

A Convex Information Relaxation for Constrained Decentralized Control Design Problems

Weixuan Lin Eilyan Bitar

Abstract—We describe a convex programming approach to the calculation of lower bounds on the minimum cost of constrained decentralized control problems with nonclassical information structures. The class of problems we consider entail the decentralized output feedback control of a linear time-varying system over a finite horizon, subject to polyhedral constraints on the state and input trajectories, and sparsity constraints on the controller's information structure. As the determination of optimal control policies for such systems is known to be computationally intractable in general, considerable effort has been made in the literature to identify efficiently computable, albeit suboptimal, feasible control policies. The construction of computationally tractable bounds on their suboptimality is the primary motivation for the techniques developed in this note. 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.

I. INTRODUCTION

In general, the design of an optimal decentralized controller amounts to an infinite-dimensional nonconvex optimization problem. The difficulty in solution derives in part from the manner in which information is shared between controllers—the so called information structure of a problem; see [2] for a survey. Considerable effort has been made to identify information structures under which the problem of decentralized control design can be recast as an equivalent convex program. For instance, partial nestedness of the information structure [3] is known to simplify the control design, as it eliminates the incentive to signal between controllers. In particular, linear control policies are guaranteed to be optimal for decentralized LQG problems with partially nested information structures [3]. Closely related notions of quadratic invariance [4] and funnel causality [5] guarantee convexity of decentralized controller synthesis, whose objective is to minimize the closed-loop norm of an LTI system.

As the tractable computation of optimal policies for the majority of decentralized control problems with nonclassical information structures remains out of reach [2], 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 [6] 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]–[9], and linear programming-based relaxations [10]. 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 [11]–[15] or quadratic invariance [16] of the expanded information structure.

The specific setting that we consider entails the decentralized output feedback control of a discrete-time, linear time-varying system over a finite horizon, subject to polyhedral constraints on the state and input trajectories. The system being controlled is partitioned into N dynamically coupled subsystems, each of which has a dedicated local controller. In this setting, the decentralization of information is expressed according to *sparsity constraints* on the information that each local controller has access to. Namely, each local controller is allowed access to the outputs of some subsystems but not others. Naturally, information constraints of this form may yield information structures that are *nonclassical* in nature, thereby making the calculation of optimal decentralized control policies computationally intractable for such systems. As a result, significant research effort has been directed towards the development of methods for the tractable calculation of constraint admissible, albeit suboptimal, policies. See [17] for an overview. The aim of this note—which serves to complement these research efforts—is the tractable evaluation of their suboptimality.

Summary of Results: In this note, 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 so-called 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 using a nonlinear change of variables akin to the Youla parameterization [18]. 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

Supported in part by NSF grant ECCS-1351621, NSF grant IIP-1632124, US DOE under the CERTS initiative, and the Simons Institute for the Theory of Computing. This technical note builds on the authors' preliminary results published as part of the 2016 IEEE American Control Conference (ACC) [1].

W. Lin (wl476@cornell.edu) and E. Bitar (eyb5@cornell.edu) are with the School of Electrical and Computer Engineering, Cornell University, Ithaca, NY, 14853, USA.

policy and the proposed lower bound is small, then one may conclude that said policy is near-optimal.

Notation: Let $\mathbf{R}$ denote the set of real numbers. Denote the transpose of a vector $x \in \mathbf{R}^n$ by $x^\top$. For any pair of vectors $x = (x_1, \dots, x_n) \in \mathbf{R}^n$ and $y = (y_1, \dots, y_m) \in \mathbf{R}^m$, we define their concatenation as $(x, y) = (x_1, \dots, x_n, y_1, \dots, y_m) \in \mathbf{R}^{n+m}$. Given a process $\{x(t)\}$ indexed by $t = 0, \dots, T-1$, we denote by $x^t = (x(0), x(1), \dots, x(t))$ its history up until and including time t . We consider block matrices throughout the paper. 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 $\text{Tr}(A)$. We denote by $\mathcal{K}_2$ a second-order cone, whose dimension will be clear from the context. Specifically, given a vector $x \in \mathbf{R}^n$, $x \in \mathcal{K}_2$ if and only if $x_1 \geq \sqrt{x_2^2 + \dots + x_n^2}$. Given a matrix A , we let $A \succeq_{\mathcal{K}_2} 0$ denote its columnwise inclusion in $\mathcal{K}_2$.

II. PROBLEM FORMULATION

A. System Model

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

$$x_i(t+1) = \sum_{j=1}^N (A_{ij}(t)x_j(t) + B_{ij}(t)u_j(t)) + G_i(t)\xi(t), \quad (1)$$

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

$$y_i(t) = \sum_{j=1}^N C_{ij}(t)x_j(t) + H_i(t)\xi(t), \quad (2)$$

for $i = 1, \dots, N$. All system matrices are assumed to be real and of compatible dimension. In the sequel, we will work with a more compact representation of the system Eqs. (1) and (2) given by

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

Here, we denote by $x(t) := (x_1(t), \dots, x_N(t)) \in \mathbf{R}^{n_x}$, $u(t) := (u_1(t), \dots, u_N(t)) \in \mathbf{R}^{n_u}$, and $y(t) := (y_1(t), \dots, y_N(t)) \in \mathbf{R}^{n_y}$ the full system state, input, and output at time t , respectively. Their dimensions are given by $n_x := \sum_{i=1}^N n_x^i$, $n_u := \sum_{i=1}^N n_u^i$, and $n_y := \sum_{i=1}^N n_y^i$, respectively. 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. (1) and (2). The system trajectories are related according to

$$x = Bu + G\xi \quad \text{and} \quad y = Cx + H\xi,$$

where x, u, ξ , and y denote the trajectories of the full system state, input, disturbance, and output, respectively.¹ We denote them by

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

¹It is straightforward to construct the block matrices $(A(t), B(t), G(t), C(t), H(t))$ from the data defining the system Eqs. (1) and (2). The specification of the matrices (B, G, C, H) can be found in Appendix V.

