

Agent-level optimal LQG control of dynamically decoupled systems with processing delays

Mruganka Kashyap

Laurent Lessard

To Appear, IEEE Conference on Decision and Control, 2020

Abstract

We consider the problem of controlling a set of dynamically decoupled plants where the plants' subcontrollers communicate with each other according to a fixed and known network topology. We assume the communication to be instantaneous but there is a fixed processing delay associated with incoming transmissions. We provide explicit closed-form expressions for the optimal decentralized controller under these communication constraints and using standard LQG assumptions for the plants and cost function. Although this problem is convex, it is challenging due to the irrationality of continuous-time delays and the decentralized information-sharing pattern. We show that the optimal subcontrollers each have an observer-regulator architecture containing LTI and FIR blocks and we characterize the signals that subcontrollers should transmit to each other across the network.

1 Introduction

When transmitting information across a network, latency can be caused either by *propagation delays*, which are due to the transmission medium and are proportional to the distance the signal must travel, or by *processing delays*, which are due to encoding, decoding, buffering, filtering, or other signal processing that must happen on either end of the transmission.

In scenarios where distances are relatively short, such as swarms of unmanned aerial vehicles (UAVs) communicating over a wireless network, propagation delays are negligible and it is reasonable to assume that latency is due entirely to processing delays. With UAVs, processing delays are often fixed and known, since they are dictated by the hardware capabilities of the UAVs and the bandwidth of the communication channel.

We consider the problem of controlling a set of N *dynamically decoupled* plants, which we refer to as *agents*. A four-agent example is depicted in Fig. 1.

M. Kashyap and L. Lessard were both with the University of Wisconsin–Madison, Madison, WI 53706, USA at the time of initial submission. M. Kashyap is now with the Department of Electrical and Computer Engineering, and L. Lessard is now with the Department of Mechanical and Industrial Engineering, both at Northeastern University, Boston, MA 02115, USA. {kashyap.mru, l.lessard}@northeastern.edu

This material is based upon work supported by the National Science Foundation under Grant No. 1710892.




Figure 1: Example of a dynamically decoupled system. The plants \mathcal{G}_i have corresponding controllers \mathcal{K}_i that share information along edges of a directed graph. Implied links are not shown, e.g., $\mathcal{K}_2 \rightarrow \mathcal{K}_1$ and $\mathcal{K}_3 \rightarrow \mathcal{K}_4$.

We make standard linear-quadratic-Gaussian (LQG) assumptions. That is, (1) the \mathcal{G}_i and \mathcal{K}_i are assumed to be continuous-time linear time-invariant (LTI) systems, (2) the exogenous noise is assumed to be Gaussian and uncorrelated between agents, and (3) the objective function is quadratic in the state and inputs. Let x_i denote the state of the agent i , and aggregate the agents' states into a global state vector x . Proceed similarly for the inputs u , measurements y , and exogenous noise w . We can represent the global dynamics compactly as in Fig. 2.




Figure 2: Classical four-block LQG problem with state-space representation of the plant. The goal is to design \mathcal{K} to minimize the \mathcal{H}_2 norm of the map from w to z . Blue matrices A have block-diagonal structure due to the dynamically decoupled assumptions.

Due to our dynamically decoupled assumption, the aggregated plant matrices A , B_1 , B_2 , C_2 , D_{21} are block-diagonal. We allow the quadratic cost to couple the states and inputs of various agents, so the cost matrices C_1 and D_{12} have no special structure.

We assume inter-controller communication happens instantaneously across the directed graph that represents the communication network (dashed lines in Fig. 1), but there is a *processing delay* of τ seconds for all incoming transmissions. We define the connectivity matrix $S \in \{0, 1\}^{N \times N}$ as $S_{ij} = 1$ if there is a directed path from controller j to controller i and $S_{ij} = 0$ otherwise. The

processing delays impose block structural constraints on \mathcal{K} . Specifically, \mathcal{K}_{ij} is: zero if $S_{ij} = 0$ (no communication), delayed by τ if $S_{ij} = 1$ and $i \neq j$ (processing delay), and unconstrained if $S_{ij} = 1$ and $i = j$ (no delay). For example, the connectivity and controller structures for the example of Fig. 1 are

$$S = \begin{bmatrix} 1 & 1 & 1 & 0 \\ 1 & 1 & 1 & 0 \\ 1 & 1 & 1 & 0 \\ 1 & 1 & 1 & 1 \end{bmatrix}, \quad \mathcal{K} \sim \begin{bmatrix} \mathcal{K}_{11} & e^{-s\tau}\mathcal{K}_{12} & e^{-s\tau}\mathcal{K}_{13} & 0 \\ e^{-s\tau}\mathcal{K}_{21} & \mathcal{K}_{22} & e^{-s\tau}\mathcal{K}_{23} & 0 \\ e^{-s\tau}\mathcal{K}_{31} & e^{-s\tau}\mathcal{K}_{32} & \mathcal{K}_{33} & 0 \\ e^{-s\tau}\mathcal{K}_{41} & e^{-s\tau}\mathcal{K}_{42} & e^{-s\tau}\mathcal{K}_{43} & \mathcal{K}_{44} \end{bmatrix}.$$

