Decentralized LQG Control of Systems with a Broadcast Architecture

Laurent Lessard[1,2]

Submitted, IEEE Conference on Decision and Control 2012

Abstract

In this paper, we consider dynamical subsystems interconnected in a broadcast architecture. In the broadcast-out case, the root node can affect several leaf nodes, but the leaf nodes do not affect any other nodes. Each subsystem is locally controlled via output feedback, and the controllers can communicate according to a structure that parallels the dynamic coupling between subsystems. Explicit state-space realizations for the optimal controllers are derived using a spectral factorization approach. An interpretation of the controller states is also provided in terms of optimal state estimators. We also address the dual broadcast-in case, where there is a single leaf node affected by multiple root nodes.

1 Introduction

In this paper, we consider a class of interconnected dynamical systems and associated controllers arranged in a *broadcast* structure. Visualizing each subsystem as a node in a directed graph, the two types of broadcast structures are illustrated in Figure 1. Each controller has access to local measurements from its associated subsystem. Additionally, a directed edge from A to B indicates:

(i) A change in the state or input of A can affect the state of B, but not vice versa.

(ii) The controller for A transmits everything it knows to the controller for B, but not vice versa.

The goal for the controllers is to optimize a global cost, which may couple the states and inputs of all subsystems. Such problems are *decentralized*, because the various controllers do not have access to the same information about the system.

Even in the case of a linear plant with Gaussian noise and a quadratic cost function, decentralization can greatly increase the difficulty of the associated controller synthesis problem. Witsenhausen showed that for a very simple instance of such a problem, linear controllers are strictly suboptimal [19]. In general, finding an optimal decentralized controller can be hard [1].

[1]L. Lessard is with the Department of Automatic Control at Lund University, Lund, Sweden. laurent.lessard@control.lth.se

[2]This work was supported by the Swedish Research Council through the LCCC Linnaeus Center

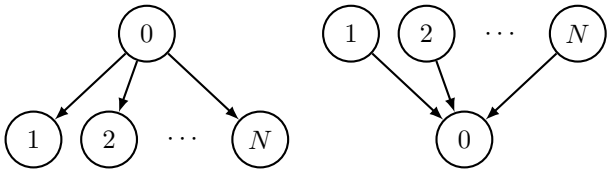

Figure 1: Diagrams representing a *broadcast-out* structure on the left and a *broadcast-in* structure on the right.

Since Witsenhausen's counterexample, efforts have largely been focused on identifying classes of decentralized control problems that can be reduced to convex programs [9, 11, 13]. When the information pattern is *partially nested* [3] as in the broadcast case considered herein, there always exists a linear optimal controller. Roughly speaking, the associated convex optimization problem can be reduced to a structured model-matching problem:

$$\begin{array}{rl} \text{minimize} & \left\|\mathcal{T}_1 + \mathcal{T}_2\mathcal{Q}\mathcal{T}_3\right\| \\ \text{subject to} & \mathcal{Q} \text{ is stable, } \mathcal{Q} \in \Lambda \end{array}$$

where $\mathcal{Q}$ and the $\mathcal{T}_i$ are matrices of proper stable transfer functions, and Λ is a subspace that characterizes the admissible controllers. For example, it may impose a sparsity constraint on $\mathcal{Q}$.

Since the design parameter $\mathcal{Q}$ is a rational function of unknown order, the resulting optimization problem is infinite-dimensional. Numerical approaches have been proposed, and examples include SDP methods [10, 14], or solution via a converging sequence of finite-dimensional problems [9]. An exact solution can be obtained by using vectorization [12], though it is only feasible for small problem instances. The methods mentioned above do not produce minimal state-space realizations in general, and perhaps most importantly, they do not reveal the structure (if any) of the optimal controller.

In the past few years, exact solutions have been found for many partially nested special cases. These solutions provide minimal state-space formulas for the optimal controller, and give great structural insight as well. For example, the controller's states can be interpreted as conditional estimates of the plant's states.

In the state-feedback case, general delayed-sharing architectures were considered in [5], and sparsity-constrained structure were solved in [15, 16, 17]. The

problem becomes considerably more complicated in the case of output feedback, and far less is known. Two special cases include:

(i) *partial output feedback* [4, 18]: leaf nodes have output feedback, internal nodes have state feedback.

