will treat as a discrete-time stochastic process. Accordingly, we denote by $w_i^p(t) \in \mathbf{R}$ and $w_i^q(t) \in \mathbf{R}$ the active and reactive power demand, respectively, at bus i and time t. It follows that the nodal active and reactive power balance equations can be expressed as

$$p_i(t) \ = \ p_i^S(t) + w_i^I(t) - w_i^p(t), \tag{6.14}$$

$$q_i(t) \ = \ q_i^I(t) - w_i^q(t), \tag{6.15}$$

where $p_i(t) \in \mathbf{R}$ and $q_i(t) \in \mathbf{R}$ denote the net active and reactive power injections, respectively, at each bus $i = 1, \ldots, n$ and time period $t = 0, \ldots, T-1$.

6.3.5 Uncertainty Model

As indicated earlier, we model the active power demand, reactive power demand, and PV active power supply as discrete-time stochastic processes. Accordingly, we associate with each bus i a *disturbance process* defined as $w_i(t) = (w_i^p(t), w_i^q(t), w_i^I(t)) \in \mathbf{R}^3$. We define the *full disturbance trajectory* as

$$w = (1, w(0), \ldots, w(T-1)) \in \mathbf{R}^{N_w}, \tag{6.16}$$

where $N_w = 1 + 3nT$ and $w(t) = (w_1(t), \ldots, w_n(t)) \in \mathbf{R}^{3n}$ for each time period t. Note that, in our specification of the disturbance trajectory w, we have included a constant scalar as its initial component. Such notational convention is made for simplifying the specification of affine control policies.

We assume that the disturbance trajectory w has support $\mathcal{W}$ that is a nonempty and compact subset of $\mathbf{R}^{N_w}$, representable by

$$\mathcal{W} = \{w \in \mathbf{R}^{N_w} \mid w_1 = 1 \text{ and } Hw \succeq_{\mathcal{K}} 0\},$$

where the matrix $H \in \mathbf{R}^{\ell \times N_w}$ is known. It follows from the compactness of $\mathcal{W}$ that the second-order moment matrix

$$M = \mathbf{E}\left[ww^{\mathsf{T}}\right],$$

is finite-valued. We assume, without loss of generality, that M is a positive definite matrix. We emphasize that our specification of the disturbance trajectory w captures a large family of disturbance processes, including those whose support can be described as the intersection of polytopes and ellipsoids.

6.4 Decentralized Control Design

6.4.1 State Space Description

In what follows, we build on the individual resource models developed in Section 6.3 to develop a discrete-time state space model describing the collective dynamics of the distribution network. We partition the system into n subsystems, where each subsystem $i \in \{1, \ldots, n\}$ encapsulates the dynamics of resources connected to bus i. For each subsystem i, we let the energy storage state $x_i(t)$ be its *state* at time t, and define its *input* according to

$$u_i(t) = \begin{bmatrix} p_i^S(t) \\ q_i^I(t) \end{bmatrix}.$$

The corresponding state equation for each subsystem i is therefore given by Eq. (6.10). We define the full system state and input at time t by $x(t) = (x_1(t), .., x_n(t)) \in \mathbf{R}^n$ and $u(t) = (u_1(t), .., u_n(t)) \in \mathbf{R}^{2n}$, respectively. The full system state equation admits the following representation

$$x(t+1) = x(t) + Bu(t).$$

Here, the matrix B is given by

$$B = I_n \otimes \begin{bmatrix} -\Delta & 0 \end{bmatrix}, \tag{6.17}$$

where $\otimes$ denotes the Kronecker product operator. The initial condition[2] and system trajectories are related according to

$$x = \mathbb{A}x(0) + \mathbb{B}u, \tag{6.18}$$

where x and u represent the *state* and *input trajectories*, respectively. They are given by

$$\begin{aligned} x &= (x(0), \ldots, x(T)) \in \mathbf{R}^{N_x}, & N_x &= n(T+1), \\ u &= (u(0), \ldots, u(T-1)) \in \mathbf{R}^{N_u}, & N_u &= 2nT. \end{aligned}$$

And the block matrices $(\mathbb{A}, \mathbb{B})$ are given by

$$\mathbb{A} = \mathbf{1}_{(T+1)\times 1} \otimes I_n, \qquad \mathbb{B} = \begin{bmatrix} 0 & & & & \\ B & 0 & & & \\ B & B & 0 & & \\ \vdots & \vdots & \ddots & \ddots & \\ \vdots & \vdots & & \ddots & 0 \\ B & B & \cdots & \cdots & B \end{bmatrix},$$

where $\mathbf{1}_{(T+1)\times 1}$ is a vector of all ones in $\mathbf{R}^{T+1}$.

6.4.2 Decentralized Control Design

The controller information structure considered in this chapter is such that each subsystem is required to determine its local control input using only its local measurements. We therefore restrict ourselves to fully decentralized

[2]Recall that the initial condition $x(0)$ is assumed fixed and known.

disturbance-feedback control policies.[3] That is to say, at each time t, the control input to each subsystem i is restricted to be of the form

$$u_i(t) = \gamma_i(w_i^t, t),$$

where $\gamma_i(\cdot, t)$ is a causal measurable function of the local disturbance history. We define the *local control policy* for subsystem i as $\gamma_i = (\gamma_i(\cdot, 0), .., \gamma_i(\cdot, T-1))$; and refer to the collection of local control policies $\gamma = (\gamma_1, .., \gamma_n)$ as the *decentralized control policy* for the system. Finally, we define Γ to be the *set of all admissible decentralized control policies*.

We consider the objective of minimizing the expected amount of active power required to meet demand over the distribution network. For the setting considered, this is technically equivalent to minimizing the expected active power losses plus the terminal energy stored in the distribution network. In a similar spirit to [9, 102], we approximate the active power loss on line $(i, j) \in \mathcal{E}$ at time period t as[4]

$$\delta p_{ij}(t) = r_{ij}\left(\frac{p_{ij}(t)^2 + q_{ij}(t)^2}{v_0(t)^2}\right).$$

By a direct substitution of the linearized branch flow Eqs. (6.5)-(6.6) into the above approximation, one can represent the total active power losses as a convex quadratic function in the input trajectory u and disturbance trajectory w. Specifically, one can construct matrices $L_u^0,\ L_w^0 \in \mathbf{R}^{2n\times 3n}$, and a positive definite

[3]For simplicity of exposition, it is assumed in this chapter that each subsystem can perfectly observe its local disturbance process. We note, however, that all of the results presented in this chapter can be immediately generalized to the setting in which each subsystem has only partial linear observations of its local disturbance process. We refer the readers to [72] for a detailed treatment of such systems.

[4]Implicit in this approximation is the assumption that the bus voltage magnitudes are uniform across the network.

diagonal matrix $\Sigma^0 \in \mathbf{R}^{2n \times 2n}$, such that

$$\sum_{(i,j)\in\mathcal{E}} \frac{r_{ij}\left(p_{ij}(t)^2 + q_{ij}(t)^2\right)}{v_0^2} = \Big(L_u^0 u(t) + L_w^0 w(t)\Big)^{\mathsf{T}} \Sigma^0 \Big(L_u^0 u(t) + L_w^0 w(t)\Big)$$

for $t = 0, \dots, T-1$. With the matrices L_u^0, L_w^0 and Σ^0 in head, we have that

$$\sum_{t=0}^{T-1} \sum_{(i,j)\in\mathcal{E}} \delta p_{ij}(t) = (L_u u + L_w w)^{\mathsf{T}} \Sigma (L_u u + L_w w). \tag{6.19}$$

where the matrices Σ, L_u, and L_w are defined according to

$$\Sigma = I_T \otimes \Sigma^0, \quad L_u = I_T \otimes L_u^0, \quad L_w = \begin{bmatrix} 0 & I_T \otimes L_w^0 \end{bmatrix}.$$

In addition, the sum of the terminal energy storage states across the network can be written as a linear function of the state trajectory x. Namely, we have that

$$\sum_{i=1}^{n} x_i(T) = c^{\mathsf{T}} x, \tag{6.20}$$

where the vector c is defined as

$$c = \begin{bmatrix} 0_{nT\times 1} \\ \mathbf{1}_{n\times 1} \end{bmatrix}.$$

Henceforth, we define the expected cost associated with a decentralized control policy $\gamma \in \Gamma$ according to

$$J(\gamma) = \mathbf{E}^{\gamma}\left[c^{\mathsf{T}} x + (L_u u + L_w w)^{\mathsf{T}} \Sigma (L_u u + L_w w)\right]. \tag{6.21}$$

Here, expectation is taken with respect to the joint distribution on (x, u, w) induced by the control policy γ.

We define the *decentralized control design* problem as

$$\begin{aligned} \text{minimize} \quad & J(\gamma) \\ \text{subject to} \quad & \gamma \in \Gamma \\ & \left. \begin{aligned} & x \in \mathcal{X},\ u \in \mathcal{U}(w) \\ & x = \mathbb{A}x(0) + \mathbb{B}u \\ & u = \gamma(w) \end{aligned} \right\} \forall w \in \mathcal{W}, \end{aligned} \tag{6.22}$$

where the decision variable is the decentralized control policy $\gamma \in \Gamma$. The set of feasible states $\mathcal{X}$ is defined according to inequality (6.11). The set of feasible control inputs $\mathcal{U}(w)$ is defined according to inequalities (6.9), (6.12), and (6.13). We let J^* denote the *optimal value* of problem (6.22).

6.5 Affine Control Design

The decentralized control design problem (6.22) amounts to an infinite-dimensional convex program, and is, in general, computationally intractable. We therefore resort to approximation by restricting the space of admissible decentralized control policies to be causal affine functions of the measured disturbance process. In addition, we approximate the feasible region of problem (6.22) from within by a polyhedral set. The combination of these two approximations enables the computation of a decentralized control policy, which is guaranteed to be feasible for problem (6.22), through solution of a finite-dimensional conic program.

6.5.1 Polyhedral Inner Approximation of Constraints

The feasible state space $\mathcal{X}$ is clearly polyhedral. The feasible input space $\mathcal{U}(w)$ is not. It can, however, be approximated from within by a polyhedral set by replacing the quadratic constraint in (6.13) with the following pair of linear constraints:

$$\left|q_i^I(t)\right| \leq \overline{q}_i^I(t). \tag{6.23}$$

Here, the deterministic constant $\overline{q}_i^I(t)$ is defined according to

$$\overline{q}_i^I(t) = \inf\left\{ \sqrt{{s_i^I}^2 - {w_i^I(t)}^2} \;\middle|\; w \in \mathcal{W} \right\}.$$

Essentially, $\overline{q}_i^I(t)$ denotes the minimum reactive power capacity that is guaranteed to be available at inverter i at time t. We provide a graphical illustration of this polyhedral inner approximation in Figure 6.2b.

Although an inner approximation of this form may appear conservative at first glance, several recent studies [62, 102] have observed such approximations to result in a small loss of performance, as measured by the objective function considered in this chapter. We corroborate these claims in Section 6.7 by developing a technique to bound the loss of optimality incurred by this inner approximation. In particular, the suboptimality incurred by such an approximation is shown to be small for the case study considered in this chapter.

Inequalities (6.9), (6.11), (6.12), and (6.23) define a collection of $m = 8nT$ linear constraints on the state, input, and disturbance trajectories. We represent them more succinctly as

$$\underline{F_x} x + \underline{F_u} u + \underline{F_w} w \leq 0, \quad \forall\, w \in \mathcal{W},$$

where it is straightforward to construct the matrices $\underline{F_x} \in \mathbf{R}^{m \times N_x}, \underline{F_u} \in \mathbf{R}^{m \times N_u}$, and $\underline{F_w} \in \mathbf{R}^{m \times N_w}$ using the given problem data. The following optimization problem is an inner approximation to the original decentralized control design problem (6.22):

$$
\begin{aligned}
\text{minimize} \quad & \mathbf{E}^{\gamma}\left[c^{\mathsf{T}}x + (L_u u + L_w w)^{\mathsf{T}}\Sigma(L_u u + L_w w)\right] \\
\text{subject to} \quad & \gamma \in \Gamma \\
& \left.\begin{aligned}
& \underline{F_x}x + \underline{F_u}u + \underline{F_w}w \leq 0 \\
& x = \mathbb{A}x(0) + \mathbb{B}u \\
& u = \gamma(w)
\end{aligned}\right\} \forall\, w \in \mathcal{W},
\end{aligned}
\tag{6.24}
$$

where the decision variable is given by γ. Although convex, problem (6.24) is an infinite-dimensional program, and is therefore computationally intractable, in general. In what follows, we refine this approximation by further restricting the space of admissible controllers to be *affine* functions of the disturbance trajectory.