In this example, the controller \mathcal{K}_3 receives (delayed) measurements from \mathcal{G}_1 and \mathcal{G}_2 , instantaneous measurement from \mathcal{G}_3 , and no information from \mathcal{G}_4 . In other words,

$$u_3(t) = f_3(y_1(0:t-\tau), y_2(0:t-\tau), y_3(0:t))$$

for some function f_3 . In the remainder of the paper, we derive the optimal structured \mathcal{K} as a function of the connectivity S , delay τ , and state-space matrices of the four-block plant (Fig. 2). We also provide a detailed and intuitive *agent-level* description of the optimal \mathcal{K}_i .

Summary of relevant literature. The sparsity and delay constraints on our plant and controller are an example of a *quadratically invariant* (QI) system [15, 16]. Therefore, the problem of finding the optimal structured linear controller amounts to solving a convex (but infinite-dimensional) optimization problem.

When the structure does not contain delays, the exact optimal controller may be computed using vectorization [17, 20]. In general, vectorization does not produce a minimal realization, nor does it elucidate structure, such as a controller-estimator separation or an interpretation of the signals communicated between subcontrollers. Explicit solutions to special cases of output-feedback have been reported, such as: triangular [9, 19], broadcast [8], or dynamically decoupled [2, 3] cases.

When delays are present, we distinguish the cases of discrete vs. continuous-time delay. In discrete time, the delay operator z^{-1} is rational, so delays may be absorbed into the plant and the problem reduces to the non-delayed case [7]. In continuous time, this reduction is not possible because the delay operator $e^{-s\tau}$ is now irrational. One approach is to use a Padé approximation for the delay [21] followed by vectorization. Alternatively, a Galerkin-style finite-dimensional approximation [14, 18] can be used.

The present work seeks an explicit solution for the delayed case that provides structure and intuition without resorting to approximations or vectorization. Some special cases have been addressed in the literature. For centralized problems with a fixed loop delay (*dead time*), a loop-shifting technique involving FIR blocks can transform the problem into an equivalent LQG problem with a finite-dimensional LTI plant [10, 13]. A similar idea was used in the discrete-time case to decompose the structure into dead time and FIR components, which can be optimized separately [6].

The aforementioned loop-shifting approach was also extended to the *adobe delay* case, where the feedback loop contains both a delayed and a non-delayed path [11, 12]. This approach was used to obtain explicit optimal controllers for bilateral teleoperation problems, where the controllers communicate across a delayed channel [1, 4]. The idea was also generalized to haptic interfaces that have two-way communication with a shared virtual environment [5]. These are special cases of the general problem we will treat, described in Section 1, where S is the all-ones matrix.

Overview. Our work generalizes the agent-level solution for the non-delayed case reported in [2] by leveraging the loop-shifting decomposition for adobe delays from [12]. The paper is organized as follows. Section 2 covers our notation and assumptions. Section 3 presents our solution to the general problem as a function of the connectivity matrix S and processing delay τ . 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. We also characterize the signals that should be transmitted between subcontrollers and show that each subcontroller has an intuitive observer-regulator structure. Finally, Section 4 shows that we recover the expected limiting cases in the limits of zero or infinite processing delay.

2 Preliminaries

Notation. State-space notation for transfer functions:

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

We let N denote the total number of agents and $[N] := \{1, \dots, N\}$. The i^{th} subsystem has state dimension n_i , input dimension m_i , and measurement dimension p_i . The global state dimension is $n := n_1 + \dots + n_N$ and similarly for m and p . The identity matrix of size k is denoted I_k . We write $\text{blkdiag}(\{X_i\})$ to denote the block-diagonal matrix formed by the blocks $\{X_1, \dots, X_n\}$ and $\text{diag}(X)$ to denote 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.

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, $\bar{i} \subseteq [N]$ denotes the *ancestors* of node i . Similarly, $\bar{\bar{i}}$ and $\underline{\underline{i}}$ denote the *strict ancestors* and *strict descendants*, respectively. We also write $s^c = [N] \setminus s$ to denote the complement. For example, in the graph of Fig. 1, we have $\underline{2} = \{1, 2, 3, 4\}$, $\bar{2} = \{1, 3, 4\}$, and $\bar{\bar{2}} = \{1, 3\}$. We also use this notation to index matrices. For example, if X is a 4×4 block matrix associated with Fig. 1, then $X_{\underline{1}\bar{2}} = [X_{11} \quad X_{13}]$.

We will require the use of specific partitions of the identity matrix. We define $I_n := \text{blkdiag}(\{I_{n_i}\})$ 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 $\tilde{E}_{n_{\bar{i}}}$ and $\tilde{E}_{n_{\underline{i}}}$ are determined by the context of use. Finally, 1_n is the $n \times 1$ matrix of 1's.