(ii) *dynamically decoupled* [4]: all nodes have output feedback, but their dynamics are fully decoupled.

The common feature shared by these special cases is that the cost function can be split into independent terms that can be separately optimized using standard techniques.

For the general case with coupled dynamics and output feedback, no such separation has been found. The first solution to such a problem appeared in [6], and it treats the case of two subsystems. The solution also reveals a new feature not present in previously reported solutions: some of the control and estimation gains are coupled, and depend on all problem parameters.

In this paper, we present a generalization of the two-player output feedback result [6, 7]. Indeed, if a broadcast structure has a single root node and single leaf node, we recover the two-player case. The approach used herein is based on spectral factorization, and is similar to the approach in [7]. However, the broadcast case presents new challenges over the two-player case. Firstly, the multiple leaf nodes are dynamically decoupled, which leads to new technicalities that we discuss in depth in Section 6. Secondly, the broadcast case requires a more complicated set of coupling conditions.

The paper is organized as follows. We state the problem in detail in Section 2, and explain new notation in Section 3. The main result is stated in Section 4, and proved in Section 8. Additional discussion regarding the solution structure is provided in Sections 5–7. Finally, we conclude in Section 9.

2 Problem Statement

We will assume a broadcast-out structure as in Figure 1. The broadcast-in case is treated separately in Section 7. All subsystems are assumed to be continuous-time and linear time-invariant. The overall dynamics can be written in state-space form as:

$$\begin{aligned} \dot{x} &= Ax + Bu + w \\ y &= Cx + v \end{aligned}$$

where w and v are random variables IID in time, and jointly Gaussian in space, with zero mean and covariance

$$\mathbf{cov}\begin{bmatrix} w \\ v \end{bmatrix} = \begin{bmatrix} W & U \\ U^{\mathsf{T}} & V \end{bmatrix}$$

We partition the state as $x = \begin{bmatrix} x_0^{\mathsf{T}} & \dots & x_N^{\mathsf{T}} \end{bmatrix}^{\mathsf{T}}$. where each part corresponds to the state of the associated subsystem. The system matrices then inherit a sparsity structure from the graph representation. In this case,

$$A = \begin{bmatrix} A_{00} & 0 & \dots & 0 \\ A_{10} & A_{11} & \dots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ A_{N0} & 0 & \dots & A_{NN} \end{bmatrix} \tag{1}$$

where only the block diagonal and the first block column are populated. For convenience, we define Λ as the set of matrices with this sparsity pattern, where the dimensions of the blocks are to be inferred by context. Note that Λ is more than a subspace, it is closed under matrix multiplication and is therefore a ring. We assume that $B \in \Lambda$ and $C \in \Lambda$ as well, and partition u and y accordingly. Thus, the input-output transfer function also has a broadcast-out structure: $C(sI-A)^{-1}B \in \Lambda$. Note that in the broadcast-in case, Λ is replaced by Λ^{T}.

The objective is to design a decentralized output-feedback control policy that minimizes the infinite-horizon quadratic cost function:

$$\text{cost} = \lim_{T\to\infty} \frac{1}{T}\, \mathbb{E} \int_0^T \begin{bmatrix} x \\ u \end{bmatrix}^{\mathsf{T}} \begin{bmatrix} Q & S \\ S^{\mathsf{T}} & R \end{bmatrix} \begin{bmatrix} x \\ u \end{bmatrix} \mathrm{d}t$$

While the noise covariance matrices W, V, and U are general, we assume in this paper that Q, R, and S have a sparsity pattern corresponding to $\Lambda^{\mathsf{T}}\Lambda$. This assumption is discussed in detail in Section 6.

The key constraint is that the controller must have the same broadcast structure as the plant. Since we are seeking a linear time-invariant controller, we will assume the state-space matrices for the controller each share

While each controller has access to different sets of measurements, it is assumed that each controller has full knowledge of the plant parameters. In other words, each controller knows A, B, C, Q, R, S, W, V, and U.