Notice that in our specification of the both the disturbance and output trajectories, ξ 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.

We close this subsection by stating a structural assumption on the system dynamics. Assumption 1, which is assumed to hold throughout the paper, 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 $[C(t)A_{s+1}^t B(s)]_{ii}$ is nonzero.

The matrix $[C(t)A_{s+1}^t B(s)]_{ii}$ refers to the $(i, i)^{\text{th}}$ block of the $N \times N$ block matrix $C(t)A_{s+1}^t B(s)$. We refer the reader to Appendix V for a definition of the matrix A_{s+1}^t .

B. Disturbance Model

We model the disturbance trajectory ξ as a random vector defined according to the probability space $(\mathbf{R}^{N_\xi}, \mathcal{B}(\mathbf{R}^{N_\xi}), \mathbf{P})$. Here, the Borel σ -algebra $\mathcal{B}(\mathbf{R}^{N_\xi})$ denotes the set of all events that are assigned probability according to the measure $\mathbf{P}$. We denote by $\mathcal{L}_n^2 := \mathcal{L}^2(\mathbf{R}^{N_\xi}, \mathcal{B}(\mathbf{R}^{N_\xi}), \mathbf{P}; \mathbf{R}^n)$ the space of all $\mathcal{B}(\mathbf{R}^{N_\xi})$ -measurable, square-integrable random vectors taking values in $\mathbf{R}^n$. With a slight abuse of notation, we occasionally use ξ to denote a realization of the random vector ξ . The following assumption on the probability distribution of the disturbance trajectory is assumed to hold throughout the paper.

Assumption 2 (Elliptically Contoured Disturbance). The disturbance trajectory ξ is assumed to have an elliptically contoured distribution. That is, there exists a vector $\mu \in \mathbf{R}^{N_\xi}$, a positive semidefinite matrix $\Sigma \in \mathbf{R}^{N_\xi \times N_\xi}$, and a scalar function g , such that the characteristic function $\varphi_{\xi-\mu}$ of the random vector $\xi - \mu$ satisfies the functional equation $\varphi_{\xi-\mu}(\theta) = g(\theta^\top \Sigma \theta)$ for every vector $\theta \in \mathbf{R}^{N_\xi}$.

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 ξ has an elliptically contoured distribution, then the conditional expectation of ξ given a subvector of ξ is affine in this subvector. And any linear transformation of ξ also follows an elliptically contoured distribution [19]. Such properties will play an integral role in the derivation of our main result in Section V.

In order to ensure the well-posedness of the problem to follow, we require that the disturbance trajectory satisfy the following conditions. We assume that the disturbance ξ has support that is an ellipsoid with a nonempty relative interior in the hyperplane $\{\xi \in \mathbf{R}^{N_\xi} \mid \xi_1 = 1\}$, and is representable by

$$\Xi = \{\xi \in \mathbf{R}^{N_\xi} \mid \xi_1 = 1 \text{ and } W\xi \succeq_{\mathcal{K}_2} 0\},$$

where $W \in \mathbf{R}^{N_\xi \times N_\xi}$. This assumption ensures that the corresponding second-order moment matrix $M := \mathbf{E}[\xi\xi^\top]$ is positive definite and finite-valued.

C. 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,

$$\left. \begin{aligned} F_x x + F_u u + F_\xi \xi + s &= 0 \\ s &\geq 0 \end{aligned} \right\} \mathbf{P}\text{-a.s.}, \quad (3)$$

where $F_x \in \mathbf{R}^{m \times N_x}$, $F_u \in \mathbf{R}^{m \times N_u}$, and $F_\xi \in \mathbf{R}^{m \times N_\xi}$. Here, $s \in \mathcal{L}_m^2$ is a slack variable that is required to be non-negative **P**-almost surely.

D. Decentralized Control Design

In this paper, we consider information structures that are specified via *sparsity constraints* on the local controllers. 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 $\mathcal{V} = \{1, \dots, N\}$ 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 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 3. 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

$$z_i(t) := \{y_j^t \mid (j, i) \in \mathcal{E}_I\}. \quad (4)$$

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

$$u_i(t) = \gamma_i(z_i(t), t), \quad (5)$$

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 $\gamma_i := (\gamma_i(\cdot, 0), \dots, \gamma_i(\cdot, T-1))$. We refer to the collection of local control policies $\gamma := (\gamma_1, \dots, \gamma_N)$ 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 $\mathcal{G}_I$.