Processing delay notation [12]. The *adobe delay* matrix $\Lambda_m^i := \text{blkdiag}(I_{m_i}, e^{-s\tau} I_{m_{\bar{i}}})$ leaves block i unchanged and imposes a delay of τ on all strict descendants of i . The completion operator $\pi_\tau\{\cdot\}$ acts on a state-space system delayed by τ and returns the FIR system with support on $[0, \tau]$ that completes it:

$$\pi_\tau \left\{ \left[\begin{array}{c|c} A & B \\ \hline C & D \end{array} \right] e^{-s\tau} \right\} := \left[\begin{array}{c|c} A & e^{-A\tau} B \\ \hline C & D \end{array} \right] - \left[\begin{array}{c|c} A & B \\ \hline C & D \end{array} \right] e^{-s\tau}.$$

We also define the function $\Gamma : (\Omega, \Lambda_m^i) \mapsto (\tilde{\Omega}, \Pi_u, \Pi_b)$, which maps a four-block plant Ω defined as

$$\Omega := \left[\begin{array}{cc|cc} \mathcal{U} & \mathcal{V} & A & B_1 & B_2 \\ \mathcal{W} & \mathcal{G} & C_1 & 0 & D_{12} \\ & & C_2 & D_{21} & 0 \end{array} \right] \quad (1)$$

and adobe delay matrix Λ_m^i to a modified plant $\tilde{\Omega}$ and FIR systems Π_u and Π_b . See Appendix A for details.

Problem statement. Given a four-block plant (1) and associated connectivity matrix S as described in Section 1 (refer to Fig. 2), the plants are dynamically decoupled but the cost function may couple the states and inputs of the different agents. It follows that \mathcal{W} and \mathcal{G} are block diagonal but \mathcal{U} and \mathcal{V} need not. Let \mathbb{S}_τ be the set of causal structured LTI controllers with processing delay τ , as described in Section 1. The problem is to find \mathcal{K} to

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

We now describe and explain the technical assumptions we make on the four-block plant Ω and connectivity S .

Riccati assumptions. Four matrices (A, B, C, D) are said to satisfy the *Riccati assumptions* [3, 12] if:

R1. $C^\top D = 0$ and $D^\top D > 0$.

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 $X \succeq 0$ satisfies $A^\top X + XA + C^\top C - XB(D^\top D)^{-1}B^\top X = 0$ with $A + BF$ Hurwitz, where $F := -(D^\top D)^{-1}B^\top X$.

Assumption 1 (System assumptions). *For the N interacting agents, we will assume the following.*

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

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

1.3. $D_{12_{\bar{i}}}^\top D_{12_{\bar{i}}} = I$ and $D_{21_{ii}} D_{21_{ii}}^\top = I$ for all $i \in [N]$.

Assumption 1.1 is an assumption of nominal stability, carried over from Kim et al. [3]. Assumption 1.2 ensures the necessary condition that the centralized LQR problem and the individual agents' estimation problems are non-degenerate. Assumption 1.3 simplifies the exposition of the delayed problem [11, 12], though the results still hold for the general case [11, Rem. 3.2].

3 Main Result

Our main result is an agent-level description of the optimal controller that solves (2). That is, we provide explicit state-space formulas for each \mathcal{K}_i and describe which signals should be transmitted and received between agents.

Theorem 2. *Consider a general instance of the structured optimal control problem described in Section 1 and suppose Assumption 1 holds. Consider the four-block sub-plant for agent i and its descendants:*

$$\Omega_i := \left[\begin{array}{cc|cc} \mathcal{U}_{:i} & \mathcal{V}_{:i} & A_{ii} & B_{1ii} & B_{2ii} \\ \mathcal{W}_{ii} & \mathcal{G}_{ii} & C_{1ii} & 0 & D_{12ii} \\ & & C_{2ii} & D_{21ii} & 0 \end{array} \right].$$

Now apply $(\tilde{\Omega}_i, \Pi_{u_i}, \Pi_{b_i}) = \Gamma(\Omega_i, \Lambda_m^i)$. Define the estimation gain \tilde{F}^i and control gain L^i by

$$(\tilde{X}^i, \tilde{F}^i) := \text{Ric}(A_{ii}, \tilde{B}_{2ii}, \tilde{C}_{1ii}, D_{12ii}), \quad (3a)$$

$$(Y^i, L^{i\top}) := \text{Ric}(A_{ii}^\top, C_{2ii}^\top, B_{1ii}^\top, D_{21ii}^\top). \quad (3b)$$