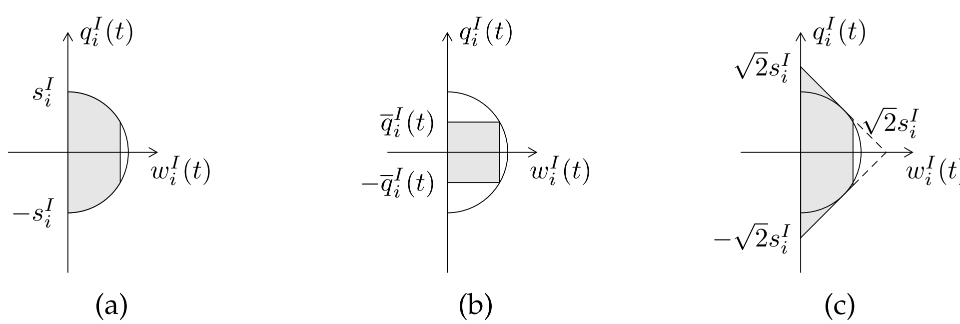

Figure 6.2: The above plots depict an inverter's range of feasible reactive power injections (in gray) at a particular time period t as specified by (a) the *original* quadratic constraints (6.13), (b) the *inner* linear constraints (6.23), and (c) the *outer* linear constraints (6.41).

6.5.2 Affine Control Design via Conic Programming

We restrict our attention to decentralized affine control policies of the form

$$u_i(t) = \overline{u}_i(t) + \sum_{s=0}^{t} Q_i(t,s) w_i(s) \tag{6.25}$$

for each subsystem $i = 1, \dots, n$ and time $t = 0, \dots, T-1$. Here, $\overline{u}_i(t) \in \mathbf{R}^2$ denotes the open loop component of the local control, and $(Q_i(t,0), \dots, Q_i(t,t))$ the collection of feedback control gains at time t. One can write the decentralized affine control policy in (6.25) more compactly as

$$u = Qw.$$

We enforce the desired information structure in Q by requiring that $Q \in S$, where S denotes the subspace of matrices that respect the information structure associated with the set of admissible decentralized control policies Γ. Specifically,

$$S = \left\{ Q \in \mathbf{R}^{N_u \times N_w} \,\middle|\, Q \in \Gamma \right\}.$$

The restriction to decentralized affine control policies gives rise to the following semi-infinite program, which stands as a more conservative inner approximation to the original decentralized control design problem (6.22).

$$\begin{aligned} \text{minimize} \quad & \mathbf{E}\left[c^\mathsf{T} x + (L_u u + L_w w)^\mathsf{T} \Sigma (L_u u + L_w w) \right] \\ \text{subject to} \quad & Q \in S \\ & \left. \begin{aligned} & \underline{F_x} x + \underline{F_u} u + \underline{F_w} w \le 0 \\ & x = \mathbb{A} x(0) + \mathbb{B} u \\ & u = Qw \end{aligned} \right\} \forall\, w \in \mathcal{W}, \end{aligned} \tag{6.26}$$

where the decision variable is given by Q. Given our assumption that the uncertainty set $\mathcal{W}$ has a conic representation, one can directly apply Proposition 3 in Chapter 3.3.2 to equivalently reformulate the semi-infinite program (6.26) as a finite-dimensional conic program.[5] In Proposition 7, we present the finite-dimensional conic reformulation of the semi-infinite program (6.26) implied by Proposition 3.

Proposition 7. The semi-infinite program (6.26) admits an equivalent reformulation as the following finite-dimensional conic program

$$\begin{aligned}
\text{minimize} \quad & \operatorname{Tr}\Bigg(\Big(Q^\intercal L_u^\intercal \Sigma L_u Q + \big(2L_w^\intercal \Sigma L_u + e_1 c^\intercal \mathbb{B}\big) Q \\
& \qquad + L_w^\intercal \Sigma L_w\Big) M\Bigg) + c^\intercal \mathbb{A} x(0) \\
\text{subject to} \quad & Q \in S \\
& Z \in \mathbf{R}^{m \times N_w}, \ \ \Pi \in \mathbf{R}^{\ell \times m}, \ \ \nu \in \mathbf{R}_+^m \\
& (\underline{F_u} + \underline{F_x}\mathbb{B})Q + \underline{F_x}\mathbb{A}x(0)e_1^\intercal + \underline{F_w} + Z = 0, \\
& Z = \nu e_1^\intercal + \Pi^\intercal H, \\
& \Pi \succeq_{\mathcal{K}^*} 0,
\end{aligned} \tag{6.27}$$

where the decision variables are given by Q, Z, Π, and ν. Let J^{in} denote the optimal value of the above program. It stands as an *upper bound* on the optimal value of the original decentralized control problem (6.22), i.e., $J^* \leq J^{\text{in}}$.

Several comments are in order. First, the specification of the conic program (6.27) relies on the probability distribution of the disturbance w only through its

[5]We note that a direct application of Proposition 3 also requires that the information structure of the underlying decentralized control problem be *partially nested*. This condition requiring partial nestedness of the information structure is trivially satisfied for the decentralized control design problem (6.22) under consideration in this chapter.

support $\mathcal{W}$ and second-order moment matrix M. Second, the conic program can be efficiently solved for a variety of cones $\mathcal{K}$, including polyhedral and second-order cones. For such cones, problem (6.27) amounts to a conic program with $O(n^2T^2)$ decision variables and $O(n^2T^2)$ constraints. It can thus be solved in time that is polynomial in the control horizon T and the number of subsystems n. Finally, assuming that the decentralized affine controller Q^* is computed at a central location, the decentralized implementation of the controller will require the communication of each local control policy to its corresponding subsystem. This entails the transmission of $3T^2 + 5T$ real numbers to each subsystem.

6.6 Fast Time-Scale Implementation of Affine Controller

In practice, the active power generated by a photovoltaic resource may fluctuate over time-scales (e.g., seconds to minutes) that are substantially shorter than the time-scale being used for control design (e.g., hourly). In what follows, we propose a method to enable the implementation of controllers designed according to Proposition 7 over more finely grained time-scales. The method we propose is simple. First, we compute an affine control policy for the original (i.e., slow) time-scale according to Proposition 7. Via a suitable rescaling of the resulting feedback control gains, we construct an affine control policy that can be implemented over a more finely grained time-scale. An attractive feature of the proposed implementation is that, under a mild assumption on the quasi-stationarity of the support of the underlying disturbance trajectory, the affine control policy we construct is guaranteed to yield state and input trajectories that are feasible on this more finely grained time-scale. In what follows, we provide a precise specification of this fast time-scale controller, and discuss its theoretical guarantees.

6.6.1 Fast Time-Scale Processes

We begin with a description of the state, input, and disturbance processes on the more finely grained time-scale by dividing each original time period t into K shorter time periods. It will be convenient to index the original time periods by t, and the more finely grained time periods by (k, t), for $k = 0, \ldots, K-1$, and $t = 0, \ldots, T$. In particular, each time period (k, t) is defined over a time interval of length Δ/K, where recall that each original time period t is of length Δ. For the remainder of this section, we will refer to the original and the more finely grained time-scales as the *slow* and *fast time-scales*, respectively.

We denote the fast time-scale state, input, and disturbance processes by $x(k, t)$, $u(k, t)$, and $w(k, t)$, respectively. We emphasize that all the fast time-scale quantities have the same units as their slow time-scale counterparts. It will prove useful to define a *slow time-scale average* of the fast time-scale disturbance process according to

$$\overline{w}(t) = \frac{1}{K}\left(\sum_{k=0}^{K-1} w\left(k, t\right)\right), \tag{6.28}$$

for $t = 0, \ldots, T-1$. We refer the reader to Fig. 6.3, which offers a graphical illustration comparing the fast time-scale disturbance process $w(k, t)$ against its slow time-scale average $\overline{w}(t)$.

We describe the evolution of the *fast time-scale state process* over each time period t according to the state equation:

$$x(k+1, t) = x(k, t) + \frac{1}{K}Bu(k, t)$$

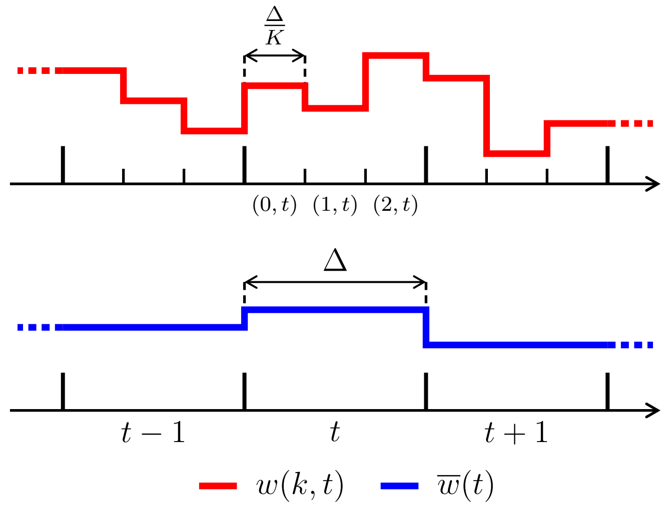

Figure 6.3: The above plots depicts (one component of) the fast time-scale disturbance process $w(k,t)$ and its slow time-scale average $\overline{w}(t)$ for $K = 3$.

for $k = 0, \ldots, K-1$. We link this process across the slow time-scale periods by enforcing the boundary conditions

$$x(K,t) = x(0,t+1)$$

for $t = 0, \ldots, T-1$. We initialize the fast time-scale state process according to $x(0,0) = x(0)$.

6.6.2 Fast Time-Scale Controller

In what follows, we construct a fast time-scale controller based on the slow time-scale controller computed according to Proposition 7. More specifically, let Q^* denote the optimal solution to problem (3.14). Since $Q^* \in S$, it follows that Q^* is a block lower-triangular matrix of the form

$$Q^* = \begin{bmatrix} \overline{u}^*(0) & Q^*(0,0) & & \\ \vdots & \vdots & \ddots & \\ \overline{u}^*(T-1) & Q^*(T-1,0) & \cdots & Q^*(T-1,T-1) \end{bmatrix},$$

where each matrix $Q^*(t,s)$ is block diagonal of the form

$$Q^*(t,s) = \begin{bmatrix} Q_1^*(t,s) & & \\ & \ddots & \\ & & Q_n^*(t,s) \end{bmatrix}$$

for each $t = 0, \ldots, T-1$ and $s = 0, \ldots, t$. Using these feedback control gains embedded in the matrix Q^*, we define the *fast time-scale control input* at period (k,t) as

$$u(k,t) = \overline{u}^*(t) + Q^*(t,t)w\,(k,t) + \sum_{s=0}^{t-1} Q^*(t,s)\overline{w}(s), \tag{6.29}$$

for all $t = 0, \ldots, T-1$ and $k = 0, \ldots, K-1$. Recall from Eq. (6.28) that $\overline{w}(t)$ denotes the average of the fast time-scale disturbance process over the period t.

6.6.3 Constraint Satisfaction Guarantees

The decentralized affine controller defined according to Eq. (6.29) is said to be *feasible* if it induces voltage, input, and state trajectories that are guaranteed to satisfy their respective constraints at the fast time-scale for all possible realizations of the fast time-scale disturbance process. That is to say, for each subsystem $i \in \{1, \ldots, n\}$, it must hold that

$$\underline{v}_i \leq v_i\,(k,t) \leq \overline{v}_i, \tag{6.30}$$

$$-\sqrt{{s_i^I}^2 - w_i^I(k,t)^2} \leq q_i^I\,(k,t) \leq \sqrt{{s_i^I}^2 - w_i^I(k,t)^2}, \tag{6.31}$$

$$\underline{p}_i^S \leq p_i^S\,(k,t) \leq \overline{p}_i^S, \tag{6.32}$$

for all time periods $t = 0, \ldots, T-1$, $k = 0, \ldots, K-1$, and

$$0 \leq x_i\,(k,t) \leq b_i, \tag{6.33}$$

for all time periods $t = 0, \ldots, T-1$, $k = 0, \ldots, K$; and all possible realizations of the fast time-scale disturbance process.

We now make a *mild* assumption on the support of the fast time-scale disturbance process, which ensures that the fast time-scale controller defined according to Eq. (6.29) is feasible.

Assumption 5 (Quasi-Stationarity)**.** We assume that

$$(1, w(k_0, 0), w(k_1, 1), \ldots, w(k_{T-1}, T-1)) \in \Xi,$$

for all $k_t \in \{0, \ldots, K-1\}$ and $t = 0, \ldots, T-1$.

Assumption 5 can be interpreted as an assumption on the *quasi-stationarity* of the support of the fast time-scale disturbance process $w(k,t)$. Moreover, Assumption 5 is reasonable, as it is always possible to construct a set Ξ such this assumption is satisfied, given a characterization of the set of all possible realizations taken by the fast time-scale disturbance process.

