

A Separation Principle for Decentralized State-Feedback Optimal Control

Laurent Lessard¹

Allerton Conference on Communication, Control, and Computing, pp. 528–534, 2013

Abstract

A cooperative control problem is considered in which dynamically decoupled subsystems must control their own states through state feedback in order to optimize a global quadratic cost. The states of the subsystems are coupled only through the cost function and correlated external disturbances. The architecture is truly decentralized; no communication between subsystems or their controllers is permitted. The main result of this paper is that the optimal decentralized controller satisfies a new separation principle that is strikingly similar to the celebrated result from centralized optimal control theory, but does not appear to follow from it. Roughly speaking, the optimal decentralized control strategy for each subsystem is the product of a static control gain and a global state estimate, and each can be separately computed.

1 Introduction

The separation principle is a classical and celebrated result in optimal control theory. When it holds, the optimal controller may be found by composing the solutions to two separate problems of control and filtering. For a recent historical perspective, see [1] and references therein.

Separation holds for linear systems with Gaussian noise and a classical information pattern [16]. Additionally, if the cost is quadratic (LQG assumptions), the optimal controller is a linear function of the conditional mean of the state. So the optimal controller takes the form

$$u^{\text{opt}} = K \mathbf{E}(x | \mathcal{Y}) \quad (1)$$

where K is a static control gain, x is the state of the plant, and \mathcal{Y} is the history of available noisy measurements. A key requirement is that of a *classical information pattern*; all decision-making stations must receive the same information and have perfect recall. Even under LQG assumptions, a nonclassical information pattern may result in a cost function with multiple local minima, and a nonlinear controller may strictly outperform any linear controller [14].

However, some structural results are known for special classes of nonclassical information patterns. If the

information pattern is *partially nested*, then there exist linear optimal strategies [2]. In particular, one-timestep delay-sharing architectures also satisfy a separation structure [15].

More recent structural results for partially nested information patterns include: poset-causal systems [12], delayed systems characterized by an underlying directed graph [3, 4] and output-feedback with sparsity [5, 7]. These examples provide *sufficient statistics* expressible in terms of conditional means. For example, the two-player output-feedback problem with triangular information structure [7] has an optimal controller of the form

$$u^{\text{opt}} = K_1 \mathbf{E}(x | \mathcal{Y}_1) + K_2 \mathbf{E}(x | \mathcal{Y}_1, \mathcal{Y}_2) \quad (2)$$

where \mathcal{Y}_1 and \mathcal{Y}_2 are the histories of two different noisy measurements. This structure is similar to that of (1) in that it provides a separation between control and estimation. However, the separation is only superficial, as the associated estimation and control problems are coupled and must be jointly solved. Structural results have also been obtained for information patterns that are not partially nested, such as *partial history sharing* [10].

In this paper, we study a problem that differs from those mentioned above in two important ways. First, the subsystems and controllers are dynamically decoupled and there is no communication or sharing of information between controllers. Second, the exogenous noise affecting the subsystems may be correlated. A state-space construction of the optimal linear controller for this problem was presented in [6] and it was found to have a state dimension that depends on the extent of the correlation in the exogenous noise. This is a curious finding, as it seems to defy the possibility of a structural result of the form (1) or (2). Indeed, a state estimate such as $\mathbf{E}(x | \mathcal{Y})$ has a state dimension equal to the dimension of x , regardless of the noise statistics.

Our main result is that the fully decentralized problem studied in [6] does indeed have a separation structure in the spirit of (1), but it uses modified control and filtering problems. This separation structure was previously unknown, and gives an explanation for why the optimal controller's state dimension depends on the noise statistics as observed in [6].

The paper is organized as follows. We begin with a formal description of the problem in Section 2, followed by main results in Section 3 and discussion in Section 4.

¹L. Lessard is with the Department of Mechanical Engineering at the University of California, Berkeley, CA 94720, USA.

lessard@berkeley.edu

In Section 5 we relate the new separation result to the work in [6], and we summarize our findings in Section 7.

2 Problem statement

We consider the simplest case of maximal decentralization; multiple subsystems are independently controlled, with no possibility of communication or influence between subsystems. The goal is to optimize a global objective that depends on both subsystems. A block diagram is shown in Figure 1. The problem considered in [6] had vector disturbances, but we will assume for now that each system is subject to a scalar disturbance. We will examine the vector case later in Section 5.

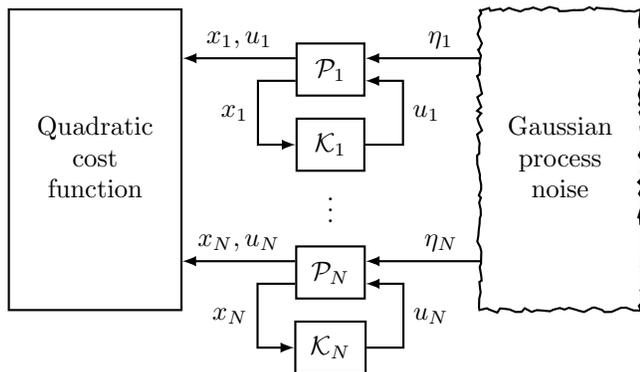


Figure 1: Diagram of the decentralized control architecture considered in this paper. The N subsystems are coupled only through the process noise and the cost function; no other communication is permitted.