A realization of the optimal controller \mathcal{K}_i for agent i that solves (2) is given by the state-space equations

$$\begin{aligned} \dot{\eta}_{i,\underline{i}} &= (A_{ii} + (LC)^i) \eta_{i,\underline{i}} + (LC)^i \sum_{k \in \bar{i}} \eta_{k,\underline{i}}(t - \tau) \\ &\quad + \tilde{B}_{2ii} \tilde{\nu}_{i,\underline{i}} - E_{n_{\bar{i}}}^\top E_{n_i} L^i (y_i + \Pi_{b_i} \nu_{i,\underline{i}}), \end{aligned} \quad (4a)$$

$$\tilde{\nu}_{i,\underline{i}} = \tilde{F}^i \eta_{i,\underline{i}}, \quad (4b)$$

$$u_i = E_{m_i}^\top E_{m_{\bar{i}}} \left(\nu_{i,\underline{i}} + \sum_{k \in \bar{i}} \tilde{\nu}_{k,\underline{i}}(t - \tau) \right), \quad (4c)$$

where $(LC)^i := E_{n_{\bar{i}}}^\top E_{n_i} L^i C_{2ii} E_{n_i}^\top E_{n_{\bar{i}}}$ and $\nu_{i,\underline{i}} := \Pi_{u_i} \tilde{\nu}_{i,\underline{i}}$.

Proof. See Appendix B. ■

We use a slight abuse of notation in (4). These are time-domain equations but Π_{b_i} and Π_{u_i} are FIR transfer matrices. Products such as $\Pi_{b_i} \nu_{i,\underline{i}}$ are to be interpreted as the time-domain output of $\Pi_{b_i}(s)$ with input $\nu_{i,\underline{i}}(t)$.

Fig. 3 shows the signals received and transmitted by each subcontroller. Fig. 4 shows a more detailed block diagram. The optimal controller is an interconnection of state-space systems, FIR blocks, and delays.




Figure 3: Block diagram representation of a local plant \mathcal{G}_i and associated controller \mathcal{K}_i (dashed red box). Controller i receives local measurements y_i and information (η, ν) from its strict ancestors \bar{i} (which incurs a processing delay $e^{-s\tau}$). The controller then computes and transmits information to its strict descendants \tilde{i} .

4 Limiting cases

We now study the behavior of our optimal controller as the processing delay τ varies. Setting $\tau = 0$, nodes belonging to the same connected component can be treated as a single node due to the instantaneous communication assumption. For example, if $\tau = 0$ in Fig. 1, nodes 1, 2, and 3 become a single node, and we are left with the equivalent two-node graph $S = \begin{pmatrix} 1 & 0 \\ 1 & 1 \end{pmatrix}$. This special case was solved in [2] and we recover this result by setting $\tau = 0$ in Theorem 2. At the other extreme, the case $\tau \rightarrow \infty$ leads to a controllers that do not communicate at all. These results are formalized in Corollary 3 below.

Corollary 3. *Consider the setting of Theorem 2.*

- 1) *If there is no processing delay ($\tau = 0$), The controller \mathcal{K}_0 that solves (2) is*

$$\dot{\eta}_{i,\underline{i}} = A_{\underline{i}\underline{i}}\eta_{i,\underline{i}} + (LC)^i \sum_{k \in \bar{i}} \eta_{k,\underline{i}} + B_{2\underline{i}\underline{i}}\nu_{i,\underline{i}} - E_{n\underline{i}}^\top E_{n_i} L^i y_i$$

$$u_i = E_{m_i}^\top \sum_{j \in \bar{i}} F_j \eta_j,$$

where $\nu_{i,\underline{i}} := F^i \eta_{i,\underline{i}}$ and $F_j := E_{m_j} F^j E_{n_j}^\top$ and the remaining parameters are defined in Theorem 2.

- 2) *If there is infinite processing delay ($\tau \rightarrow \infty$), the controller \mathcal{K}_∞ that solves (2) is*

$$\dot{\eta}_{i,i} = A_{ii}\eta_{i,i} + L^i C_{2ii}\eta_{i,i} + B_{2ii}u_i - L^i y_i \quad (6a)$$

$$u_i = F_\infty^i \eta_{i,i}, \quad (6b)$$

where $(X_\infty^i, F_\infty^i) := \text{Ric}(A_{ii}, B_{2ii}, C_{1,i}, D_{12,i})$.

Proof. When $\tau \rightarrow 0$, we have $\Pi_{u_i} = I$ and $\Pi_{b_i} = 0$. Moreover, $\tilde{\Omega}_i = \Omega_i$ and $\tilde{F}^i = F^i$. When $\tau \rightarrow \infty$, we also have $\Pi_{u_i} \rightarrow I$ and $\Pi_{b_i} \rightarrow 0$ but this time, $\tilde{B}_{2\underline{i}\underline{i}} \rightarrow \text{diag}(B_{2\underline{i}\underline{i}})$. The corresponding control gain \tilde{F}^i reduces to an augmented nominal gain $E_{m_i} F_\infty^i E_{n_i}^\top$ after elimination of the uncontrollable (\underline{i}) and unobservable (\underline{i}^c) modes to obtain (6). Letting $\tau \rightarrow \infty$ is equivalent to solving N separate LQG problems (the global cost matrices are block-diagonal). ■

The effect of changing the processing delay τ is illustrated in the block diagram of Fig. 4. When τ varies, only the green blocks are affected. The remaining parts of the controller can be designed without knowing τ .

