Information Structures in Optimal Decentralized Control

Aditya Mahajan, Nuno C. Martins, Michael C. Rotkowitz and Serdar Yüksel

Abstract— This tutorial paper provides a comprehensive characterization of information structures in team decision problems and their impact on the tractability of team optimization. Solution methods for team decision problems are presented in various settings where the discussion is structured in two foci: The first is centered on solution methods for stochastic teams admitting state-space formulations. The second focus is on norm-optimal control for linear plants under information constraints.

I. INTRODUCTION

A decentralized control system has structural properties that may restrict the storage, processing and the dissemination of information in the feedback loop. An example would be a large-scale power-grid with multiple coupled generators and consumers. Here, the decision center at each generator has access to only local measurements and based on these it must regulate variables that may impact the entire grid.

Decentralization may be imposed as a way to achieve scalability, by constraining the connectivity of the underlying communication network and the computational complexity of the control algorithms. New methods for the analysis of decentralized systems have also been proposed that focus on the detrimental effects that real communication networks may introduce, such as bit-rate limits [1] and packet losses [2].

As illustrated in [3], [4], the information structure may determine the tractability of optimal decentralized control problems, and research on this topic remains quite active ever since Witsenhausen [5] and Ho [6] considered a considerable investigation of information structures in the context of team decision theory.

A. Paper organization

This paper focuses on the impact of information structures on the tractability of computing optimal decentralized controllers. It is centered on two subclasses of problems. The

first pertains to the design of optimal team decision systems specified using the state-space formalism, while the second discusses the design of optimal policies when the system and controller are represented in linear operatorial form and decentralization is specified via sparsity constraints.

The paper has the following three main Sections:

- **Section II:** Introduces the main concepts and definitions that allow the classification of information structures.
- **Section III:** Focuses on optimal team decision problems in state-space form.
- **Section IV:** Discusses the design of optimal controllers when the decentralization constraint is imposed via subspace constraints.

II. BASIC CONCEPTS AND DEFINITIONS

A decentralized control system may either be sequential or non-sequential. In a sequential system, the decision makers (DMs) act according to an order that is specified before the system starts running; while in a non-sequential system the DMs act in an order that depends on the realization of the system uncertainty and the control actions of other DMs. It is much more difficult to analyze and formulate a well-posed optimal control problem for non-sequential systems because we need to ensure that it is causal and deadlock free. Here, we restrict attention to sequential systems. We start with an abstract model of decentralized control systems, called the *intrinsic model*, that was proposed by Witsenhausen [7]. We present a version of the intrinsic model that is restricted to sequential systems. Later we define state-space and input-output models that are also restricted to sequential systems. We refer the reader to [7]–[11] for more details on non-sequential systems.

According to the instrinsic model, any (finite horizon) sequentail team problem can be characterized by a tuple $\big((\Omega, \mathcal{F}), N, \{(\mathbb{U}^i, \mathcal{U}^i), i = 1, \ldots, N\}, \{\mathcal{J}^i, i = 1, \ldots, N\}\big)$ or equivalently by a tuple $\big((\Omega, \mathcal{F}), N, \{(\mathbb{U}^i, \mathcal{U}^i), i = 1, \ldots, N\}, \{(\mathbb{I}^i, \mathcal{I}^i), i = 1, \ldots, N\}\big)$ where

- $(\Omega, \mathcal{F})$ is a measurable space representing all the uncertainty in the system. The realization of this uncertainty is called the *primitive variable* of the system. Ω denotes all possible realizations of the primitive random variable and $\mathcal{F}$ is a sigma-algebra over Ω.
- N denotes the number of decision makers (DMs) in the system. Each DM takes only one action. If the system has a control station that takes multiple actions over time, it is modeled as a collection of DMs, one for each time instant.
- $\{(\mathbb{U}^i, \mathcal{U}^i), i = 1, \ldots, N\}$ is a collection of measurable spaces representing the action space for each DM. The

control action u^i of DM i takes value in $\mathbb{U}^i$ and $\mathcal{U}^i$ is a sigma-algebra over $\mathbb{U}^i$.

- $\{\mathcal{J}^i, i = 1, \ldots, N\}$ is a collection of sets in $\mathcal{F}$ and represents the *information* available to a DM to take an action. Sometimes it is useful to assume that the information is available in terms of an explicit observation that takes values in a measurable space $(\mathbb{I}^i, \mathcal{I}^i)$. Such an observation is generated by a measurable observation function from $\Omega \times \mathbb{U}^1 \times \cdots \times \mathbb{U}^{i-1}$ to $\mathbb{I}^i$. The collection $\{\mathcal{J}^i, i = 1, \ldots, N\}$ or $\{(\mathbb{I}^i, \mathcal{I}^i), i = 1, \ldots, N\}$ is called the *information structure* of the system.
- A *control strategy* (also called a *control policy* or *design*) of a decentralized control system is given by a collection $\{\gamma^i, i = 1, \ldots, N\}$ of functions where $\gamma^i : (\mathbb{I}^i, \mathcal{I}^i) \to (\mathbb{U}^i, \mathcal{U}^i)$ (or equivalently, $\gamma^i : (\Omega, \mathcal{J}^i) \to (\mathbb{U}^i, \mathcal{U}^i)$).

Although, there are different ways to define a control objective of a decentralized system, we focus on minimizing a loss function. Other performance measures include minimizing regret, minimizing risk, ensuring safety, and ensuring stability. We will assume that we are given a probability measure P on $(\Omega, \mathbb{F})$ and a real-valued loss function ℓ on $(\Omega \times \mathbb{U}^1 \times \cdots \times \mathbb{U}^N, \mathcal{F} \otimes \mathcal{U}^1 \otimes \cdots \otimes \mathcal{U}^N) =: (\mathbb{H}, \mathcal{H})$. Any choice $\gamma = (\gamma^1, \ldots, \gamma^n)$ of the control strategy induces a probability measure P^γ on $(\mathbb{H}, \mathcal{H})$. We define the performance $J(\gamma)$ of a strategy as the expected loss (under probability measure P^γ), i.e.,

$$J(\gamma) = \mathcal{E}^\gamma[\ell(\omega, u^1, \ldots, u^n)]$$

where ω is the primitive variable (or the primitive *random* variable, since a measure is specified) and u^i is the control action of DM i.

As an example, consider the following model of a system with two decision makers which is taken from [12]. Let $\Omega = \{\omega_1, \omega_2, \omega_3\}$, $\mathcal{F}$ be the power set of Ω. Let the action space be $\mathbb{U}^1 = \{U(\text{up}), D(\text{down})\}$, $\mathbb{U}^2 = \{L(\text{left}), R(\text{right})\}$, and $\mathcal{U}^1$ and $\mathcal{U}^2$ be the power sets of $\mathbb{U}^1$ and U^2 respectively. Let the information fields $\mathcal{J}^1 = \{\emptyset, \{\omega_1\}, \{\omega_2, \omega_3\}, \Omega\}$ and $\mathcal{J}^2 = \{\emptyset, \{\omega_1, \omega_2\}, \{\omega_3\}, \Omega\}$. (This information corresponds to the non-identical imperfect (quantized) measurement setting considered in [12]).

Suppose the probability measure P is given by $P(\omega_i) = p_i, i = 1, 2, 3$ and $p_1 = p_2 = 0.3$, $p_3 = 0.4$, and the loss function $\ell(\omega, u^1, u^2)$ is given by

		u^2	
u^1		L	R
	U	1	0
	D	3	1

$\omega: \quad \omega_1 \leftrightarrow 0.3$

		u^2	
u^1		L	R
	U	2	3
	D	2	1

$\omega_2 \leftrightarrow 0.3$

		u^2	
u^1		L	R
	U	1	2
	D	0	2

$\omega_3 \leftrightarrow 0.4$

For the above model, the unique optimal control strategy is given by

$$\gamma^{1,*}(y^1) = \begin{cases} U, & y^1 \in \{\omega_1\} \\ D, & \text{else} \end{cases}$$

$$\gamma^{2,*}(y^2) = \begin{cases} R, & y^2 \in \{\omega_1, \omega_2\} \\ L, & \text{else} \end{cases}$$

A solution to the generalized sequential decentralized stochastic control problem is very difficult. Most of the work in the literature has concentrated on identifying solution techniques for specific subclasses. Typically, these subclasses are characterized on the basis of the information structure of the system. We describe the most common classification below.

A. Static and dynamic information structures

The simplest, and at first glance, the most critical, distinction is between static and dynamic information structures. An information structure is called *static* if $\mathcal{J}^i$ is $\mathcal{F}$ measurable for all i, i.e., the observation of all DMs depends only on the primitive random variable (and not on the control actions of others). Systems that don't have static information structure are said to have *dynamic* information structure. In such systems, some DMs influence the observations of others through their actions. The authors of [13], [14] studied the general properties of static information structures. One of their strongest results is that when the primitive random variable is Gaussian, observation functions are linear, and the loss function is quadratic—i.e., the static LQG case—affine control strategies are optimal. The result relies on the convexity of the problem, and is explained later in Section III.

Witsenhausen [15] showed that *any* dynamic decentralized control system can be converted to a static decentralized control system by an appropriate change of measures. However, very little is known regarding the solution of a non-LQG static system; hence, the above transformation is not practically useful.

B. Classical, quasiclassical and nonclassical information structures

Centralized control systems are a special case of decentralized control systems; their characterizing feature is centralization of information, i.e., any DM knows the information available to *all* the DMs that acted before it, or formally, $\mathcal{J}^i \subseteq \mathcal{J}^{i+1}$ for all i. Such information structures are called *classical*.

A decentralized system is called *quasiclassical* or *partially nested* if the following condition holds: whenever DM i can influence DM j, then DM j must know the observations of DM i, or more formally, $\mathcal{J}^i \subseteq \mathcal{J}^j$. Ho and Chu [16] showed that when the primitive random variable is Gaussian, observation functions are linear, and the loss function is quadratic—i.e., the partially nested LQG case—affine control laws are optimal. The result relies on showing that an invertible linear transformation can convert the partially nested LQG system into a static LQG system and then use the results of the static LQG case. Not much is known regarding other subclasses (non-LQG systems) of quasiclassical or partially nested teams.

Any information structure that is not classical or quasiclassical is called *nonclassical*.

Since classical information structures are effectively equivalent to centralized control systems, they can be solved using

techniques from centralized stochastic control. For example, for state space models, we can use dynamic programming to find optimal strategies (see Sec. III.D); for input-output models, we can use convex parametrizations to find optimal control strategies (see Sec. IV). In general, these techniques do not work for nonclassical information structures. Many researchers have investigated specific sub-classes of nonclassical information structures, both for the state-space models and for the input-output models. We present some of these results for these models in the Sec. IV. Further results and examples on information structures are available in [12].

III. STOCHASTIC TEAM PROBLEMS IN STATE-SPACE FORM

In a state space model, we assume that the decentralized control system has a *state* x_t that is evolving with time. The evolution of the state is controlled by the actions of the control stations. We assume that the system has N control stations where each control station i chooses a control action u_t^i at time t. The system runs in discrete time, either for finite or infinite horizon.