Of interest is the characterization of control policies, which solve the following constrained *decentralized control design problem*:

$$\begin{aligned} & \text{minimize} \quad \mathbf{E} \left[x^\top R_x x + u^\top R_u u \right] \\ & \text{subject to} \quad \gamma \in \Gamma(\mathcal{G}_I), \quad s \in \mathcal{L}_m^2 \\ & \quad \left. \begin{aligned} F_x x + F_u u + F_\xi \xi + s &= 0 \\ x &= Bu + G\xi \\ y &= Cx + H\xi \\ u &= \gamma(y) \\ s &\geq 0 \end{aligned} \right\} \quad \mathbf{P}\text{-a.s.} \end{aligned} \quad (6)$$

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 positive semidefinite. We denote the *optimal value* of problem (6) by $J^*(\mathcal{G}_I)$, where we have made explicit the dependence of the optimal value of problem (6) on the underlying information graph $\mathcal{G}_I$. In general, the decentralized control design problem (6) amounts to an infinite-dimensional, nonconvex optimization problem with neither analytical nor computationally efficient solution available at present time [2], [20], [21]. As a result, considerable effort has been directed towards the development of methods that enable the tractable calculation of feasible control policies [17]. Although these methods are known to produce decentralized controllers that perform well empirically, they are suboptimal in general; and the question as to how far from optimal these policies are remains unanswered. The primary objective of this note is the development of tractable computational methods to estimate their suboptimality.

III. PRELIMINARIES

In what follows, we describe how to equivalently reformulate the decentralized control design problem (6) as a static team problem [3] 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.

A. Nonlinear Youla Parameterization

Define the nonlinear Youla parameterization of the decentralized control policy $\gamma \in \Gamma(\mathcal{G}_I)$ as

$$\phi := \gamma \circ (I - CB\gamma)^{-1}. \quad (7)$$

Note that the map $I - CB\gamma : \mathbf{R}^{N_y} \rightarrow \mathbf{R}^{N_y}$ 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

$$\phi(\eta) = \gamma(y) \quad (8)$$

for every disturbance trajectory $\xi \in \Xi$. Here, η denotes the so-called *purified output trajectory* defined by $\eta := P\xi$, where the matrix $P \in \mathbf{R}^{N_y \times N_\xi}$ is given by $P := CG + H$. Note that Eq. (8) 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 (6) can be equivalently reformulated as a static team problem by applying the nonlinear change of variables in (7). This yields the following optimization problem:

$$\begin{aligned} & \text{minimize} \quad \mathbf{E} \left[x^\top R_x x + u^\top R_u u \right] \\ & \text{subject to} \quad \phi \in \Phi(\mathcal{G}_I), \quad s \in \mathcal{L}_m^2 \\ & \quad \left. \begin{aligned} F_x x + F_u u + F_\xi \xi + s &= 0 \\ x &= Bu + G\xi \\ \eta &= P\xi \\ u &= \phi(\eta) \\ s &\geq 0 \end{aligned} \right\} \quad \mathbf{P}\text{-a.s.} \end{aligned} \quad (9)$$

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 (9) is in the set of Youla parameters $\Phi(\mathcal{G}_I)$. In particular, problem (9) is a convex program if and only if the set $\Phi(\mathcal{G}_I)$ is convex.

B. Convexity under Partially Nested Information Structures

In what follows, we show that the static team problem (9) 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 [22].

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 $[C(t)A_{s+1}^t B(s)]_{kj} \neq 0$.

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 $t = 0, \dots, T-1$.

We denote by $\text{PN}(\Theta)$ the set of information graphs that are partially nested with respect to the system Θ . The information structure defined by $\mathcal{G}_I$ is said to be *nonclassical* if $\mathcal{G}_I \notin \text{PN}(\Theta)$. 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 [3], [22], [23], which applies to the setting in which controllers are subject to both delay and sparsity constraints on information sharing. Equipped with this definition, we state the following result, which shows 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 1, as it directly follows from existing arguments in [24, Thm. 1] and [25, Cor. 7].

Lemma 1. 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 1 implies Ho and Chu's classical result [3, 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 1 that the reformulated decentralized control problem in (9) is convex if and only if the underlying information structure is partially nested.²

IV. A CONVEX INFORMATION RELAXATION

In what follows, we consider systems with nonclassical information structures, and address the question as to how one might convexify the corresponding decentralized control design problems 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:

$$\underset{\mathcal{G} \supseteq \mathcal{G}_I}{\text{maximize}} \quad J^*(\mathcal{G}) \quad \text{subject to} \quad \mathcal{G} \in \text{PN}(\Theta). \quad (10)$$

Recall that $J^*(\mathcal{G})$ is the optimal value of the decentralized control design problem (6) given an information graph $\mathcal{G}$. Also, note that any feasible solution to problem (10) is required to 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 (10).

²We note that the convexity result in Lemma 1 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 (9) can be computed via the solution of a convex program when the information structure is nonclassical. See [11], [15], [26], [27] for recent advances.

Definition 3 (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 4 (Transitive Closure). The *transitive closure* of a directed graph $\mathcal{G} = (\mathcal{V}, \mathcal{E})$ is defined as the directed graph $\bar{\mathcal{G}} = (\mathcal{V}, \bar{\mathcal{E}})$, where $(i, j) \in \bar{\mathcal{E}}$ 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 [28]. Equipped with these definitions, we state the following result, which provides a 'closed-form' solution to problem (10).

Theorem 1 (Information Relaxation). An optimal solution to (10) is given by $\bar{\mathcal{G}}_P(\Theta, \mathcal{G}_I)$, the *transitive closure of the precedence graph*.