5 Conclusion

In this work, we considered an optimal control problem where dynamically decoupled agents communicate over a network and incur processing delays to receive transmissions from neighboring controllers. We described the structure of the optimal control strategies at the level of individual agents (Fig. 4), which shows which signals should be transmitted and which parts of the controller depend on the processing delay.

A Definition of the Γ function

The input matrices B_2 and D_{12} of Ω defined in (1) are partitioned according to the blocks of Λ_m^i . Therefore, $B_2 = [B_{2_0} \ B_{2_\tau}]$, where the two blocks have delays of 0 and τ , respectively. We partition $D_{12} = [D_{12_0} \ D_{12_\tau}]$ similarly. Now define the Hamiltonian

$$H = \begin{bmatrix} H_{11} & H_{12} \\ H_{21} & H_{22} \end{bmatrix} := \begin{bmatrix} A - B_{2_0} D_{12_0}^\top C_1 & -B_{2_0} B_{2_0}^\top \\ -C_1^\top P_\tau C_1 & -A^\top + C_1^\top D_{12_0} B_{2_0}^\top \end{bmatrix},$$

where $P_0 := D_{12_0} D_{12_0}^\top$ and $P_\tau := I - P_0$, and define its symplectic matrix exponential as $\Sigma := e^{H\tau}$. Define the modified matrices corresponding to B_2 and C_1 as

$$\begin{aligned} \tilde{B}_2 &:= [B_{2_0} \ \Sigma_{12}^\top C_1^\top D_{12_\tau} + \Sigma_{22}^\top B_{2_\tau}] \\ \tilde{C}_1 &:= (P_\tau C_1 + P_0 C_1 \Sigma_{22}^\top - D_{12_0} B_{2_0}^\top \Sigma_{21}^\top) \Sigma_{22}^{-\top}, \end{aligned}$$

where the Σ_{ij} are partitioned the same was as the H_{ij} . The modified four-block plant output by Γ is then

$$\tilde{\Omega} := \begin{bmatrix} \tilde{U} & \tilde{V} \\ \tilde{W} & \tilde{G} \end{bmatrix} := \left[\begin{array}{c|cc} A & B_1 & \tilde{B}_2 \\ \tilde{C}_1 & 0 & D_{12} \\ \hline C_2 & D_{21} & 0 \end{array} \right], \quad (7)$$

where $\tilde{B}_2 = [B_{2_0} \ \tilde{B}_{2_\tau}]$ is partitioned the same way as the original B_2 . Finally, define the FIR systems

$$\begin{bmatrix} \tilde{\Pi}_u \\ \tilde{\Pi}_b \end{bmatrix} := \pi_\tau \left\{ \left[\begin{array}{cc|c} H_{11} & H_{12} & B_{2_\tau} \\ H_{21} & H_{22} & -C_1^\top D_{12_\tau} \\ \hline D_{12_0}^\top C_1 & B_{2_0}^\top & 0 \\ C_2 & 0 & 0 \end{array} \right] e^{-s\tau} \right\}$$

and Γ outputs $\Pi_u := \begin{bmatrix} I & \tilde{\Pi}_u \\ 0 & I \end{bmatrix}$ and $\Pi_b := [0 \ \tilde{\Pi}_b]$.

B Proof of Theorem 2

A roadmap for the proof is shown in Fig. 5. We begin in the top left block, which is (2), then we follow the arrows, each of which is explained in the following paragraphs.

Step 1. Assumption 1.1 together with quadratic invariance [15, 16] ensures that all stabilizing controllers are

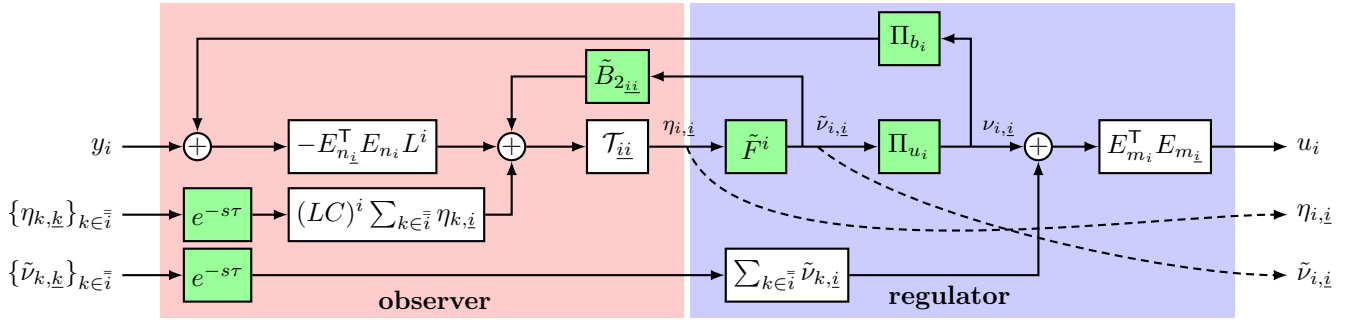


Figure 4: Expanded diagram of the optimal controller shown in Fig. 3 that reveals the observer-regulator structure of \mathcal{K}_i from Theorem 2. The green blocks represent blocks that depend on the processing delay τ . Agent i updates its local innovation $\eta_{i,i}$ and computes its local input u_i using the local measurement y_i and the delayed innovations $\eta_{k,k}$ and partial inputs $\tilde{v}_{k,k}$ from strict ancestors $k \in \bar{i}$. For simplicity, we used the following notation for the observer dynamics: $\mathcal{T}_{ii} := (sI - A_{ii} - (LC)^i)^{-1}$, where $(LC)^i := E_{n_i}^T E_{n_i} L^i C_{2ii} E_{n_i} E_{n_i}^T$.