Let $\mathbb{X}$ denote the space of realizations of the state x_t, and $\mathbb{U}^i$ denote the space of realization of control actions u_t^i. Let $\mathcal{T}$ denote the set of time for which the system runs.

The initial state x_1 is a random variable and the state of the system evolves as

$$x_{t+1} = f_t(x_t, u_t^1, \ldots, u_t^N; w_t^0), \quad t \in \mathcal{T}, \qquad (1)$$

where $\{w_t^0, t \in \mathcal{T}\}$ is an independent noise process that is also independent of x_1.

We assume that each control station i observes the following at time t

$$y_t^i = g_t^i(x_t, w_t^i), \qquad (2)$$

where $\{w_t^i, t \in \mathcal{T}\}$ are measurement noise processes that are independent across time, independent of each other, and independent of $\{w_t^0, t \in \mathcal{T}\}$ and x_1.

The above evolution does not completely describe the dynamic control system, because we have not specified the data available at each control station. In general, the data I_t^i available at control station i at time t will be a function of all the past system variables $\{x_{[1,t]}, \mathbf{y}_{[1,t]}, \mathbf{u}_{[1,t-1]}, \mathbf{w}_{[1,t]}\}$, i.e.,

$$I_t^i = \eta_t^i(x_{[1,t]}, \mathbf{y}_{[1,t]}, \mathbf{u}_{[1,t-1]}, \mathbf{w}_{[1,t]}), \qquad (3)$$

where we use the notation $\mathbf{u} = \{u^1, \ldots, u^N\}$ and $x_{[1,t]} = \{x_1, \ldots, x_t\}$. The collection $\{I_t^i, i = 1, \ldots, N, t \in \mathcal{T}\}$ comprises the observation variables which generate the information structure in the system.

When $\mathcal{T}$ is finite, say equal to $\{1, \ldots, T\}$, the above model is a special case of the sequential intrinsic model presented above. The set $\{x_1, w_t^0, w_t^1, \ldots, w_t^N, t \in \mathcal{T}\}$ denotes the primitive random variable with probability measure given by the product measure of the marginal probabilities; the system has $N \times T$ DMs, one for each control station at each time. DM (i, t) observes I_t^i and chooses u_t^i. The information sub-fields $\mathcal{J}^k$ are determined by $\{\eta_t^i, i = 1, \ldots, N, t \in \mathcal{T}\}$.

Some important information structures are

1) *Complete information sharing:* In complete information sharing, each DM has access to present and past measurements and past actions of all DMs. Such a system is equivalent to a centralized system.

$$I_t^i = \{\mathbf{y}_{[1,t]}, \mathbf{u}_{[1,t-1]}\}, t \in \mathcal{T}.$$

2) *Complete measurement sharing:* In complete measurement sharing, each DM has access to the present and past measurements of all DMs. Note that past control actions are not shared.

$$I_t^i = \{y_{[1,t]}\}, \quad t \in \mathcal{T}.$$

3) *Delayed information sharing:* In delayed information sharing, each DM has access to n-step delayed measurements and control actions of all DMs.

$$I_t^i = \begin{cases} \{y_{[t-n+1,t]}^i, u_{[t-n+1,t-1]}^i \mathbf{y}_{[1,t-n]}, \mathbf{u}_{[1,t-n]}\}, & t > n \\ \{y_{[1,t]}^i, u_{[1,t-1]}^i\}, & t \leq n \end{cases}$$

$$(4)$$

4) *Delayed measurement sharing:* In delayed measurement sharing, each DM has access to n-step delayed measurements of all DMs. Note that control actions are not shared.

$$I_t^i = \begin{cases} \{y_{[t-n+1,t]}^i, u_{[1,t-1]}^i, \mathbf{y}_{[1,t-n]}\}, & t > n \\ \{y_{[1,t]}^i, u_{[1,t-1]}^i\}, & t \leq n \end{cases}$$

5) *Delayed control sharing*: In delayed control sharing, each DM has access to n-step delayed control actions of all DMs. Note that measurements are not shared.

$$I_t^i = \begin{cases} \{y_{[1,t]}^i, u_{[t-n+1,t-1]}^i, \mathbf{u}_{[1,t-n]}\}, & t > n \\ \{y_{[1,t]}^i, u_{[1,t-1]}^i\}, & t \leq n \end{cases}$$

6) *Periodic information sharing:* In periodic information sharing, the DMs share their measurements and control periodically after every k time steps. No information is shared at other time instants.

$$I_t^i = \begin{cases} \{y_{[\lfloor t/k \rfloor k, t]}^i, u_{[\lfloor t/k \rfloor k, t]}^i, \mathbf{y}_{[1, \lfloor t/k \rfloor k]}, \mathbf{u}_{[1, \lfloor t/k \rfloor k]}\}, \\ \qquad\qquad\qquad\qquad\qquad\qquad t \geq k \\ \{y_{[1,t]}^i, u_{[1,t-1]}^i\}, \qquad\qquad\qquad t < k \end{cases}$$

7) *Completely decentralized information:* In a completely decentralized system, no data is shared between the DMs.

$$I_t^i = \{y_{[1,t]}^i, u_{[1,t-1]}^i\}, \quad t \in \mathcal{T}.$$

In all the information structures given above, each DM has perfect recall (PR), that is, each DM has full memory of its past information. In general, a DM need not have perfect recall. For example, a DM may only have access to its current observation, in which case the information structure is

$$I_t^i = \{y_t^i\}, \quad t \in \mathcal{T}. \qquad (5)$$

To complete the description of the team problem, we have to specify the loss function. We will assume that the loss function is of an additive form:

$$\ell(x_{[1,T]}, \mathbf{u}_{[1,T]}) = \sum_{t \in \mathcal{T}} c(x_t, \mathbf{u}_t) \qquad (6)$$

where each term in the summation is known as the *incremental* (or *stagewise*) *loss*.

The objective is to choose control laws γ_t^i such that $u_t^i = \gamma_t^i(I_t^i)$ so as to minimize the expected loss (6). In the sequel, we will denote the set of all measurable control laws γ_t^i under the given information structure by Γ_t^i.

A. Solutions to Static Teams

Let $J(\gamma) := \mathcal{E}[c(\omega, \gamma^1(\eta^1(\omega)), \ldots, \gamma^N(\eta^N(\omega)))]$. We say that a policy γ^* is person by person optimal if

$$J(\gamma^*) \le J(\gamma^{*1}, \ldots, \gamma^{*(k-1)}, \beta, \gamma^{*(k+1)}, \ldots),$$
$$\beta \in \Gamma^k, \quad k = 1, 2, \ldots, N. \quad (7)$$

A policy γ^* is optimal if

$$J(\gamma^*) \le J(\gamma), \quad \text{for all } \gamma \in \Gamma.$$

It has been observed by Radner [13] and Krainak et al [17] that a static team problem with a loss (or cost) function $c(x, \cdot)$ which is (i) continuously differentiable in the actions and (ii) strictly convex in the actions, admits an optimal policy that satisfies local stationarity conditions. An important application of the above result is the following static Linear Quadratic Gaussian Problem: Consider a two-controller system evolving in $\mathbb{R}^n$ with the following description: Let x_1 be Gaussian and $x_2 = Ax_1 + B^1 u_1^1 + B^2 u_1^2 + w_1$

$$y_1^1 = C^1 x_1 + v_1^1,$$
$$y_1^2 = C^2 x_1 + v_1^2,$$

with w, v^1, v^2 zero-mean, i.i.d. disturbances. For $\rho_1, \rho_2 > 0$, let the goal be the minimization of

$$J(\gamma^1, \gamma^2) = \mathcal{E}\left[||x_1||_2^2 + \rho_1 ||u_1^1||_2^2 + \rho_2 ||u_1^2||_2^2 + ||x_2||_2^2 \right] \quad (8)$$

over the control policies of the form:

$$u_t^i = \mu_t^i(y_1^i), \quad i = 1, 2.$$

For this problem, the cost function is convex in the actions of the decision makers, and moreover, it is continuously differentiable. Furthermore, linear policies are person by person optimal since linear policies adopted by the other decision makers reduce the problem to a standard Linear Quadratic Gaussian cost optimization problem with partial, Gaussian observations. Hence, the solution to this problem is affine. This remarkable observation allows one to show that optimal team policies are affine.

The above analysis is not applicable to stochastic dynamic team problems with nonclassical information as we will see in the next subsection. As discussed above, nonclassical information structure (*IS*) arises if a Decision Maker (DM) i's action affects the information available to another DM j, who however does not have access to the information available to DM i based on which her action was constructed. Another way, perhaps mathematically more precise way of stating this is that the information sigma field of agent DM j is dependent explicitly on the policy (decision rule, or control law) of DM i.

In team problems with partially nested information, one talks about *precedence relationships* among agents: an agent **A**i *preceeds* another agent **A**j (or **A**i *communicates* to **A**j), if the former agent's actions affect the information of the latter, in which case (to be partially nested) **A**j has to have the information based on which the action-generating policy of **A**i was constructed. Under quasi-classical information, LQG stochastic team problems are tractable by conversion into an equivalent static team problem, of the type discussed. The team-optimal solution under this new (static) *IS* can then be expressed in terms of the original *IS*. Examples of such an indirect derivation for dynamic teams with quasi-classical information are given in several papers, such as [16], [18], [19]), see also [12].

B. Signaling and its effect on lack of convexity

What makes a large number of problems possessing the nonclassical information structure difficult is the fact that signaling is present: Signaling is the policy of communication through control actions. Under signaling, the decision makers apply their actions to affect the information available at the other decision makers. In this case, the control policies induce a probabilistic map (hence, a channel or a stochastic kernel) from the exogenous random variable space to the observation space of the signaled decision makers. For the nonclassical case, the problem thus also features an information transmission aspect, and the signaling decision maker's objective also includes the design of an optimal measurement channel. In [20], it has been established that, an optimal control problem is concave on the space of information structures, viewed as stochastic kernels (see also [12]). Hence, convexity hardly holds when there is a signaling incentive.

To make this important issue more explicit, let us consider the following example from [21]. Consider a two-controller system evolving in $\mathbb{R}^n$:

$$x_{t+1} = Ax_t + B^1 u_t^1 + B^2 u_t^2 + w_t,$$
$$y_t^1 = C^1 x_t + v_t^1,$$
$$y_t^2 = C^2 x_t + v_t^2,$$

where w, v^1, v^2 are zero-mean, i.i.d. disturbances, and A, B^1, B^2, C^1, C^2 matrices of appropriate dimensions. For $\rho_1, \rho_2 > 0$, let the objective be the minimization of the cost functional be a generalization of (8)

$$J = \mathcal{E}\left[\left(\sum_{t=1}^{T} |x_t|^2 + \rho_1 |u_t^1|^2 + \rho_2 |u_t^2|^2 \right) + ||x_T||^2 \right]$$

over control policies of the form:

$$u_t^i = \mu_t^i(y_{[0,t]}^i, u_{[0,t-1]}^i), \quad i = 1, 2; \quad t = 0, 1, \ldots, T - 1.$$

For a multi-stage problem (say with $T = 2$), unlike $T = 1$ in (8), the cost is in general no-longer *convex* in the action variables of the controllers acting in the first stage $t = 0$. This is because these actions might affect the estimation quality of the other controller in the future stages, if one DM can signal information to the other DM in one stage. We note that this

condition is equivalent to $C^1 A^l B^2 \neq 0$ or $C^2 A^l B^1 \neq 0$ with $l+1$ denoting the delay in signaling with $l = 0$ in the problem considered. In particular, if the controller is allowed to apply a randomized policy (for example by possibly using private random information that it has from the past realizations), this induces a conditional probability measure (channel) from the external variables and the initial state of the system to the observation variables at the other decision maker. The optimization problem, as such, is not jointly convex in such policies, and as such finding a fixed point to the stationarity conditions in the optimal policies does not necessarily lead to the conclusion that such policies are optimal. Even when restricted to linear policies, the problem is not convex in general.