Each \mathcal{P}_i is a continuous-time linear dynamical system with state $x_i \in \mathbb{R}^{n_i}$, input $u_i \in \mathbb{R}^{m_i}$, and scalar disturbance $\eta_i \in \mathbb{R}$. The state-space equations are as follows

$$\mathcal{P}_i: \quad \dot{x}_i = A_i x_i + B_i u_i + h_i \eta_i \quad \text{for } i = 1, \dots, N \quad (3)$$

For each plant, $A_i \in \mathbb{R}^{n_i \times n_i}$ is the state transition matrix, $B_i \in \mathbb{R}^{n_i \times m_i}$ is the controlled input matrix, and $h_i \in \mathbb{R}^{n_i \times 1}$ is the exogenous input vector. The noise terms η_1, \dots, η_N are assumed to be zero-mean, jointly Gaussian, and IID in time. We characterize this property via the $N \times N$ covariance matrix

$$\Sigma := \text{cov} \begin{bmatrix} \eta_1 \\ \eta_2 \\ \vdots \\ \eta_N \end{bmatrix} = \begin{bmatrix} 1 & \sigma_{12} & \dots & \sigma_{1N} \\ \sigma_{21} & 1 & \dots & \sigma_{2N} \\ \vdots & \vdots & \ddots & \vdots \\ \sigma_{N1} & \sigma_{N2} & \dots & 1 \end{bmatrix} \succ 0 \quad (4)$$

The assumption that $\mathbf{E}(\eta_i^2) = 1$ is not restrictive, as we may simply scale each equation so that each η_i has unit variance. By concatenating $x = [x_1^\top \dots x_N^\top]^\top$ and similarly for u and η , we can write one compact equation that describes the global dynamics.

$$\dot{x} = Ax + Bu + H\eta \quad (5)$$

where $A = \text{diag}(A_1, \dots, A_N)$ and similarly for B and H . The objective is to find a *fully decentralized* stabilizing control policy that minimizes the following infinite-horizon LQR cost:

$$\mathcal{J} = \lim_{T \rightarrow \infty} \mathbf{E} \int_0^T \begin{bmatrix} x \\ u \end{bmatrix}^\top \begin{bmatrix} Q & S \\ S^\top & R \end{bmatrix} \begin{bmatrix} x \\ u \end{bmatrix} dt \quad (6)$$

Note that Q , R , and S may be full matrices in general, so the cost couples the states and inputs of the various subsystems. We assume $R \succ 0$ and $Q - SR^{-1}S^\top \succeq 0$ as is standard, to ensure strict convexity of the cost.

We seek fully decentralized policies $u_i = \gamma_i(\mathcal{X}_i)$ that depend on the state histories of the respective plants. Naturally, a necessary condition for the existence of a stabilizing controller is that (A_i, B_i) be stabilizable for each i . We assume without loss of generality that each A_i is Hurwitz, for if it were not, we could perform a loop-shifting change of coordinates by writing $u_i = F_i x_i + \tilde{u}_i$, where F_i is any matrix such that $A_i + B_i F_i$ is Hurwitz.

No communication is permitted between the subsystems or their controllers. Although the plants (3) are dynamically decoupled, the performance of their controllers is coupled through the noise (4) and cost function (6).

3 Main result

The classical separation principle for LQG problems states that the optimal policy has the form $u = K\hat{x}$.

1. K is a solution to the state-feedback version of the problem. In other words, if we had perfect state information, the optimal policy would be $u = Kx$. This is also known as the LQR controller.
2. \hat{x} is the estimate of the state that minimizes the mean-squared error conditioned on the available measurements. This is equal to the conditional expectation $\hat{x} = \mathbf{E}(x | \mathcal{Y})$, and is also known as the Kalman-Bucy filter.

The main reason for calling this a separation principle is that finding the optimal controller is reduced to finding the solutions to two simpler problems that can be solved separately. Another reason is that it effectively gives a separation between estimation and control. Indeed, the control gain K does not depend on noise covariances, nor does the optimal estimate \hat{x} depend on the cost function.

The main result of this paper is a separation principle for the decentralized problem described in Section 2. As in the classical case, it leads to an optimal controller of the form $u = K\hat{x}$, where the control gain K and optimal estimate \hat{x} can be computed separately by solving separate simpler problems. There is one important difference however. In this new separation principle, K *does* depend on the noise covariance, and \hat{x} *does* depend on the cost function. In other words, even though K and

\hat{x} can be computed separately, they each depend on all parameters of the problem.

There are two ingredients to the new separation principle, and each is a well-understood classical problem. We now describe these simpler problems and explain some notation used for each one.

Centralized LQR control. If we relax the constraint that u_i measures only x_i , and instead allow u_i to measure the global state x , the problem described by (3)–(6) is a standard centralized LQR problem. The solution is given by $u = Kx$ where

$$K = -R^{-1}(B^T P + S^T), \quad \text{and}$$

$$P = \text{Ric} \begin{bmatrix} A - BR^{-1}S^T & -B^T R^{-1}B \\ -(Q - SR^{-1}S^T) & -(A - BR^{-1}S^T)^T \end{bmatrix}$$

where $\text{Ric}(H)$ denotes the unique stabilizing solution to the algebraic Riccati equation with Hamiltonian matrix H . See [17] for more details. As a convention, we will summarize the above solution using the compact notation $K = \text{LQR}(A, B, Q, R, S)$.