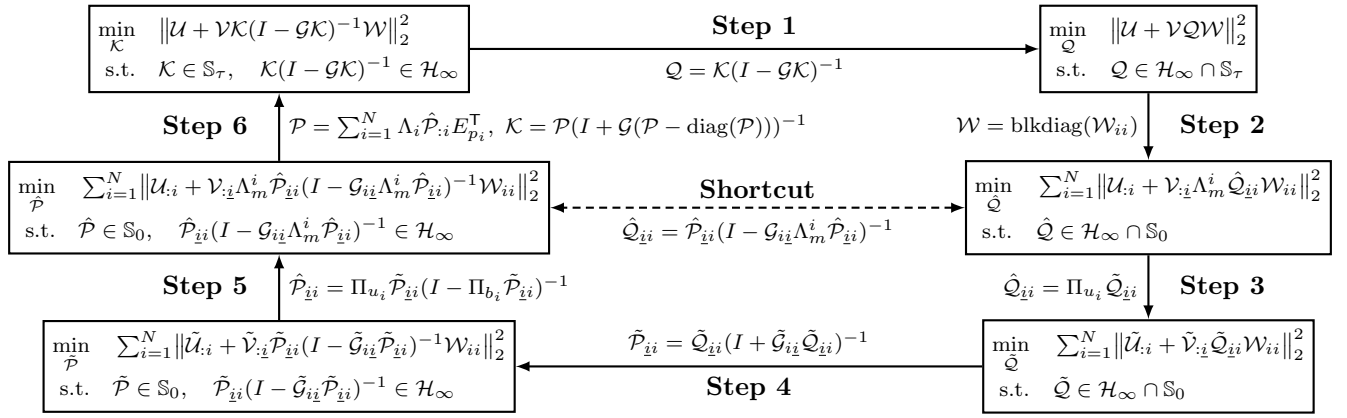


Figure 5: Sequence of transformations used in the proof of Theorem 2. The proof starts in the upper-left box, which is our original structured optimization problem (2) and follows the arrows clockwise.

parameterized by $\mathcal{Q} = \mathcal{K}(I - \mathcal{G}\mathcal{K})^{-1}$, where the parameter \mathcal{Q} is stable and has the same structure as \mathcal{K} . This leads to a convex model-matching problem in \mathcal{Q} .

Step 2. Note that \mathcal{W} is block-diagonal and $\mathcal{Q}_{ii} = \Lambda_m^i \hat{\mathcal{Q}}_{ii}$, where Λ_m^i is an adobe delay matrix (defined in Section 2), and $\hat{\mathcal{Q}}$ is delay-free. Thus, we may separate the cost by its block columns and optimize each summand separately. A similar separation technique was leveraged in [2, 3] to solve the non-delayed version of this problem.

Step 3. We now apply the loop-shifting result of et al. [12, Thm. 1], which we state below as Lemma 4. This result transforms an optimal control problem containing an adobe delay Λ_m^i to one with no delays. The original controller can then be recovered via a transformation involving FIR blocks.

Lemma 4 ([12, Thm. 1]). *Consider the following structured optimal control problem subject to Assumption 1 similar to (2), but with a simpler structure: there are two control inputs, and the second one is delayed by τ . In other words, we would like to solve*

$$\underset{\Lambda_m \mathcal{K} \text{ stabilizes } \Omega}{\text{minimize}} \quad \|\mathcal{U} + \mathcal{V}\Lambda_m \mathcal{K}(I - \mathcal{G}\Lambda_m \mathcal{K})^{-1}\mathcal{W}\|_2^2$$

where $\Lambda_m := \text{blkdiag}(I, e^{-s\tau}I)$. Let $(\tilde{\Omega}, \Pi_u, \Pi_b) = \Gamma(\Omega, \Lambda_m)$, where Γ is defined in Section 2 and Appendix A. Then, $\Lambda_m \mathcal{K}$ stabilizes Ω if and only if $\tilde{\mathcal{K}}$ stabilizes $\tilde{\Omega}$, where \mathcal{K} and $\tilde{\mathcal{K}}$ are related via the bijective transformation $\mathcal{K} = \Pi_u \tilde{\mathcal{K}}(I - \Pi_b \tilde{\mathcal{K}})^{-1}$.

