Explicit Agent-Level Optimal Cooperative Controllers for Dynamically Decoupled Systems with Output Feedback

Mruganka Kashyap   Laurent Lessard

Abstract

We consider a dynamically decoupled network of agents each with a local output-feedback controller. We assume each agent is a node in a directed acyclic graph and the controllers share information along the edges in order to cooperatively optimize a global objective. We develop explicit state-space formulations for the jointly optimal networked controllers that highlight the role of the graph structure. Specifically, we provide generically minimal agent-level implementations of the local controllers along with intuitive interpretations of their states and the information that should be transmitted between controllers.

1   Introduction

Dynamically decoupled systems are examples of multi-agent systems where each agent has independent dynamics but the agents' controllers share information via a communication network to achieve a common goal. Dynamically decoupled architectures occur naturally in the context of swarms of unmanned aerial vehicles (UAVs). For example, a swarm of UAVs might be deployed to survey an uncharted region or to optimize geographic coverage to combat a forest fire.
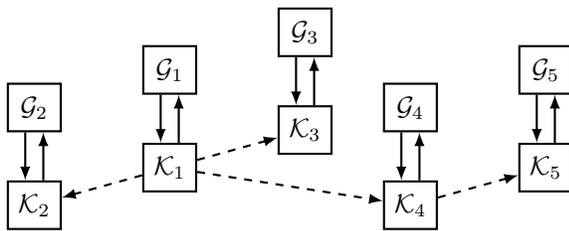

Figure 1: Example of a dynamically decoupled system. The plants $\mathcal{G}_i$ each have controllers $\mathcal{K}_i$ that share information instantaneously along a directed acyclic graph. The edge $\mathcal{K}_1 \rightarrow \mathcal{K}_5$ is implied but not shown.

In this paper, we abstract the communication network as a directed acyclic graph (DAG). An example with five

M. Kashyap and L. Lessard are with the Department of Electrical and Computer Engineering, at the University of Wisconsin–Madison, Madison, WI 53706, USA. L. Lessard is also with the Wisconsin Institute for Discovery at the University of Wisconsin–Madison. {mkashyap2,laurent.lessard}@wisc.edu

nodes is shown in Figure 1. Each plant $\mathcal{G}_i$ is assumed to obey state-space linear dynamics of the form

$$\dot{x}_i = A_{ii}x_i + B_{1_{ii}}w_i + B_{2_{ii}}u_i \tag{1a}$$

$$y_i = C_{2_{ii}}x_i + D_{21_{ii}}w_i, \tag{1b}$$

where $x_i(t) \in \mathbb{R}^{n_i}$, $u_i(t) \in \mathbb{R}^{m_i}$, and $y_i(t) \in \mathbb{R}^{p_i}$ are the state, input, and measurement for each agent, and $w_i(t) \in \mathbb{R}^{q_i}$ is exogenous standard Gaussian noise, assumed to be independent across agents and time.

The global cooperative task is to optimize a standard infinite-horizon quadratic objective that couples the states and inputs of all agents. Specifically, if we form the global state by aggregating agents $\{1, \dots, N\}$ as $x^{\mathsf{T}} := \begin{bmatrix} x_1^{\mathsf{T}} & \cdots & x_N^{\mathsf{T}} \end{bmatrix}$ and similarly for the aggregated u, y, and w, the cost function is of the form

$$\mathcal{J} = \lim_{T \to \infty} \frac{1}{T} \, \mathbf{E} \int_0^T \left\| C_1 x(t) + D_{12} u(t) \right\|^2 \mathrm{d}t, \tag{2}$$

where the expectation $\mathbf{E}$ is taken with respect to the exogenous noise w. If each controller $\mathcal{K}_i$ could measure the past history of all global measurements y, then finding the $\mathcal{K}_i$ that optimizes $\mathcal{J}$ in (2) would be a standard Linear Quadratic Gaussian (LQG) problem [20].

What makes our problem more difficult is that each $\mathcal{K}_i$ only has access to the measurement histories of its *ancestor* nodes in the DAG. In the example of Figure 1, $\mathcal{K}_1$ only measures y_1, $\mathcal{K}_3$ measures $\{y_1, y_3\}$, $\mathcal{K}_5$ measures $\{y_1, y_4, y_5\}$, etc. All measurements are instantaneous.

Contributions. We present an *agent-level* solution to the dynamically decoupled output-feedback problem. Specifically, we provide state-space realizations for the individual $\mathcal{K}_i$. Moreover, we provide intuitive interpretations for the controller states that highlight the role of the underlying graph structure and characterize the signals that should be transmitted between controllers. This work builds on [3, 4], which reported a state-space realization of the aggregated optimal controller but no agent-level details. Our results are presented in continuous time with an infinite-horizon cost but can easily be generalized to discrete time and/or a finite-horizon cost.

Literature review. The elegant solutions found for optimal centralized control do not generally carry over to the decentralized case. It was shown for example that under LQG assumptions, linear compensators can

be strictly suboptimal [18]. Several related decentralized control problems were shown to be intractable [1].

Tractability is closely related to how information is shared between the different controllers. For LQG problems with *partially nested* information, the optimal controller is linear [2]. When the control structure is *quadratically invariant* with respect to the plant, the problem of finding the optimal linear controller can be cast as a convex optimization over a set of stable Youla parameters [9, 10]. Under mild assumptions, convexity implies the existence of a unique optimal controller and efficient numerical methods for finding it. Numerical approaches include vectorization [11, 17] and Galerkin-style finite-dimensional approximation [8, 12].

The optimal decentralized controller often carries a rich structure, similar to the estimator-controller separation structure of classical centralized LQG control [19]. Such structure is lost in the aforementioned numerical approaches, which prompted recent efforts to solve these problems in an *explicit* way that uncovers the structure and intuition reminiscent of the centralized theory.