Decentralized estimation. Suppose two systems are described by the following state-space equations.

$$\begin{aligned} \dot{x}_1 &= A_1 x_1 + B_1 u_1 + w_1 \\ \dot{x}_2 &= A_2 x_2 + B_2 u_2 + w_2 \end{aligned}$$

where the A_i are stable and both u_1 and u_2 are known to the estimator. The noises w_1 and w_2 are zero-mean jointly Gaussian random variables with covariance

$$W := \text{cov} \begin{bmatrix} w_1 \\ w_2 \end{bmatrix} = \begin{bmatrix} W_{11} & W_{12} \\ W_{21} & W_{22} \end{bmatrix} \quad (7)$$

Then $\mathbf{E}(x_2 | \mathcal{X}_1)$ is given by the output of the following dynamical estimator

$$\begin{aligned} \dot{q} &= A_2 q + (A_2 L - L A_1) x_1 + B_2 u_2 - L B_1 u_1 \\ \mathbf{E}(x_2 | \mathcal{X}_1) &= q + L x_1 \end{aligned}$$

where L is any solution to $LW_{11} = W_{21}$. Note that such an L exists even when W_{11} is not invertible. The estimator error covariance is given by

$$\text{cov}(x_2 - \mathbf{E}(x_2 | \mathcal{X}_1)) = Y$$

where Y is a solution to the continuous-time Lyapunov equation $A_2 Y + Y A_2^T + W_{22} - L W_{11} L^T = 0$. The matrix A_2 is stable by assumption, so the Lyapunov equation has a unique solution and it satisfies $Y \succeq 0$. See [8] for a reference on how to construct such observers.

We now show how these two ingredients are assembled to give a solution to the decentralized control problem described in Section 2. What follows is our main result, a new separation principle for decentralized control.

Theorem 1. Consider the decentralized control problem described by (3)–(6), and perform these two steps.

- (i) Partition Q, R, S into $N \times N$ block matrices that conform to the dimensions n_i and m_i . Therefore, $Q_{ij} \in \mathbb{R}^{n_i \times n_j}$, $R_{ij} \in \mathbb{R}^{m_i \times m_j}$ and $S_{ij} \in \mathbb{R}^{n_i \times m_j}$. Now define the blockwise Hadamard product

$$Q \circ \Sigma := \begin{bmatrix} Q_{11} & \sigma_{12} Q_{12} & \dots & \sigma_{1N} Q_{1N} \\ \sigma_{21} Q_{21} & Q_{22} & \dots & \sigma_{2N} Q_{2N} \\ \vdots & \vdots & \ddots & \vdots \\ \sigma_{N1} Q_{N1} & \sigma_{N2} Q_{N2} & \dots & Q_{NN} \end{bmatrix}$$

and similarly for $R \circ \Sigma$ and $S \circ \Sigma$. Finally, define

$$K^\sigma := \text{LQR}(A, B, (Q \circ \Sigma), (R \circ \Sigma), (S \circ \Sigma))$$

Hence, K^σ is the solution to a centralized LQR problem where the cost matrices have been weighted.

- (ii) For each $i = 1, \dots, N$, consider the following system defined by the state z^i

$$\begin{aligned} \dot{z}^i &= A z^i + B u^i + h \eta_i \\ u^i &= K^\sigma z^i \end{aligned} \quad (8)$$

where h is the vertical concatenation of h_1, \dots, h_N . This is similar to (5), except η_i is driving every state of the system, not just state i . Now define \hat{x}^i as the decentralized estimate $\hat{x}^i := \mathbf{E}(z^i | \mathcal{X}_i)$.

The optimal decentralized controller is given by

$$u_i = K_i^\sigma \hat{x}^i$$

where K_i^σ is the i^{th} block-row of K^σ .

Proof. See Section 6. ■

4 Discussion

The result of Theorem 1 bears a striking resemblance to the centralized separation principle in terms of structure, but there are some important differences.

Firstly, the control gain K^σ depends on the noise covariance (4). This is different from the centralized case, where the control gain is independent of the noise covariance. Indeed, the centralized LQR problem is *certainly equivalent*, so the optimal controller remains the same even if noise is removed entirely and we consider a deterministic formulation of the problem. This feature is clearly not present in the decentralized LQR problem considered herein.

Secondly, the estimators actually depend on the cost function as well. Indeed, in estimating the global states, each assumes the control policy of the other agents is a function of K^σ , and this depends on the cost matrices

Q , R , and S . So, even though Theorem 1 expressed the optimal decentralized controller as a product of the solutions to a control and an estimation problem that can be solved separately, these problems each depend on all parameters of the original problem.

It is worth noting that the estimates $\hat{x}^i = \mathbf{E}(z^i | \mathcal{X}_i)$ are always exact. In other words, the covariance of the estimation error is zero; $Y = \mathbf{cov}(z^i - \mathbf{E}(z^i | \mathcal{X}_i)) = 0$. This is intuitive because the same η_i drives the entire state. So measuring x_i is equivalent to measuring any other component of x . One could perhaps call the \hat{x}^i *inferences* rather than *estimates*.