Transfer matrices $\tilde{\mathcal{U}}, \tilde{\mathcal{V}}, \mathcal{W}$ are rational with realizations given in (7), so we have N separate standard model-matching problems [22, §14.5] in $\tilde{\mathcal{Q}}_{ii}$, with solutions

$$\tilde{\mathcal{Q}}_{ii} = \begin{bmatrix} A_{ii} + \tilde{B}_{2ii} \tilde{F}^i & -\tilde{B}_{2ii} \tilde{F}^i & 0 \\ 0 & A_{ii} + (LC)^i & E_{n_i}^T E_{n_i} L^i \\ \tilde{F}^i & -\tilde{F}^i & 0 \end{bmatrix}, \quad (8)$$

where \tilde{F}^i and L^i are defined in (3).

The rest of the proof consists of algebraic substitutions and simplifications to transform $\tilde{\mathcal{Q}}_{ii}$ back into \mathcal{K} .

Step 4. Applying the inverse transformation from Step 1 to each sub-problem, we obtain N separate control problems in the variables $\tilde{\mathcal{P}}_{ii} = \tilde{\mathcal{Q}}_{ii}(I + \tilde{\mathcal{G}}_{ii} \tilde{\mathcal{Q}}_{ii})^{-1}$.

Step 5. Given the solutions $\tilde{\mathcal{P}}_{ii}$, we can invert the transformation in Lemma 4 to obtain the adobe-delayed controller $\hat{\mathcal{P}}_{ii} = \Pi_{u_i} \tilde{\mathcal{P}}_{ii}(I - \Pi_{b_i} \tilde{\mathcal{P}}_{ii})^{-1}$.

Step 6. By comparing terms, we have (see shortcut arrow in Fig. 5): $Q_{ii} = \Lambda_m^i \hat{Q}_{ii} = \Lambda_m^i \hat{P}_{ii} (I - \mathcal{G}_{ii} \Lambda_m^i \mathcal{P}_{ii})^{-1}$. Inverting this equation yields $\Lambda_m^i \hat{P}_{ii} = Q_{ii} (I + \mathcal{G}_{ii} Q_{ii})^{-1}$. Now zero-pad and horizontally concatenate the $\Lambda_m^i \hat{P}_{ii}$ to obtain $\mathcal{P} = \sum_{i=1}^N \Lambda_i \hat{P}_{ii} E_{p_i}^T = Q(\text{diag}(I + \mathcal{G}Q))^{-1}$. Together with $Q = \mathcal{K}(I - \mathcal{G}\mathcal{K})^{-1}$ from Step 1, eliminate Q to obtain (see also [3, Lemma 9])

$$\mathcal{K} = \mathcal{P}(I + \mathcal{G}(\mathcal{P} - \text{diag}(\mathcal{P})))^{-1}. \quad (9)$$

Carrying the \tilde{Q}_{ii} found in (8) through the transformations in steps 4–6, and simplifying, we obtain

$$\mathcal{P} = -\bar{\mathbf{1}}_m^T \bar{\Lambda} \bar{\Pi}_u \tilde{F} (sI - \bar{A} - \bar{L}\bar{C} - \tilde{B}\tilde{F} + \bar{L}\bar{\Pi}_b \tilde{F})^{-1} \bar{L} \bar{\mathbf{1}}_p.$$

Here, $\bar{\Pi}_u$, $\bar{\Pi}_b$, $\bar{\Lambda}$, \tilde{F} , \bar{L} are block-diagonal matrices (zero-padded if necessary) made from $\{\Pi_{u_i}\}$, $\{\Pi_{b_i}\}$, $\{\Lambda_m^i\}$, $\{\tilde{F}^i\}$, $\{L^i\}$, respectively, and $\bar{\mathbf{1}}_n := \mathbf{1}_N \otimes I_n$, where \otimes denotes the Kronecker product. To compute \mathcal{K} and extract subcontrollers \mathcal{K}_i , we follow similar steps to the non-delayed case [2, Thm. 4]. Substituting the expression for \mathcal{P} into (9) and further simplifying, we obtain

$$\mathcal{K} = -\bar{\mathbf{1}}_m^T \bar{\Lambda} \bar{\Pi}_u \tilde{F} (sI - \bar{A} - \bar{L}\bar{C} - \tilde{B}\tilde{F} + \bar{L}\bar{\Pi}_b \tilde{F} - \bar{L} \bar{\mathbf{1}}_p C_2 (sI - A)^{-1} B_2 \bar{\mathbf{1}}_{\Delta m}^T \bar{\Lambda} \bar{\Pi}_u \tilde{F})^{-1} \bar{L} \bar{\mathbf{1}}_p, \quad (10)$$

where $\bar{\mathbf{1}}_{\Delta m} := \bar{\mathbf{1}}_m - \text{diag}(\bar{\mathbf{1}}_m)$. Define the controller state variable η and input variable u as follows:

$$\begin{aligned} \eta &= -(sI - \bar{A} - \bar{L}\bar{C} - \tilde{B}\tilde{F} + \bar{L}\bar{\Pi}_b \tilde{F} \\ &\quad - \bar{L} \bar{\mathbf{1}}_p C_2 (sI - A)^{-1} B_2 \bar{\mathbf{1}}_{\Delta m}^T \bar{\Lambda} \bar{\Pi}_u \tilde{F})^{-1} \bar{L} \bar{\mathbf{1}}_p y, \\ u &= \bar{\mathbf{1}}_m^T \bar{\Lambda} \bar{\Pi}_u \tilde{F} \eta. \end{aligned}$$