Theorem 1 implies the following lower bound on the optimal value of the original decentralized control problem (6):

$$J^*(\bar{\mathcal{G}}_P(\Theta, \mathcal{G}_I)) \leq J^*(\mathcal{G}_I). \quad (11)$$

Moreover, this lower bound can be computed via the solution of a convex infinite-dimensional optimization problem (9). In Theorem 2, we provide a finite-dimensional relaxation of problem (9) 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 $\bar{\mathcal{G}}_P(\Theta, \mathcal{G}_I)$ yields a partially nested information structure—a result that was originally shown in [12]. 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 1 (Tightness of the Relaxation). We also note that the information relaxation in Theorem 1 is *tight*. That is, $J^*(\bar{\mathcal{G}}_P(\Theta, \mathcal{G}_I)) = J^*(\mathcal{G}_I)$ 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 [13], [15], [26, Sec. 3.5] for a rigorous explication of such claims. It can also be shown that the lower bound (11) is achieved by nonclassical LQG control problems that satisfy the so-called *substitutability condition*. See [11, Sec. 3] for a formal proof of this claim.

In Lemma 2, we present an alternative characterization of partially nested information structures that will prove useful in the proof of Theorem 1.

Lemma 2. $\mathcal{G} \in \text{PN}(\Theta)$ if and only if $\mathcal{G} = \bar{\mathcal{G}}_P(\Theta, \mathcal{G})$.

The graph theoretic fixed-point condition in Lemma 2 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 2 is closely related to the graph theoretic necessary and sufficient condition for quadratic invariance presented in [29], 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 2: The proof of the “if” direction is straightforward, 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 $(j, i) \in \mathcal{E}_P(\Theta, \mathcal{G})$ implies that $(j, i) \in \mathcal{E}$. Note that $(j, i) \in \mathcal{E}_P(\Theta, \mathcal{G})$ 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 $v_1, \dots, v_m \in \mathcal{V}$ that satisfy $j = v_1 \prec v_2 \prec \dots \prec v_m = i$. Since $\mathcal{G} \in \text{PN}(\Theta)$, it also holds that $z_{v_1}(t) \subseteq z_{v_2}(t) \subseteq \dots \subseteq z_{v_m}(t)$ 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 3, 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 2 showing that any graph, which is feasible for problem (10), 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 2 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 $J^*(\overline{\mathcal{G}_P(\Theta, \mathcal{G}_I)}) \geq J^*(\mathcal{G})$ for every graph $\mathcal{G}$ that is feasible for problem (10). 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 (10). We previously showed in the proof of Lemma 2 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

$$\overline{\mathcal{G}_P(\Theta, \mathcal{G}_I)} = \mathcal{G}_P(\Theta, \overline{\mathcal{G}_P(\Theta, \mathcal{G}_I)}). \quad (12)$$

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 $(i, j) \in \mathcal{E}_P(\Theta, \mathcal{G}_I)$. It follows from (12) and Lemma 2 that $\overline{\mathcal{G}_P(\Theta, \mathcal{G}_I)} \in \text{PN}(\Theta)$, which completes the proof. ■

V. A DUAL APPROACH TO CONSTRAINT RELAXATION

The information relaxation developed in Section IV provides a convex programming relaxation of the original decentralized control design problem (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 [30]–[32] 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 (6).

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

A. Restriction to Affine Dual Control Policies

The derivation of our lower bound centers on a partial Lagrangian relaxation of problem (6). We do so by introducing a *dual control*

policy $v \in \mathcal{L}_m^2$, 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 (6):

$$\begin{aligned} & \text{minimize} \quad \sup_{v \in \mathcal{L}_m^2} \mathbf{E} \left[x^\top R_x x + u^\top R_u u \right. \\ & \quad \left. + v^\top (F_x x + F_u u + F_\xi \xi + s) \right] \\ & \text{subject to} \quad \gamma \in \Gamma(\mathcal{G}_I), \quad s \in \mathcal{L}_m^2 \\ & \quad \left. \begin{aligned} x &= Bu + G\xi \\ \eta &= P\xi \\ u &= \gamma(\eta) \\ s &\geq 0 \end{aligned} \right\} \quad \mathbf{P}\text{-a.s.} \end{aligned} \quad (13)$$

In presenting the equivalent min-max reformulation of problem (6), we have used the fact that problem (6) is equivalent to problem (9); and Lemma 1, 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 (13), we restrict ourselves to dual control policies that are affine in the disturbance trajectory, i.e., $v = V\xi$ for some $V \in \mathbf{R}^{m \times N_\xi}$. With this restriction, it is possible to derive a closed-form solution for the inner maximization in problem (13). This yields another minimization problem, whose optimal value stands as a lower bound on that of problem (13). We have the following result, which clarifies this claim.