In the case where $\Sigma = \mathbf{1}\mathbf{1}^\top$, the matrix of all-ones, all noise signals are perfectly correlated. Therefore, all η_i are equal and the global estimates \hat{x}^i are each identically equal to the global state x . Also, $Q \circ \Sigma = Q$ and similarly for R and S . Therefore, $K^\sigma = K = \text{LQR}(A, B, Q, R, S)$, and we recover the optimal centralized policy $u = Kx$.

In the case where $\Sigma = I$, the noise signals are independent. Also, $Q \circ \Sigma = \text{diag}(Q_{11}, \dots, Q_{NN})$ and similarly for R and S . We therefore obtain the optimal greedy gains defined by $K_i := \text{LQR}(A_i, B_i, Q_{ii}, R_{ii}, S_{ii})$. In other words, $K^\sigma = \text{diag}(K_1, \dots, K_N)$ and we recover the optimal greedy policy $u_i = K_i x_i$.

The cases discussed above show that in the extremes where the noise is perfectly correlated or independent, we recover the optimal centralized policy or the optimal greedy one, respectively. While these edge cases correspond to classical problems, the intermediate values of Σ do not. The key fact revealed in Theorem 1 is that a unifying separation principle holds for all Σ .

5 Extension to the vector noise case

The case of two subsystems. We begin with the case $N = 2$, where the scalar noises η_1, η_2 are replaced by vector noises. This problem was solved explicitly in [6], and the optimal controller was found to have a state dimension that depends on the noise covariance W . We will now show that under an appropriate change of coordinates, this vector-noise problem can be reduced to decoupled instances of the scalar noise problem described by (3)–(6). The vector-noise problem for $N = 2$ is described by the following equations.

$$\begin{bmatrix} \dot{x}_1 \\ \dot{x}_2 \end{bmatrix} = \begin{bmatrix} A_1 & 0 \\ 0 & A_2 \end{bmatrix} \begin{bmatrix} x_1 \\ x_2 \end{bmatrix} + \begin{bmatrix} B_1 & 0 \\ 0 & B_2 \end{bmatrix} \begin{bmatrix} u_1 \\ u_2 \end{bmatrix} + \begin{bmatrix} w_1 \\ w_2 \end{bmatrix} \quad (9)$$

with a cost function of the form (6). We assume each state equation has been scaled such that

$$W := \mathbf{cov} \begin{bmatrix} w_1 \\ w_2 \end{bmatrix} = \begin{bmatrix} I & X \\ X^\top & I \end{bmatrix} \quad (10)$$

We further assume that each input has been shifted using $u_i = \tilde{u}_i + K_i x_i$ where $K_i = \text{LQR}(A_i, B_i, Q_{ii}, R_{ii}, S_{ii})$.

This ensures that the A_i in (9) are stable. Note that in these new coordinates, $\text{LQR}(A_i, B_i, Q_{ii}, R_{ii}, S_{ii}) = 0$.

The vector noise w may be decomposed into a sum involving scalar noise terms. To see why, taking the full singular value decomposition (SVD) of X . Namely, write

$$X = [U_1 \quad \bar{U}_1] \begin{bmatrix} \hat{\Sigma} & 0 \\ 0 & 0 \end{bmatrix} [U_2 \quad \bar{U}_2]^\top \quad (11)$$

where $\hat{\Sigma} = \text{diag}(\sigma_1, \dots, \sigma_r)$, and $\sigma_1 \geq \dots \geq \sigma_r > 0$ are the singular values of X , and $r := \text{rank } X$. Note that the various matrices have sizes $U_i \in \mathbb{R}^{n_i \times r}$, $\bar{U}_i \in \mathbb{R}^{n_i \times (n_i - r)}$. The following lemma gives the required decomposition.

Lemma 2. *Suppose that (w_1, w_2) are jointly Gaussian random vectors with a covariance matrix given by (10), and the SVD of X is given by (11). Define the scalar noises $\eta_i^1, \dots, \eta_i^r, v_i$ for $i = 1, 2$ using the formulas*

$$\begin{aligned} \eta_i^k &:= e_k^\top U_i^\top w_i & \text{for } k = 1, \dots, r & & \text{for } i = 1, 2 \\ v_i &:= \bar{U}_i^\top w_i \end{aligned}$$

Then we have the properties

- 1) $\left\{ \begin{bmatrix} \eta_1^1 \\ \eta_2^1 \end{bmatrix}, \dots, \begin{bmatrix} \eta_1^r \\ \eta_2^r \end{bmatrix}, v_1, v_2 \right\}$ are independent
- 2) $\mathbf{cov} \begin{bmatrix} \eta_1^k \\ \eta_2^k \end{bmatrix} = \begin{bmatrix} 1 & \sigma_k \\ \sigma_k & 1 \end{bmatrix}$ for $k = 1, \dots, r$
- 3) $\mathbf{cov} v_1 = I$ and $\mathbf{cov} v_2 = I$.