Further simplify using the identities $\bar{L} \bar{\mathbf{1}}_p C_2 = \bar{L} \bar{C} \bar{\mathbf{1}}_n$ and $(sI - A)^{-1} B_2 \bar{\mathbf{1}}_{\Delta m}^T \bar{\Lambda} = \bar{\mathbf{1}}_{\Delta n}^T \bar{\Lambda} (sI - \bar{A})^{-1} \tilde{B}$ and split the state equations into their agent-level components η_i , and we obtain the reduced realization (4), as required. ■

References

- [1] J. H. Cho and M. Kristalny. On the H^2 decentralized controller synthesis for delayed bilateral teleoperation systems. *IFAC Proceedings*, 45(22):393–398, 2012.
- [2] M. Kashyap and L. Lessard. Explicit agent-level optimal cooperative controllers for dynamically decoupled systems with output feedback. In *IEEE Conference on Decision and Control*, pages 8254–8259, 2019.
- [3] J.-H. Kim and S. Lall. Explicit solutions to separable problems in optimal cooperative control. *IEEE Trans. Autom. Control*, 60(5):1304–1319, 2015.
- [4] M. Kristalny and J. H. Cho. On the decentralized H_2 optimal control of bilateral teleoperation systems with time delays. In *IEEE Conference on Decision and Control*, pages 6908–6914, 2012.
- [5] M. Kristalny and J. H. Cho. Decentralized h^2 optimal control of haptic interfaces for a shared virtual environment. In *IEEE Conference on Decision and Control*, pages 5204–5209, 2013.
- [6] A. Lamperski and J. C. Doyle. The \mathcal{H}_2 control problem for quadratically invariant systems with delays. *IEEE Trans. Autom. Control*, 60(7):1945–1950, 2015.
- [7] A. Lamperski and L. Lessard. Optimal decentralized state-feedback control with sparsity and delays. *Automatica*, 58:143–151, 2015.
- [8] L. Lessard. Decentralized LQG control of systems with a broadcast architecture. In *IEEE Conference on Decision and Control*, pages 6241–6246, 2012.
- [9] L. Lessard and S. Lall. Optimal control of two-player systems with output feedback. *IEEE Trans. Autom. Control*, 60(8):2129–2144, 2015.
- [10] L. Mirkin. On the extraction of dead-time controllers and estimators from delay-free parametrizations. *IEEE Trans. Autom. Control*, 48(4):543–553, 2003.
- [11] L. Mirkin, Z. J. Palmor, and D. Shneiderman. Dead-time compensation for systems with multiple i/o delays: A loop-shifting approach. *IEEE Trans. Autom. Control*, 56(11):2542–2554, 2011.
- [12] L. Mirkin, Z. J. Palmor, and D. Shneiderman. H^2 optimization for systems with adobe input delays: A loop shifting approach. *Automatica*, 48(8):1722–1728, 2012.
- [13] L. Mirkin and N. Raskin. Every stabilizing dead-time controller has an observer–predictor–based structure. *Automatica*, 39(10):1747–1754, 2003.
- [14] X. Qi, M. V. Salapaka, P. G. Voulgaris, and M. Khammash. Structured optimal and robust control with multiple criteria: A convex solution. *IEEE Trans. Autom. Control*, 49(10):1623–1640, 2004.
- [15] M. Rotkowitz, R. Cogill, and S. Lall. Convexity of optimal control over networks with delays and arbitrary topology. *International Journal of Systems, Control and Communications*, 2(1-3):30–54, 2010.
- [16] M. Rotkowitz and S. Lall. A characterization of convex problems in decentralized control. *IEEE Trans. Autom. Control*, 50(12):1984–1996, 2005.
- [17] 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.
- [18] C. W. Scherer. Structured finite-dimensional controller design by convex optimization. *Linear Algebra and its Applications*, 351–352:639–669, 2002.
- [19] T. Tanaka and P. A. Parrilo. Optimal output feedback architecture for triangular LQG problems. In *American Control Conference*, pages 5730–5735, 2014.
- [20] A. S. M. Vamsi and N. Elia. Optimal distributed controllers realizable over arbitrary networks. *IEEE Trans. Autom. Control*, 61(1):129–144, 2016.
- [21] J. Yan and S. E. Salcudean. Teleoperation controller design using H_∞ -optimization with application to motion-scaling. *IEEE Trans. Control Syst. Technol.*, 4(3):244–258, 1996.
- [22] K. Zhou, J. C. Doyle, and K. Glover. *Robust and optimal control*, volume 40. Prentice Hall, New Jersey, 1996.