C. LQG Problems Exhibiting Non-Classical Information Structure

In the following, we discuss two important LQG type problems with nonclassical information.

1) Transmission of Gaussian Source over a Gaussian Channel under Various Information Structures: Consider the transmission of a Gaussian source over a Gaussian relay channel, as depicted in Figure 1.
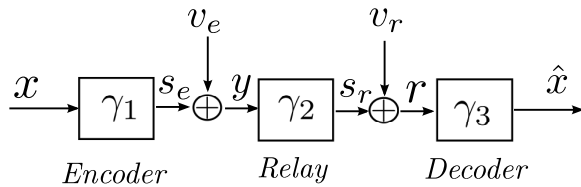

Fig. 1. Gaussian relay channel.

We wish to minimize $\mathcal{E}[(x - \hat{x})^2]$ over encoder and relay encoder policies. We assume that x is Gaussian with zero mean and variance σ_x^2. The encoder mapping satisfies, $s_e = \gamma^1(x)$ such that,

$$\mathcal{E}[s_e^2] \leq P_S. \tag{9}$$

The transmitted signal s_e is then observed in noise by the relay node as $y = s_e + v_e$, where v_e is a zero-mean independent Gaussian noise of variance N_e. The relay node applies a measurable mapping γ^2 on the received signal to produce s_r under the following average relay power constraint,

$$\mathcal{E}[s_r^2] \leq P_R. \tag{10}$$

The signal s_r is then transmitted over a Gaussian channel. Accordingly the destination node receives $z = s_r + v_r$, where $\{v_r\}$ is zero mean white Gaussian noise with variance N_r. The decoder generates $\hat{x} = g(z)$.

Some facts for this problem are as follows.

- If the relay is restricted to be linear, the optimal encoder is linear from information theoretic arguments see [12]
- If the encoder is restricted to be linear, the best relay is linear [22].
- The problem is non-convex when the encoders are viewed as stochastic kernels [20]. Hence, person by

person-optimality above does not imply optimality of linear policies. For further discussion, see [23] and [12]. This implies that the person-by-person optimal encoding policies do not guarantee team optimality. Even under linear policies, the problem is not convex (see [24]).

- Linear policies are not globally optimal [23]. If there are more that two relays, linear policies are not optimal [25]. Both references mentioned have obtained counterexamples.

2) Witsenhausen's Counterexample and the Generalized Gaussian test channel: Consider the following two-stage stochastic control problem with non-classical information structure:

$$x_1 = x_0 + u_1, \qquad x_2 = x_1 - u_2$$

and with the measurement equations

$$y_1 = x_0, \qquad y_2 = x_1 + v$$

and control policies

$$u_1 = \gamma_1(y_1), \qquad u_2 = \gamma_2(y_2),$$

with cost function now written as

$$Q(\gamma_1, \gamma_2) = \mathcal{E}[k^2 u_1^2 + x_2^2].$$

where k is a positive real parameter, while x_0 and v are zero mean independent Gaussian random variables with variance σ^2 and 1, respectively. This is Witsenhausen's counterexample [4]. For this problem, Witsenhausen established that a solution exists (we note that Wu and Verdu provided an alternative proof using tools from Transportation theory [26]), and established that an optimal policy is non-linear.

Under a different cost selection, the formulation above can also be used to cast the well known Gaussian test channel, which admits an optimal affine solution. An interesting study on the optimality of affine policies for different quadratic costs is provided in [27], and a concrete example is solved in [12].

Remark 3.1: As mentioned earlier, [15] showed that it is possible to transform any dynamic team problem into a static team problem. The static reduction of the Witsenhausen's counterexample is a two controller static team where the observations y^1 and y^2 of the two controllers are *independent* zero-mean Gaussian random variables with variance σ^2 and 1, respectively. See [15] for the exact correspondence between the static formulation and the original example.

Remark 3.2: In a class of quadratic cost minimization problems, one poses the problem not as an expectation minimization but as a min-max optimization where nature acts as the maximizer and the controllers act as the minimizers for cost functions of the form

$$\inf_{\{\gamma^i\}} \sup_{\omega} J(\gamma, \omega),$$

with J being a quadratic function, γ^i denoting controller policies and ω a disturbance with norm constraints. Linear policies are optimal for a large class of such settings in both encoder-decoder design as well as controller design (see

different setups in for example [28], [29], [30], [31], [32]). The proof of such results typically use the fact that, such min-max problems can be converted to a quadratic optimization problem by a re-parametrization of the cost function. More on this connection will be presented in the next section.

Remark 3.3: There is a large class of problems where signaling is *performance-irrelevant*, see [21], [33], [12] for a class of problems where partial nestedness does not hold, yet, optimal solutions are linear. Hence, one can look beyond partial nestedness and exploit the *measure dependency of information structures* to look for settings where optimal policies may be linear and more importantly, the optimization problem may be convex.

D. Dynamic Programming Approach to Team Problems and Limited Information Sharing

For centralized stochastic control, dynamic programming (and its specialization for LQG systems—Ricatti equations) provide a systematic methodology to synthesize optimal controllers. Thus, a natural queston is whether the dynamic programming principle extends to decentralized stochastic control as well.

In order to attempt to answer the above question, let us first briefly describe how dynamic programming works in centralized systems. Consider the model of Section III for $N = 1$ control stations and $I_t = (y_{[1,t]}, u_{[1,t-1]})$. Notice that this system has classical information structure. The main conceptual difficulties in the optimal design of such a system are:

D1) The domain I_t of the control law γ_t is increasing with time; consequently, optimal control laws are harder to search for and harder to implement as time increases.

D2) The optimal control problem is a functional optimization problem where we need to find a control policy $(\gamma_t, t \in \mathcal{T})$ to minimize the expected total loss.

The theory of centralized stochastic control overcomes these difficulties by identifying an *information state* π_t at each time with the following properties:

P1) π_t is a function of the information I_t;

P2) π_{t+1} is a function of π_t and new information (u_t, y_{t+1});

P3) π_t is a sufficient statistic for predicting the future observations, i.e., $\mathcal{P}(y_{t+1}|I_t) = \mathcal{P}(y_{t+1}|\pi_t)$.

P4) π_t is a sufficient statistic for performance evaluation, i.e., $\mathcal{E}[c(x_t, u_t)|I_t] = \mathcal{E}[c(x_t, u_t)|\pi_t]$.

For any information state satisfying the above properties, the centralized stochastic control theory provides the following results:

1) **Structure of optimal controllers.** Restricting attention to control laws of the form $u_t = \gamma_t(\pi_t)$ does not entail any loss of optimality.

2) **Dynamic programming decomposition.** Recursively define:[1]

$$V_t(\pi_t) = \inf_{u_t \in \mathbb{U}_t} \mathcal{E}[c(x_t, u_t) + V_{t+1}(\pi_{t+1}) \mid \pi_t, u_t]. \quad (11)$$

[1] When $\mathcal{T}$ is finite, initialize $V_{|\mathcal{T}|+1}(\cdot) = 0$; otherwise $V_t(\cdot) = V_{t+1}(\cdot)$ and (11) reduces to a fixed-point equation.

Then, if the infimum above is achieved, the $\arg\min$ at time t gives the optimal control action when the information state is π_t.

The above results hold for any choice of information state. In general, a system may have more than one information state, and the "best" choice of information state is model dependent. Some examples for the choice of information state are:

1) *Markov decision process (MDP).* When the controller observes the state x_t of the system perfectly, i.e., when $y_t = x_t$, then $\pi_t = x_t$ is an information state.

2) *Partially observable Markov decision process (POMDP).* For the general centralized control system described above, the belief state $\pi_t(\cdot) = \mathcal{P}(x_t = \cdot|I_t)$ is an information state.

An appropriate information state overcomes the conceptual difficulties (D1) and (D2) described above. Let $\mathbb{B}_t$ denote the space of all possible values of π_t. If $\mathbb{B}_t$ is time-invariant, then difficulty (D1) is resolved. If $\mathbb{B}_t$ is finite, we need to solve $\sum_{t \in \mathcal{T}} |\mathbb{B}_t|$ parameteric optimization equations to solve the dynamic program of (11), thereby resolving difficulty (D2). Even if $\mathbb{B}_t$ is infinite, the dynamic program of (11) can be solved efficiently either exactly (e.g., LQG systems) or approximately (e.g., POMDPs).

Now lets come back to the question of extending the dynamic programming principle to decentralized stochastic control systems. Decentralized stochastic problems belong to NEXP complexity class [34]. So, in general, no efficient solution algorithm is possible. Nonetheless, it is possible to develop a dynamic programming decomposition for specific information structures. We present three such generic approaches.

1) The person-by-person approach: The person-by-person approach is motivated by techniques for computing the Nash equilibrium of a game in policy space and works when some controllers have perfect recall.

The approach proceeds as follows. Pick a controller, say i, that has perfect recall and *arbitrarily* fix the control policies γ^{-i} of all other controllers. Now consider the sub-problem of optimally choosing the best policy at controller i. Since the policies of all other controllers are fixed and controller i has perfect recall, this optimal control sub-problem at controller i is a centralized stochastic control problem. Let π_t^i be any information state for controller i at time t. Then, the structural results of centralized stochastic control described above show that a policy of the form $u_t^i = \gamma_t^i(\pi_t^i)$ performs as well as (or, in other words, dominates) any other control policy at controller i. If the *form* of π_t^i does not depend on the policies γ^{-i} for other controllers, then the policy $u_t^i = \gamma_t^i(\pi_t^i)$ is optimal for every choice of γ^{-i}; and hence is globally optimal. Thus, the person-by-person approach is used to identify the structure of globally optimal policies.

As an example, consider the decentralized sequential hypothesis testing problem described in [35]. Let $H \in \{0, 1\}$ be a binary hypothesis and two sensors observe noisy versions of H. At each time, a sensor can either stop and declare an

estimate $\hat{H}^i$, or continue to take additional measurements. Each measurement costs c and when both sensors have declared an estimate, a cost $\ell(H, H^1, H^2)$ is incurred.