3 Notation

Matrix Manipulation. Before we state our main result, we require some new notation in order to efficiently manipulate structured matrices. Superscripts will be used to denote submatrices or subvectors. For example, if A is defined as in (1), then for $k = 1, \dots, N$ we define:

$$A^{0k} = \begin{bmatrix} A_{00} & 0 \\ A_{k0} & A_{kk} \end{bmatrix} \qquad \text{and} \qquad x^{0k} = \begin{bmatrix} x_0 \\ x_k \end{bmatrix}$$

Note that blocks are indexed starting from 0. Finally, we define the E matrices as operators that pick out certain block-columns. Specifically, they are identity matrices with certain columns removed. We indicate the columns by using subscripts. For example, if $N = 2$,

$$E_0 = \begin{bmatrix} I \\ 0 \\ 0 \end{bmatrix} \quad \text{such that} \quad AE_0 = \begin{bmatrix} A_{00} \\ A_{10} \\ A_{20} \end{bmatrix}$$

The dimensions of the identity matrices within the blocks of E matrices may vary depending on the context. For

example, If there are three states with dimension n_i and three inputs with dimensions m_i, E_0 will have dimension $(n_0+n_1+n_2)\times n_0$ in the expression AE_0, but will have dimension $(m_0+m_1+m_2)\times m_0$ in the expression BE_0.

Expected Values. We write $\mathbb{E}(x\,|\,y)$ as shorthand to denote the expected value of $x(t)$ conditioned on the past history of measurements $y(\tau)$ for $0\le\tau\le t$. This notation is often combined with superscripting. For example, $\mathbb{E}(x_0\,|\,y^{02})$ denotes the expected value of $x_0(t)$ conditioned on the past history of y_0 and y_2.

Riccati Equations. Consider the continuous-time algebraic Riccati equation:

$$A^{\mathsf{T}}X + XA - (XB+S)R^{-1}(XB+S)^{\mathsf{T}} + Q = 0 \quad (2)$$

We say that X is a *stabilizing* solution to (2) if $A+BK$ is Hurwitz, where $K=-R^{-1}(B^{\mathsf{T}}X+S^{\mathsf{T}})$. In such a case, we use the notation:

$$(X,K) = \mathbf{care}(A,B,Q,R,S)$$

For cases of interest in this paper, this notation is always well-defined, because stabilizing solutions will always exist and be unique.

4 Main Result

The main result is an explicit state-space representation of the optimal controller with broadcast-out sparsity structure. Our assumptions are as follows:

A1) (A_{kk},B_{kk}) is stabilizable for $k=1,\dots,N$.

A2) (C_{kk},A_{kk}) is detectable for $k=1,\dots,N$.

A3) $R>0$ and $V>0$.

A4) $\begin{bmatrix}A-j\omega I & B\end{bmatrix}^{\mathsf{T}}\begin{bmatrix}A-j\omega I & B\end{bmatrix}+\begin{bmatrix}Q & S\\ S^{\mathsf{T}} & R\end{bmatrix}>0$ for all $\omega\in\mathbb{R}$.

A5) $\begin{bmatrix}A-j\omega I\\ C\end{bmatrix}\begin{bmatrix}A-j\omega I\\ C\end{bmatrix}^{\mathsf{T}}+\begin{bmatrix}W & U\\ U^{\mathsf{T}} & V\end{bmatrix}>0$ for all $\omega\in\mathbb{R}$.

A6) Q, R, and S have sparsity pattern $\Lambda^{\mathsf{T}}\Lambda$.

The first two assumptions are necessary and sufficient for the existence of a stabilizing controller with the required sparsity pattern [9]. Note that these requirements are stronger than requiring that (C,A,B) be stabilizable and detectable. Assumptions *A3*–*A5* are standard assumptions used in $\mathcal{H}_2$ optimal control to guarantee that the problem has a unique optimal controller which is linear, time-invariant, and finite-dimensional. See for example [20]. Assumption *A6* is discussed in detail in Section 6.

The first step is to compute solutions to several Riccati equations. For the root node, we have the pair:

$$(X,K) = \mathbf{care}(A,B,Q,R,S) \quad (3)$$

$$(\tilde{Y}_0,M_0^{\mathsf{T}}) = \mathbf{care}(A_{00}^{\mathsf{T}},C_{00}^{\mathsf{T}},W_{00},V_{00},U_{00}) \quad (4)$$