Proposition 8 (Fast Time-Scale Feasibility)**.** Let Assumption 5 hold. The fast time-scale controller specified according to Equation (6.29) is feasible.

Proposition 8 reveals that the *slow time-scale* controller computed according to Proposition 3 can be implemented as a feasible *fast time-scale* controller. We provide the proof of Proposition 8 in the next subsection.

6.6.4 Proof of Proposition 8

Let the system control input be specified according to Eq. (6.29). The proof consists of two parts. In Part 1, we show that for any realization of the fast time-scale disturbance process, the input constraints specified in inequalities (6.30)-(6.32) are all satisfied. In Part 2, we show that for any realization of the fast time-scale disturbance process, the state constraint specified in inequality (6.33) is satisfied.

Part 1: We will only show that for any realization of the fast time-scale disturbance process, the voltage magnitude constraint specified in inequality (6.30) is satisfied. The proof of the satisfaction of the input constraints specified in inequalities (6.31) and (6.32) is analogous. It is thus omitted for the sake of brevity.

It will be convenient to work with the vector of squared voltage magnitudes $v\left(k,t\right)^2 = (v_1(k,t)^2, \ldots, v_n(k,t)^2)$ for the remainder of the proof. We will show that

$$\underline{v}^2 \leq v\left(k,t\right)^2 \leq \overline{v}^2,$$

where $\underline{v}^2 = (\underline{v}_1^2, \ldots, \underline{v}_n^2)$, and $\overline{v}^2 = (\overline{v}_1^2, \ldots, \overline{v}_n^2)$. It follows from the linearized branch flow model (6.8) that for each time period (k,t), the vector of squared voltage magnitudes is given by

$$v\left(k,t\right)^2 = V_u u\left(k,t\right) + V_w w\left(k,t\right) + v_0^2 \mathbf{1},$$

where the matrices V_u and V_w are defined according to

$$\begin{aligned} V_u &= R \otimes \begin{bmatrix} 1 & 0 \end{bmatrix} + X \otimes \begin{bmatrix} 0 & 1 \end{bmatrix}, \\ V_w &= R \otimes \begin{bmatrix} -1 & 0 & 1 \end{bmatrix} - X \otimes \begin{bmatrix} 0 & 1 & 0 \end{bmatrix}. \end{aligned}$$

Given the specification of the fast time-scale control input $u\left(k,t\right)$ according to

Eq. (6.29), we have that

$$
\begin{aligned}
v\left(k,t\right)^{2} =& V_u \left(\overline{u}^*(t) + \sum_{s=0}^{t-1} Q^*(t,s)\overline{w}(s) \right) \\
&+ \left(V_u Q^*(t,t) + V_w\right) w\left(k,t\right) + v_0^2 \mathbf{1}.
\end{aligned} \tag{6.34}
$$

Given Assumption 5 and the convexity of the set Ξ, it holds that

$$
(1, \overline{w}(0), .., \overline{w}(t-1), w(k,t), .., w(k,T-1)) \in \mathcal{W}. \tag{6.35}
$$

Condition (6.35), in combination with the guaranteed feasibility of the control policy Q^* for the original *slow time-scale* decentralized control design problem (3.6), implies that

$$
\begin{aligned}
\underline{v}^2 \leq & V_u \left(\overline{u}^*(t) + \sum_{s=0}^{t-1} Q^*(t,s)\overline{w}(s) \right) \\
&+ \left(V_u Q^*(t,t) + V_w\right) w\left(k,t\right) + v_0^2 \mathbf{1} \leq \overline{v}^2.
\end{aligned}
$$

It immediately follows from Eq. (6.34) that $\underline{v}^2 \leq v(k,t)^2 \leq \overline{v}^2$. This completes Part 1 of the proof.

Part 2: We show that for any realization of the fast time-scale disturbance process, the system state satisfies $0 \leq x(k,t) \leq b$ for $t = 0, \ldots, T-1$, $k = 0, \ldots, K$. Here, $b = (b_1, \ldots, b_n)$ is the vector of energy storage capacities. We fix an arbitrary realization of the fast time-scale disturbance process throughout this part of the proof.

We first consider the case of $k = 0$, and show that $0 \leq x(0,t) \leq b$ for $t =$

$0, \ldots, T$. It holds that

$$x(0,t) = x(0,0) + \frac{1}{K} B \left(\sum_{s=0}^{t-1} \sum_{\ell=0}^{K-1} u(\ell, s) \right) \tag{6.36}$$

$$= x(0) + B \left(\sum_{s=0}^{t-1} \left(\overline{u}^*(s) + \sum_{r=0}^{s} Q^*(s,r)\overline{w}(r) \right) \right), \tag{6.37}$$

where Eq. (6.37) follows from the specification of the fast time-scale system control input according to Eq. (6.29).

Condition (6.35), in combination with the guaranteed feasibility of the control policy Q^* for the slow time-scale decentralized control design problem (3.6), implies that

$$0 \leq x(0) + B \left(\sum_{s=0}^{t-1} \left(\overline{u}^*(s) + \sum_{r=0}^{s} Q^*(s,r)\overline{w}(r) \right) \right) \leq b.$$

It follows from Eq. (6.37) that $0 \leq x(0,t) \leq b$ for $t = 0, \ldots, T$. The enforcement of the boundary condition of the fast time-scale state equation requires that $x(K,t) = x(0,t+1)$ for $t = 0, \ldots, T-1$. This, in combination with the above inequality, implies that $0 \leq x(k,t) \leq b$ for $t = 0, \ldots, T-1$, and $k = 0$ and K.

Next, we show that $0 \leq x(k,t) \leq b$ for $k = 1, \ldots, K-1$, $t = 0, \ldots, T-1$. We first write $x(k,t)$ as

$$x(k,t) = x(0,t) + \frac{k}{K} B \left(\frac{1}{k} \sum_{\ell'=0}^{k-1} u(\ell', t) \right)$$

$$= \frac{K-k}{K} x(0,t) + \frac{k}{K} \left(x(0,t) + B \left(\frac{1}{k} \sum_{\ell'=0}^{k-1} u(\ell', t) \right) \right),$$

where $x(0,t)$ is specified according to Eq. (6.36). Recall that we previously established that $0 \leq x(0,t) \leq b$. Thus, to show that $0 \leq x(k,t) \leq b$, it suffices to show that

$$0 \leq x(0,t) + B \left(\frac{1}{k} \sum_{\ell'=0}^{k-1} u(\ell', t) \right) \leq b. \tag{6.38}$$

First notice that under the fast time-scale control policy specified by Eq. (6.29), we have that

$$\frac{1}{k}\sum_{\ell'=0}^{k-1} u(\ell', t) = \overline{u}^*(t) + Q^*(t,t)\widetilde{w}(k,t) + \sum_{s'=0}^{t-1} Q^*(t,s')\overline{w}(s'), \tag{6.39}$$

where the vector $\widetilde{w}(k,t)$ is defined according to

$$\widetilde{w}(k,t) = \frac{1}{k}\sum_{\ell'=0}^{k-1} w\left(\ell', t\right).$$

Given Assumption 5 and the convexity of $\mathcal{W}$, it holds that

$$(1, \overline{w}(0), .., \overline{w}(t-1), \widetilde{w}(k,t), w\left(0, t+1\right), .., w\left(0, T-1\right)) \in \mathcal{W}. \tag{6.40}$$

Condition (6.40), in combination with the guaranteed feasibility of the control policy Q^* for the slow time-scale decentralized control design problem (3.6), implies that

$$0 \leq x(0,t) + B\left(\overline{u}^*(t) + Q^*(t,t)\widetilde{w}(k,t) + \sum_{s'=0}^{t-1} Q^*(t,s')\overline{w}(s')\right) \leq b,$$

where $x(0,t)$ is specified according to Eq. (6.37). It follows from Eq. (6.39) that inequality (6.38) is satisfied. This completes Part 2 of the proof.

6.7 Lower Bounds

The restriction to affine policies computed according to Proposition 7 may result in the loss of optimality with respect to the original decentralized control design problem. In this section, we develop a tractable method to bound this loss of optimality via the solution of a conic programming relaxation–the optimal value of which is guaranteed to stand as a lower bound on the optimal value of the original decentralized control design problem (6.22). With such a lower bound in hand, one can estimate the suboptimality incurred by any feasible decentralized control policy.

6.7.1 Polyhedral Outer Approximation of Constraints

As an initial step in the derivation of this relaxation, we construct a polyhedral outer approximation of the feasible region of problem (6.22). Specifically, the quadratic constraint in (6.13) can be relaxed to the following pair of linear constraints:

$$\left|q_i^I(t)\right| \leq \sqrt{2}s_i^I - w_i^I(t). \tag{6.41}$$

We provide a graphical illustration of this polyhedral outer approximation in Figure 6.2c.

Inequalities (6.9), (6.11), (6.12), and (6.41) define a collection of m linear constraints on the state, input, and disturbance trajectories. We represent them more succinctly as

$$\overline{F_x}x + \overline{F_u}u + \overline{F_w}w \leq 0, \quad \forall\, w \in \mathcal{W},$$

where it is straightforward to construct the matrices $\overline{F_x} \in \mathbf{R}^{m \times N_x}, \overline{F_u} \in \mathbf{R}^{m \times N_u}$, and $\overline{F_w} \in \mathbf{R}^{m \times N_w}$ using the given problem data. The following optimization problem is an outer approximation to the original decentralized control design problem (6.22):

$$\begin{array}{ll} \text{minimize} & \mathbf{E}^{\gamma}\left[c^{\mathsf{T}}x + (L_u u + L_w w)^{\mathsf{T}}\Sigma(L_u u + L_w w)\right] \\ \text{subject to} & \gamma \in \Gamma \\ & \left.\begin{array}{l} \overline{F_x}x + \overline{F_u}u + \overline{F_w}w \leq 0 \\ x = \mathbb{A}x(0) + \mathbb{B}u \\ u = \gamma(w) \end{array}\right\} \forall\, w \in \mathcal{W}, \end{array} \tag{6.42}$$

where the decision variable is given by γ.

6.7.2 Lower Bounds via Conic Programming

Problem (6.42) is, in general, computationally intractable due to the infinite-dimensionality of its decision space. In what follows, we further relax problem (6.42) to a finite-dimensional conic program via an application of the constraint relaxation technique we developed in Chapter 5.3. We first require an additional assumption on the probability distribution of the disturbance trajectory w.

Assumption 6 (Disturbance Process)**.** There exist matrices $H_i^t \in \mathbf{R}^{3n(t+1)\times(1+3(t+1))}$ and $H^t \in \mathbf{R}^{N_w\times(1+3n(t+1))}$ such that

$$\mathbf{E}\left[w^t \,\middle|\, w_i^t\right] = H_i^t \begin{bmatrix} 1 \\ w_i^t \end{bmatrix} \quad \text{and} \quad \mathbf{E}\left[w \,\middle|\, w^t\right] = H^t \begin{bmatrix} 1 \\ w^t \end{bmatrix}$$

almost surely, for all time periods $t = 0, \ldots, T-1$ and subsystems $i = 1, \ldots, n$.

Although Assumption 6 may appear restrictive, it was shown in [54] to hold for a large family of distributions. In particular, Assumption 6 holds for all disturbance processes that possess elliptically contoured distributions. In Appendix B, we provide a formal definition of elliptically contoured distributions, and discuss its properties that are useful in optimal control problems. It is also straightforward to show that Assumption 6 is satisfied by any disturbance process for which the random vectors $w_i(t)$ ($i = 1, \ldots, n$, $t = 0, \ldots, T-1$) are mutually independent.

With Assumption 6 in hand, a direct application[6] of Theorem 2 in Chapter 5.3 yields a conic programming relaxation of problem (6.42). Its optimal value stands as a lower bound on the optimal value of the original decentralized control design problem (6.22).

Proposition 9. Consider the following finite-dimensional conic program:

$$\begin{aligned}
\text{minimize} \quad & \mathrm{Tr}\Big(\Big(Q^\top L_u^\top \Sigma L_u Q + \big(2L_w^\top \Sigma L_u + e_1 c^\top \mathbb{B}\big) Q \\
& \qquad + L_w^\top \Sigma L_w\Big) M\Big) + c^\top \mathbb{A} x(0) \\
\text{subject to} \quad & Q \in S, \;\; Z \in \mathbf{R}^{m \times N_w} \\
& (\overline{F_u} + \overline{F_x}\mathbb{B})Q + \overline{F_x}\mathbb{A}x(0)e_1^\top + \overline{F_w} + Z = 0, \\
& HMZ^\top \succeq_{\mathcal{K}} 0, \\
& e_1^\top M Z^\top \geq 0,
\end{aligned} \tag{6.43}$$

where the decision variables are given by Q and Z. Let J^{out} denote the optimal value of the above program. If Assumption 6 holds, then $J^{\mathrm{out}} \leq J^*$.