Explicit solutions for problems that are both partially nested and quadratically invariant have appeared for a variety of special cases. For example, state-feedback problems over general graphs where agent dynamics are coupled according to the graph structure [5, 14, 15]. Some output-feedback structures have also been solved, including two-player [7], broadcast [6], triangular [16], and dynamically decoupled [3, 4] cases. In the aforementioned works, explicit solutions typically consist of a state-space realization for the structured map $y \mapsto u$ between the global measurements y and the global inputs u.

In this paper, we look at the dynamically decoupled case and drill down even further. We provide an *agent-level* solution, which is an explicit state-space realization of the optimal controllers for each of the agents that is both minimal and easily interpretable. Section 2 covers notation, formal assumptions, and other preliminaries. Section 3 presents our alternative form for the optimal controller, followed by interpretations of the controller states and signals in Section 4 and the explicit agent-level structure in Section 5. We conclude in Section 6.

2 Preliminaries

Notation. A square matrix is *Hurwitz* if all eigenvalues have a strictly negative real part. $\mathcal{R}_p$ is the set of real rational proper transfer matrices and $\mathcal{RH}_\infty$ is the set of real-rational proper transfer matrices with no poles on the imaginary axis and no unstable poles. Similarly, $\mathcal{RH}_2$ is the set of real-rational strictly proper transfer matrices with no poles on the imaginary axis and no unstable poles. We use the state-space notation

$$\mathcal{G}(s) = \left[\begin{array}{c|c} A & B \\ \hline C & D \end{array} \right] = D + C(sI - A)^{-1}B.$$

If the realization above is stabilizable and detectable, then $\mathcal{G} \in \mathcal{RH}_\infty$ if and only if A is Hurwitz. The $\mathcal{H}_2$ norm of $\mathcal{G}$ is $\|\mathcal{G}\|_2^2 = \frac{1}{2\pi} \int_{-\infty}^{\infty} \mathrm{tr}(\mathcal{G}^*(jw)\mathcal{G}(jw)) \, dw$. All transfer matrices are represented using calligraphic symbols.

We use N to denote the total number of agents and $[N] := \{1, \ldots, N\}$. The global state dimension is $n := n_1 + \cdots + n_N$ and similarly for m and p. I_k will represent an identity matrix of size k. blkdiag$(\{X_i\})$ is the block-diagonal matrix formed by the blocks $\{X_1, \ldots, X_n\}$ and diag(X) is the block-diagonal matrix formed by the diagonal blocks of X. The zeros used throughout are matrix or vector zeros and their sizes are determined from context. The symbol $\star$ represents a possible nonzero value in a matrix with dimensions inferred from context.

For $i \in [N]$, we write $\underline{i} \subseteq [N]$ to denote the *descendants* of node i, i.e., the set of nodes j such that there is a directed path from i to j. Likewise, we write $\bar{i} \subseteq [N]$ to denote the *ancestors* of node i. Similarly, $\bar{\bar{i}}$ and $\underline{\underline{i}}$ denote the *strict ancestors* and *strict descendants*, respectively. For example, in the graph of Figure 1, $\underline{1} = [5]$, $\underline{4} = \{4, 5\}$, and $\bar{\bar{5}} = \{1, 4\}$. We also use this notation as a means of indexing matrices. For example, if X is a $N \times N$ block matrix, then we can write: $X_{2\underline{4}} = \begin{bmatrix} X_{24} & X_{25} \end{bmatrix}$.

Finally, we will require the use of specific partitions of the identity matrix. $I_n := \mathrm{blkdiag}(\{I_{n_i}\})$ and for each agent $i \in [N]$, we define $E_{n_i} := (I_n)_{:i}$ (the i^{th} block column of I_n). Similar to the descendant and ancestor definitions, $n_{\underline{i}} = \sum_{k \in \underline{i}} n_k$ and $n_{\bar{i}} = \sum_{k \in \bar{i}} n_k$. The dimensions of $E_{n_{\bar{i}}}$ and $E_{n_{\underline{i}}}$ are determined by the context of use. 1_n is the $n \times 1$ matrix of 1's. The definitions above also hold for other sets of indices, for example if we replace n by m or p.

Problem statement. Aggregating (1) across agents, we can form global dynamics where x, u, w, and y are the global state, input, exogenous noise, and measured output, respectively. We also define the regulated output $z = C_1 x + D_{12} u$. The transfer matrix $(w, u) \mapsto (z, y)$ is

$$\begin{bmatrix} \mathcal{U} & \mathcal{V} \\ \mathcal{W} & \mathcal{G} \end{bmatrix} := \left[\begin{array}{c|cc} A & B_1 & B_2 \\ \hline C_1 & 0 & D_{12} \\ C_2 & D_{21} & 0 \end{array} \right].$$

The dynamically decoupled structure of (1) implies that $A := \mathrm{blkdiag}(\{A_{ii}\})$ and similarly for B_1, B_2, C_2, D_{21}. The cost matrices C_1 and D_{12} are not assumed to have any special structure. It follows that $\mathcal{W}$ and $\mathcal{G}$ are block diagonal but $\mathcal{U}$ and $\mathcal{V}$ need not be.

The problem is to find the optimal controller $\mathcal{K}$ that solves the optimization problem

$$\begin{aligned} \underset{\mathcal{K} \text{ stabilizes } \mathcal{G}}{\text{minimize}} \quad & \left\| \mathcal{U} + \mathcal{V}\mathcal{K}(I - \mathcal{G}\mathcal{K})^{-1}\mathcal{W} \right\|_2^2 \\ \text{subject to} \quad & \mathcal{K} \in \mathbb{S}_{\mathcal{R}_p}. \end{aligned} \quad (3)$$

We assume the DAG characterizing the information-sharing pattern has an adjacency matrix $S \in \{0, 1\}^{N \times N}$, and $\mathbb{S}_{\mathcal{R}_p}$ is the set of proper matrix transfer functions with a block-sparsity pattern corresponding to S.