Following a person-by-person approach, we get that $\pi_t^1 = \mathcal{P}(H = 0|y_{[1,t]}^i)$ is an information state for controller i. Since the form of π_t^1 does not depend on the control policies of the other sensor, restricting attention to control policies of the form $u_t^i = \gamma_t^i(\pi_t^i)$ does not entail any loss of optimality. Furthermore, by following a procedure similar to the Wald sequential hypothesis testing, we can show that the optimal control policy γ_t^i is of the threshold type, i.e., there exists numbers $\underline{\tau}_t^i, \bar{\tau}_t^i \in [0, 1]$ with $\underline{\tau}_t^i \le \bar{\tau}_t^i$ such that if $\pi_t^i \in [0, \underline{\tau}_t^i)$, then it is optimal to stop and declare 1; if $\pi_t^i \in (\bar{\tau}_t^i, 1]$, then it is optimal to stop and declare 0; otherwise, it is optimal to take another measurement. (See [35] for details). Thus, the person-by-person approach identifies the structure of the optimal controller.

When all controllers have perfect recall and the information state at each controller belongs to a time-invariant space, then the person-by-person approach also gives coupled dynamic programs. If these coupled dynamic programs have a fixed point solution, then the resulting policies are person-by-person optimal.

The main idea for finding such coupled dynamic programs is the following. Pick a controller, say i, and arbitrarily fix the controller policy γ^{-i} of all other controllers. Write the dynamic program to find the optimal policy at controller i. This dynamic program determines the *best response* of controller i to the policies γ^{-i}. Write similar dynamic programs for all controllers. Thus, we end up with N coupled dynamic programs, one for each controller. These coupled dynamic programs have a fixed point if for every i, γ^i is the best response to γ^{-i}.

We can use an orthogonal search to find the fixed point of such coupled dynamic programs. Arbitrarily initialize the control policies $\gamma^{(0)}$ for all controllers. Pick a controller i, and use the dynamic program for controller i to find the best response $\gamma^{*,i}$ to $\gamma^{(0),-i}$. Set $\gamma^{(1)}$ such that $\gamma^{(1),i} = \gamma^{*,i}$ and the rest of the components are same as in $\gamma^{(0)}$. Repeat the above process by picking some other controller j. Continue until a fixed-point is reached. By construction, such a fixed point determines a person-by-person optimal policy. Such an approach was used in [35] to find person-by-person optimal strategies for the decentralized sequential hypothesis testing problem described above.

In summary, the person-by-person approach identifies structural properties of globally optimal control policies and provides a search method to find person-by-person optimal control policies. This method has been used to identify globally optimal strategies for specific information structures (e.g., stochastically nested information structures [21] and broadcast information structures [36]) and for various applications (e.g., real-time communication [37]–[42], decentralized hypothesis testing and quickest change detection [35], [43]–[49], and networked control systems [50], [51]). The person-by-person approach has also been used to identify person-by-person optimal control strategies for

specific information structures (e.g., control sharing information structure [52]). As discussed earlier, a decentralized control problem need not be convex. Thus, a person-by-person optimal policy found using the above approach need not be globally optimal. However, if the problem is convex (e.g., an LQG system that is partially nested), then the approach gives a globally optimal solution.

2) The designer's approach: The designer's approach was proposed in [53] (where it was called the standard form) and refined in [54], [55]. The main idea behind the designer's approach is that although the dynamic team problem is informationally decentralized, it is a centralized planning problem that may be solved from the viewpoint of a system designer that (centrally) chooses the *control laws* of all the controllers *before the system starts running*. Since the control laws are picked before the system starts running, no data is observed by the system designer. We can obtain a dynamic programming decomposition of this centralized planning problem by identifying an appropriate information state for the designer.

For the system described in Section III, the planning problem is a POMDP. Thus, based on the results for POMDPs, the information state is given by $\mathcal{P}(x_t, I_t^1, \ldots, I_t^n)$. In general, a dynamic program based on such an information state is not useful because the space of realizations of I_t^i is increasing with time. Nonetheless, the designer's approach gives useful results for specific information structures, as is illustrated by the following example.

Decentralized control with finite memory controllers: Consider the model of Section III in which no controller has perfect recall. A simple example of this case is $I_t^i = \{y_t^i, y_{t-1}^i\}$. Based on the designer's approach, the information state for this system is $\mathcal{P}(x_t, \mathbf{y}_t, \mathbf{y}_{t-1})$, which is computable from $\pi_t = \mathcal{P}(x_t, \mathbf{y}_{t-1})$. Based on this information state, we get a dynamic program characterized by the following recursion:

$$V_t(\pi_t) = \inf_{\boldsymbol{\gamma}_t} \mathcal{E}[c(x_t, \mathbf{u}_t) + V_{t+1}(\pi_{t+1}) \mid \pi_t, \boldsymbol{\gamma}_t].$$

The minimizer at time t, if it exists, gives the optimal control laws $\boldsymbol{\gamma}_t$ when the information state is π_t. Such an approach was used for one controller with finite memory in [56] and for two controllers with finite memory in [54].

Notice that the information state defined above just depends on the past choices of control laws $\boldsymbol{\gamma}_{[1,t-1]}$. Hence, the evolution of the information state is deterministic and the optimal control policies are obtained using a two step approach. In the first step, proceed backwards in time and determine the value function $V_t(\cdot)$ and the corresponding minimizer. In the second step, start from the initial value of the information state π_1 and proceeds forward in time as follows. From the result of the first step, find the optimal control law $\boldsymbol{\gamma}_1$ corresponding to π_1. Based on π_1 and $\boldsymbol{\gamma}_1$, determining π_2 and continue the above process until all control laws $\boldsymbol{\gamma}_t$, $t \in \mathcal{T}$ are determined.

The designer's approach can also be used in tandem with the person-by-person approach as is illustrated by the

following example. Such a tandem approach has been used in various applications including real-time communication [38], [55], [57], decentralized hypothesis testing [47], and networked control systems [50], [51].

Two agent team with one finite-memory controller: Consider a two agent ($N = 2$), where one controller has perfect recall with $I_t^1 = (y_{[1,t]}^1, u_{[1,t-1]}^1)$ and the second controller has finite memory, say $I_t^2 = (y_t^2, y_{t-1}^2)$. (A more general form of this system was investigated in [54]). Directly using the designer's approach on this system will give an information state $\mathcal{P}(x_t, y_{[1,t]}^1, u_{[1,t-1]}^1, y_{[t-1,t]}^2)$ which does not take values in a time invariant space. We can find a more useful information state if we first use the person-by-person approach on controller 1, which shows that $\mathcal{P}(x_t, y_{[t-1,t]}^2 \mid I_t^1)$, which is computable from $\theta_t = \mathcal{P}(x_t, y_{t-1}^2 \mid I_t^1)$, is an information state for controller 1. Therefore, based on the structural results, restricting attention to controller 1 of the form $u_t^1 = \gamma_t^1(\theta_t)$ does not entail any loss of optimality.

Now, restrict attention to controller 1 of the form $u_t^1 = \gamma_t^1(\theta_t)$ and use the designer's approach. Based on the designer's approach, we get that $\mathcal{P}(x_t, \theta_t, y_{[t-1,t]}^2)$, which is computable from $\pi_t = \mathcal{P}(x_t, \theta_t, y_{t-1}^2)$, is an information state. Notice that π_t takes value in a time invariant space. Based on this information state, we get the following dynamic program. Define

$$V_t(\pi_t) = \inf_{\gamma_t} \mathcal{E}[c(x_t, \mathbf{u}_t) + V_{t+1}(\pi_{t+1}) \mid \pi_t, \gamma_t].$$

The minimizer at time t, if it exists, gives the optimal control laws γ_t when the information state is π_t. Although the above information state looks formidable (it is a probability measure on a probability measure), the above dynamic program provides a means of synthesizing approximately optimal control laws.

3) The common information approach: The common information approach was proposed in [58]–[60] and works for decentralized control problems in which the controllers sequentially share information with each other e.g., the delayed information/measurement/control sharing information structure and the periodic sharing information structures defined earlier. The general class of such information structures is called partial history sharing information structure [60].

The main idea behind the common information approach is the following. Given any information structure, we can split the information available at the controllers into *common information* $C_t = \bigcap_{s \geq t} \bigcap_{i=1}^n I_s^i$ and *local information* $L_t^i = I_t^i \setminus C_t$, $i = 1, \ldots, N$. If all controllers have perfect recall, then we can simply define common information as $C_t = \bigcap_{i=1}^n I_t^i$. By construction, common information is increasing with time, i.e., $C_t \subseteq C_{t+1}$.

The common information approach consists of the following five steps. The first step is to formulate a centralized *coordinated system* from the viewpoint of a *coordinator* that observes the common information C_t and chooses *prescriptions* $\varphi_t = (\varphi_t^1, \ldots, \varphi_t^N)$, where φ_t^i maps the local information L_t^i to control action u_t^i, $i = 1, \ldots, N$. In this coordinated systems, the controllers simply take the prescription φ_t^i given

to them by the coordinator, and use it to generate control action $u_t^i = \varphi_t^i(L_t^i)$. The second step is to show that the coordinated system is a centralized POMDP. The third step is to identify an information state for the coordinator and use it to identify the structure of optimal control laws and write a dynamic programming decomposition. The fourth step is to show that the coordinated system is equivalent to the original system, i.e., any policy of the coordinated system is implementable in the original system, and vice versa, in such a manner that both policies result in identical realization of all system variables. The fifth and the final step is to use the result of the fourth step to translate the structural result and the dynamic programming decomposition of the third step to the original system.

For the general system defined above, the information state is given by $\mathcal{P}(x_t, L_t^1, \ldots, L_t^N \mid C_t)$. In general, the dynamic programming decomposition based on this information state is not useful because the local information L_t^i is increasing with time. Nonetheless, the common information approach gives useful results for specific information structures, as is illustrated by the following example.

The delayed sharing information structure: Consider the model of Section III with a n-step delayed sharing information structure, i.e., I_t is given by (4). This model was proposed in [61], the structural result and dynamic programming decomposition for $n = 1$ was given in [62], and two structural results and dynamic programming decomposition for the general n was given in [63]. We show how this model can be solved using the common information approach.

In this model, the common information is $C_t = (\mathbf{y}_{[1,t-n]}, \mathbf{u}_{[1,t-n]})$, while the local information is $L_t^i = (y_{[t-n+1,t]}^i, u_{[t-n+1,t]}^i)$. Based on the common information approach, we get that $\pi_t = \mathcal{P}(x_t, \mathbf{y}_{[t-n+1,t]}, \mathbf{u}_{[t-n+1,t]})$ is an information state for the coordinator. In addition, we get the following:

1) *Structure of optimal controller.* In the coordinated system, restricting attention to coordination strategies of the form $\varphi_t = d_t(\pi_t)$ does not entail any loss of optimality. Since $u_t^i = \varphi_t^i(L_t^i)$, the structural result for the coordinator implies that in the original system, restricting attention to control strategies of the form $u_t^i = \gamma_t^i(L_t^i, \pi_t)$ does not entail any loss of optimality.

2) *Dynamic programming decomposition.* In the coordinated system, define

$$V_t(\pi_t) = \inf_{\varphi_t} \mathcal{E}[c(x_t, \mathbf{u}_t) + V_{t+1}(\pi_{t+1}) \mid \pi_t, \varphi_t].$$

The minimizer at time t, if it exists, gives the optimal prescription φ_t when the information state is π_t. The corresponding optimal control policy in the original system is given by $\gamma_t^i(\cdot, \pi_t) = \varphi_t^i(\pi_t)(\cdot)$.