Proposition 1. The optimal value of the following problem is a lower bound on the optimal value of problem (13):

$$\begin{aligned} & \text{minimize} \quad \sup_{V \in \mathbf{R}^{m \times N_\xi}} \mathbf{E} \left[x^\top R_x x + u^\top R_u u \right. \\ & \quad \left. + \xi^\top V^\top (F_x x + F_u u + F_\xi \xi + s) \right] \\ & \text{subject to} \quad \gamma \in \Gamma(\mathcal{G}_I), \quad s \in \mathcal{L}_m^2 \\ & \quad \left. \begin{aligned} x &= Bu + G\xi \\ \eta &= P\xi \\ u &= \gamma(\eta) \\ s &\geq 0 \end{aligned} \right\} \quad \mathbf{P}\text{-a.s.} \end{aligned} \quad (14)$$

Moreover, the optimal value of problem (14) equals that of the following optimization problem:

$$\begin{aligned} & \text{minimize} \quad \mathbf{E} \left[x^\top R_x x + u^\top R_u u \right] \\ & \text{subject to} \quad \gamma \in \Gamma(\mathcal{G}_I), \quad s \in \mathcal{L}_m^2 \\ & \quad \mathbf{E} \left[(F_x x + F_u u + F_\xi \xi + s) \xi^\top \right] = 0 \\ & \quad \left. \begin{aligned} x &= Bu + G\xi \\ \eta &= P\xi \\ u &= \gamma(\eta) \\ s &\geq 0 \end{aligned} \right\} \quad \mathbf{P}\text{-a.s.} \end{aligned} \quad (15)$$

Proof of Proposition 1: The fact that the optimal value of problem (14) lower bounds that of (13) is straightforward, since any dual affine control policy $v = V\xi$ is feasible for the inner maximization problem in (13). To see that the optimal values of problem (14) and (15) are equal, we note that

$$\begin{aligned} & \sup_{V \in \mathbf{R}^{m \times N_\xi}} \mathbf{E} \left[\xi^\top V^\top (F_x x + F_u u + F_\xi \xi + s) \right] \\ &= \sup_{V \in \mathbf{R}^{m \times N_\xi}} \mathbf{E} \left[\text{Tr} \left(V^\top (F_x x + F_u u + F_\xi \xi + s) \xi^\top \right) \right] \\ &= \begin{cases} 0, & \text{if } \mathbf{E} \left[(F_x x + F_u u + F_\xi \xi + s) \xi^\top \right] = 0, \\ +\infty, & \text{otherwise.} \end{cases} \end{aligned}$$

■

B. Relaxation to a Finite-dimensional Conic Program

Problem (15) 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 finite-dimensional conic program. Before proceeding, we present a formal definition of the subspace of causal affine controllers respecting the information structure defined by $\mathcal{G}_I$.

Definition 5. 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 (15) 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 [30], [32]. We provide a proof of Theorem 2 in Appendix I, which extends these techniques to accommodate the added complexity of decentralized information constraints on the controller.

Theorem 2. Let Assumption 2 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 (6):

$$\begin{aligned} & \text{minimize} \quad \text{Tr} \left(P^T Q^T R Q P M + 2G^T R_x B Q P M + G^T R_x G M \right) \\ & \text{subject to} \quad Q \in S(\mathcal{G}_I), \quad Z \in \mathbf{R}^{m \times N_\xi} \\ & \quad (F_u + F_x B) Q P + F_x G + F_\xi + Z = 0, \\ & \quad W M Z^T \succeq_{\mathcal{K}_2} 0, \\ & \quad e_1^T M Z^T \geq 0, \end{aligned} \quad (16)$$

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

Let $J^d(\mathcal{G}_I)$ denote the optimal value of the finite-dimensional conic program (16). 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 (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 (6).

APPENDIX I 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 (15). Consider the following result, which we prove in Appendix II.

Lemma 3. Let Assumption 2 hold. For each $s \in \mathcal{L}_m^2$, there exists a matrix $Z \in \mathbf{R}^{m \times N_\xi}$ that satisfies

$$Z M = \mathbf{E}[s \xi^T]. \quad (17)$$

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

$$Q P M = \mathbf{E}[u \xi^T], \quad (18)$$

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

With Lemma 3 in hand, we obtain an equivalent reformulation of problem (15) as the following optimization problem—via the introduction of the finite-dimensional decision variables Z and Q through the constraints (17) and (18), respectively.

$$\begin{aligned} & \text{minimize} \quad \mathbf{E}[u^T R u] + \text{Tr} \left(2G^T R_x B Q P M + G^T R_x G M \right) \\ & \text{subject to} \quad \gamma \in \Gamma(\mathcal{G}_I), \quad s \in \mathcal{L}_m^2, \quad Q \in S(\mathcal{G}_I), \quad Z \in \mathbf{R}^{m \times N_\xi} \\ & \quad (F_u + F_x B) Q P M + F_x G M + F_\xi M + Z M = 0 \\ & \quad Q P M = \mathbf{E}[u \xi^T] \\ & \quad Z M = \mathbf{E}[s \xi^T] \\ & \quad \left. \begin{aligned} \eta &= P \xi \\ u &= \gamma(\eta) \\ s &\geq 0 \end{aligned} \right\} \quad \mathbf{P}\text{-a.s.} \end{aligned} \quad (19)$$

where $R = R_u + B^T R_x B$.

We now introduce two technical Lemmas that permit us to construct a finite-dimensional relaxation of problem (19).

Lemma 4. 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:

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

We omit the proof of Lemma 4, as it is an immediate corollary of [32, Lem. 4.5]. A direct application of Lemma 4 yields the following equivalent reformulation of problem (19) as:

$$\begin{aligned} & \text{minimize} \quad \text{Tr} \left(P^T Q^T R Q P M + 2G^T R_x B Q P M + G^T R_x G M \right) \\ & \text{subject to} \quad s \in \mathcal{L}_m^2, \quad Q \in S(\mathcal{G}_I), \quad Z \in \mathbf{R}^{m \times N_\xi} \\ & \quad (F_u + F_x B) Q P + F_x G + F_\xi + Z = 0 \\ & \quad Z M = \mathbf{E}[s \xi^T] \\ & \quad s \geq 0 \quad \mathbf{P}\text{-a.s.} \end{aligned} \quad (20)$$

Note that, in reformulating problem (19), we have eliminated the second-order moment matrix M from the equality constraint $(F_u + F_x B) Q P M + F_x G M + F_\xi M + Z M = 0$, as M is assumed to be positive definite, and, therefore, invertible.