Namely, $\mathbb{S}_{\mathcal{R}_{\mathrm{p}}} := \{\mathcal{K} \in \mathcal{R}_{\mathrm{p}}^{m \times p} \mid \mathcal{K}_{ij} = 0 \text{ if } S_{ij} = 0\}$. Note that the sub-blocks of $\mathcal{K}$ satisfy $\mathcal{K}_{ij} \in \mathcal{R}_{\mathrm{p}}^{m_i \times p_j}$. The optimal controller in this case is known to be linear and finite-dimensional [3]. So there is no loss of generality in assuming $\mathcal{K} \in \mathbb{S}_{\mathcal{R}_{\mathrm{p}}}$ in (3). In the sequel, we make several simplifying assumptions regarding the adjacency matrix and the system matrices.

Assumption 1 (**Information sharing structure**). *Information is shared instantaneously between agent i and its descendants $\underline{i}$, so there is an implicit edge from i to each $j \in \underline{i}$ (the graph is transitively closed). Thus, any directed cycle can be treated as a single node and the graph becomes a DAG. By relabeling the nodes according to the parital ordering induced by the DAG, the adjacency matrix $S \in \mathbb{R}^{N \times N}$ becomes lower triangular. S is also invertible because each node has a self-loop.*

Assumption 1 is critical in ensuring that the problem is partially nested [2] and quadratically invariant [10, 11], so the optimal controller is linear and finite-dimensional.

Riccati assumptions. Matrices (A, B, C, D) are said to satisfy the *Riccati assumptions* [3] if:

R1. $C^{\mathsf{T}} D = 0$ and $D^{\mathsf{T}} D > 0$.

R2. (A, B) is stabilizable.

R3. $\begin{bmatrix} A - j\omega I & B \\ C & D \end{bmatrix}$ is full column rank for all $\omega \in \mathbb{R}$.

If the Riccati assumptions hold, there is a unique stabilizing solution to the associated algebraic Riccati equation, which we write as $(X, F) = \text{Ric}(A, B, C, D)$. Thus $X \succ 0$ satisfies $A^{\mathsf{T}} X + X A + C^{\mathsf{T}} C - X B (D^{\mathsf{T}} D)^{-1} B^{\mathsf{T}} X = 0$ with $A + BF$ Hurwitz, where $F := -(D^{\mathsf{T}} D)^{-1} B^{\mathsf{T}} X$.

Assumption 2 (**System assumptions**). *For the N interacting agents, we will assume the following.*

2.1. *As described in Section 2, the system is dynamically decoupled, $A := \text{blkdiag}(\{A_{ii}\})$ and similarly for B_1, B_2, C_2, and D_{21}. The cost matrices C_1 and D_{12} are not assumed to have any special structure.*

2.2. *A_{ii} is Hurwitz for all $i \in [N]$, i.e., A is Hurwitz.*

2.3. *The Riccati assumptions hold for (A, B_2, C_1, D_{12}) and for $(A_{ii}^{\mathsf{T}}, C_{2_{ii}}^{\mathsf{T}}, B_{1_{ii}}^{\mathsf{T}}, D_{21_{ii}}^{\mathsf{T}})$ for all $i \in [N]$.*

Assumption 2.2 is an assumption of nominal stability, carried over from Kim and Lall [3]. Assumption 2.3 ensures the necessary condition that the centralized LQR problem and the individual agents' estimation problems are non-degenerate.

3 Main Results

Kim and Lall [3] provide an explicit state-space formula for the controller that optimizes (3) under Assumptions 1–2. This will be the starting point for our work.

Lemma 3 ([3, Cor. 14]). *Suppose Assumptions 1–2 hold. The solution to (3) is $\mathcal{K} = \mathcal{P}(I + \mathcal{GP} - \text{diag}(\mathcal{GP}))^{-1}$, where the i^{th} column of $\mathcal{P} \in \mathbb{S}_{\mathcal{R}_{\mathrm{p}}}$ is given by*

$$\mathcal{P}_{\underline{i}i} = \left[\begin{array}{c|c} A_{\underline{ii}} + B_{2_{\underline{ii}}} F^i + \begin{bmatrix} L^i \\ 0 \end{bmatrix} C_{2_{i\underline{i}}} & \begin{bmatrix} -L^i \\ 0 \end{bmatrix} \\ \hline F^i & 0 \end{array} \right] \qquad (4)$$

and we define the control and estimation gains as

$$(X^i, F^i) := \text{Ric}(A_{\underline{ii}}, B_{2_{\underline{ii}}}, C_{1_{\cdot\underline{i}}}, D_{12_{\cdot\underline{i}}}) \qquad (5a)$$

$$(Y^i, L^{i\mathsf{T}}) := \text{Ric}(A_{ii}^{\mathsf{T}}, C_{2_{ii}}^{\mathsf{T}}, B_{1_{ii}}^{\mathsf{T}}, D_{21_{ii}}^{\mathsf{T}}). \qquad (5b)$$

Lemma 3 provides a state-space realization of the optimal controller. Our first main result is an even more explicit formula for the optimal controller $\mathcal{K}$ that highlights the duality between estimation and control. We will also see in Section 5 that this new formula is more readily implementable on an agent-level basis.