Given Assumption 6, the conic program (6.43) can be used to evaluate the performance of any feasible control policy. Namely, a policy $\gamma \in \Gamma$ is close to optimal (for a given problem instance) if $J(\gamma)$ is close to J^{out}. Additionally, Propositions 7 and 9 imply that the optimal value of the original decentralized

[6]While Assumption 6 differs from the assumption that Theorem 2 requires, it is straightforward to generalize the proof of Theorem 2 to the setting in which Assumption 6 is satisfied. We omit the details of the generalized proof, as it mirrors that of Theorem 2.

control problem (6.22) satisfies

$$J^{\text{out}} \leq J^* \leq J^{\text{in}}.$$

Therefore, a small gap between J^{in} and J^{out} implies that decentralized affine control policies are close to optimal for the underlying problem instance. Finally, we note that the conic program (6.43) can be efficiently solved for a variety of cones $\mathcal{K}$, including polyhedral and second-order cones. For such cones, problem (6.43) amounts to a conic program with $O(nT^2)$ decision variables and $O(nT)$ constraints. It can thus be solved in time that is polynomial in the control horizon T and the number of subsystems n.

6.8 Case Study

We consider the control of distributed energy resources in a 12 kV radial distribution feeder depicted in Fig. 6.4. The distribution feeder considered in this chapter is similar in structure to the network considered in [119]. Apart from the substation, the distribution feeder consists of $n = 14$ buses. We operate the system over a finite time horizon of $T = 24$ hours, beginning at twelve o'clock (midnight).

6.8.1 System Description

We assume that only buses 4 and 8 have storage devices and PV inverters installed. All PV inverters are assumed to have an identical active power capacity, which we denote by by θ (MW). As for demand, we assume that only buses 3, 4, 5, 13, and 14 have loads; and these loads are assumed to have identical distributions. We specify their mean active and reactive power trajectories in Fig. 6.5. In order to ensure that Assumption 6 is satisfied, we assume that the random vectors $w_i(t)$ ($i = 1, \ldots, n$, $t = 0, \ldots, T-1$) are mutually independent. In addition, we assume that the random variables $w_i^p(t)$, $w_i^q(t)$, and $w_i^I(t)$ are mutually independent for each bus i and time t. Recall that Assumption 6 is necessary only for the calculation of the performance bound specified in Proposition 9.

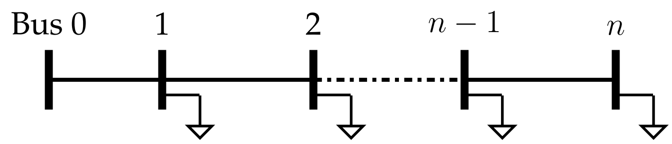

Figure 6.4: Schematic diagram of a 12 kV radial distribution feeder with $n+1$ buses.

In Table 6.1, we present additional notation pertinent to this section. In Table 6.2, we specify the parameter values of the distribution network, storage devices, PV inverters, and load.

Table 6.1: Additional notation.

Notation	**Description**
θ	Active power capacity of each PV inverter.
$\mu_i^p(t)$	Mean active power demand at bus i and time t .
$\mu_i^q(t)$	Mean reactive power demand at bus i and time t.
$\mu_i^I(t)$	Mean active power supply from PV inverter i at time t.
Uni $[a, b]$	Uniform distribution on $[a, b]$.

6.8.2 Discussion

We begin by examining the performance of the decentralized controller proposed in this chapter. In Fig. 6.6, we plot both the upper and lower bounds

Table 6.2: Specification of system data.

Distribution network	
Base voltage magnitude	12 kV
Substation voltage magnitude	$v_0 = 1$ (per-unit)
Impedance on line $(i,j) \in \mathcal{E}$	$r_{ij} = 0.466,\ x_{ij} = 0.733$ (Ω)
Voltage magnitude constraints	$\underline{v} = 0.95 \cdot \mathbf{1},\ \overline{v} = 1.05 \cdot \mathbf{1}$ (per-unit)
Storage at bus $i \in \{4, 8\}$	
Energy capacity	$b_i = 0.5$ (MWh)
Power capacity	$\underline{p}_i^S = -0.2,\ \overline{p}_i^S = 0.2$ (MW)
Initial condition	$x_i(0) = 0$ (MWh)
PV inverter at bus $i \in \{4, 8\}$	
Apparent power capacity	$s_i^I = 1.25\theta$ (MVA)
Active power supply	$w_i^I(t) \sim$ Uni $[0, 2\mu_i^I(t)]$ (MW)
Mean active power supply	$\mu_i^I(t) = \theta \cdot \max\left\{0.5 \sin\left(\frac{t-6}{12}\pi\right), 0\right\}$
Load at bus $i \in \{3, 4, 5, 13, 14\}$	
Active power demand	$w_i^p(t) \sim$ Uni $[0.7\mu_i^p(t),\ 1.3\mu_i^p(t)]$ (MW)
Reactive power demand	$w_i^q(t) \sim$ Uni $[0.7\mu_i^q(t),\ 1.3\mu_i^q(t)]$ (Mvar)

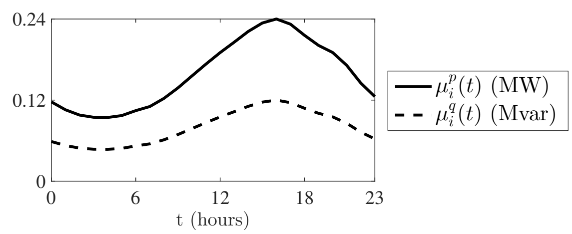

Figure 6.5: Buses i = 3, 4, 5, 13, and 14 are assumed to be identical in terms of their mean load trajectories. The above figure depicts the mean active power and reactive power demand trajectories at these buses. Both trajectories are scaled versions of the load profile DOM-S/M on 07/01/2016 from Southern California Edison [2].

on the optimal value J^* of the decentralized control design problem (6.22), as a function of the PV inverter active power capacity θ. Recall that J^{in} measures the cost incurred by the decentralized affine control policy computed according to Proposition 7. Notice that, at low PV penetration levels (i.e., for low values of θ), the upper and lower bounds nearly coincide. This indicates that the decentralized affine control policy is nearly optimal for the original decentralized control design problem. More interestingly, at high PV penetration levels (i.e., for high values of θ), the gap between the upper and lower bounds remains small. This reveals that decentralized affine control policies persist in being close to optimal for the system considered, despite the presence of large and unpredictable fluctuations in PV active power generation. Therefore, for the system under consideration, there is little additional value to be had in the design of more sophisticated (nonlinear) control policies.

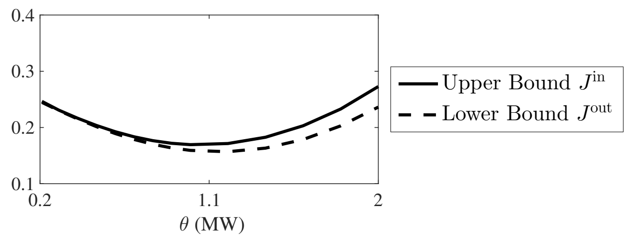

Figure 6.6: This figure depicts the upper and lower bounds, J^{in} and J^{out}, respectively, on the optimal value of the decentralized control design problem J^* (measured in MWh) as a function of the PV inverter active power capacity θ.

In Fig. 6.7, we illustrate the behavior of input and state trajectories generated by the decentralized affine controller computed according to Proposition 7. We consider the case of high PV penetration at a level of $\theta = 4$ MW. In the first and third columns of Fig. 6.7, we plot several independent realizations of disturbance, input, and state trajectories associated with bus 4 and 8, respectively. In the second and fourth columns, we plot the corresponding empirical confidence intervals.[7] First, notice that both the sequence of reactive power injections from PV inverters and the sequence of active power extractions from storage exhibit large fluctuations during daytime hours. These fluctuations are due in large part to the underlying variability in the active power supplied by the PV

[7]The empirical confidence intervals were estimated using 3×10^5 independent realizations of the disturbance, input, and state trajectories.

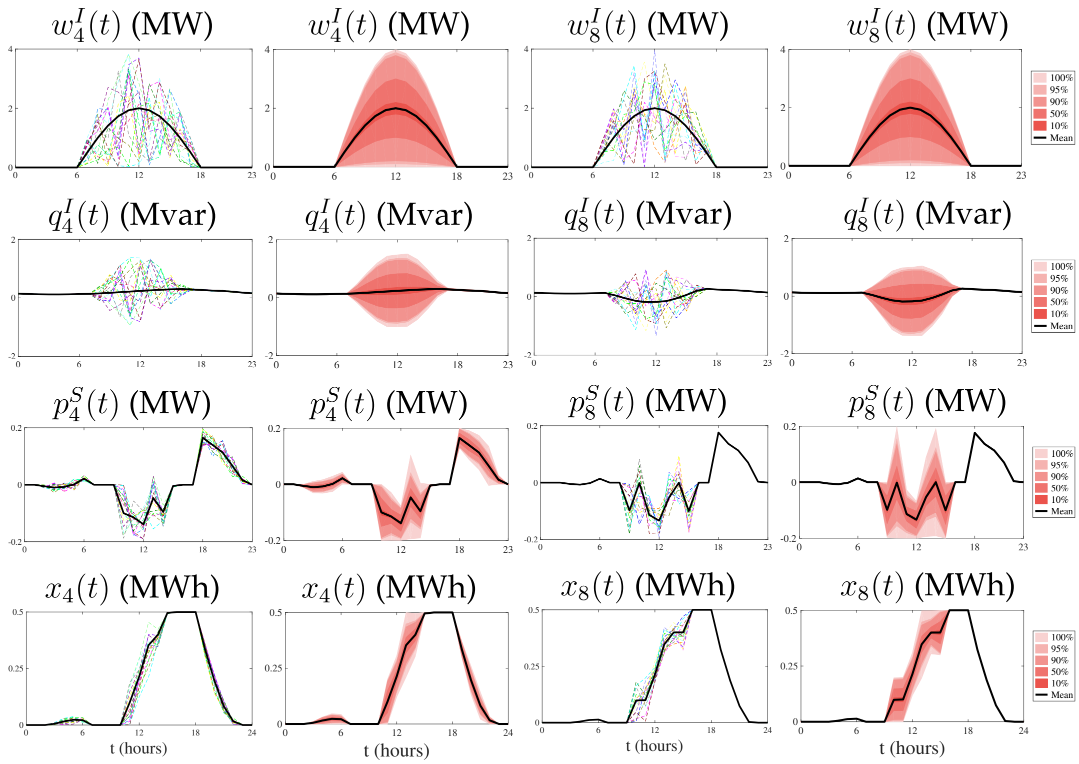

(a) System trajectories and their confidence intervals at bus 4.

(b) System trajectories and their confidence intervals at bus 8.

Figure 6.7: The figures in the first and third columns plot independent realizations of disturbance, input, and state trajectories associated with bus 4 and 8, respectively. The dashed colored lines represent the trajectory realizations, while the solid black lines denote the mean trajectories. The figures in the second and fourth columns depict the empirical confidence intervals associated with these trajectories. They were estimated using 3×10^5 independent realizations of the disturbance trajectories.

resources. In particular, a large excess of active power supply from PV can manifest in overvoltage in the distribution network. In order to ensure that voltage magnitude constraints are not violated, the proposed control policy induces reactive power injections from PV inverters that are negatively correlated with their own active power supply. Clearly, in the absence of such a feedback control mechanism, certain realizations of the disturbance trajectory would have resulted in the violation of the voltage magnitude constraints at certain buses in the distribution system.