The above result is similar to the first structural result and dynamic programming derived in [63].

The common information approach helps in deriving structural results and dynamic programs that cannot be derived using the person-by-person approach or the designer's approach. A case in point is the results for delayed sharing

information structure derived above. If the common information is $C_t = \emptyset$, then the common information approach collapses to the designer's approach.

Variations of this approach have been used for various information structures, including delayed state sharing [64], partially nested systems with common past [65], teams with sequential partitions [66], coupled subsystems with control sharing [67], periodic sharing information structure [68], and belief sharing information structure [21].

IV. Norm-Optimal Design

In this section, we discuss methods for incorporating most types of information structure into a modern control framework. We focus here on problems where both the system to be controlled and the controllers under consideration are all linear (and, where applicable, time-invariant), except where otherwse noted. The number of controllers / decision makers, as well as the amount of times that each one has to act, may be finite or infinite, and the systems may be discrete-time or continuous-time. We mainly focus on problems where there is an infinite horizon and where stability issues thus need to be considered as well.

In IV-A, we describe the framework; in particular, introducing the idea of an information constraint, and explaining how optimal decentralized control problems may be addressed just like standard (centralized) optimal control problems, attempting to minimize a closed-loop norm, but subject to a constraint on the controller to be designed. In the simplest case, when certain controllers / decision makers can access some measurements but not others, as discussed in several examples above, this is captured by enforcing a sparsity constraint on the controller. In IV-B, we review how to parametrize all of the stabilizing controllers for centralized problems, in particular, the Youla-Kucera parametrization.

In IV-C we introduce a condition that allows the information constraint on the controller to be incorporated in such a way that finding the optimal controller subject to the information constraint can be cast as a convex optimization problem. This holds regardless of which closed-loop norm one wishes to minimize. The condition, called quadratic invariance, is an algebraic condition relating the system (plant) to the constraint, typically providing a simple test for convexity, and turns out to be closely related to the partially nested condition described above, developed for classifying linear optimality in LQG problems. In IV-D, we discuss a new result allowing these ideas to be applied more broadly, and allowing the optimization problem to be handled together with the problem of stabilization. In IV-E, armed with our test for convexity, we revisit some of the examples that were used to motivate the constraint framework, and determine when those problems are amenable to convex synthesis.

In IV-F, we discuss what happens when this approach is applied to problems which are not quadratically invariant, and perfectly decentralized problems in particular. In IV-G, we briefly discuss a related result for nonlinear systems.

A. Framework and Setup

We introduce a framework for designing optimal controllers for LTI systems, subject to decentralized information constraints.

a) Standard LTI framework: We first review a standard framework for centralized control synthesis.
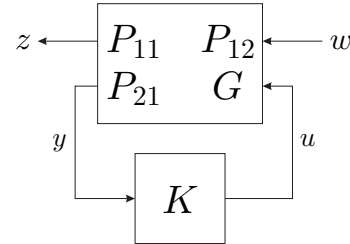

Fig. 2.   Standard LTI feedback control framework

Figure 2 represents a standard design framework often used in modern control theory. The signal w represents the vector of exogenous inputs, those the designer has no control over, such as wind gusts if one is considering an example in aerospace, and z represents everything the designer would like to keep small, which would typically include deviations from a desired state or trajectory, or a measure of control effort, for example. The signal y represents the vector of measurements that the controller K has access to, and u is the vector of inputs from the controller that is fed back into the plant. The plant is subdivided into four blocks which map w and u into z and y. The block which maps the controller input u to the measurements y is simply referred to as G, since it corresponds to the plant of classical control analysis, and so that we can later refer to its subdivisions without any ambiguity. Note that the four parts of the plant can, and often are, chosen such that all or parts of the vectors y and u are repeated in the vector z that we are trying to keep small, and such that parts of the vector u are repeated in y, to which the controller has access.

The design objective is to construct a controller K to keep a measure of the size of the mapping from w to z, known as the *closed-loop map*, as small as possible. There are many ways one can measure the size of a mapping, and thus this basic setup underpins much of modern (linear) controls including $\mathcal{H}_2$-control and $\mathcal{H}_\infty$-control. The choice of $\mathcal{H}_2$ leads to a stochastic interpretation, as that system norm is equivalent to the 2-norm of the vector z if w was a standard normal vector, and so the cost is equivalent to that of an associated LQG problem. In this framework, a decentralized information structure may be viewed as a constraint on the structure of the controller K, as now illustrated by examples.

b) Information constraint: We now illustrate why, in this framework, decentralization may be simply encapsulated as a constraint that the controller lies in a particular subspace. We focus now on the controller K, and on G, the part of the plant which takes the vector of control inputs u to the measurements y.

The diagram in Figure 3 represents three different subsystems, each of which may effect its neighbors, and each
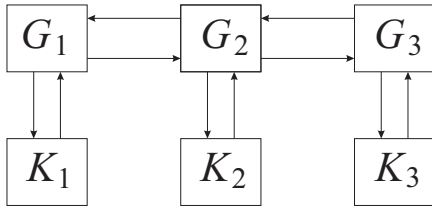
Fig. 3.   Perfectly decentralized control

of which has its own controller, which only has access to measurements coming from its own subsystem. In this case, if we look at the system as a whole, we need to design a controller K that can be written as

$$\begin{bmatrix} u_1 \\ u_2 \\ u_3 \end{bmatrix} = \underbrace{\begin{bmatrix} K_1 & 0 & 0 \\ 0 & K_2 & 0 \\ 0 & 0 & K_3 \end{bmatrix}}_{K} \begin{bmatrix} y_1 \\ y_2 \\ y_3 \end{bmatrix}$$

since each controller input may only depend upon the measurement from its corresponding subsystem. In other words, we need to design the best possible K which is block diagonal. The overall problem can be viewed as minimizing the size of the closed-loop map subject to the additional constraint that $K \in S$, where S is the set of all block diagonal controllers. This concept readily extends to any type of structural constraint we may need to impose in formulating an optimal control problem for controller synthesis. For instance, if in the above example, each controller were able to share information with its neighbors, then we would end up with a constraint set S which is tri-diagonal. In general, if the controller that generates u_i is unable to see y_j then the ijth entry of the controller must be set to zero.
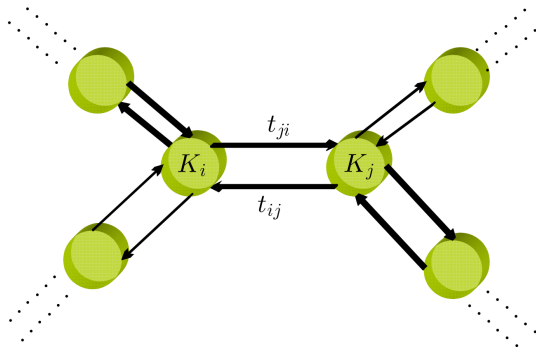

Fig. 4.   Network with delays

If controllers were instead allowed to communicate with each other, but with some delays, this too could be reflected in another constraint set S. This situation is represented in Figure 4, where the controller that generates u_i, which regulates subsystem i, can see the information from another subsystem j only after a transmission delay of t_{ij}. In this case, if we look at the system as a whole, we need to design

a controller K that can be written as

$$\begin{bmatrix} u_1 \\ u_2 \\ u_3 \end{bmatrix} = \underbrace{\begin{bmatrix} D_{t_{11}}\tilde{K}_{11} & D_{t_{12}}\tilde{K}_{12} & D_{t_{13}}\tilde{K}_{13} \\ D_{t_{21}}\tilde{K}_{21} & D_{t_{22}}\tilde{K}_{22} & D_{t_{23}}\tilde{K}_{23} \\ D_{t_{31}}\tilde{K}_{31} & D_{t_{32}}\tilde{K}_{32} & D_{t_{33}}\tilde{K}_{33} \end{bmatrix}}_{K} \begin{bmatrix} y_1 \\ y_2 \\ y_3 \end{bmatrix}$$

where each $D_{t_{ij}}$ realizes the corresponding delay t_{ij}, and each $\tilde{K}_{ij}$ represents a factor of the controller that we are free to design. Notice that, according to the representation above, the controller responsible for u_i must wait the prescribed delay t_{ij} until it can access measurements from subsystem j.

The set S above is called the *information constraint*, as it captures the information available to various parts of the controller. This includes the notion of information structures as defined earlier in Section III, as the constraint can be set to encapsulate which measurements and/or control inputs (and from which subsystems and from which times) are available to the controllers of which subsystems at which times.

The overarching point is that the objective of decentralized control may be considered to be the minimization of a closed-loop map subject to an information constraint $K \in S$. The approach is extremely broad, as it seamlessly incorporates any type of decentralization, any control objective, and heterogeneous subsystems.

c) Problem formulation: The mapping from w to z that we wish to keep small in Figure 2, the closed-loop map, can be written as $f(P, K) = P_{11} + P_{12}K(I - GK)^{-1}P_{21}$. The problem that we would like to address may then be formulated as:

$$\begin{aligned} \text{minimize} \quad & \|f(P, K)\| \\ \text{subject to} \quad & K \text{ stabilizes } P \\ & K \in S \end{aligned} \tag{12}$$

The norm ($\|\cdot\|$) is any appropriate system norm, chosen based on the particular performance objectives, which could be the $\mathcal{H}_2$-norm or $\mathcal{H}_\infty$-norm, as briefly described earlier. The information constraint S is the subspace of admissible controllers that encapsulates the decentralized nature of the system, as exemplified above. Choosing P and S, along with the $\mathcal{H}_2$-norm, to correspond to a given LQG problem with information structure is discussed in [69]. The stabilization constraint is needed in the most typical case where the signals lie in extended spaces and the plant and controller are rational proper systems whose interconnections may thus be unstable. It may not be necessary, or another technical condition may be necessary such as the invertibility of $(I - GK)$, for other spaces of interest, such as Banach spaces with bounded linear operators [70], [71].

B. Youla-Kucera Parametrization of Stabilizing Controllers

If the plant to be controlled is stable, we could use the following change of variables

$$Q = -K(I - GK)^{-1} \iff K = -Q(I - GQ)^{-1} \tag{13}$$

and then allowing the new parameter Q to be stable is equivalent to the controller K stabilizing the plant P, and the set of all achievable closed-loop maps (ignoring the information constraint) is then given as

$$\{P_{11} - P_{12}QP_{21} \mid Q \text{ stable}\}. \tag{14}$$

This is generalized by the Youla-Kucera or YJBK parametrization [72], which uses a doubly coprime factorization of G over $\mathcal{RH}_\infty$, that is, $M_l, N_l, X_l, Y_l, M_r, N_r, X_r, Y_r \in \mathcal{RH}_\infty$ such that $G = N_r M_r^{-1} = M_l^{-1} N_l$ and

$$\begin{bmatrix} X_l & -Y_l \\ -N_l & M_l \end{bmatrix} \begin{bmatrix} M_r & Y_r \\ N_r & X_r \end{bmatrix} = \begin{bmatrix} I & 0 \\ 0 & I \end{bmatrix}. \tag{15}$$

Then the set of all stabilizing controllers is given by

$$\begin{aligned} \{K \in \mathcal{R}_p & \mid K \text{ stabilizes } G\} \\ = & \Big\{ (Y_r - M_r Q)(X_r - N_r Q)^{-1} \mid \\ & X_r - N_r Q \text{ is invertible}, Q \in \mathcal{RH}_\infty \Big\} \\ = & \Big\{ (X_l - QN_l)^{-1}(Y_l - QM_l) \mid \\ & X_l - QN_l \text{ is invertible}, Q \in \mathcal{RH}_\infty \Big\}. \end{aligned} \tag{16}$$

Thus allowing the new (Youla) parameter Q to vary over all stable systems is still equivalent to considering all stabilizing controllers K, and the set of all achievable closed-loop maps is then given by

$$\{T_1 - T_2 Q T_3 \mid Q \in \mathcal{RH}_\infty\} \tag{17}$$

where T_1, T_2, T_3 are other stable systems which are functions of the coprime factors and the generalized plant parameters.