Theorem 4. *Suppose Assumptions 1–2 hold. A realization for the optimal controller that solves (3) is*

$$\mathcal{K} = \left[\begin{array}{c|c} \bar{A} + \bar{L}\bar{C} + \bar{B}\bar{S}_m \bar{F} \bar{S}_n^{-1} & -\bar{L}\bar{\mathbf{1}}_p \\ \hline \bar{\mathbf{1}}_m^{\mathsf{T}} \bar{F} \bar{S}_n^{-1} & 0 \end{array} \right], \qquad (6)$$

where we have defined the new symbols

$$\bar{A} := I_N \otimes A, \qquad \bar{B} := I_N \otimes B_2, \qquad \bar{C} := I_N \otimes C_2,$$
$$\bar{S}_m := S \otimes I_m, \qquad \bar{S}_n := S \otimes I_n, \qquad \bar{S}_p := S \otimes I_p,$$
$$\bar{\mathbf{1}}_m := 1_N \otimes I_m, \qquad \bar{\mathbf{1}}_p := 1_N \otimes I_p.$$

For all $i \in [N]$, with F^i and $L^{i\mathsf{T}}$ from (5), we also defined

$$L^{\bar{i}} := \text{blkdiag}(\{L^j\}_{j \in \bar{i}})$$
$$\bar{L} := \text{blkdiag}(\{L_i\}) \quad \text{where: } L_i := E_{n_i} L^{\bar{i}} E_{p_{\bar{i}}}^{\mathsf{T}}$$
$$\bar{F} := \text{blkdiag}(\{F_i\}) \quad \text{where: } F_i := E_{m_{\underline{i}}} F^i E_{n_{\underline{i}}}^{\mathsf{T}}.$$

Proof. See Appendix A for a complete proof. ∎

Remark 5. *The realization (6) of Theorem 4 is not minimal. Indeed, the gain matrices L_i and F_i have been padded with zeros so that they have compatible dimensions. We have $\bar{A} \in \mathbb{R}^{nN \times nN}$, where recall N is the number of agents and $n := n_1 + \cdots + n_N$ is the global aggregate state dimension of all agents. We use this non-minimal form to simplify algebraic manipulations; we present a generically minimal realization of (6) in Section 5.*

Applying the state transformation $x \mapsto \bar{S}_n x$ to (6), we obtain a dual realization for the optimal controller.

Corollary 6. *Suppose Assumptions 1–2 hold. A realization for the optimal controller that solves (3) is*

$$\mathcal{K} = \left[\begin{array}{c|c} \bar{A} + \bar{S}_n^{-1} \bar{L} \bar{S}_p \bar{C} + \bar{B}\bar{F} & -\bar{S}_n^{-1}\bar{L}\bar{\mathbf{1}}_p \\ \hline \bar{\mathbf{1}}_m^{\mathsf{T}} \bar{F} & 0 \end{array} \right], \qquad (7)$$

where the notation is the same as in Theorem 4.

4 Interpreting the optimal controller

State interpretation.   Based on Theorem 4, a realization of the controller (calling the state ξ) is

$$\dot{\xi} = (\bar{A} + \bar{L}\bar{C} + \bar{B}\bar{S}_m\bar{F}\bar{S}_n^{-1})\xi - \bar{L}\bar{\mathbf{1}}_p y \qquad (8a)$$

$$u = \bar{\mathbf{1}}_m^\mathsf{T}\bar{F}\bar{S}_n^{-1}\xi. \qquad (8b)$$

We call ξ the *state* coordinates. Informally, ξ_i can be interpreted as agent i's best estimate of the global state x given the available information $y_{\bar{i}}$. Likewise, a realization from Corollary 6 with state η is

$$\dot{\eta} = (\bar{A} + \bar{S}_n^{-1}\bar{L}\bar{S}_p\bar{C} + \bar{B}\bar{F})\eta - \bar{S}_n^{-1}\bar{L}\bar{\mathbf{1}}_p y \qquad (9a)$$

$$u = \bar{\mathbf{1}}_m^\mathsf{T}\bar{F}\eta. \qquad (9b)$$

We call η the *innovation* coordinates. Informally, η captures improvements in state estimates as information is aggregated along the DAG. A block-diagram of the global controller is shown in Figure 2.

Relationship to posets.   The work by Shah and Parrilo [13, 14] considers a different version of this decentralized control problem but obtains a similar optimal structure. The authors consider agents communicating on a DAG, except all agents use *state feedback* rather than output feedback, and A is not required to be block-diagonal; it may have a block-sparsity pattern conforming to that of the DAG. The resulting controller for this problem has a controller-estimator structure very similar to the one in Figure 2. What the state-feedback coupled case and the output-feedback decoupled case share in common is the notion of *separability* [3]. The more general output-feedback case with dynamic coupling is *not* separable and the control and estimation gains are generally coupled in that case [7]. Nevertheless, the fact that a state-feedback and output-feedback variants of the problem share a similar optimal controller structure suggests that there may exist an even broader class of control problems sharing this interesting structure.

5 Agent-level implementation

In this section, we further refine the global controller structure of Figure 2 to obtain agent-level implementations of the optimal controller.

The realizations (8) and (9) describe the optimal controller dynamics in terms of the global state ξ or η. Define the partial input $\nu := \bar{S}_m\bar{F}\bar{S}_n^{-1}\xi$, and further isolate the dynamics of each individual agent by writing the i^{th} block-component of (8) as

$$\dot{\xi}_i = (A + L_iC_2)\xi_i + B_2\nu_i - L_i y \qquad (10a)$$

$$u_i = E_{m_i}^\mathsf{T}\nu_i = \nu_{i,i} \qquad (10b)$$

$$\nu_i = \sum_{j\in\bar{i}} F_j(\bar{S}_n^{-1}\xi)_j. \qquad (10c)$$

Note that (10) can be implemented by agent i because updates only require information from ancestors

$\bar{i}$. Specifically, each ν_i is a function of $\xi_{\bar{i}}$ alone. Moreover, (10a) does not depend on the full y, but rather only $y_{\bar{i}}$, since $L_i := E_{n_{\bar{i}}}L^{\bar{i}}E_{p_{\bar{i}}}^\mathsf{T}$ (defined in Theorem 4).