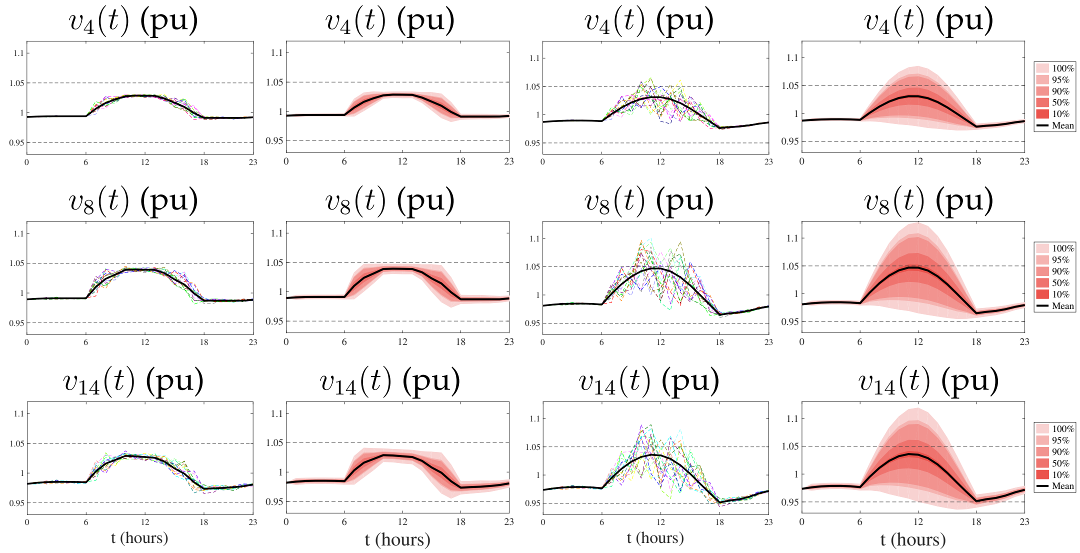

(a) Bus voltages in a *controlled distribution system*.

(b) Bus voltages in an *uncontrolled distribution system*.

Figure 6.8: The above figures depict independent realizations of bus voltage magnitude trajectories and their empirical confidence intervals for (a) a *controlled distribution system* operated under the decentralized affine controller, and (b) an *uncontrolled distribution system*. The empirical confidence intervals were estimated using 3×10^5 independent realizations of the disturbance trajectories. The dashed black lines indicate the range of allowable voltage magnitudes.

In Fig. 6.8, we illustrate the effectiveness of the proposed decentralized affine controller in maintaining bus voltage magnitudes within their allowable range. In particular, we compare the behavior of bus voltage magnitudes that occur in the distribution system *with* and *without control*. In the first and second columns of Fig. 6.8, we illustrate the behavior of voltage magnitude trajectories that materialize in the *controlled distribution system* operated under the proposed decentralized affine controller. In the third and fourth columns of Fig. 6.8, we illustrate the behavior of voltage magnitude trajectories that materialize in the *uncontrolled distribution system*, i.e., under the control policy $\gamma = 0$. Notice that, in the absence of control, the distribution system may realize bus voltage magnitudes that substantially deviate from their allowable range. In particular, the distribution system appears to suffer from overvoltage when there is an overabundance of active power supply from PV, and undervoltage during hours of peak demand. However, when operated under the decentralized affine controller, the distribution system is guaranteed to satisfy the bus voltage magnitude constraints for any possible realization of the disturbance trajectory.

6.9 Conclusion

There are several interesting directions for future work. For example, one potential drawback of the approach considered in this chapter is the explicit reliance of the control policy on the entire disturbance history. Such dependency may result in the computational intractability of calculating control policies for problems with a long horizons T. Accordingly, it will be of interest to extend the techniques developed in this chapter to accommodate fixed-memory constraints on the control policy.

It is also worth noting that the class of controllers considered in this chapter are *fully decentralized*, in that explicit communication between subsystems is not permitted. It would be of theoretical and practical interest to investigate the extent to which the introduction of additional communication links between subsystems might improve system performance. In particular, it would be of interest to explore the problem of designing a communication topology between subsystems, in order to minimize the optimal control cost, subject to a constraint on the maximum number of allowable communication links. While such problems are inherently combinatorial in nature, it is conceivable that regularization techniques, similar to those proposed in [71, 79], might yield good approximations.

Finally, all of our results rely on the assumption that the distribution system is three-phase balanced. Such an assumption will not always hold in practice. It would be of interest to extend the techniques in this chapter to accommodate the possibility of imbalance in three-phase distribution systems.

CHAPTER 7
CONCLUSION

The decentralized control of constrained linear systems with nonclassical information structures amounts to an infinite-dimensional nonconvex program that is, in general, computationally intractable. In this dissertation, we provide computationally tractable methods to the calculation of feasible decentralized control policies and the estimation of their suboptimality via the solution of finite-dimensional convex programs. Our results provide a systematic approach to the calculation of feasible decentralized control policies that are affine in the state history for decentralized control design problems with arbitrary information structures and arbitrary polyhedral constraints on the state and input trajectories. Additionally, our results also enable the estimation of the suboptimality of feasible decentralized control policies for the general family of decentralized control problems considered in this dissertation. In what follows, we conclude the dissertation with a discussion on potential directions for future work.

Control Design over a Long Time Horizon. Our results in the dissertation enable the calculation of decentralized control policies that parameterize the input trajectory as a causal function of the disturbance trajectory using convex optimization methods. As a consequence of such a parameterization of control policy, the number of decision variables in the resulting convex problem is quadratic in the time horizon T. Consequently, an important drawback of our control design approach is that it does not provide a computationally tractable method for designing decentralized controllers that operate over a long time horizon. In order to enable the design of controllers that operate over a long time horizon, one possible approach is to design the controller in a hierarchical fashion—that is to say, we formulate the decentralized control design problem over a slow time-scale, and impose an additional constraint that the implementation of the decentralized controller over a fast time-scale would yield state and input trajectories that are guaranteed to be feasible. In Chapter 6, we provide an initial result of such a flavor under the additional assumption that the A matrix of the system be equal to the identity matrix.

Unknown System Model. All results developed in this dissertation rely on the assumption that the system model is fixed and known. However, as we discussed in Section 6.2, such an assumption might not hold in the control of practical large-scale dynamical systems. As a result, one might need to consider an alternative formulation of the decentralized control design problem, in which the system matrices are assumed to belong to a bounded parameter set that is known a priori. For such systems, one needs to specify a decentralized control policy that is guaranteed to induce feasible state and input trajectories for all system matrices belonging to the parameter set and all realizations of the disturbance trajectory. More importantly, the control of such systems naturally entails a tradeoff between exploration and exploitation, as one needs to constantly perturb the control inputs to learn the underlying system matrices.

A Dual Approach to Information Relaxation. The information relaxation we construct in Chapter 5 can be thought of as a *primal* approach to the construction of information relaxations, in the sense that it only entails the expansion of the information that each controller has access to. An important drawback to this approach is that it does not penalize the use of the extra information that is introduced in the information relaxation, which might manifest in the looseness of the resulting performance lower bound. In [24, 39, 111], the authors provide a *dual approach* to information relaxation in the construction of lower bounds on the optimal values of centralized Markov decision processes. Such an approach essentially entails the dualization of the causality constraint on decision making—that is, we obtain a lower bound by relaxing the causality constraint and penalizing the use of future information. The question as to whether one could leverage on such a dual approach to reduce the looseness of the lower bound we derive in Chapter 5 is an interesting direction for future investigation.

APPENDIX A

CONVEX OPTIMIZATION

In this chapter, we review several results in convex optimization that are useful for problems in optimal control.

A.1 Conic Linear Programs

A large family of convex programs can be described as conic linear programs. In what follows, we define a conic linear program in a way that is analogous to [80], and identify families of convex program that can be formulated as conic linear programs.

Let $\mathbf{E}$ be a Euclidean space equipped with the inner product $\langle\cdot,\cdot\rangle$. A set $\mathcal{K} \subseteq \mathbf{E}$ is called a *cone* if for each $x \in \mathcal{K}$ and each $\theta \in \mathbf{R}_+$, we have that $\theta x \in \mathcal{K}$. A cone $\mathcal{K}$ is said to be *pointed* if $\mathcal{K} \cap -\mathcal{K} = \{0\}$. We say that a cone $\mathcal{K}$ is proper if it has a nonempty interior and is convex, closed, and pointed. Additionally, the *dual cone* associated with the cone $\mathcal{K}$ is defined as

$$\mathcal{K}^* := \{y \in \mathbf{E} \,|\, \langle x, y\rangle \geq 0 \;\forall x \in \mathcal{K}\}.$$

Let $\mathcal{K}$ be a proper cone. A conic linear program in its most general form is given by

$$\begin{aligned} \text{minimize} \quad & \langle c, x\rangle \\ \text{subject to} \quad & x \in \mathbf{E} \\ & \langle A_i, x\rangle = b_i,\ i = 1, \ldots, m \\ & x \in \mathcal{K}. \end{aligned} \tag{A.1}$$

And the dual to problem (A.1) is given by

$$\begin{aligned}
\text{maximize} \quad & b^\mathsf{T} y \\
\text{subject to} \quad & z \in \mathbf{E},\ y \in \mathbf{R}^m \\
& z + \sum_{i=1}^{m} A_i y_i = c \\
& z \in \mathcal{K}^*.
\end{aligned} \tag{A.2}$$

Examples of conic linear programs include:

1. *Linear programs*, where the Euclidean space $\mathbf{E} = \mathbf{R}^n$, the inner product is given by $\langle x, y \rangle = x^\mathsf{T} y$, and the cone $\mathcal{K} = \mathbf{R}^n_+$.
2. *Second-order cone programs*, where the Euclidean space $\mathbf{E} = \mathbf{R}^n$, the inner product is given by $\langle x, y \rangle = x^\mathsf{T} y$, and the cone $\mathcal{K}$ is the second order cone:
$$\mathcal{K} = \left\{ x \in \mathbf{R}^n \;\middle|\; x_1 \geq \sqrt{\sum_{i=2}^{n} x_i^2} \right\}.$$
3. *Semidefinite programs*, where the Euclidean space $\mathbf{E} = \mathbf{S}^n$ (i.e., the space of symmetric n-by-n matrices), the inner product is given by $\langle X, Y \rangle = \operatorname{Tr}\left(X^\mathsf{T} Y\right)$, and the cone $\mathcal{K} = \mathbf{S}^n_+$ (i.e., the cone of positive semidefinite matrices in $\mathbf{S}^n$).

A.2 Finite-dimensional Reformulation of Robust Linear Constraints

Consider a robust linear program of the form

$$\begin{aligned} \text{minimize} \quad & c^\mathsf{T} x \\ \text{subject to} \quad & x \in \mathbf{R}^n \\ & (A_i \xi)^\mathsf{T} x \le b_i \quad \forall \xi \in \Xi,\ i = 1, \ldots, m, \end{aligned} \tag{A.3}$$

where the uncertainty set $\Xi \subseteq \mathbf{R}^d$ is assumed to be convex and compact.

Problem (A.3) is a semi-infinite program, as its feasible solutions are required to satisfy an infinite number of linear constraints. Under the additional assumption that the uncertainty set Ξ admits a conic representation, problem (A.3) admits an equivalent reformulation as a finite-dimensional conic program. Specifically, assume that the uncertainty set Ξ admits a representation of the form

$$\Xi := \{\xi \in \mathbf{R}^d | W\xi - g \in \mathcal{K}\},$$

where $\mathcal{K}$ is assumed to be a proper cone in $\mathbf{R}^p$, and the matrix W and the vector g are both assumed to be known. It follows from the weak duality of conic linear programs that robust linear constraints of the form $\xi^\mathsf{T} y \le \alpha$ for all $\xi \in \Xi$ admits the following inner approximation as a finite-dimensional cone constraint. Additionally, such an inner approximation is an equivalent reformulation is the uncertainty set Ξ admits a nonempty interior—a consequence of the strong duality of conic linear programs that have strictly feasible solutions.

Lemma 22. We have that $\xi^\mathsf{T} y \le \alpha$ for all $\xi \in \Xi$ if there exists $x \in \mathcal{K}^*$ that satisfies

$$g^\mathsf{T} x + \alpha \le 0 \quad \text{and} \quad y + W^\mathsf{T} x = 0.$$

Additionally, the converse is also true if the uncertainty set Ξ admits a nonempty interior.

It follows from Lemma 22 that the robust linear program (A.3) admits the following inner approximation as a finite-dimensional conic program:

$$\begin{array}{ll} \text{minimize} & c^\mathsf{T} x \\ \text{subject to} & x \in \mathbf{R}^n, \\ & y^i \in \mathcal{K}^*, \; i = 1, \ldots, m \\ & g^\mathsf{T} y^i + b_i \leq 0, \; i = 1, \ldots, m \\ & A_i^\mathsf{T} x + W^\mathsf{T} y^i = 0, \; i = 1, \ldots, m. \end{array} \tag{A.4}$$

Additionally, problem (A.4) is an equivalent reformulation of problem (A.3) if the uncertainty set Ξ admits a nonempty interior.

A.3 The S-procedure

The S-procedure provides a method of verifying the satisfaction a quadratic inequality constraint given the satisfaction of a finite number of quadratic inequality constraints. We have the following proposition from [23].