We see that these parametrizations allow the set of achievable closed-loop maps to be expressed as an affine function of a stable parameter, and thus allow our objective function in our main problem (12) to be cast as a convex function of that parameter. However, the information constraint $K \in S$ will typically not be simple to express in the new parameter, and this will ruin the convexity of the optimization problem.

C. Quadratic Invariance

We have seen that we can employ a change of variables that will make our objective convex, but that will generally cause the information constraint to no longer be affine. We thus seek to characterize problems for which the information constraint may be written as an affine constraint in the Youla parameter, such that a convex reformulation of our main problem will result.

The following property, first introduced in [70], provides that characterization.

Definition 4.1: The set S is **quadratically invariant** with respect to G if

$$KGK \in S \qquad \text{for all} \quad K \in S.$$

In other words, given any admissible controller K, the composition KGK has to be admissible as well. For problems where both are well-defined, this is shown to be equivalent to the partially nested condition [69] discussed in previous sections. When this condition holds, it follows that a controller being admissible is further equivalent to the linear-fractional transformation we encountered earlier lying in the constraint set [71], [73]:

$$K \in S \quad \Longleftrightarrow \quad K(I - GK)^{-1} \in S \tag{18}$$

Quadratic invariance is not only necessary and sufficient for the above equivalence, but also for the linear-fractional transformation of the admissible set to be any convex set [74]. We can see immediately from (13) that for the stable case this results in the equivalence of enforcing the information constraint on the controller or on the new parameter:

$$K \in S \quad \Longleftrightarrow \quad Q \in S \tag{19}$$

and it can be shown that when the plant is unstable, as long as it can be stabilized by a stable controller (known as strong stabilizability), another change of variables can be made such that this equivalence still holds [73].

Thus when the information constraint S is quadratically invariant with respect to the plant G, the optimal decentralized control problem (12) may be recast as the following:

$$\begin{aligned} \text{minimize} \quad & \|T_1 - T_2 Q T_3\| \\ \text{subject to} \quad & Q \text{ stable} \\ & Q \in S \end{aligned} \tag{20}$$

which is a convex optimization problem.

This leaves the question of how to find that initial stabilizing controller, itself a longstanding difficult problem, as well as whether the stabilization and optimization need to be handled separately, or whether they can be handled simultaneously. These questions are answered in the following subsection.

D. A Parametrization of Information-Constrained Controllers

Recent results show [75] that when quadratic invariance holds, an arbitrary coprime factorization, which is always available, can be used to characterize the information constraint on the controller as the following constraint on the Youla parameter:

$$K \in S \quad \Longleftrightarrow \quad M_r Q M_l - M_r Y_l \in S \tag{21}$$

While the constraint on Q is not identical to the information constraint, as it was in the previous cases we discussed, it is still an affine constraint.

The problem of finding a stabilizing controller is then reduced to finding a $Q \in \mathcal{RH}_\infty$ such that this constraint is satisfied, and moreover, if this cannot be done, then we know that the plant is not stabilizable with the given information constraint. It shown in [75] that the existence of such a Q is equivalent to the feasibility of an exact model matching problem [76], with stability restrictions [77], for which there are tractable solution methods [78]. If the associated exact model-matching problem is feasible then we can use the resulting Q to recover the stabilizing controller via (16). If

this controller is stable then we can use the previous results to write the optimal decentralized control problem as (20). When it is not, similar results can be obtained by using a parametrization discussed in [79], with a different convex problem resulting in a convex parametrization that is distinct from Youla's [80].

Instead of this two-step compensation scheme, finding a stabilizing controller first, then using this to parametrize all of the stabilizing controllers and to find the optimal one, we can also utilize this new equivalence (21) to address both issues at once.

We can now, after finding any doubly coprime factorization, recast the optimal decentralized control problem as follows:

$$
\begin{aligned}
\text{minimize} \quad & \|T_1 - T_2 Q T_3\| \\
\text{subject to} \quad & Q \text{ stable} \\
& M_r Q M_l - M_r Y_l \in S
\end{aligned}
\tag{22}
$$

This is still a convex optimization problem in the Youla parameter Q. The solution allows us to recover the optimal decentralized controller via (16), and the problem is infeasible if and only if the plant is not stabilizable with the given information constraint.

E. Examples

This subsection looks at particular classes of information constraints to see when this quadratic invariance condition holds, to identify those decentralized problems which are amenable to convex synthesis. We see that this algebraic condition often has intuitive interpretations for specific classes of problems.

1) Structural Constraints: We first look at structural constraints, or sparsity constraints, where each sub-controller can see the measurements from some subsystems but not from others. This structure can be represented with a binary matrix K^{bin}. For instance, $K^{\text{bin}}_{kl} = 1$ if the kth control input u_k is allowed to be a function of the lth measurement y_l, and $K^{\text{bin}}_{kl} = 0$ if it cannot see that measurement. The information constraint S is then the set of all controllers which have the structure prescribed by K^{bin}; that is, all of the controllers such that none of the sub-controllers use information which they cannot see.

A binary matrix G^{bin} can similarly be used to give the structure of the plant. For instance, $G^{\text{bin}}_{ij} = 1$ if G_{ij} is non-zero and the ith measurement y_i is affected by the jth control input u_j, and $G^{\text{bin}}_{ij} = 0$ if it is unaffected by that input. Given this representation of the structure of the plant and the controller constraints, we have the following result:

S is quadratically invariant with respect to G if and only if

$$
K^{\text{bin}}_{ki} G^{\text{bin}}_{ij} K^{\text{bin}}_{jl} (1 - K^{\text{bin}}_{kl}) = 0 \qquad \text{for all } i,j,k,l. \tag{23}
$$

Figure 5 illustrates this condition. The condition in (23) requires that, for arbitrary i, j, k, l, if the three blocks on the bottom are all non-zero (or allowed to be chosen non-zero), then the top block must be allowed to be non-zero as well. In other words, if there is an indirect connection from a
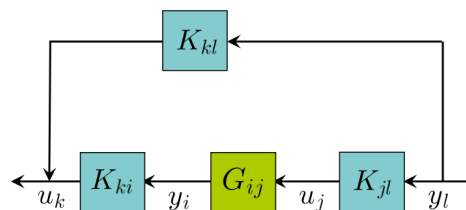

Fig. 5. Structural quadratic invariance

measurement to a control input, then there has to be a direct connection as well.

One subclass of problems for which this condition holds is where the plant is lower triangular, and the controller constraints are such that if a given part of the controller is allowed to be nonzero, then all of the parts of the controller below it must be active as well, sometimes called skyline structure [70]. This can be viewed as the manifestation in this framework of the aforementioned classical structures. Other classes of problems previously shown to admit convex solutions [81] satisfy the condition developed here as well.

When this condition is met, the problem is quadratically invariant, and we can recast our optimal decentralized control problem as the convex optimization problem in (20) or (22).

2) Symmetry: We briefly consider the problem of symmetric synthesis. Suppose that we need to design the best symmetric controller; that is, the best controller such that $K_{kl} = K_{lk}$ for all k, l, and that the information constraint S is the set of all such symmetric controllers. If the plant is also symmetric; that is, if $G_{ij} = G_{ji}$ for all i, j, then KGK is symmetric for any symmetric K. Thus, $KGK \in S$ for all $K \in S$, the problem is quadratically invariant, and the optimal symmetric control problem may be recast as (20) or (22).

3) Delays: We now return to the problem of Figure 4, where we have multiple nodes/subsystems, each with its own controller, and each subsystem i can see the information from another subsystem j after a transmission delay of t_{ij}.

We similarly consider that the inputs to a given subsystem j may affect other subsystems after some delay, and denote the amount of time after which it may affect another subsystem i by the propagation delay p_{ij}.

The overall problem of controlling such a network with propagation delays, with controllers that may communicate with transmission delays, is depicted in Figure 6.

When this problem is tested for quadratic invariance, one first finds that the following condition is necessary and sufficient:

$$
t_{ki} + p_{ij} + t_{jl} \geq t_{kl} \qquad \text{for all } i, j, k, l \tag{24}
$$

This is reminiscent of condition (23) for structural constraints, as it similarly requires that any direct path from y_l to u_k must be at least as fast as any indirect path through the plant. This condition can be further reduced to a very simple intuitive condition [82], as long as we may assume that the transmission delays themselves satisfy the triangle
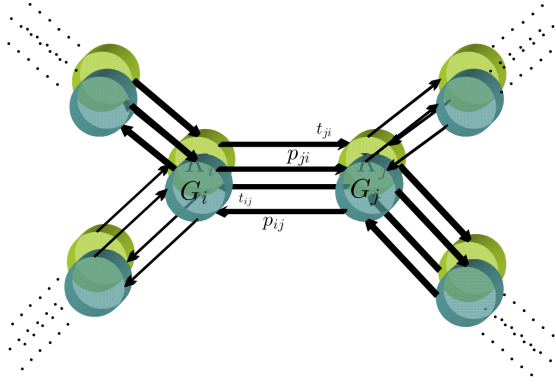
Fig. 6.   Network with delays

inequality; that is

$$t_{ki} + t_{ij} \ \geq \ t_{kj} \qquad \text{for all } k, i, j. \tag{25}$$

This is typically a very reasonable assumption, as it states that information is transmitted between nodes in the quickest manner available through the network. If the inequality failed for some k, j, one would want to reroute the transmissions from j to k along the faster route such that the inequality would then hold.

If the triangle inequality among transmissions does hold, then condition (24), and thus quadratic invariance, is reduced to simply:

$$p_{ij} \ \geq \ t_{ij} \qquad \text{for all } i, j. \tag{26}$$

In other words, for any pair of nodes, information needs to be transmitted faster than the dynamics propagate. When this simple condition holds, the problem is quadratically invariant, and the optimal decentralized control problem may be recast as the convex problem (20) or (22).

This very intuitive result has a counterintuitive complement when one considers computational delays as well. Suppose now that the ith controller cannot use a measurement from the jth subsystem until a 'pure' transmission delay of $\tilde{t}_{ij}$, representing the time it takes to send the information from one subsystem to the other, as well as a computational delay of c_i, representing the time it takes to process the information once it is received.