Furthermore, (w_1, w_2) may be decomposed as follows.

$$\begin{bmatrix} w_1 \\ w_2 \end{bmatrix} = \sum_{k=1}^r \begin{bmatrix} U_1 e_k & 0 \\ 0 & U_2 e_k \end{bmatrix} \begin{bmatrix} \eta_1^k \\ \eta_2^k \end{bmatrix} + \begin{bmatrix} \bar{U}_1 & 0 \\ 0 & \bar{U}_2 \end{bmatrix} \begin{bmatrix} v_1 \\ v_2 \end{bmatrix}$$

Proof. Expand (w_1, w_2) using (10)–(11) and the orthogonality of the $[U_i \quad \bar{U}_i]$ matrices to obtain

$$w_i = \sum_{k=1}^r U_i e_k (e_k^\top U_i^\top w_i) + \bar{U}_i (\bar{U}_i^\top w_i)$$

which is precisely the desired decomposition. The three properties follow from direct evaluation of the appropriate covariances. \blacksquare

We now apply Lemma 2 to separate the state equations (9) into independent parts. Since (x_1, u_1) is a linear function of w_1 , we can consider the components that depend on each of the terms in the decomposition. Namely, we can write

$$\begin{bmatrix} x \\ u \end{bmatrix} = \begin{bmatrix} \zeta^1 \\ \varphi^1 \end{bmatrix} + \dots + \begin{bmatrix} \zeta^r \\ \varphi^r \end{bmatrix} + \begin{bmatrix} \bar{\zeta} \\ \bar{\varphi} \end{bmatrix}$$

where (ζ^k, φ^k) is a linear function of (η_1^k, η_2^k) , and $(\bar{\zeta}, \bar{\varphi})$ is a linear function of (v_1, v_2) . Furthermore, the terms in

the above decomposition are independent. We therefore have the dynamics

$$\begin{aligned}\zeta^k &= A\zeta^k + B\varphi^k + \begin{bmatrix} U_1 e_k & 0 \\ 0 & U_2 e_k \end{bmatrix} \begin{bmatrix} \eta_1^k \\ \eta_2^k \end{bmatrix} \quad \text{for } k = 1, \dots, r \\ \dot{\zeta} &= A\bar{\zeta} + B\bar{\varphi} + \begin{bmatrix} \bar{U}_1 & 0 \\ 0 & \bar{U}_2 \end{bmatrix} \begin{bmatrix} v_1 \\ v_2 \end{bmatrix}\end{aligned}$$

Each such equation is a decentralized LQR problem of the type described in (3)–(6). The cost functions decouple as well, so we may apply Theorem 1 for each of the $r + 1$ equations above. Thus, the optimal controller for the $N = 2$ case with vector noise can be written as the linear combination

$$u_i = \sum_{k=1}^r K_i^{\sigma_k} (\hat{x}^k)^i$$

where each $K_i^{\sigma_k}$ is generated by the corresponding σ_k according to the definition in the statement of Theorem 1. There is no term corresponding to the ζ_i component of the decomposition because the associated gain would be LQR($A_i, B_i, Q_{ii}, R_{ii}, S_{ii}$), which is zero due to the normalization discussed in the beginning of this section.

Recall that $(\hat{x}^k)^i$ is the estimate of ζ^k given ζ_i^k , where both noise inputs are set to η_i^k . As discussed in Section 4, this estimate is an exact inference. Furthermore, measuring ζ_i^k is equivalent to measuring η_i^k , which is in turn equivalent to measuring w_i , and therefore x_i . Thus, $(\hat{x}^k)^i = \mathbf{E}(\zeta^k | \mathcal{X}_i)$, where

$$\begin{aligned}\dot{x}_i &= A_i x_i + B_i u_i + w_i \\ \dot{\zeta}^k &= A\zeta^k + BK^{\sigma_k} \zeta^k + \begin{bmatrix} U_1 e_k \\ U_2 e_k \end{bmatrix} \eta_i^k\end{aligned}$$

Applying the decentralized estimation result of Section 3 to evaluate each of the expectations, we obtain

$$\begin{aligned}\dot{q}^k &= (A + BK^{\sigma_k})q^k - L_i^k B_i u_i + (AL_i^k - L_i^k A_i)x_i \\ u_i &= \sum_{k=1}^r K_i^{\sigma_k} (q^k + L_i^k x_i)\end{aligned} \quad \text{for } k = 1, \dots, r$$

where $L_i^k := [U_1^T \ U_2^T]^T e_k e_k^T U_i^T$. This is precisely the result derived in [6]. In general, each state q^k has dimension $n_1 + n_2$, so the overall state dimension of the controller $\mathcal{K}_i : x_i \rightarrow u_i$ is $r(n_1 + n_2)$.

Therefore, the case $N = 2$ with vector noise is nothing more than r separate instances of the case $N = 2$ with scalar noise, and each such instance satisfies the separation principle of Theorem 1. As discussed in [6], the vector noise problem can be formulated in such a way that it has a unique solution, and it can be shown that the optimal controller with state dimension $r(n_1 + n_2)$ is *generically minimal*. That is, there is a choice of problem parameters such that the optimal controller has a minimal realization with this state dimension. This means that the interpretation of the states in terms of separate filters is necessary; one may need each of these r estimates to accurately represent the optimal controller.

The case of N subsystems. The case $N \geq 3$ with vector noises is substantially more complicated than the case $N = 2$ described above. The reason is that the decomposition of Lemma 2 does not hold for more than two subsystems. For example, if $N = 3$, the noise covariance has the general form

$$W = \text{cov} \begin{bmatrix} w_1 \\ w_2 \\ w_3 \end{bmatrix} = \begin{bmatrix} I & X_{12} & X_{13} \\ X_{12}^T & I & X_{23} \\ X_{13}^T & X_{23}^T & I \end{bmatrix}$$

If the X_{ij} happen to be diagonal, then the problem splits into separate instances of the $N = 3$ scalar-noise case. However, the technique used in Lemma 2 can only diagonalize one of the X_{ij} matrices at a time, so a new type of decomposition will be required to address the case $N \geq 3$.