Proposition 10. Let $f_i : \mathbf{R}^n \to \mathbf{R}$ be a quadratic function of the form

$$f_i(x) := x^\mathsf{T} A_i x + 2 b_i^\mathsf{T} x + c_i,$$

for $i = 0, \ldots, p$, where $A_i = A_i^\mathsf{T}$ for $i = 0, \ldots, p$. We have that $f_i(x) \geq 0$ for $i = 1, \ldots, p$ implies $f_0(x) \geq 0$ if there exist $\tau_1, \ldots, \tau_p \geq 0$ such that

$$\begin{bmatrix} A_0 & b_0 \\ b_0^\mathsf{T} & c_0 \end{bmatrix} - \sum_{i=1}^{p} \tau_i \begin{bmatrix} A_i & b_i \\ b_i^\mathsf{T} & c_i \end{bmatrix} \succeq 0.$$

Moreover, if $p = 1$ and there exists $x \in \mathbf{R}^n$ such that $f_1(x) > 0$, then the converse of the above statement is also true.

A.4 The Schur Complement

The Schur complement provides a method of converting nonlinear (convex) matrix inequalities into linear matrix inequalities. Specifically, consider a block matrix A given by

$$A = \begin{bmatrix} Q & S \\ S^\mathsf{T} & R \end{bmatrix}, \tag{A.5}$$

where $Q \in \mathbf{R}^{n\times n}$ and $R \in \mathbf{R}^{m\times m}$ are both symmetric. If the matrix Q is invertible, then the *Schur complement* of the block Q in the block matrix A is given by

$$A/Q = R - S^\mathsf{T} Q^{-1} S.$$

Additionally, if the matrix R is invertible, then the Schur complement of the block R in the block matrix A is given by

$$A/R = Q - S R^{-1} S^\mathsf{T}.$$

Building on the concept of Schur complement, the Schur's lemma provides necessary and sufficient conditions under which the block matrix A is positive definite.

Lemma 23 (Schur's Lemma)**.** Let the block matrix A be given by Eq. (A.5). Both of the following statements are true

(i) $A \succ 0$ if and only if $Q \succ 0$ and $A/Q \succ 0$.

(ii) $A \succ 0$ if and only if $R \succ 0$ and $A/R \succ 0$.

Additionally, Schur's Lemma can be generalized to provide necessary and sufficient conditions under which the block matrix A is positive semidefinite. Specifically, the *generalized Schur complement* (cf. [117, Chapter 0]) of the block Q in the block matrix A is defined as $R - S^{\mathsf{T}} Q^{\dagger} S$, where $Q^{\dagger}$ denotes the Moore-Penrose pseudoinverse of the matrix Q. And the generalized Schur complement of the block R can be defined in an analogous fashion. We have the following lemma, which provides a necessary and sufficient condition under which $A \succeq 0$ using the concept of generalized Schur complement.

Lemma 24. Let the block matrix A be given by Eq. (A.5). Both of the following statements are true

(i) $A \succeq 0$ if and only if $R \succeq 0$, $Q - SR^{\dagger}S^{\mathsf{T}} \succ 0$ and $S(I - RR^{\dagger}) = 0$.

(ii) $A \succeq 0$ if and only if $Q \succeq 0$, $R - S^{\mathsf{T}}Q^{\dagger}S \succ 0$ and $S^{\mathsf{T}}(I - QQ^{\dagger}) = 0$.

APPENDIX B

ELLIPTICALLY CONTOURED DISTRIBUTIONS

Elliptically contoured distributions are considered to be a natural generalization to Gaussian distributions [48, 63, 73], as they inherit many properties of Gaussian distributions that are useful in problems in optimal control. As a consequence, many results in optimal control that are established under the assumption of Gaussian disturbance can be generalized to problems whose disturbances have elliptically contoured distributions [30]. In this chapter, we briefly review the definition and main properties of elliptically contoured distributions. Additionally, we illustrate how classical results that are established under the assumption of Gaussian disturbance might be generalized to problems with elliptically contoured disturbances via the investigation of a linear-quadratic (LQ) control problem whose disturbance has an elliptically contoured distribution.

B.1 Definition and Properties

The notion of elliptically contoured distributions was introduced in [63] and thoroughly studied in [25] and [43]. In what follows, we provide a formal definition of elliptically contoured distributions, and state their properties that are pertinent to the technical results established in this dissertation.

Definition 7 (Elliptically Contoured Disturbance)**.** A random vector ξ taking value in $\mathbf{R}^n$ is said to have an elliptically contoured distribution if there exists a vector $\mu \in \mathbf{R}^n$, a symmetric positive semidefinite matrix $\Sigma \in \mathbf{R}^{n\times n}$, and a scalar function g, such that the characteristic function $\varphi_{\xi-\mu}$ of the random vector $\xi - \mu$ satisfies the functional equation $\varphi_{\xi-\mu}(\theta) = g(\theta^\mathsf{T}\Sigma\theta)$ for every vector $\theta \in \mathbf{R}^n$.

The family of elliptically contoured distributions is broad. It includes the multivariate Gaussian distribution, multivariate t-distribution, their truncated versions, and uniform distributions on ellipsoids. Such a family of distributions are considered to be natural generalizations of multivariate Gaussian distributions, as they inherit several important properties of Gaussian distributions that are useful in problems in optimal control. We review such properties as follows.

Lemma 25. Let the random vector ξ be distributed according to an elliptically contoured distribution. It follows that the random vector $A\xi + b$ has an elliptically contoured distribution for each $A \in \mathbf{R}^{m\times n}$ and $b \in \mathbf{R}^m$.

Lemma 25 shows that any affine transformation of a random vector with an elliptically contoured distribution remains a random vector with an elliptically contoured distribution. Apparently, this is a generalization to the property that any affine transformation of a Gaussian random vector still has a Gaussian distribution. In the following proposition, we show that the conditional expectation of a random vector with an elliptically contoured distribution given its subvector is affine in this subvector.

Proposition 11 (Conditional Distribution)**.** Let $\xi = (\xi_1, \xi_2)$ be a random vector that has an elliptically contoured distribution and finite-valued second moment matrix. Additionally, assume that the mean μ and covariance matrix Σ of the random vector ξ are given by

$$\mu = \begin{bmatrix} \mu_1 \\ \mu_2 \end{bmatrix} \quad \Sigma = \begin{bmatrix} \Sigma_{11} & \Sigma_{12} \\ \Sigma_{12}^{\mathsf{T}} & \Sigma_{22} \end{bmatrix},$$

where $\xi_1, \mu_1 \in \mathbf{R}^{n_1}$ and $\Sigma_{11} \in \mathbf{R}^{n_1}$ for an integer $n_1 < n$. We have that

(i) The conditional probability distribution of the random vector ξ_1 given the random vector ξ_2 is an elliptically contoured distribution.

(ii) The conditional expectation of the random vector ξ_1 given the random vector ξ_2 is

$$\mathbf{E}[\xi_1|\xi_2] = \mu_1 + \Sigma_{12}\Sigma_{22}^{\dagger}(\xi_2 - \mu_2),$$

where $\Sigma_{22}^{\dagger}$ denotes the Moore-Penrose pseudoinverse of the matrix Σ_{22}.

Lemma 25 and Proposition 11 both reveal that properties of Gaussian distributions that are useful in optimal control problems also hold for random vectors with elliptically contoured distributions. This naturally implies the possibility that results in optimal control that are established under the assumption of Gaussian noise might be generalized to the setting in which the noise follows an elliptically contoured distribution. In what follows, we provide an example of such a generalization via the investigation of a LQ control problem with elliptically contoured noise.

B.2 Linear Quadratic Control with Elliptically Contoured Noise

In what follows, we consider a variant of the linear-quadratic-Gaussian (LQG) problem, in which the process and measurement noise is distributed according to an elliptically contoured distribution. We show that all properties that the optimal control policy of the LQG problem satisfies are also satisfied by the optimal control policy of a linear quadratic control problem with elliptically contoured noise. Note that in order to simplify notation, we adopt the notation convention in [20], in which subscripts denote the time indices.

Consider a discrete-time, linear time-varying system whose dynamics is described according to

$$x_{t+1} = A_t x_t + B_t u_t + w_t,$$

where $x_t \in \mathbf{R}^n$, $u_t \in \mathbf{R}^m$, and $w_t \in \mathbf{R}^n$ denotes the system state, input, and process noise at each time t. The initial state x_0, and the process noise w_t at each time t are both assumed to be zero-mean random vectors, whose second moment matrices are given by $S := \mathbf{E}[x_0 x_0^\mathsf{T}]$ and $M_t := \mathbf{E}[w_t w_t^\mathsf{T}]$, respectively. The system operates over a finite time horizon indexed by $t = 0, \ldots, T-1$. We denote by $y_t \in \mathbf{R}^p$ the measured output at each time t. It is given by

$$y_t = C_t x_t + v_t,$$

where $v_t \in \mathbf{R}^p$ denotes the measurement noise at time t. We assume that v_t is a zero-mean random vector at each time t, whose second moment matrix $N_t := \mathbf{E}[v_t v_t^\mathsf{T}]$ is assumed to be positive definite. The control input at each time t is specified as a function of the entire history of outputs up until and including time t. That is, the control input at time t is specified as

$$u_t = \gamma_t(y^t),$$

where the function $\gamma_t(\cdot)$ is a measurable function of its argument. Our objective is to specify a control policy $\gamma := (\gamma_0, \ldots, \gamma_{T-1})$ which solves the following linear quadratic control problem:

$$\begin{array}{ll} \text{minimize} & \mathbf{E}\left[x_T^\mathsf{T} Q_T x_T + \sum_{t=0}^{T-1} \left(x_t^\mathsf{T} Q_t x_t + u_t^\mathsf{T} R_t u_t \right) \right] \\ \text{subject to} & u_t = \gamma_t(y^t), \quad t = 0, \ldots, T-1 \\ & x_{t+1} = A_t x_t + B_t u_t + w_t, \quad t = 0, \ldots, T-1 \\ & y_t = C_t x_t + v_t, \quad t = 0, \ldots, T-1, \end{array} \tag{B.1}$$

where the cost matrices Q_t and R_t are assumed to be symmetric positive semidefinite and symmetric positive definite, respectively, for each time t. Our characterization of the optimal policy for problem (B.1) requires the following assumption on the process and measurement noise.

Assumption 7. We assume that the random vector $(x_0, w_0, v_0, w_1, v_1, \ldots, w_{T-1}, v_{T-1})$ has an elliptically contoured distribution with zero mean and covariance matrix

$$\Sigma = \operatorname{diag}(S, M_0, N_0, M_1, N_1, \ldots, M_{T-1}, N_{T-1}).$$

In other words, Assumption 7 requires that the random vectors $x_0, w_0, v_0, w_1, v_1, \ldots, w_{T-1}, v_{T-1}$ be mutually uncorrelated, and that their joint probability distribution be elliptically contoured. We have the following result, which shows that all properties that the optimal control policy of the LQG problem satisfies are satisfied by the optimal control policy for problem (B.1).

Theorem 3. Let Assumption 7 hold. It follows that the optimal control policy for problem (B.1) satisfies the following properties

(i) The optimal control input at each time t is linear in the conditional expectation of the system state at time t given the history of outputs y^t:

$$u_t = \gamma^*(y^t, t) = L_t \mathbf{E}[x_t | y^t].$$

(ii) The feedback control gain L_t is given by

$$L_t = -\left(R_t + B_t^\mathsf{T} K_{t+1} B_t\right)^{-1} B_t^\mathsf{T} K_{t+1} A_t,$$

where the matrix K_t can be computed recursively backwards using the Riccati equation:

$$\begin{aligned} P_t &= A_t^\mathsf{T} K_{t+1} B_t \left(R_t + B_t^\mathsf{T} K_{t+1} B_t\right)^{-1} B_t^\mathsf{T} K_{t+1} A_t \\ K_t &= A_t^\mathsf{T} K_{t+1} A_t - P_t + Q_t \end{aligned}$$

for $t = 0, \dots, T-1$ with initial condition given by $K_T = Q_T$.

(iii) The conditional expectation of the system state at time t can be calculated using the Kalman filter. That is, the

$$\begin{aligned}\mathbf{E}[x_t | y^t] =& (A_{t-1} + B_{t-1}L_{t-1})\mathbf{E}[x_{t-1}|y^{t-1}] \\ &+ \Sigma_{t|t} C_t^\mathsf{T} N_t^{-1} \left(y_t - C_t(A_{t-1} + B_{t-1}L_{t-1})\mathbf{E}[x_{t-1}|y^{t-1}]\right)\end{aligned}$$

where the matrices $\Sigma_{t|t}$ can be computed recursively according to

$$\begin{aligned}\Sigma_{t|t} &= \Sigma_{t|t-1} - \Sigma_{t|t-1} C_t^\mathsf{T} \left(C_t \Sigma_{t|t-1} C_t^\mathsf{T} + N_t \right)^{-1} C_t \Sigma_{t|t-1} \\ \Sigma_{t|t-1} &= A_{t-1}\Sigma_{t-1|t-1}A_{t-1}^\mathsf{T} + M_{t-1},\end{aligned}$$

for $t = 0, \dots, T-1$, with initial conditions given by $\Sigma_{0|-1} = S$ and $\mathbf{E}[x_{-1}|y^{-1}] = 0$.