We will now subdivide (10) to obtain agent-level update equations for each of the agents. It turns out agent i only needs to store $\xi_{i,\underline{i}}$, which are the components of ξ_i corresponding to the descendants of i.

By isolating these components, (10a) reduces to

$$\dot{\xi}_{i,\underline{i}} = A_{\underline{i}\underline{i}}\xi_{i,\underline{i}} + B_{2\underline{i}\underline{i}}\nu_{i,\underline{i}} - E_{n_{\underline{i}}}^\mathsf{T}E_{n_i}L^i(y_i - C_{2ii}\xi_{i,i}). \quad (11a)$$

Express (10c) in a more convenient form by expanding: $\nu_i = \sum_{j\in\bar{i}} F_j(\bar{S}_n^{-1}\xi)_j = \sum_{j\in\bar{i}} F_j \sum_{k\in\bar{j}} \bar{S}_n^{-1}(j,k)\xi_k$. Now define the summand $\tilde{\nu}_j := F_j \sum_{k\in\bar{j}} \bar{S}_n^{-1}(j,k)\xi_k$. Since $F_i := E_{m_i}F^iE_{n_{\underline{i}}}^\mathsf{T}$ (defined in Theorem 4), we can write ν_i compactly as: $\nu_i = \sum_{k\in\bar{i}} E_{m_k}\tilde{\nu}_{k,\underline{k}}$, which leads to the following expression for $\nu_{i,\underline{i}}$.

$$\nu_{i,\underline{i}} = \tilde{\nu}_{i,\underline{i}} + E_{m_{\underline{i}}}^\mathsf{T}\sum_{k\in\bar{\bar{i}}} E_{m_{\underline{k}}}\tilde{\nu}_{k,\underline{k}}. \qquad (11b)$$

Furthermore, each $\tilde{\nu}_{i,\underline{i}}$ is computable from the information available to agent i. To see why, compute:

$$\tilde{\nu}_{i,\underline{i}} = E_{m_{\underline{i}}}^\mathsf{T}F_i\sum_{k\in\bar{i}}\bar{S}_n^{-1}(i,k)\xi_k = F^i\sum_{k\in\bar{i}}\bar{S}_n^{-1}(i,k)\xi_{k,\underline{i}}$$

$$= F^i\sum_{k\in\bar{i}}\bar{S}_n^{-1}(i,k)E_{n_{\underline{i}}}^\mathsf{T}E_{n_{\underline{k}}}\xi_{k,\underline{k}}. \quad (11c)$$

In summary, the optimal agent-level controller for agent i maintains a state $\xi_{i,\underline{i}}$ and uses update equations (11) and (10b). Agent i receives $\{\xi_{k,\underline{k}}\}_{k\in\bar{\bar{i}}}$ and $\{\tilde{\nu}_{k,\underline{k}}\}_{k\in\bar{\bar{i}}}$ from its strict ancestors $\bar{\bar{i}}$, and transmits $\xi_{i,\underline{i}}$ and $\tilde{\nu}_{i,\underline{i}}$ to its strict descendants $\underline{\underline{i}}$. A block diagram of the agent-level structure is shown in Figure 3.

A similar structure can be derived for the coordinate choice of Corollary 6, which we omit.

Minimal realization.   As pointed out in Remark 5 and above, the realizations obtained in Theorem 4 and Corollary 6 are not minimal. The agent-level implementation encapsulates that each node i doesn't transmit information about the $[N]\setminus\underline{i}$ nodes. These correspond to the unobservable modes that can be removed. Therefore, the reduced realizations are:

$$\mathcal{K} = \left[\begin{array}{c|c} \bar{E}_{\underline{n}}^\mathsf{T}(\bar{A} + \bar{L}\bar{C} + \bar{B}\bar{S}_m\bar{F}\bar{S}_n^{-1})\bar{E}_{\underline{n}} & -\bar{E}_{\underline{n}}^\mathsf{T}\bar{L}\bar{\mathbf{1}}_p \\ \hline \bar{\mathbf{1}}_m^\mathsf{T}\bar{F}\bar{S}_n^{-1}\bar{E}_{\underline{n}} & 0 \end{array}\right]$$

$$\mathcal{K} = \left[\begin{array}{c|c} \bar{E}_{\underline{n}}^\mathsf{T}(\bar{A} + \bar{S}_n^{-1}\bar{L}\bar{S}_p\bar{C} + \bar{B}\bar{F})\bar{E}_{\underline{n}} & -\bar{E}_{\underline{n}}^\mathsf{T}\bar{S}_n^{-1}\bar{L}\bar{\mathbf{1}}_p \\ \hline \bar{\mathbf{1}}_m^\mathsf{T}\bar{F}\bar{E}_{\underline{n}} & 0 \end{array}\right],$$