6 Proof of Theorem 1

We prove the result by first explicitly constructing a state-space realization of the optimal controller using an approach similar to the one used in [6], and then showing that the solution satisfies the required separation structure. We begin by converting the problem (3)–(6) into model-matching form. Define \mathcal{F} and \mathcal{G} as

$$\mathcal{F} = \left[\begin{array}{c|c} A & H \\ \hline C & 0 \end{array} \right] \quad \text{and} \quad \mathcal{G} = \left[\begin{array}{c|c} A & B \\ \hline C & D \end{array} \right] \quad (12)$$

where C and D are defined such that

$$[C \ D]^T [C \ D] = \begin{bmatrix} Q & S \\ S^T & R \end{bmatrix} \quad (13)$$

Now consider the following decentralized model-matching problem with one-sided dynamics.

$$\begin{aligned}\min \quad & \left\| \mathcal{F}\Sigma^{1/2} + \mathcal{G} \begin{bmatrix} \mathcal{Q}_1 & & 0 \\ & \ddots & \\ 0 & & \mathcal{Q}_N \end{bmatrix} H\Sigma^{1/2} \right\|_{\mathcal{H}_2} \\ \text{s.t.} \quad & \mathcal{Q}_1, \dots, \mathcal{Q}_N \in \mathcal{RH}_2\end{aligned} \quad (14)$$

We then have the following result.

Lemma 3. *Suppose a solution to (14) is given by*

$$\mathcal{Q}_i = \left[\begin{array}{c|c} \hat{A}_i & \hat{B}_i \\ \hline \hat{C}_i & 0 \end{array} \right] \quad \text{for } i = 1, \dots, N \quad (15)$$

A solution to the decentralized LQR problem described by (3)–(6) is given by $u_i = \mathcal{K}_i x_i$ for $i = 1, \dots, N$, where

$$\mathcal{K}_i = \left[\begin{array}{c|c} \hat{A}_i - \hat{B}_i B_i \hat{C}_i & (\hat{A}_i - \hat{B}_i B_i \hat{C}_i) \hat{B}_i - \hat{B}_i A_i \\ \hline \hat{C}_i & \hat{C}_i \hat{B}_i \end{array} \right] \quad (16)$$

Proof. The LQR problem described by (3)–(6) is quadratically invariant [11], and the associated plant is

stable. Therefore, the set of stabilizing controllers is parameterized by $\mathcal{K} = \bar{\mathcal{Q}}(I + (sI - A)^{-1}B\bar{\mathcal{Q}})^{-1}$, where $\bar{\mathcal{Q}} \in \mathcal{RH}_\infty$ and $\bar{\mathcal{Q}}$ is block-diagonal [17]. The closed-loop map $e \rightarrow z$ where e is IID Gaussian with identity covariance and $z = Cx + Du$ is the cost functional is

$$\begin{aligned} & \mathcal{F}\Sigma^{1/2} + \mathcal{G}\mathcal{K}(I - (sI - A)^{-1}B\mathcal{K})^{-1}(sI - A)^{-1}H\Sigma^{1/2} \\ &= \mathcal{F}\Sigma^{1/2} + \mathcal{G}\bar{\mathcal{Q}}(sI - A)^{-1}H\Sigma^{1/2} \end{aligned}$$

Let $\mathcal{Q} := \bar{\mathcal{Q}}(sI - A)^{-1}$. Whenever $\bar{\mathcal{Q}} \in \mathcal{RH}_\infty$, we have $\mathcal{Q} \in \mathcal{RH}_2$. Similarly, if

$$\begin{aligned} \mathcal{Q} &= \left[\begin{array}{c|c} \bar{A} & \bar{B} \\ \bar{C} & 0 \end{array} \right] \in \mathcal{RH}_2 \quad \text{then} \\ \bar{\mathcal{Q}} &= \left[\begin{array}{c|c} \bar{A} & \bar{A}\bar{B} - \bar{B}A \\ \bar{C} & \bar{C}\bar{B} \end{array} \right] \in \mathcal{RH}_\infty \end{aligned}$$

Furthermore, \mathcal{Q} is block-diagonal if and only if $\bar{\mathcal{Q}}$ is block-diagonal. Therefore, the set of closed-loop maps may be equivalently parameterized by $\mathcal{F}\Sigma^{1/2} + \mathcal{G}\mathcal{Q}H\Sigma^{1/2}$, where $\mathcal{Q} \in \mathcal{RH}_2$ and is block-diagonal. Inverting the maps $\mathcal{Q} \rightarrow \bar{\mathcal{Q}} \rightarrow \mathcal{K}$ yields the desired result. ■

We will therefore concentrate our efforts on solving the decentralized model-matching problem (14). Viewed as an optimization problem in the Hilbert space \mathcal{H}_2 , the optimality condition for (14) is given by [9]

$$\mathcal{G}^\sim \mathcal{F}\Sigma H^\top + \mathcal{G}^\sim \mathcal{G}\mathcal{Q}H\Sigma H^\top \in \begin{bmatrix} \mathcal{H}_2^\perp & \mathcal{L}_2 \\ & \ddots \\ \mathcal{L}_2 & \mathcal{H}_2^\perp \end{bmatrix}$$