While intuition might suggest that these two quantities would end up being added and then replacing the right-hand side of equation (26), if we now assume that the pure transmission delays satisfy the triangle inequality, the condition for quadratic invariance becomes:

$$p_{ij} + c_j \ \geq \ \tilde{t}_{ij} \qquad \text{for all } i, j \tag{27}$$

with the computational delay on the other side of the inequality.

This shows that, regardless of computational delay, if information can be transmitted faster than dynamics propagate, then the optimal decentralized control problem can be reformulated as a convex optimization problem. If we consider a problem with multiple aerial vehicles, for example, where

dynamics between any pair of subsystems will propagate at the speed of sound, this tells us that transmissions just have to be faster than that threshold for the optimal control problem to be recast as (20) or (22).

The results of this section have also been extended to spatio-temporal systems [83], where plants and controllers can act across a continuum, and results very similar to the above still hold. When applied to the special case of spatially invariant systems, the controller still needs to be able to receive information faster than the plant can propagate its inputs over any given distance, analogous to (26), and the triangle inequality (25) discussed above becomes a condition that the support function imposed on the controller needs to be subadditive. This includes funnel causal systems, developed in the study of convexity for these problems [84].

F. Perfectly Decentralized Control

We now revisit the problem of Figure 3, where each controller can only use the measurements from its own subsystem, and thus the information constraint is block diagonal. This problem is never quadratically invariant, and will never satisfy condition (23), except for the case where the subsystems do not affect one another; that is, except for the case where G is block diagonal as well.

In all other cases where subsystems may have some affect on others, we thus cannot parametrize all of the admissible stabilizing controllers in a convex fashion, and cannot cast the optimal decentralized control problem as a convex problem such as in (20). However, a Youla parametrization can similarly be used, and while (19) does not hold, as the information constraint on the controller is not equivalent to enforcing it on the Youla parameter, it is instead equivalent to a quadratic equality constraint on the parameter [85]:

$$K \in S \quad \Longleftrightarrow \quad W_2 + QW_4 - W_1 Q - QW_3 Q \ = \ 0 \tag{28}$$

for stable operators W_1, W_2, W_3, W_4. When returning to the optimal decentralized control problem, this equality constraint replaces the final $Q \in S$ constraint of (20) or $M_r Q M_l - M_r Y_l \in S$ constraint of (22). The problem is no longer convex due to the quadratic term, but the overall difficulty is transformed to one well-understood type of constraint, for which many methods exist to approximate optimal solutions. Another approach inspired on Youla's is given in [86] for the fully decentralized case, where the lack of convexity is readily recognizable from certain nonlinear constraints imposed on the parameters.

Other structural constraints, which are neither block diagonal nor quadratically invariant, can be similarly parametrized by first converting them to a perfectly decentralized problem [87]. One can also find the closest constraint set which is quadratically invariant to get approximate solutions, or the closest quadratically invariant subset or superset to obtain bounds on the solution [88].

G. Nonlinear Decentralized Controller Parametrization

The parametrization and optimization results discussed thus far assume that the operators, both the plant to be

controlled and the possible controllers that we may design for it, are all linear, and when applicable, time-invariant as well. A similar convex parametrization of stabilizing decentralized controllers exists even when the systems are possibly nonlinear and possibly time-varying (NLTV) [89]. The condition allowing for the parametrization then becomes

$$K_1(I \pm GK_2) \in S \qquad \text{for all} \quad K_1, K_2 \in S.$$

When the plant is stable, the (finite gain) stabilizing controllers may be parametrized similarly to (14) [90], and when the plant is unstable, the stabilizing controllers may typically be parametrized similarly to (17) [91]. Similar to quadratic invariance, the above condition then yields the equivalence of the controller and the feedback map satisfying the information constraint (18), which then gives the equivalence of the controller and the parameter satisfying the constraint as in (19). The convex parametrization of all causal stabilizing decentralized controllers then results, analogous to the linear case with quadratic invariance.

While this condition may appear quite different from quadratic invariance, they actually both reduce to the same conditions when one considers the classes of sparsity constraints or delay constraints, and so these results extend to all of the cases covered in Sections IV-E.1 and IV-E.3.

REFERENCES

[1] G. N. Nair, F. Fagnani, S. Zampieri, and R. J. Evans, "Feedback Control Under Data Rate Constraints: An Overview," *Proceedings of the Ieee*, vol. 95, no. 1, pp. 108–137, 2007.

[2] J. P. Hespanha, P. Naghshtabrizi, and Y. Xu, "A Survey of Recent Results in Networked Control Systems," *Proceedings of the IEEE*, vol. 95, no. 1, pp. 138–162, 2007.

[3] C. Papadimitriou and J. Tsitsiklis, "1985 24th IEEE Conference on Decision and Control," in *1985 24th IEEE Conference on Decision and Control*. IEEE, 1985, pp. 1099–1103.

[4] H. Witsenhausen, "A Counterexample in Stochastic Optimum Control," *Siam Journal Of Control And Optimization*, vol. 6, pp. 131–147, Dec. 1968.

[5] H. S. Witsenhausen, "Separation of estimation and control for discrete time systems," *Proceedings of the IEEE*, vol. 59, pp. 1557–1566, November 1971.

[6] Y.-C. Ho, "Team decision theory and information structures," in *Proceedings of the IEEE*, 1980, pp. 644–654.

[7] H. Witsenhausen, "The Intrinsic Model for Discrete Stochastic Control: Some Open Problems," in *Lecture Notes in Economics and Mathematical Systems*, A. Bensoussan and J. L. Lions, Eds. Springer Verlag, Jun. 1975, pp. 322–335.

[8] M. S. Andersland, "Decoupling Non-Sequential Stochastic Control Problems," in *American Control Conference, 1991*, 1991, pp. 2665–2669.

[9] M. Andersland and D. Teneketzis, "Information Structures, Causality, and Nonsequential Stochastic Control II: Design-Dependent Properties," *Siam Journal Of Control And Optimization*, vol. 32, pp. 1726–1751, Jun. 1994.

[10] M. S. Andersland and D. Teneketzis, "Information Structures, Causality, and Nonsequential Stochastic Control I: Design-Independent Properties ," *Siam Journal Of Control And Optimization*, vol. 30, pp. 1447–1475, Jun. 1992.

[11] D. Teneketzis, "On information structures and nonsequential stochastic control," *CWI Quarterly*, vol. 9, pp. 241–261, May 1996.

[12] S. Yüksel and T. Başar, *Stochastic Networked Control Systems: Stabilization and Optimization under Information Constraints*. (under review, Springer-Birkhäuser), Sep. 2012.

[13] R. Radner, "Team decision problems," *The Annals of Mathematical Statistics*, 1962.

[14] J. Marschack and R. Radner, *Economic Theory Of Teams*. New Haven: Yale University Press, Feb. 1972.

[15] H. S. Witsenhausen, "Equivalent Stochastic Control Problems," *Mathematics of Control, Signals and Systems (Springer Verlag)*, vol. 1, pp. 3–11, Jan. 1988.

[16] Y.-C. Ho and K.-C. Chu, "Team decision theory and information structures in optimal control problems–Part I," *Automatic Control, IEEE Transactions on*, vol. 17, no. 1, pp. 15–22, 1972.

[17] J. Krainak, J. Speyer, and S. Marcus, "Static team problems–Part I: Sufficient conditions and the exponential cost criterion," *IEEE Transactions on Automatic Control*, vol. 27, no. 4, pp. 839–848, Aug. 1982.

[18] N. Sandell and M. Athans, "Solution of some nonclassical LQG stochastic decision problems," *IEEE Transactions on Automatic Control*, vol. 19, no. 2, pp. 108–116, Apr. 1974.

[19] T. Yoshikawa, "Decomposition of Dynamic Team Decision Problems," *IEEE Transactions on Automatic Control*, vol. 23, no. 4, pp. 627–632, 1978.

[20] S. Yüksel and T. Linder, "On optimization and convergence of observation channels and quantizers in stochastic control," in *Communication, Control, and Computing (Allerton), 2010 48th Annual Allerton Conference on*, 2010, pp. 286–293.

[21] S. Yüksel, "Stochastic Nestedness and the Belief Sharing Information Pattern," *IEEE Transactions on Automatic Control*, vol. 54, no. 12, pp. 2773–2786, 2009.

[22] R. Bansal and T. Başar, "Solutions to a class of linear-quadratic-Gaussian (LQG) stochastic team problems with nonclassical information," *Systems & Control Letters*, vol. 9, pp. 125–130, 1987.

[23] A. A. Zaidi, S. Yüksel, T. J. Oechtering, and M. Skoglund, "On optimal policies for control and estimation over a Gaussian relay channel," in *Decision and Control and European Control Conference (CDC-ECC), 2011 50th IEEE Conference on*, 2011, pp. 5720–5725.

[24] A. Özçelikkale, S. Yüksel, and H. M. Ozaktas, "Unitary Precoding and Basis Dependency of MMSE Performance for Gaussian Erasure Channels," *arXiv.org*, vol. cs.IT, Nov. 2011.

[25] G. M. Lipsa and N. Martins, "Optimal memoryless control in Gaussian noise: A simple counterexample," *Automatica*, 2011.

[26] Y. Wu and S. Verdu, "Witsenhausen's counterexample: A view from optimal transport theory," in *Decision and Control and European Control Conference (CDC-ECC), 2011 50th IEEE Conference on*, 2011, pp. 5732–5737.

[27] R. Bansal and T. Başar, "Stochastic team problems with nonclassical information revisited: When is an affine law optimal?" *IEEE Transactions Automat. Contr.*, vol. 32, pp. 554–559, 1987.

[28] T. Başar, "The Gaussian test channel with an intelligent jammer," *IEEE Transactions on Information Theory*, vol. 54, no. 29, pp. 152–157, 1983.

[29] T. Başar and M. Mintz, "Minimax estimation under generalized quadratic loss," in *Decision and Control, 1971 IEEE Conference on*, 1971, pp. 456–461.

[30] ——, "On the existence of linear saddle-point strategies for a two-person zero-sum stochastic game," in *Decision and Control, 1972 and 11th Symposium on Adaptive Processes. Proceedings of the 1972 IEEE Conference on*, 1972, pp. 188–192.

[31] M. Rotkowitz, "Linear Controllers are Uniformly Optimal for the Witsenhausen Counterexample," in *Decision and Control, 2006 45th IEEE Conference on*, 2006, pp. 553–558.

[32] A. Gattami, B. M. Bernhardsson, and A. Rantzer, "Robust Team Decision Theory," *IEEE Transactions on Automatic Control*, vol. 57, no. 3, pp. 794–798, 2012.

[33] A. Mahajan and S. Yüksel, "Measure and cost dependent properties of information structures," in *American Control Conference (ACC), 2010*, 2010, pp. 6397–6402.

[34] D. S. Bernstein, S. Zilberstein, and N. Immerman, "The complexity of decentralized control of markov decision processes," in *Proceedings of the 16th International Conference on Uncertainty in Artificial Intelligence (UAI)*, Stanford, CA, Jun. 2000, pp. 32–27.