Theorem 3 provides a concrete example showing that results in optimal control that are established under the assumption of Gaussian noise can be generalized to the setting in which the noise is distributed according to an elliptically contoured distributions. Additionally, we note that out of the three properties stated in Theorem 3, properties (i) and (ii) (i.e., the separation principle and the optimality of certainty equivalence) hold for a much larger family of LQ control problems. Specifically, it was shown in [3] that for arbitrary LQ control problems, in which the process noise and the measurement noise might correlate across time, the separation principle holds, and the optimal control policy is certainty equivalent. However, Assumption 7 is required for the satisfaction of property (iii), as the ability to recursively estimate the system state using the Kalman filter follows from a combination of Lemma 25 and Proposition 11.

BIBLIOGRAPHY

[1] *IEEE 1547 Standard for Interconnecting Distributed Resources with Electric Power Systems*. Available at: http://grouper.ieee.org/groups/scc21/1547/1547_index.html.

[2] *SCE 2016 Static Load Profiles*. Available at: https://www.sce.com/wps/portal/home/regulatory/load-profiles/2016%20static%20load%20profiles/.

[3] Hajime Akashi and Kazuo Nose. On certainty equivalence of stochastic optimal control problem. *International Journal of Control*, 21(5):857–863, 1975.

[4] Rajeev Alur and Thomas A Henzinger. Reactive modules. *Formal methods in system design*, 15(1):7–48, 1999.

[5] Seyed Mohammad Asghari and Ashutosh Nayyar. Dynamic teams and decentralized control problems with substitutable actions. *IEEE Transactions on Automatic Control*, 62(10):5302–5309, 2017.

[6] Marco Baglietto, Thomas Parisini, and Riccardo Zoppoli. Numerical solutions to the witsenhausen counterexample by approximating networks. *IEEE Transactions on Automatic Control*, 46(9):1471–1477, 2001.

[7] Rajesh Bansal and Tamer Basar. Stochastic teams with nonclassical information revisited: When is an affine law optimal? In *1986 American Control Conference*, pages 45–50. IEEE, 1986.

[8] Rajesh Bansal and Tamer Başar. Solutions to a class of linear-quadratic-gaussian (lqg) stochastic team problems with nonclassical information. *Systems & Control Letters*, 9(2):125–130, 1987.

[9] Mesut E Baran and Felix F Wu. Network reconfiguration in distribution systems for loss reduction and load balancing. *Power Delivery, IEEE Transactions on*, 4(2):1401–1407, 1989.

[10] Mesut E Baran and Felix F Wu. Optimal capacitor placement on radial distribution systems. *Power Delivery, IEEE Transactions on*, 4(1):725–734, 1989.

[11] Mesut E Baran and Felix F Wu. Optimal sizing of capacitors placed on a radial distribution system. *Power Delivery, IEEE Transactions on*, 4(1):735–743, 1989.

[12] João Carlos Alves Barata and Mahir Saleh Hussein. The Moore–Penrose pseudoinverse: A tutorial review of the theory. *Brazilian Journal of Physics*, 42(1-2):146–165, 2012.

[13] Tamer Basar. Variations on the theme of the witsenhausen counterexample. In *Decision and Control, 2008. CDC 2008. 47th IEEE Conference on*, pages 1614–1619. IEEE, 2008.

[14] Mohammadhafez Bazrafshan and Nikolaos Gatsis. Decentralized stochastic programming for real and reactive power management in distribution systems. In *Smart Grid Communications (SmartGridComm), 2014 IEEE International Conference on*, pages 218–223. IEEE, 2014.

[15] Alberto Bemporad, Maurice Heemels, and Mikael Johansson. *Networked control systems*, volume 406. Springer, 2010.

[16] Aharon Ben-Tal, Stephen Boyd, and Arkadi Nemirovski. Control of uncertainty-affected discrete time linear systems via convex programming. Technical report, Minerva Optimization Center, Technion, Haifa, Israel, 2005.

[17] Aharon Ben-Tal, Stephen Boyd, and Arkadi Nemirovski. Extending scope of robust optimization: Comprehensive robust counterparts of uncertain problems. *Mathematical Programming*, 107(1-2):63–89, 2006.

[18] Aharon Ben-Tal, Laurent El Ghaoui, and Arkadi Nemirovski. *Robust optimization*. Princeton University Press, 2009.

[19] Toby Berger. Rate-distortion theory. *Wiley Encyclopedia of Telecommunications*, 2003.

[20] Dimitri P Bertsekas. *Dynamic programming and optimal control*, volume 1. Athena scientific Belmont, MA, 1995.

[21] Dimitris Bertsimas, Dan Andrei Iancu, Pablo Parrilo, et al. A hierarchy of near-optimal policies for multistage adaptive optimization. *Automatic Control, IEEE Transactions on*, 56(12):2809–2824, 2011.

[22] Saverio Bolognani and Sandro Zampieri. A distributed control strategy for reactive power compensation in smart microgrids. *Automatic Control, IEEE Transactions on*, 58(11):2818–2833, 2013.

[23] Stephen Boyd, Laurent El Ghaoui, Eric Feron, and Venkataramanan Balakrishnan. *Linear matrix inequalities in system and control theory*, volume 15. Siam, 1994.

[24] David B Brown, James E Smith, and Peng Sun. Information relaxations and duality in stochastic dynamic programs. *Operations research*, 58(4-part-1):785–801, 2010.

[25] Stamatis Cambanis, Steel Huang, and Gordon Simons. On the theory of elliptically contoured distributions. *Journal of Multivariate Analysis*, 11(3):368–385, 1981.

[26] Eduardo Camponogara, Dong Jia, Bruce H Krogh, and Sarosh Talukdar. Distributed model predictive control. *Control Systems, IEEE*, 22(1):44–52, 2002.

[27] Michael Caramanis and Justin M Foster. Management of electric vehicle charging to mitigate renewable generation intermittency and distribution network congestion. In *Decision and Control, 2009 held jointly with the 2009 28th Chinese Control Conference. CDC/CCC 2009. Proceedings of the 48th IEEE Conference on*, pages 4717–4722. Ieee, 2009.

[28] Lin Cheng, Yao Chang, and Renle Huang. Mitigating voltage problem in distribution system with distributed solar generation using electric vehicles. *Sustainable Energy, IEEE Transactions on*, 6(4):1475–1484, 2015.

[29] Kai Ching Chu. Team decision theory and information structures in optimal control problems–Part II. *IEEE Transactions on Automatic Control*, 17(1):22–28, 1972.

[30] K'ai Ching Chu. Estimation and decision for linear-systems with elliptical random processes. *IEEE Transactions on Automatic Control*, (5):499–505, 1973.

[31] Kristien Clement, Edwin Haesen, and Johan Driesen. Coordinated charging of multiple plug-in hybrid electric vehicles in residential distribution grids. In *Power Systems Conference and Exposition, 2009. PSCE'09. IEEE/PES*, pages 1–7. IEEE, 2009.

[32] Emiliano Dall'Anese, Sairaj V Dhople, and Georgios B Giannakis. Optimal dispatch of photovoltaic inverters in residential distribution systems. *Sustainable Energy, IEEE Transactions on*, 5(2):487–497, 2014.

[33] Emiliano Dall'Anese, Hao Zhu, and Georgios B Giannakis. Distributed optimal power flow for smart microgrids. *Smart Grid, IEEE Transactions on*, 4(3):1464–1475, 2013.

[34] Emiliano DallAnese, Sairaj V Dhople, Brian B Johnson, and Georgios B Giannakis. Decentralized optimal dispatch of photovoltaic inverters in residential distribution systems. *Energy Conversion, IEEE Transactions on*, 29(4):957–967, 2014.

[35] Emiliano DallAnese, Sairaj V Dhople, Brian B Johnson, and Georgios B Giannakis. Optimal dispatch of residential photovoltaic inverters under forecasting uncertainties. *IEEE Journal of Photovoltaics*, 5(1):350–359, 2015.

[36] Georgios Darivianakis, Angelos Georghiou, and John Lygeros. Decentralized decision making for networks of uncertain systems. *arXiv preprint arXiv:1803.07660*, 2018.

[37] Sara Deilami, Amir S Masoum, Paul S Moses, and Mohammad AS Masoum. Real-time coordination of plug-in electric vehicle charging in smart grids to minimize power losses and improve voltage profile. *Smart Grid, IEEE Transactions on*, 2(3):456–467, 2011.

[38] Deepjyoti Deka, Scott Backhaus, and Michael Chertkov. Structure learning in power distribution networks. *IEEE Transactions on Control of Network Systems*, 5(3):1061–1074, 2017.

[39] V Desai, V Farias, and C Moallemi. Bounds for markov decision processes. *Reinforcement Learning and Approximate Dynamic Programming for Feedback Control,(FL Lewis, D. Liu, eds.)*, pages 452–473, 2011.

[40] William B Dunbar. Distributed receding horizon control of dynamically coupled nonlinear systems. *IEEE Transactions on Automatic Control*, 52(7):1249–1263, 2007.

[41] William B Dunbar and Richard M Murray. Distributed receding horizon control for multi-vehicle formation stabilization. *Automatica*, 42(4):549–558, 2006.

[42] C Durieu, E Walter, and B Polyak. Multi-input multi-output ellipsoidal state bounding. *Journal of optimization theory and applications*, 111(2):273–303, 2001.

[43] Kai Wang Fang. *Symmetric multivariate and related distributions*. Chapman and Hall/CRC, 2018.

[44] Marcello Farina and Riccardo Scattolini. Distributed predictive control: A non-cooperative algorithm with neighbor-to-neighbor communication for linear systems. *Automatica*, 48(6):1088–1096, 2012.

[45] Masoud Farivar, Lijun Chen, and Steven Low. Equilibrium and dynamics of local voltage control in distribution systems. In *Decision and Control (CDC), 2013 IEEE 52nd Annual Conference on*, pages 4329–4334. IEEE, 2013.

[46] Masoud Farivar, Russell Neal, Christopher Clarke, and Steven Low. Optimal inverter var control in distribution systems with high pv penetration. In *Power and Energy Society General Meeting, 2012 IEEE*, pages 1–7. IEEE, 2012.

[47] Paulo DF Ferreira, Pedro Carvalho, Luis AFM Ferreira, and Marija D Ilic. Distributed energy resources integration challenges in low-voltage networks: Voltage control limitations and risk of cascading. *Sustainable Energy, IEEE Transactions on*, 4(1):82–88, 2013.

[48] Gabriel Frahm. *Generalized elliptical distributions: theory and applications*. PhD thesis, Universität zu Köln, 2004.

[49] Lingwen Gan, Ufuk Topcu, and Steven H Low. Optimal decentralized protocol for electric vehicle charging. *Power Systems, IEEE Transactions on*, 28(2):940–951, 2013.

[50] Ather Gattami. *Optimal decisions with limited information*. PhD thesis, Lund University, 2007.

[51] Angelos Georghiou, Wolfram Wiesemann, and Daniel Kuhn. Generalized decision rule approximations for stochastic programming via liftings. *Mathematical Programming*, 152(1-2):301–338, 2015.

[52] Pulkit Grover, Se Yong Park, and Anant Sahai. Approximately optimal solutions to the finite-dimensional Witsenhausen counterexample. *IEEE Transactions on Automatic Control*, 58(9):2189–2204, 2013.

[53] Pulkit Grover and Anant Sahai. Witsenhausen's counterexample as assisted interference suppression. *International Journal of Systems, Control and Communications*, 2(1-3):197–237, 2010.

[54] Michael J Hadjiyiannis, Paul J Goulart, and Daniel Kuhn. An efficient method to estimate the suboptimality of affine controllers. *Automatic Control, IEEE Transactions on*, 56(12):2841–2853, 2011.

[55] Bernardo Hernandez, Pablo Baldivieso, and Paul Trodden. Distributed mpc: Guaranteeing global stability from locally designed tubes. *IFAC-PapersOnLine*, 50(1):11829–11834, 2017.

[56] Yu Chi Ho and Kai-hing Chu. Team decision theory and information structures in optimal control problems–part i. *Automatic Control, IEEE Transactions on*, 17(1):15–22, 1972.

[57] Yu-Chi Ho and KC Chu. Information structure in dynamic multi-person control problems. *Automatica*, 10(4):341–351, 1974.