Where $\mathcal{Q} := \text{diag}(\mathcal{Q}_1, \dots, \mathcal{Q}_N)$. Extract the N block-diagonal entries, and let E_i be the $(n_1 + \dots + n_N) \times n_i$ matrix such that $I = [E_1 \ \dots \ E_N]$. Similarly, let e_i be the $N \times n$ matrix such that $I = [e_1 \ \dots \ e_N]$. Using the fact that H is block-diagonal, we obtain

$$E_i^\top \mathcal{G}^\sim \mathcal{F}\Sigma e_i + E_i^\top \mathcal{G}^\sim \mathcal{G}\mathcal{Q}H\Sigma e_i \in \mathcal{H}_2^\perp$$

Expanding, and making use of the definition (7)

$$E_i^\top \mathcal{G}^\sim \mathcal{F}\Sigma e_i + E_i^\top \mathcal{G}^\sim \mathcal{G} \begin{bmatrix} \mathcal{Q}_1 h_1 \sigma_{1i} \\ \vdots \\ \mathcal{Q}_N h_N \sigma_{Ni} \end{bmatrix} \in \mathcal{H}_2^\perp$$

Commuting the scalars σ_{ji} ,

$$E_i^\top \mathcal{G}^\sim \mathcal{F}\Sigma e_i + E_i^\top \mathcal{G}^\sim \mathcal{G} \begin{bmatrix} \sigma_{1i} I & & \\ & \ddots & \\ & & \sigma_{Ni} I \end{bmatrix} \begin{bmatrix} \mathcal{Q}_1 h_1 \\ \vdots \\ \mathcal{Q}_N h_N \end{bmatrix} \in \mathcal{H}_2^\perp$$

Where the identity matrices are sized according to the corresponding \mathcal{Q}_i blocks. Stacking these equations for $i = 1, \dots, N$, we obtain

$$\left((\mathcal{G}^\sim \mathcal{F}) \circ \Sigma \right) \mathbf{1} + \left((\mathcal{G}^\sim \mathcal{G}) \circ \Sigma \right) \begin{bmatrix} \mathcal{Q}_1 h_1 \\ \vdots \\ \mathcal{Q}_N h_N \end{bmatrix} \in \mathcal{H}_2^\perp \quad (17)$$

where \circ is the blockwise Hadamard product defined in the theorem statement, $\mathbf{1}$ is the vector of all-ones, and we made use of the fact that $\Sigma = \Sigma^\top$. This is a centralized LQR problem in model-matching form! It remains to find a joint realization for the blocks. Specifically, we would like to find some $\bar{\mathcal{F}}$ and $\bar{\mathcal{G}}$ such that

$$\bar{\mathcal{G}}^\sim \left[\bar{\mathcal{F}} \quad \bar{\mathcal{G}} \right] = \left[\left((\mathcal{G}^\sim \mathcal{F}) \circ \Sigma \right) \mathbf{1} \quad (\mathcal{G}^\sim \mathcal{G}) \circ \Sigma \right] \quad (18)$$

This will ensure that (17) is the optimality condition for the LQR problem determined by $\bar{\mathcal{F}}$ and $\bar{\mathcal{G}}$. Substituting the definitions (12)–(13), we obtain

$$\begin{aligned} \mathcal{G}^\sim \mathcal{F} &= \left[\begin{array}{cc|c} -A^\top & Q & 0 \\ 0 & A & H \\ -B^\top & S^\top & 0 \end{array} \right] \quad \text{and} \\ \mathcal{G}^\sim \mathcal{G} &= \left[\begin{array}{cc|c} -A^\top & Q & S \\ 0 & A & B \\ -B^\top & S^\top & R \end{array} \right] \end{aligned}$$

Making use of the fact that A , B , and H are block-diagonal, we find after some algebra that

$$\begin{aligned} (\mathcal{G}^\sim \mathcal{F}) \circ \Sigma &= \left[\begin{array}{cc|c} -A^\top & Q \circ \Sigma & 0 \\ 0 & A & H \\ -B^\top & S^\top \circ \Sigma & 0 \end{array} \right] \quad \text{and} \\ (\mathcal{G}^\sim \mathcal{G}) \circ \Sigma &= \left[\begin{array}{cc|c} -A^\top & Q \circ \Sigma & S \circ \Sigma \\ 0 & A & B \\ -B^\top & S^\top \circ \Sigma & R \end{array} \right] \end{aligned}$$

Now, define \bar{C} and \bar{D} such that

$$[\bar{C} \quad \bar{D}]^\top [\bar{C} \quad \bar{D}] = \begin{bmatrix} Q \circ \Sigma & S \circ \Sigma \\ S^\top \circ \Sigma & R \circ \Sigma \end{bmatrix} \quad (19)$$