Lemma 5 provides a conic relaxation of the constraints in problem (20) involving the infinite-dimensional decision variable $s \in \mathcal{L}_m^2$. We provide a proof of this technical Lemma in Appendix III.

Lemma 5. If $s \in \mathcal{L}_m^2$ and $Z \in \mathbf{R}^{m \times N_\xi}$ satisfy $Z M = \mathbf{E}[s \xi^T]$ and $s \geq 0$ $\mathbf{P}$ -a.s., then $W M Z^T \succeq_{\mathcal{K}_2} 0$ and $e_1^T M Z^T \geq 0$.

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

$$(16) \underset{(a)}{\leq} (20) \underset{(b)}{=} (19) \underset{(c)}{=} (15) \underset{(d)}{\leq} (13) \underset{(e)}{=} (9) \underset{(f)}{=} (6)$$

Inequality (a) follows from Lemma 5, which implies that problem (16) is a relaxation of problem (20). Equality (b) follows from Lemma 4. Equality (c) follows from Lemma 3. Inequality (d) follows from Proposition 1. Finally, Equality (e) follows from Lemma 1, as the assumption of a partially nested information structure implies equivalence between the optimal values of problems (13) and (9). The equivalence between (9) and (6) is argued in Section III-A.

APPENDIX II PROOF OF LEMMA 3

This proof extends arguments originally developed in [32, 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 $s \in \mathcal{L}_m^2$. The matrix M is invertible, since it is assumed to be positive definite. Setting $Z = \mathbf{E}[s\xi^\top]M^{-1}$ yields the desired result in (17).

Proof of the second part: We first introduce the notion of a truncation operator. Given a nonempty set of indices $J \subseteq \{1, \dots, N_y\}$ and an arbitrary vector $x \in \mathbf{R}^{N_y}$, we define $x_J \in \mathbf{R}^{|J|}$ to be the subvector of x , whose entries are given by x_j for all $j \in J$. The entries of x_J are ordered in ascending order of their indices. For example, if $J = \{1, 3\}$, then $x_J = (x_1, x_3)$. 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. (18). The proof of Lemma 6 is in Appendix IV.

Lemma 6. Let Assumption 2 hold. Let $z \in \mathcal{L}_1^2$ be random variable that is a (possibly nonlinear) function of the random vector $\eta_J = \Pi_J \eta$, where $\{1\} \subseteq J \subseteq \{1, \dots, N_y\}$ is a given index set. Then, there exists another random variable $\tilde{z} \in \mathcal{L}_1^2$, which is an *affine*³ function of η_J , and satisfies $\mathbf{E}[\tilde{z}\eta^\top] = \mathbf{E}[z\eta^\top]$.

Stated in other words, Lemma 6 asserts the existence of a vector $q \in \mathbf{R}^{N_y}$ that satisfies

$$\mathbf{E}[q^\top \eta \eta^\top] = \mathbf{E}[z \eta^\top], \quad (21)$$

where the vector q respects the sparsity pattern encoded by the index set J , i.e., $q = \Pi_J^\top \Pi_J q$. It follows that one can apply Lemma 6 to each row of the matrix $\mathbf{E}[u\eta^\top]$ to establish the existence of a matrix $Q \in S(\mathcal{G}_I)$ that satisfies

$$\mathbf{E}[Q\eta\eta^\top] = \mathbf{E}[u\eta^\top]. \quad (22)$$

Consider a matrix $Q \in S(\mathcal{G}_I)$ that satisfies Eq. (22). We complete the proof by showing that this matrix also satisfies Eq. (18). First note that the combination of Assumption 2 and [19, Thm. 1] implies that the random vector $(\xi, \eta) = (\xi, P\xi)$ has an elliptically contoured distribution. Hence, it follows from [19, Cor. 5] that the conditional expectation of ξ 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

$$\mathbf{E}[\xi|\eta] = L\eta \quad \mathbf{P}\text{-a.s.} \quad (23)$$

for some matrix $L \in \mathbf{R}^{N_\xi \times N_y}$. It holds that

$$\mathbf{E}[u\xi^\top] = \mathbf{E}[\mathbf{E}[u\xi^\top|\eta]] = \mathbf{E}[u\eta^\top]L^\top = Q\mathbf{E}[\eta\eta^\top]L^\top.$$

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. (23); and the third equality follows from Lemma 6. It also holds that

$$\mathbf{E}[\eta\eta^\top]L^\top = \mathbf{E}[\eta\mathbf{E}[\xi^\top|\eta]] = \mathbf{E}[\eta\xi^\top] = \mathbf{E}[\eta\xi^\top] = PM,$$

which completes the proof.

³As 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.