where $\bar{E}_{\underline{n}} = \mathrm{blkdiag}(\{E_{n_{\underline{i}}}\})$. In both realizations above, the total number of states is $\sum_{i=1}^N n_{\underline{i}}$ and the state transition matrix has block-sparsity conforming to S.
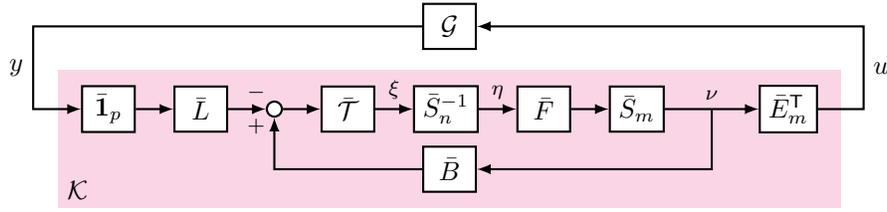
Figure 2: Block diagram representation of the global plant $\mathcal{G}$ and global controller $\mathcal{K}$ from Theorem 4. Here, the (block-diagonal) estimator dynamics are contained in $\bar{\mathcal{T}} := (sI - \bar{A} - \bar{L}\bar{C})^{-1}$ and the remaining blocks of $\mathcal{K}$ are static.
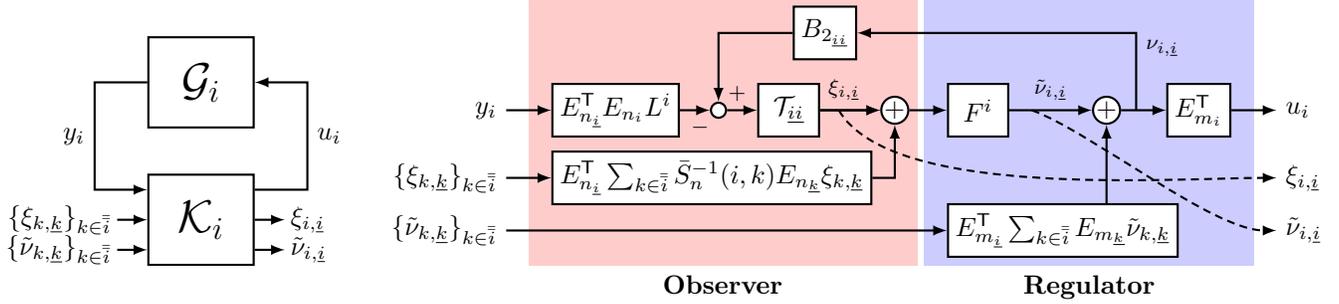

Figure 3: Block diagram representation of a local plant $\mathcal{G}_i$ and associated controller $\mathcal{K}_i$ (left) and the Observer-Regulator structure of $\mathcal{K}_i$ (right) from (11). Agent i updates its local state $\xi_{i,\underline{i}}$ and computes its local input u_i using the local measurement y_i and the states $\xi_{k,\underline{k}}$ and partial inputs $\tilde{v}_{k,\underline{k}}$ from strict ancestors $k \in \bar{\bar{i}}$. For simplicity, we used the following notation for the estimator dynamics: $\mathcal{T}_{\underline{ii}} := (sI - A_{\underline{ii}} - E_{n_{\underline{i}}}^{\mathsf{T}} E_{n_i} L^i C_{2_{ii}} E_{n_i}^{\mathsf{T}} E_{n_{\underline{i}}})^{-1}$.

6 Conclusions

In this work, our starting point was the explicit solution to the output-feedback dynamically decoupled cooperative control problem found by Kim and Lall [3]. We obtained an equivalent and compact realization of the optimal controller, which enabled us to derive agent-level implementations for the optimal control laws and to characterize precisely what each agent needs to receive, compute, and transmit.

A Proof of Theorem 4

In order to manipulate the expression for $\mathcal{K}$ in Lemma 3, we begin by finding a realization for $\mathcal{P}$ by aggregating the individual $\mathcal{P}_{ii}$ given in (4). For ease of exposition, we augment the states of each $\mathcal{P}_{ii}$ by adding modes that are simultaneously uncontrollable and unobservable. Specifically, we augment each A_{ii} to A and likewise with $B_{2_{ii}}$ to B_2 and $C_{2_{ii}}$ to C_2. For $i \in [N]$, we also augment the F^i by zero-padding to obtain $F_i := E_{m_{\underline{i}}} F^i E_{n_{\underline{i}}}^{\mathsf{T}}$. The estimation gains are similarly zero-padded to obtain $\tilde{L} := \mathrm{blkdiag}(\{E_{n_i} L^i E_{p_i}^{\mathsf{T}}\})$. So finally, we obtain

$$\mathcal{P} = \left[\begin{array}{c|c} \bar{A} + \bar{B}\bar{F} + \tilde{L}\bar{C} & -\tilde{L}\bar{1}_p \\ \hline \bar{1}_m^{\mathsf{T}}\bar{F} & 0 \end{array} \right]. \quad (12)$$

The input and output matrices, $-\tilde{L}\bar{1}_p$ and $\bar{1}_m^{\mathsf{T}}\bar{F}$ respectively, also arose by zero-padding so that all the modes we added to $\mathcal{P}$ are both uncontrollable and unobservable.

Now consider the related transfer matrix

$$\mathcal{P}_2 := \left[\begin{array}{c|c} \bar{A} + \bar{B}\bar{F} + L_x\bar{C} & -\tilde{L}\bar{1}_p \\ \hline \bar{1}_m^{\mathsf{T}}\bar{F} & 0 \end{array} \right], \quad (13)$$

where $L_x := \bar{S}_n^{-1}\bar{L}(\bar{S}_p - \bar{1}_p(\bar{1}_p^{\mathsf{T}} - \mathrm{diag}(\bar{1}_p^{\mathsf{T}})))$. We will now prove that $\mathcal{P}(s) = \mathcal{P}_2(s)$, which can be done by explicitly computing the transfer matrices and performing appropriate algebraic manipulations. Beginning with $\mathcal{P}_2$,

$$\mathcal{P}_2(s) = -\bar{1}_m^{\mathsf{T}}\bar{F}\big(s\bar{I}_n - \bar{A} - \bar{B}\bar{F} - \bar{S}_n^{-1}\bar{L}\bar{S}_p\bar{C}$$
$$+ \bar{S}_n^{-1}\bar{L}\bar{1}_p\bar{1}_p^{\mathsf{T}}\bar{C} - \bar{S}_n^{-1}\bar{L}\bar{1}_p\,\mathrm{diag}(\bar{1}_p^{\mathsf{T}})\bar{C}\big)^{-1}\tilde{L}\bar{1}_p, \quad (14)$$

where $\bar{I}_n := I_N \otimes I_n = I_{Nn}$. By the sparsity pattern of $\tilde{L}$ and $\bar{S}_n$, we can verify that $\bar{S}_n\tilde{L}\bar{1}_p = \bar{L}\bar{1}_p$. Therefore,

$$\mathcal{P}_2(s) = -\bar{1}_m^{\mathsf{T}}\bar{F}\big(s\bar{I}_n - \bar{A} - \bar{B}\bar{F} - \bar{S}_n^{-1}\bar{L}\bar{S}_p\bar{C}$$
$$+ \tilde{L}\bar{1}_p\bar{1}_p^{\mathsf{T}}\bar{C} - \tilde{L}\bar{1}_p\,\mathrm{diag}(\bar{1}_p^{\mathsf{T}})\bar{C}\big)^{-1}\tilde{L}\bar{1}_p.$$