This new cost matrix satisfies the required definiteness constraints, see for example [13]. It is straightforward to check that

$$\bar{\mathcal{F}} = \begin{bmatrix} A & h \\ \bar{C} & 0 \end{bmatrix} \quad \text{and} \quad \bar{\mathcal{G}} = \begin{bmatrix} A & B \\ \bar{C} & \bar{D} \end{bmatrix}$$

satisfy (18), where $h := H\mathbf{1}$ is the vector obtained by stacking h_1, \dots, h_N . The solution to this centralized model-matching problem is therefore

$$\begin{bmatrix} \mathcal{Q}_1 h_1 \\ \vdots \\ \mathcal{Q}_N h_N \end{bmatrix} = \begin{bmatrix} A + BK^\sigma & h \\ K^\sigma & 0 \end{bmatrix} \quad (20)$$

with K^σ defined in the theorem statement. Extracting the i^{th} block-row, one possible solution is given by

$$\mathcal{Q}_i = \begin{bmatrix} A + BK^\sigma & h h_i^\dagger \\ E_i^\top K^\sigma & 0 \end{bmatrix}$$

Where $h_i^\dagger = (h_i^\top h_i)^{-1} h_i^\top$. Note that this solution is not unique, and there are many possible choices of \mathcal{Q}_i that

satisfy (20). Now apply Lemma 3 to obtain a realization for \mathcal{K}_i . The state-space equations are

$$\begin{aligned} \dot{q}^i &= (A + BK^\sigma - hh_i^\dagger B_i E_i^\top K^\sigma) q^i \\ &\quad + ((A + BK^\sigma - hh_i^\dagger B_i E_i^\top K^\sigma) hh_i^\dagger - hh_i^\dagger A_i) x_i \\ u_i &= E_i^\top K^\sigma q^i + E_i^\top K^\sigma hh_i^\dagger x_i \end{aligned}$$

Define $L := hh_i^\dagger$ and $\xi^i := q^i + Lx_i$. The equations then simplify to

$$\begin{aligned} \dot{q}^i &= Aq^i + (AL - LA_i)x_i + BK^\sigma \xi^i - LB_i u_i \\ \xi^i &= q^i + Lx_i \\ u_i &= E_i^\top K^\sigma \xi^i \end{aligned}$$

Upon comparing these equations with the decentralized estimation result from Section 3, we conclude that $\xi^i = \mathbf{E}(z^i | \mathcal{X}_i)$, where z^i satisfies

$$\begin{aligned} \dot{z}^i &= Az^i + Bu^i + v^i \\ u^i &= K^\sigma z^i \end{aligned} \tag{21}$$

and v^i has noise statistics that satisfy

$$L \mathbf{E}(h_i \eta_i^2 h_i^\top) = \mathbf{E}(v^i \eta_i h_i^\top)$$

Substituting the definition $L := hh_i^\dagger$ and using the fact that $\mathbf{E}(\eta_i^2) = 1$, we deduce that $v^i = h\eta_i$. Consequently, (21) becomes (8), as required. This completes the proof of Theorem 1. ■

7 Summary

In this paper, we showed that the N -subsystem fully decentralized LQR problem with coupled cost and correlated scalar noise satisfies a new separation principle reminiscent of the classical separation principle for LQG control. We also showed that the $N = 2$ case with correlated vector noise solved in [6] can be reduced to separate instances of the scalar-noise case.

The estimation and separation structure of the general N -subsystem case with correlated vector noise remains an open problem, as our decomposition approach used for $N = 2$ does not appear to generalize to $N \geq 3$.

References

[1] T. Georgiou and A. Lindquist. The separation principle in stochastic control, redux. *IEEE Transactions on Automatic Control*, 58(10):2481–2494, 2013.

[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] A. Lamperski and J. C. Doyle. Dynamic programming solutions for decentralized state-feedback LQG problems with output feedback. In *American Control Conference*, pages 6322–6327, 2012.

[4] A. Lamperski and L. Lessard. Optimal state-feedback control under sparsity and delay constraints. In *3rd IFAC Workshop on Distributed Estimation and Control in Networked Systems*, pages 204–209, 2012.

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

[6] L. Lessard. Optimal control of a fully decentralized quadratic regulator. In *Allerton Conference on Communication, Control, and Computing*, pages 48–54, 2012.

[7] L. Lessard and S. Lall. A state-space solution to the two-player decentralized optimal control problem. In *Allerton Conference on Communication, Control, and Computing*, pages 1559–1564. IEEE, 2011.

[8] D. Luenberger. An introduction to observers. *IEEE Transactions on Automatic Control*, 16(6):596–602, 1971.

[9] D. Luenberger. *Optimization by vector space methods*. Wiley-Interscience, 1997.

[10] A. Nayyar, A. Mahajan, and D. Teneketzis. Decentralized stochastic control with partial history sharing: A common information approach. *IEEE Transactions on Automatic Control*, 58(7):1644–1658, 2013.

[11] M. Rotkowitz and S. Lall. A characterization of convex problems in decentralized control. *IEEE Transactions on Automatic Control*, 51(2):274–286, 2006.

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

[13] B. Wang and F. Zhang. Schur complements and matrix inequalities of hadamard products. *Linear and Multilinear Algebra*, 43(1-3):315–326, 1997.

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

[15] H. S. Witsenhausen. Separation of estimation and control for discrete time systems. *Proceedings of the IEEE*, 59(11):1557–1566, 1971.

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

[17] K. Zhou, J. Doyle, and K. Glover. *Robust and optimal control*. Prentice-Hall, 1995.