[58] International Electrotechnical Commission et al. *Efficient Electrical Energy Transmission and Distribution*. IEC, 2007.

[59] Irwin Mark Jacobs and JM Wozencraft. Principles of communication engineering. 1965.

[60] Dong Jia and Bruce H Krogh. Distributed model predictive control. In *American Control Conference, 2001. Proceedings of the 2001*, volume 4, pages 2767–2772. IEEE, 2001.

[61] Sharu Theresa Jose and Ankur A Kulkarni. A linear programming relaxation for stochastic control problems with non-classical information patterns. In *Decision and Control (CDC), 2015 IEEE 54th Annual Conference on*, pages 5743–5748. IEEE, 2015.

[62] Vassilis Kekatos, Gang Wang, Antonio J Conejo, and Georgios B Giannakis. Stochastic reactive power management in microgrids with renewables. *Power Systems, IEEE Transactions on*, 30(6):3386–3395, 2015.

[63] Douglas Kelker. Distribution theory of spherical distributions and a location-scale parameter generalization. *Sankhyā: The Indian Journal of Statistics, Series A*, pages 419–430, 1970.

[64] TamáS Keviczky, Francesco Borrelli, and Gary J Balas. Decentralized receding horizon control for large scale dynamically decoupled systems. *Automatica*, 42(12):2105–2115, 2006.

[65] Daniel Kuhn, Wolfram Wiesemann, and Angelos Georghiou. Primal and dual linear decision rules in stochastic and robust optimization. *Mathematical Programming*, 130(1):177–209, 2011.

[66] Ankur Kulkarni, Todd P Coleman, et al. An optimizer's approach to stochastic control problems with nonclassical information structures. *Automatic Control, IEEE Transactions on*, 60(4):937–949, 2015.

[67] Marta Kwiatkowska, Gethin Norman, David Parker, and Hongyang Qu. Assume-guarantee verification for probabilistic systems. In *International Conference on Tools and Algorithms for the Construction and Analysis of Systems*, pages 23–37. Springer, 2010.

[68] Jonathan T Lee, Edward Lau, and Yu-Chi Ho. The witsenhausen counterexample: A hierarchical search approach for nonconvex optimization problems. *IEEE Transactions on Automatic Control*, 46(3):382–397, 2001.

[69] Laurent Lessard and Sanjay Lall. Quadratic invariance is necessary and sufficient for convexity. In *American Control Conference (ACC), 2011*, pages 5360–5362. IEEE, 2011.

[70] Na Li, Guannan Qu, and Munther Dahleh. Real-time decentralized voltage control in distribution networks. In *Communication, Control, and Computing (Allerton), 2014 52nd Annual Allerton Conference on*, pages 582–588. IEEE, 2014.

[71] Fu Lin, Mohammad Fardad, and Mihailo R Jovanovic. Design of optimal sparse feedback gains via the alternating direction method of multipliers. *Automatic Control, IEEE Transactions on*, 58(9):2426–2431, 2013.

[72] Weixuan Lin and Eilyan Bitar. Performance bounds for robust decentralized control. In *American Control Conference (ACC), 2016*, pages 4323–4330. IEEE, 2016.

[73] JK Lindsey. Multivariate elliptically contoured distributions for repeated measurements. *Biometrics*, 55(4):1277–1280, 1999.

[74] Emmanouil Loukarakis, Chris J Dent, and Janusz W Bialek. Decentral-

ized multi-period economic dispatch for real-time flexible demand management. *Power Systems, IEEE Transactions on*, 31(1):672–684, 2016.

[75] Sergio Lucia, Markus Kögel, and Rolf Findeisen. Contract-based predictive control of distributed systems with plug and play capabilities. *IFAC-PapersOnLine*, 48(23):205–211, 2015.

[76] Aditya Mahajan, Nuno C Martins, Michael C Rotkowitz, and Serdar Yuksel. Information structures in optimal decentralized control. In *Decision and Control (CDC), 2012 IEEE 51st Annual Conference on*, pages 1291–1306. IEEE, 2012.

[77] Aditya Mahajan and Serdar Yuksel. Measure and cost dependent properties of information structures. In *American Control Conference (ACC), 2010*, pages 6397–6402. IEEE, 2010.

[78] Francesco Marra, GY Yang, Yehia Tarek Fawzy, Chresten Traeholt, Esben Larsen, Rodrigo Garcia-Valle, and M Moller Jensen. Improvement of local voltage in feeders with photovoltaic using electric vehicles. *Power Systems, IEEE Transactions on*, 28(3):3515–3516, 2013.

[79] Nikolai Matni and Venkat Chandrasekaran. Regularization for design. In *Decision and Control (CDC), 2014 IEEE 53rd Annual Conference on*, pages 1111–1118. IEEE, 2014.

[80] Hans D Mittelmann. The state-of-the-art in conic optimization software. In *Handbook on Semidefinite, Conic and Polynomial Optimization*, pages 671–686. Springer, 2012.

[81] Roy Radner. Team decision problems. *The Annals of Mathematical Statistics*, 33(3):857–881, 1962.

[82] Arthur Richards and Jonathan How. A decentralized algorithm for robust constrained model predictive control. In *American Control Conference, 2004. Proceedings of the 2004*, volume 5, pages 4261–4266. IEEE, 2004.

[83] Arthur Richards and Jonathan P How. Robust distributed model predictive control. *International Journal of control*, 80(9):1517–1531, 2007.

[84] Stefano Riverso and Giancarlo Ferrari-Trecate. Tube-based distributed control of linear constrained systems. *Automatica*, 48(11):2860–2865, 2012.

[85] Brett Robbins, Christoforos N Hadjicostis, Alejandro D Domínguez-García, et al. A two-stage distributed architecture for voltage control in power distribution systems. *Power Systems, IEEE Transactions on*, 28(2):1470–1482, 2013.

[86] Michael Rotkowitz. On information structures, convexity, and linear optimality. In *Decision and Control, 2008. CDC 2008. 47th IEEE Conference on*, pages 1642–1647. IEEE, 2008.

[87] Michael C Rotkowitz and Nuno C Martins. On the nearest quadratically invariant information constraint. *Automatic Control, IEEE Transactions on*, 57(5):1314–1319, 2012.

[88] Anant Sahai and Pullkit Grover. Demystifying the witsenhausen counterexample [ask the experts]. *IEEE Control Systems Magazine*, 30(6):20–24, 2010.

[89] Nils R Sandell Jr, Pravin Varaiya, Michael Athans, and Michael G Safonov. Survey of decentralized control methods for large scale systems. *Automatic Control, IEEE Transactions on*, 23(2):108–128, 1978.

[90] Joelle Skaf and Stephen P Boyd. Design of affine controllers via convex optimization. *Automatic Control, IEEE Transactions on*, 55(11):2476–2487, 2010.

[91] R Srikant and T Basar. Asymptotic solutions to weakly coupled stochastic teams with nonclassical information. *IEEE transactions on Automatic Control*, 37(2):163–173, 1992.

[92] Charlotte Striebel. Sufficient statistics in the optimum control of stochastic systems. *Journal of Mathematical Analysis and Applications*, 12(3):576–592, 1965.

[93] Vignesh Subramanian, Laura Brink, Nikunj Jain, Kailas Vodrahalli, Akhil Jalan, Nikhil Shinde, and Anant Sahai. Some new numeric results concerning the witsenhausen counterexample. In *2018 56th Annual Allerton Conference on Communication, Control, and Computing (Allerton)*, pages 413–420. IEEE, 2018.

[94] Petr Sulc, Scott Backhaus, and Michael Chertkov. Optimal distributed control of reactive power via the alternating direction method of multipliers. *Energy Conversion, IEEE Transactions on*, 29(4):968–977, 2014.

[95] John Swigart and Sanjay Lall. A graph-theoretic approach to distributed control over networks. In *Decision and Control, 2009 held jointly with the 2009 28th Chinese Control Conference. CDC/CCC 2009. Proceedings of the 48th IEEE Conference on*, pages 5409–5414. IEEE, 2009.

[96] Joshua Traube, Fenglong Lu, Dragan Maksimovic, Joseph Mossoba, Matthew Kromer, Peter Faill, Stan Katz, Bogdan Borowy, Steve Nichols, and Leo Casey. Mitigation of solar irradiance intermittency in photovoltaic power systems with integrated electric-vehicle charging functionality. *Power Electronics, IEEE Transactions on*, 28(6):3058–3067, 2013.

[97] PA Trodden and AG Richards. Cooperative tube-based distributed mpc for linear uncertain systems coupled via constraints. In *Distributed Model Predictive Control Made Easy*, pages 57–72. Springer, 2014.

[98] Paul Trodden and Arthur Richards. Robust distributed model predictive control using tubes. In *American Control Conference, 2006*, pages 6–pp. IEEE, 2006.

[99] Paul Trodden and Arthur Richards. Distributed model predictive control of linear systems with persistent disturbances. *International Journal of Control*, 83(8):1653–1663, 2010.

[100] Paul A Trodden and Jose Maria Maestre. Distributed predictive control with minimization of mutual disturbances. *Automatica*, 77:31–43, 2017.

[101] John N Tsitsiklis and Michael Athans. On the complexity of decentralized decision making and detection problems. *Automatic Control, IEEE Transactions on*, 30(5):440–446, 1985.

[102] Konstantin Turitsyn, Scott Backhaus, Michael Chertkov, et al. Options for control of reactive power by distributed photovoltaic generators. *Proceedings of the IEEE*, 99(6):1063–1073, 2011.

[103] US Department of Energy (DoE). Advanced distribution management system program, multi-year program plan, 2016-2020. Technical report, 2016. Available at: `https://www.smartgrid.gov/files/ADMS_MYPP_20160824_DRAFT.pdf`.

[104] DH Van Hessem and OH Bosgra. A conic reformulation of model predictive control including bounded and stochastic disturbances under state and input constraints. In *Decision and Control, 2002, Proceedings of the 41st IEEE Conference on*, volume 4, pages 4643–4648. IEEE, 2002.

[105] Aswin N Venkat, James B Rawlings, and Stephen J Wright. Stability and optimality of distributed model predictive control. In *Decision and Control, 2005 and 2005 European Control Conference. CDC-ECC'05. 44th IEEE Conference on*, pages 6680–6685. IEEE, 2005.

[106] Maria Vrakopoulou, Johanna L Mathieu, and Göran Andersson. Stochastic optimal power flow with uncertain reserves from demand response. In *2014 47th Hawaii International Conference on System Sciences*, pages 2353–2362. IEEE, 2014.

[107] Chen Wang, Chong-Jin Ong, and Melvyn Sim. Model predictive control using segregated disturbance feedback. *Automatic Control, IEEE Transactions on*, 55(4):831–840, 2010.

[108] Stephen Warshall. A theorem on boolean matrices. *Journal of the ACM (JACM)*, 9(1):11–12, 1962.

[109] Hans S Witsenhausen. A counterexample in stochastic optimum control. *SIAM Journal on Control*, 6(1):131–147, 1968.

[110] AD Wyner. Another look at the coding theorem of information theorya tutorial. *Proceedings of the IEEE*, 58(6):894–913, 1970.

[111] Fan Ye and Enlu Zhou. Information relaxation and dual formulation of controlled markov diffusions. 2013.

[112] Dante C Youla, Hamid Jabr, Joseph J Bongiorno Jr, et al. Modern wiener-hopf design of optimal controllers–part ii: The multivariable case. *Automatic Control, IEEE Transactions on*, 21(3):319–338, 1976.

[113] Serdar Yüksel. Stochastic nestedness and the belief sharing information pattern. *IEEE Transactions on Automatic Control*, 54(12):2773–2786, 2009.

[114] Serdar Yüksel and Tamer Basar. *Stochastic networked control systems: Stabilization and optimization under information constraints*. Springer Science & Business Media, 2013.

[115] Serdar Yüksel and Naci Saldi. Convex analysis in decentralized stochastic control, strategic measures, and optimal solutions. *SIAM Journal on Control and Optimization*, 55(1):1–28, 2017.

[116] Baosen Zhang, Albert Lam, Alejandro D Domínguez-García, and David

Tse. An optimal and distributed method for voltage regulation in power distribution systems. *Power Systems, IEEE Transactions on*, 30(4):1714–1726, 2015.

[117] Fuzhen Zhang. *The Schur complement and its applications*, volume 4. Springer Science & Business Media, 2006.

[118] Xiaojing Zhang, Maryam Kamgarpour, Angelos Georghiou, Paul Goulart, and John Lygeros. Robust optimal control with adjustable uncertainty sets. *Automatica*, 75:249–259, 2017.

[119] Hao Zhu and Hao Jan Liu. Fast local voltage control under limited reactive power: Optimality and stability analysis. *arXiv preprint arXiv:1510.07013*, 2015.