Now, $\tilde{L}\bar{1}_p\,\mathrm{diag}(\bar{1}_p^{\mathsf{T}})\bar{C} = \tilde{L}\bar{C}$. Defining $Z_1 := (s\bar{I}_n - \bar{A} - \bar{B}\bar{F} - \tilde{L}\bar{C})$ and $Z_2 := -\bar{S}_n^{-1}\bar{L}\bar{S}_p\bar{C} + \tilde{L}\bar{1}_p\bar{1}_p^{\mathsf{T}}\bar{C}$,

$$\mathcal{P}_2(s) = -\bar{1}_m^{\mathsf{T}}\bar{F}\,(Z_1 + Z_2)^{-1}\tilde{L}\bar{1}_p.$$

Applying the Woodbury matrix identity,

$$\mathcal{P}_2(s) = -\bar{1}_m^{\mathsf{T}}\bar{F}Z_1^{-1}\tilde{L}\bar{1}_p + M = \mathcal{P}(s) + \mathcal{M}(s), \quad (15)$$

where $\mathcal{M}(s) := \bar{1}_m^{\mathsf{T}}\bar{F}Z_1^{-1}Z_2(\bar{I}_n + Z_1^{-1}Z_2)^{-1}Z_1^{-1}\tilde{L}\bar{1}_p$. The matrix Z_1^{-1} is block-diagonal, where each $Z_{1_{ii}}$ is also block-diagonal. Partitioning the rows and columns according to observable and unobservable modes, respectively, we obtain the following block-sparsity patterns:

5

$Z_1^{-1} \sim \left(\begin{smallmatrix} \star & 0 \\ 0 & \star \end{smallmatrix}\right)$, $Z_2 \sim \left(\begin{smallmatrix} 0 & \star \\ 0 & \star \end{smallmatrix}\right)$, $\tilde{L}\bar{1}_p \sim \left(\begin{smallmatrix} \star \\ 0 \end{smallmatrix}\right)$, and $\bar{1}_m^\mathsf{T}\bar{F} \sim \left(\begin{smallmatrix} \star & 0 \end{smallmatrix}\right)$. Substituting into the definition of $\mathcal{M}$, we conclude that $\mathcal{M}(s) = 0$ and so from (15), we have $\mathcal{P}_2(s) = \mathcal{P}(s)$.

Next, we substitute this expression for $\mathcal{P}_2$ into the expression for $\mathcal{K}$ in Lemma 3, one term at a time.

$$\tilde{\mathcal{P}} := \mathcal{P} - \operatorname{diag}(\mathcal{P}) = \left[\begin{array}{c|c} \bar{A} + \bar{B}\bar{F} + L_x\bar{C} & -\tilde{L}\bar{1}_p \\ \hline F_\Delta & 0 \end{array}\right],$$

where $F_\Delta := \bar{1}_m^\mathsf{T}\bar{F} - \operatorname{diag}(\bar{1}_m^\mathsf{T}\bar{F}) = (\bar{1}_m^\mathsf{T} - \operatorname{diag}(\bar{1}_m^\mathsf{T}))\bar{F}$. We evaluate, $(I + \mathcal{G}\mathcal{P} - \operatorname{diag}(\mathcal{G}\mathcal{P}))^{-1} = (I + \mathcal{G}\tilde{\mathcal{P}})^{-1} =$

$$\left[\begin{array}{cc|c} A & B_2 F_\Delta & 0 \\ \tilde{L}\bar{1}_p C_2 & \bar{A} + \bar{B}\bar{F} + L_x\bar{C} & -\tilde{L}\bar{1}_p \\ \hline -C_2 & 0 & I \end{array}\right].$$

Finally, $\mathcal{K} = \mathcal{P}(I + \mathcal{G}\tilde{\mathcal{P}})^{-1} =$

$$\left[\begin{array}{ccc|c} \bar{A} + \bar{B}\bar{F} + L_x\bar{C} & \tilde{L}\bar{1}_p C_2 & 0 & -\tilde{L}\bar{1}_p \\ 0 & A & B_2 F_\Delta & 0 \\ 0 & \tilde{L}\bar{1}_p C_2 & \bar{A} + \bar{B}\bar{F} + L_x\bar{C} & -\tilde{L}\bar{1}_p \\ \hline \bar{1}_m^\mathsf{T}\bar{F} & 0 & 0 & 0 \end{array}\right].$$

Applying the similarity transform with $T_1 := \left(\begin{smallmatrix} \bar{I}_n & 0 & \bar{I}_n \\ 0 & I_n & 0 \\ 0 & 0 & \bar{I}_n \end{smallmatrix}\right)$ reveals an uncontrollable mode, which we eliminate:

$$\mathcal{K} = \left[\begin{array}{cc|c} A & B_2 F_\Delta & 0 \\ \tilde{L}\bar{1}_p C_2 & \bar{A} + \bar{B}\bar{F} + L_x\bar{C} & -\tilde{L}\bar{1}_p \\ \hline 0 & \bar{1}_m^\mathsf{T}\bar{F} & 0 \end{array}\right].$$