APPENDIX III PROOF OF LEMMA 5

It follows from the symmetry of the matrix M that $MZ^\top = (ZM)^\top = \mathbf{E}[\xi s^\top]$. It, therefore, holds that

$$e_1^\top MZ^\top = e_1^\top \mathbf{E}[\xi s^\top] = \mathbf{E}[e_1^\top \xi s^\top] = \mathbf{E}[s^\top] \geq 0.$$

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

$$WE[s_i\xi] \succeq_{\mathcal{K}_2} 0, \quad \text{for } i = 1, \dots, m,$$

where $s_i \in \mathcal{L}_1^2$ is the i^{th} element of the random vector s . By definition, we have that $W\xi \succeq_{\mathcal{K}_2} 0$ for all $\xi \in \Xi$. Also, since $s_i \geq 0$ almost surely, we have that $W(s_i\xi) \succeq_{\mathcal{K}_2} 0$ almost surely. It follows from the convexity of the second-order cone that $WE[s_i\xi] \succeq_{\mathcal{K}_2} 0$.

APPENDIX IV PROOF OF LEMMA 6

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

$$r^\top := \mathbf{E}[z\eta_J^\top] (P_J M P_J^\top)^\dagger, \quad (24)$$

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

$$\mathbf{E}[z\eta_J^\top] = r^\top P_J M P_J^\top. \quad (25)$$

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

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

The second and the last equalities both follow from the fact that $\eta_J = \Pi_J P\xi = P_J\xi$. 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 [33, Prop. 3.2] that $\Psi^\top (\Psi \Psi^\top)^\dagger = \Psi^\dagger$. The sixth equality follows from the fact [33, Prop. 3.1] that $\Psi^\dagger \Psi \Psi^\top = \Psi^\top$. It follows that the vector r satisfies Eq. (25).

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

$$\mathbf{E}[\tilde{z}\eta^\top] = \mathbf{E}[z\eta^\top]. \quad (26)$$

First recall that the random vector (ξ, η) is shown to have an elliptically contoured distribution in Appendix II. Hence, it follows from [19, Cor. 5] 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 $L_J \in \mathbf{R}^{N_y \times |J|}$, such that

$$\mathbf{E}[\eta|\eta_J] = L_J \eta_J \quad \mathbf{P}\text{-a.s.} \quad (27)$$

It holds that

$$\mathbf{E}[z\eta^\top] = \mathbf{E}[\mathbf{E}[z\eta^\top|\eta_J]] = \mathbf{E}[z\eta_J^\top] L_J^\top = r^\top P_J M P_J^\top L_J^\top$$

Here, the first equality follows from the law of iterated expectations. The second equality follows from a combination of Eq. (27) and the assumption that z is a function of η_J . The third equality follows from Eq. (25). It also holds that

$$\begin{aligned} r^\top P_J M P_J^\top L_J^\top &= r^\top \mathbf{E}[\eta_J \eta_J^\top] L_J^\top = r^\top \mathbf{E}[\mathbf{E}[\eta_J \eta_J^\top | \eta_J]] \\ &= r^\top \mathbf{E}[\eta_J \eta_J^\top] = \mathbf{E}[\tilde{z} \eta^\top], \end{aligned}$$

which completes the proof.

APPENDIX V MATRIX DEFINITIONS

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

$$\begin{aligned} B &:= \begin{bmatrix} 0 & & & & \\ A_1^1 B(0) & 0 & & & \\ A_1^2 B(0) & A_2^2 B(1) & 0 & & \\ \vdots & & & \ddots & \\ \vdots & & & & 0 \\ A_1^T B(0) & A_2^T B(1) & \cdots & \cdots & A_T^T B(T-1) \end{bmatrix} \\ G &:= \begin{bmatrix} A_0^0 x(0) & & & & \\ A_0^1 x(0) & A_1^1 G(0) & & & \\ A_0^2 x(0) & A_1^2 G(0) & A_2^2 G(1) & & \\ \vdots & \vdots & & \ddots & \\ A_0^T x(0) & A_1^T G(0) & A_2^T G(1) & \cdots & A_T^T G(T-1) \end{bmatrix} \\ C &:= \begin{bmatrix} 0 & & & & \\ C(0) & 0 & & & \\ & \ddots & & \ddots & \\ & & C(T-1) & 0 \end{bmatrix} \\ H &:= \text{diag}(1, H(0), \dots, H(T-1)), \end{aligned}$$