[35] D. Teneketzis and Y. Ho, "The decentralized Wald problem," *Information and Computation (formerly Information and Control)*, vol. 73, no. 1, pp. 23–44, Apr. 1987.

[36] J. Wu and S. Lall, "A dynamic programming algorithm for decentralized markov decision processes with a broadcast structure," in *Proceedings of the 49th IEEE Conference on Decision and Control*, 2010, pp. 6143–6148.

[37] H. S. Witsenhausen, "On the structure of real-time source coders," *Bell System Technical Journal*, vol. 58, no. 6, pp. 1437–1451, July-August 1979.

[38] J. C. Walrand and P. Varaiya, "Optimal causal coding-decoding problems," *IEEE Trans. Inf. Theory*, vol. 29, no. 6, pp. 814–820, Nov. 1983.

[39] D. Teneketzis, "On the structure of optimal real-time encoders and decoders in noisy communication," *IEEE Trans. Inf. Theory*, pp. 4017–4035, Sep. 2006.

[40] A. Nayyar and D. Teneketzis, "On the structure of real-time encoders and decoders in a multi-terminal communication system," *IEEE Trans. Inf. Theory*, vol. 57, no. 9, pp. 6196–6214, Sep. 2011.

[41] Y. Kapsi and N. Merhav, "Structure theorem for real-time lossy source encoders and memory-limited decoders with side information," in *International Symposium on Information Theory*, 2010.

[42] S. Yüksel, "On optimal causal codng of partially observed markov sources in single and multi-terminal settings," *accepted for publication in IEEE Trans. Inf. Theory*, 2012.

[43] R. R. Tenney and N. R. Sandell Jr., "Detection with distributed sensors," *IEEE Trans. Aerospace Electron. Systems*, vol. AES-17, no. 4, pp. 501–510, Jul. 1981.

[44] J. N. Tsitsiklis, "Decentralized detection," in *Advances in Statistical Signal Processing*. JAI Press, 1993, pp. 297–344.

[45] V. V. Veeravalli, T. Başar, and H. V. Poor, "Decentralized sequential detection with a fusion center performing the sequential test," *IEEE Transactions on Information Theory*, vol. 39, no. 2, pp. 433–442, 1993.

[46] ——, "Decentralized sequential detection with a sensors performing the sequential test," *Mathematics of Control, Signals, and Systems*, vol. 7, no. 4, pp. 292–305, 1994.

[47] A. Nayyar and D. Teneketzis, "Sequential problems in decentralized detection with communication," *IEEE Trans. Inf. Theory*, vol. 57, no. 8, pp. 5410–5435, Aug. 2011.

[48] D. Teneketzis and P. Varaiya, "The decentralized quickest detection problem," *IEEE Trans. Autom. Control*, vol. AC-29, no. 7, pp. 641–644, Jul. 1984.

[49] V. V. Veeravalli, "Decentralized quickest change detection," *IEEE Trans. Inf. Theory*, vol. 47, no. 4, pp. 1657–1665, Apr. 2001.

[50] J. C. Walrand and P. Varaiya, "Causal coding and control of Markov chains," *System and Control Letters*, vol. 3, pp. 189–192, 1983.

[51] A. Mahajan and D. Teneketzis, "Optimal performance of networked control systems with non-classical information structures," *SIAM Journal of Control and Optimization*, vol. 48, no. 3, pp. 1377–1404, May 2009.

[52] M. Aoki, "On decentralized linear stochastic control problems with quadratic cost," *IEEE Trans. Autom. Control*, vol. 18, no. 3, pp. 243 – 250, Jun. 1973.

[53] H. S. Witsenhausen, "A standard form for sequential stochastic control," *Mathematical Systems Theory*, vol. 7, no. 1, pp. 5–11, 1973.

[54] A. Mahajan, "Sequential decomposition of sequential dynamic teams: applications to real-time communication and networked control systems," Ph.D. dissertation, University of Michigan, Ann Arbor, MI, Sep. 2008.

[55] A. Mahajan and D. Teneketzis, "Optimal design of sequential real-time communication systems," *IEEE Trans. Inf. Theory*, vol. 55, no. 11, pp. 5317–5338, Nov. 2009.

[56] N. R. Sandell, Jr., "Control of finite-state, finite-memory stochastic systems," Ph.D. Thesis, Massachussets Institute of Technology, Cambridge, MA, 1974.

[57] A. Mahajan and D. Teneketzis, "On the design of globally optimal communication strategies for real-time noisy communication systems with noisy feedback," *IEEE J. Sel. Areas Commun.*, vol. 26, no. 4, pp. 580–595, May 2008.

[58] A. Mahajan, A. Nayyar, and D. Teneketzis, "Identifying tractable decentralized control problems on the basis of information structure," in *Proc. 46th Annual Allerton Conf. Communication, Control, and Computing*, Monticello, IL, Sep. 2008, pp. 1440–1449.

[59] A. Nayyar, "Sequential decision making in decentralized systems," Ph.D. dissertation, University of Michigan, Ann Arbor, MI, 2011.

[60] A. Nayyar, A. Mahajan, and D. Teneketzis, "Dynamic programming for multi-controller stochastic control with partial information sharing: A common-information approach," 2011, submitted to IEEE Trans. Autom. Control.

[61] H. S. Witsenhausen, "Separation of estimation and control for discrete time systems," *Proc. IEEE*, vol. 59, no. 11, pp. 1557–1566, Nov. 1971.

[62] P. Varaiya and J. Walrand, "On delayed sharing patterns," *IEEE Trans. Autom. Control*, vol. 23, no. 3, pp. 443–445, 1978.

[63] A. Nayyar, A. Mahajan, and D. Teneketzis, "Optimal control strategies in delayed sharing information structures," *IEEE Trans. Autom. Control*, vol. 56, no. 7, pp. 1606–1620, Jul. 2011.

[64] M. Aicardi, F. Davoli, and R. Minciardi, "Decentralized optimal control of Markov chains with a common past information set," *IEEE Trans. Autom. Control*, vol. 32, no. 11, pp. 1028–1031, 1987.

[65] G. Casalino, F. Davoli, R. Minciardi, P. Puliafito, and R. Zoppoli, "Partially nested information structures with a common past," *IEEE Trans. Autom. Control*, vol. 29, no. 9, pp. 846–850, Sep. 1984.

[66] T. Yoshikawa, "Decomposition of dynamic team decision problems," *IEEE Trans. Autom. Control*, vol. 23, no. 4, pp. 627–632, Aug. 1978.

[67] A. Mahajan, "Optimal decentralized control of coupled subsystems with control sharing," in *Proc. 50th IEEE Conf. Decision and Control and European Control Conf. (CDC-ECC)*, Orlando, FL, Dec. 2011, pp. 5726–5731.

[68] J. M. Ooi, S. M. Verbout, J. T. Ludwig, and G. W. Wornell, "A separation theorem for periodic sharing information patterns in decentralized control," *IEEE Trans. Autom. Control*, vol. 42, no. 11, pp. 1546–1550, Nov. 1997.

[69] M. Rotkowitz, "On information structures, convexity, and linear optimality," in *Proc. IEEE Conference on Decision and Control*, 2008, pp. 1642–1647.

[70] M. Rotkowitz and S. Lall, "Decentralized control information structures preserved under feedback," in *Proc. IEEE Conference on Decision and Control*, December 2002, pp. 569–575.

[71] ——, "Affine controller parameterization for decentralized control over Banach spaces," *IEEE Transactions on Automatic Control*, vol. 51, no. 9, pp. 1497–1500, September 2006.

[72] D. Youla, H. Jabr, and J. B. Jr., "Modern Wiener-Hopf design of optimal controllers: part II," *IEEE Transactions on Automatic Control*, vol. 21, no. 3, pp. 319–338, 1976.

[73] 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, February 2006.

[74] L. Lessard and S. Lall, "Quadratic invariance is necessary and sufficient for convexity," in *Proc. American Control Conference*, 2011, pp. 5360–5362.

[75] S. Sabau and N. C. Martins, "Stabilizability and norm-optimal control design subject to sparsity constraints," *arXiv:1209.1123v1*, September 2012.

[76] W. A. Wolovich, "The application of state feedback invariants to exact model-matching," in *5th Annual Princeton Conference on Information Sci. Syst., Princeton, NJ*, 1971.

[77] Z. Gao and P. J. Antsaklis, "On stable solutions of one-sided and 2-sided model-matching problems," *IEEE Trans. Automatic Control*, vol. 34, no. 9, pp. 978–982, 1989.

[78] D. Chu and P. V. Dooren, "A novel numerical method for exact model-matching problem with stability," *Automatica*, vol. 42, pp. 1697–1704, 2006.

[79] K. Mori and K. Abe, "Relationship between standard control problems and model-matching problems without coprime factorizability," *IEEE Trans. Automatic Control*, vol. 49, no. 2, pp. 230–233, 2004.

[80] S. Sabau and N. C. Martins, "A convex parameterization of all stabilizing controllers for non–strongly stabilizable plants under quadratically invariant sparsity constraints," in *Proceedings of the American Control Conference*, June 2009, pp. 878–883.

[81] X. Qi, M. Salapaka, P. Voulgaris, and M. Khammash, "Structured optimal and robust control with multiple criteria: A convex solution," *IEEE Transactions on Automatic Control*, vol. 49, no. 10, pp. 1623–1640, 2004.

[82] M. Rotkowitz, R. Cogill, and S. Lall, "A simple condition for the convexity of optimal control over networks with delays," in *Proc. IEEE Conference on Decision and Control*, 2005, pp. 6686–6691.

[83] ——, "Convexity of optimal control over networks with delays and arbitrary topology," *International Journal of Systems, Control, and Communications*, vol. 2, pp. 30–54, 2010.

[84] B. Bamieh and P. G. Voulgaris, "A convex characterization of distributed control problems in spatially invariant systems with communications constraints," *Systems and Control Letters*, vol. 54, no. 6, pp. 575–583, June 2005.

[85] V. Manousiouthakis, "On the parametrization of all stabilizing decentralized controllers," *Systems and Control Letters*, vol. 21, no. 5, pp. 397–403, 1993.

[86] R. A. Date and J. H. Chow, "Decentralized stable factors and a

parameterization of decentralized controllers," *IEEE Trans. Automatic Control*, vol. 39, pp. 347–351, 1994.

[87] M. Rotkowitz, "Parametrization of all stabilizing controllers subject to any structural constraint," in *Proc. IEEE Conference on Decision and Control*, December 2010, pp. 108 – 113.

[88] M. Rotkowitz and N. Martins, "On the nearest quadratically invariant constraint," *IEEE Transactions on Automatic Control*, vol. 57, no. 5, pp. 1314–1319, 2012.

[89] M. Rotkowitz, "Information structures preserved under nonlinear time-varying feedback," in *Proc. American Control Conference*, 2006, pp. 4207–4212.

[90] C. A. Desoer and R. W. Liu, "Global parametrization of feedback systems with nonlinear plants," *Systems and Control Letters*, vol. 1, no. 4, pp. 249–251, 1982.

[91] V. Anantharam and C. A. Desoer, "On the stabilization of nonlinear systems," *IEEE Transactions on Automatic Control*, vol. 29, no. 6, pp. 569–572, 1984.