Applying the similarity transform $T_2 := \left(\begin{smallmatrix} I_n & 0 \\ 0 & \bar{S}_n^{-1} \end{smallmatrix}\right)$ and using the relationships $\tilde{L}\bar{1}_p C_2 = \tilde{L}\bar{C}\bar{1}_n$, $\bar{S}_n\bar{A} = \bar{A}\bar{S}_n$, $\bar{S}_n\bar{B} = \bar{B}\bar{S}_m$, and $\bar{C}\bar{S}_n^{-1} = \bar{S}_p^{-1}\bar{C}$, we obtain

$$\mathcal{K} = \left[\begin{array}{cc|c} A & B_2(\bar{1}_m^\mathsf{T} - \operatorname{diag}(\bar{1}_m^\mathsf{T}))\bar{F}\bar{S}_n^{-1} & 0 \\ \bar{S}_n\tilde{L}\bar{C}\bar{1}_n & \bar{A} + \bar{B}\bar{S}_m\bar{F}\bar{S}_n^{-1} + \bar{S}_n L_x\bar{S}_p^{-1}\bar{C} & -\bar{S}_n\tilde{L}\bar{1}_p \\ \hline 0 & \bar{1}_m^\mathsf{T}\bar{F}\bar{S}_n^{-1} & 0 \end{array}\right].$$

Finally, we apply $T_3 := \left(\begin{smallmatrix} I_n & (\bar{1}_n^\mathsf{T} - \operatorname{diag}(\bar{1}_n^\mathsf{T}))\bar{S}_n^{-1} \\ 0 & \bar{I}_n \end{smallmatrix}\right)$ and use the relationship $\bar{S}_n\tilde{L}\bar{1}_p = \bar{L}\bar{1}_p$, eliminate the uncontrollable states, and obtain a reduced realization of the controller:

$$\mathcal{K} = \left[\begin{array}{c|c} \bar{A} + \bar{L}\bar{C} + \bar{B}\bar{S}_m\bar{F}\bar{S}_n^{-1} & -\bar{L}\bar{1}_p \\ \hline \bar{1}_m^\mathsf{T}\bar{F}\bar{S}_n^{-1} & 0 \end{array}\right]. \qquad \blacksquare$$

References

[1] V. D. Blondel and J. N. Tsitsiklis. A survey of computational complexity results in systems and control. *Automatica*, 36(9):1249–1274, 2000.

[2] Y.-C. Ho and K.-C. Chu. Team decision theory and information structures in optimal control problems—Part I. *IEEE Transactions on Automatic Control*, 17(1):15–22, 1972.

[3] J.-H. Kim and S. Lall. Explicit solutions to separable problems in optimal cooperative control. *IEEE Transactions on Automatic Control*, 60(5):1304–1319, 2015.

[4] J.-H. Kim, S. Lall, and C.-K. Ryoo. Optimal cooperative control of dynamically decoupled systems. In *IEEE Conference on Decision and Control*, pages 4852–4857, 2012.

[5] A. Lamperski and L. Lessard. Optimal decentralized state-feedback control with sparsity and delays. *Automatica*, 58:143–151, 2015.

[6] L. Lessard. Decentralized LQG control of systems with a broadcast architecture. In *IEEE Conference on Decision and Control*, pages 6241–6246, 2012.

[7] L. Lessard and S. Lall. Optimal control of two-player systems with output feedback. *IEEE Transactions on Automatic Control*, 60(8):2129–2144, 2015.

[8] X. Qi, M. V. Salapaka, P. G. Voulgaris, and M. Khammash. Structured optimal and robust control with multiple criteria: A convex solution. *IEEE Transactions on Automatic Control*, 49(10):1623–1640, 2004.

[9] M. Rotkowitz and S. Lall. Decentralized control information structures preserved under feedback. In *IEEE Conference on Decision and Control*, pages 569–575, 2002.

[10] M. Rotkowitz and S. Lall. A characterization of convex problems in decentralized control. *IEEE Transactions on Automatic Control*, 50(12):1984–1996, 2005.

[11] M. Rotkowitz and S. Lall. Convexification of optimal decentralized control without a stabilizing controller. In *International Symposium on Mathematical Theory of Networks and Systems*, pages 1496–1499, 2006.

[12] C. W. Scherer. Structured finite-dimensional controller design by convex optimization. *Linear Algebra and its Applications*, 351–352:639–669, 2002.

[13] P. Shah and P. A. Parrilo. An optimal controller architecture for poset-causal systems. In *IEEE Conference on Decision and Control*, pages 5522–5528, 2011.

[14] P. Shah and P. A. Parrilo. $\mathcal{H}_2$-optimal decentralized control over posets: A state-space solution for state-feedback. *IEEE Transactions on Automatic Control*, 58(12):3084–3096, 2013.

[15] J. Swigart and S. Lall. Optimal controller synthesis for decentralized systems over graphs via spectral factorization. *IEEE Transactions on Automatic Control*, 59(9):2311–2323, 2014.

[16] T. Tanaka and P. A. Parrilo. Optimal output feedback architecture for triangular LQG problems. In *American Control Conference*, pages 5730–5735, 2014.

[17] A. S. M. Vamsi and N. Elia. Optimal distributed controllers realizable over arbitrary networks. *IEEE Transactions on Automatic Control*, 61(1):129–144, 2016.

[18] H. S. Witsenhausen. A counterexample in stochastic optimum control. *SIAM Journal on Control*, 6(1):131–147, 1968.

[19] W. Wonham. On the separation theorem of stochastic control. *SIAM Journal on Control*, 6(2):312–326, 1968.

[20] K. Zhou, J. C. Doyle, and K. Glover. *Robust and optimal control*, volume 40. Prentice Hall, New Jersey, 1996.