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

REFERENCES

- [1] W. Lin and E. Bitar, "Performance bounds for robust decentralized control," in *American Control Conference (ACC)*, 2016. IEEE, 2016, pp. 4323–4330.
- [2] A. Mahajan, N. C. Martins, M. C. Rotkowitz, and S. Yüksel, "Information structures in optimal decentralized control," in *Decision and Control (CDC), 2012 IEEE 51st Annual Conference on*. IEEE, 2012, pp. 1291–1306.
- [3] Y. C. Ho and K. C. Chu, "Team decision theory and information structures in optimal control problems—Part I," *IEEE Transactions on Automatic Control*, vol. 17, no. 1, pp. 15–22, 1972.
- [4] M. Rotkowitz and S. Lall, "A characterization of convex problems in decentralized control," *IEEE Transactions on Automatic Control*, vol. 51, no. 2, pp. 274–286, 2006.
- [5] B. Bamieh and P. G. Voulgaris, "A convex characterization of distributed control problems in spatially invariant systems with communication constraints," *Systems & Control Letters*, vol. 54, no. 6, pp. 575–583, 2005.
- [6] H. S. Witsenhausen, "A counterexample in stochastic optimum control," *SIAM Journal on Control*, vol. 6, no. 1, pp. 131–147, 1968.
- [7] R. Bansal and T. Basar, "Stochastic teams with nonclassical information revisited: When is an affine law optimal?" in *American Control Conference, 1986*. IEEE, 1986, pp. 45–50.
- [8] P. Grover, S. Y. Park, and A. Sahai, "Approximately optimal solutions to the finite-dimensional Witsenhausen counterexample," *IEEE Transactions on Automatic Control*, vol. 58, no. 9, pp. 2189–2204, 2013.
- [9] A. Kulkarni, T. P. Coleman *et al.*, "An optimizer's approach to stochastic control problems with nonclassical information structures," *IEEE Transactions on Automatic Control*, vol. 60, no. 4, pp. 937–949, 2015.
- [10] S. T. Jose and A. 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*. IEEE, 2015, pp. 5743–5748.
- [11] S. M. Asghari and A. Nayyar, "Dynamic teams and decentralized control problems with substitutable actions," *IEEE Transactions on Automatic Control*, vol. 62, no. 10, pp. 5302–5309, 2017.
- [12] K. C. Chu, "Team decision theory and information structures in optimal control problems—Part II," *IEEE Transactions on Automatic Control*, vol. 17, no. 1, pp. 22–28, 1972.
- [13] A. Mahajan and S. Yüksel, "Measure and cost dependent properties of information structures," in *American Control Conference (ACC), 2010*. IEEE, 2010, pp. 6397–6402.
- [14] R. Srikant and T. Basar, "Asymptotic solutions to weakly coupled stochastic teams with nonclassical information," *IEEE transactions on Automatic Control*, vol. 37, no. 2, pp. 163–173, 1992.
- [15] S. Yüksel, "Stochastic nestedness and the belief sharing information pattern," *IEEE Transactions on Automatic Control*, vol. 54, no. 12, pp. 2773–2786, 2009.
- [16] M. C. Rotkowitz and N. C. Martins, "On the nearest quadratically invariant information constraint," *IEEE Transactions on Automatic Control*, vol. 57, no. 5, pp. 1314–1319, 2012.
- [17] A. Bemporad and D. Barcelli, "Decentralized model predictive control," in *Networked control systems*. Springer, 2010, pp. 149–178.
- [18] D. C. Youla, H. Jabr, J. J. Bongiorno Jr *et al.*, "Modern Wiener-Hopf design of optimal controllers—part II: The multivariable case," *IEEE Transactions on Automatic Control*, vol. 21, no. 3, pp. 319–338, 1976.
- [19] S. Cambanis, S. Huang, and G. Simons, "On the theory of elliptically contoured distributions," *Journal of Multivariate Analysis*, vol. 11, no. 3, pp. 368–385, 1981.
- [20] N. R. Sandell Jr, P. Varaiya, M. Athans, and M. G. Safonov, "Survey of decentralized control methods for large scale systems," *IEEE Transactions on Automatic Control*, vol. 23, no. 2, pp. 108–128, 1978.
- [21] J. N. Tsitsiklis and M. Athans, "On the complexity of decentralized decision making and detection problems," *IEEE Transactions on Automatic Control*, vol. 30, no. 5, pp. 440–446, 1985.
- [22] Y. C. Ho and K. C. Chu, "Information structure in dynamic multi-person control problems," *Automatica*, vol. 10, no. 4, pp. 341–351, 1974.
- [23] A. Gattami, "Optimal decisions with limited information," Ph.D. dissertation, Lund University, 2007.
- [24] M. Rotkowitz, "On information structures, convexity, and linear optimality," in *Decision and Control, 2008. CDC 2008. 47th IEEE Conference on*. IEEE, 2008, pp. 1642–1647.
- [25] L. Lessard and S. Lall, "Quadratic invariance is necessary and sufficient for convexity," in *American Control Conference (ACC), 2011*. IEEE, 2011, pp. 5360–5362.
- [26] S. Yüksel and T. Basar, *Stochastic networked control systems: Stabilization and optimization under information constraints*. Springer Science & Business Media, 2013.
- [27] S. Yüksel and N. Saldi, "Convex analysis in decentralized stochastic control, strategic measures, and optimal solutions," *SIAM Journal on Control and Optimization*, vol. 55, no. 1, pp. 1–28, 2017.
- [28] S. Warshall, "A theorem on boolean matrices," *Journal of the ACM (JACM)*, vol. 9, no. 1, pp. 11–12, 1962.
- [29] J. Swigart and S. 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*. IEEE, 2009, pp. 5409–5414.
- [30] D. Kuhn, W. Wiesemann, and A. Georgiou, "Primal and dual linear decision rules in stochastic and robust optimization," *Mathematical Programming*, vol. 130, no. 1, pp. 177–209, 2011.
- [31] A. Georgiou, W. Wiesemann, and D. Kuhn, "Generalized decision rule approximations for stochastic programming via liftings," *Mathematical Programming*, vol. 152, no. 1–2, pp. 301–338, 2015.
- [32] M. J. Hadjiyiannis, P. J. Goulart, and D. Kuhn, "An efficient method to estimate the suboptimality of affine controllers," *IEEE Transactions on Automatic Control*, vol. 56, no. 12, pp. 2841–2853, 2011.
- [33] J. C. A. Barata and M. S. Hussein, "The Moore–Penrose pseudoinverse: A tutorial review of the theory," *Brazilian Journal of Physics*, vol. 42, no. 1–2, pp. 146–165, 2012.