For the k^{th} leaf node, we have the pair:

$$(\tilde{X}_k,J_k) = \mathbf{care}(A_{kk},B_{kk},Q_{kk},R_{kk},S_{kk}) \quad (5)$$

$$(Y_{0k},L_{0k}^{\mathsf{T}}) = \mathbf{care}\big((A^{0k})^{\mathsf{T}},(C^{0k})^{\mathsf{T}},W^{0k},V^{0k},U^{0k}\big) \quad (6)$$

Under Assumptions *A1*–*A5*, (3)–(6) have unique stabilizing solutions. The next step is to solve a system of simultaneous linear equations in the variables Ψ_k and Φ_k. There are two sets of equations, each indexed over $k=1,\dots,N$:

$$\begin{aligned}&(A_{kk}+B_{kk}J_k)\Psi_k+\Psi_k(A_{00}+M_0C_{00})^{\mathsf{T}}\\ &\qquad -B_{kk}R_{kk}^{-1}\big(B_{kk}^{\mathsf{T}}\Phi_k+S_{0k}^{\mathsf{T}}\big)(\tilde{Y}_0-Y_0)\\ &\qquad +\big(A_{k0}\tilde{Y}_0+U_{k0}M_0^{\mathsf{T}}+W_{k0}-B_{kk}J_kY_{0k}\big)=0\end{aligned} \quad (7)$$

$$\begin{aligned}&(A_{kk}+B_{kk}J_k)^{\mathsf{T}}\Phi_k+\Phi_k(A_{00}+M_0C_{00})\\ &\qquad -\sum_{i=1}^{N}(\delta_{ik}\tilde{X}_k-X_{ki})\big(\Psi_iC_{00}^{\mathsf{T}}+U_{i0}\big)V_{00}^{-1}C_{00}\\ &\qquad +\big(\tilde{X}_kA_{k0}+J_kS_{0k}^{\mathsf{T}}+Q_{k0}-X_{k0}M_0C_{00}\big)=0\end{aligned} \quad (8)$$

One can verify that (7)–(8) are $n_0(n_1+\cdots+n_N)$ linear equations in as many unknowns. While they appear complicated because they are coupled, they are in principle much easier to solve than the Riccati equation for X, which is quadratic and involves more variables. Finally, define the new gains for $K=1,\dots,N$:

$$\hat{K}_k = -R_{kk}^{-1}\left(B_{kk}^{\mathsf{T}}\Phi_k+S_{0k}^{\mathsf{T}}\right) \quad (9)$$

$$\hat{L}_k = -\left(\Psi_kC_{00}^{\mathsf{T}}+U_{k0}\right)V_{00}^{-1} \quad (10)$$

Theorem 1. *The optimal broadcast-out controller for the plant described in Section 2 is given by the following state-space equations. The controller for the root node has state ζ, with dimension $n_0+\cdots+n_N$:*

$$\begin{aligned}\dot{\zeta}_0 &= A_{00}\zeta_0+B_{00}u_0-M\big(y_0-C_{00}\zeta_0\big)\\ \dot{\zeta}_k &= A_{k0}\zeta_0+A_{kk}\zeta_k+B_{k0}u_0-\hat{L}_k\big(y_0-C_{00}\zeta_0\big)\\ u_0 &= E_0^{\mathsf{T}}K\zeta\end{aligned}$$

The controller for the k^{th} leaf node requires a new state ξ_k, which has state dimension n_0+n_k:

$$\begin{aligned}\dot{\xi}_k &= A^{0k}\xi_k+B^{0k}u^{0k}-L_{0k}\big(y^{0k}-C^{0k}\xi_k\big)\\ u_k &= E_k^{\mathsf{T}}K\zeta+\hat{K}_k(\xi_0-\zeta_0)+J_k(\xi_k-\zeta_k)\end{aligned}$$

Note that each of the leaf nodes also depends on ζ. This does not violate the broadcast constraint, since each leaf node has access to y_0 and can therefore compute ζ. An outline of the proof of Theorem 1 is provided in Section 8.

5 Estimation Structure

The solution provided by Theorem 1 is an explicit state-space realization of the optimal decentralized controller. The following theorem gives a physical interpretation of its states.

Theorem 2. *The states of the optimal broadcast-out controller described by Theorem 1 have the following interpretations:*

(i) $\zeta_0 = \mathbb{E}(x_0 \,|\, y_0)$

(ii) $\zeta_k = \mathbb{E}(x_k \,|\, y_0)$ *for* $k = 1, \dots, N$

(iii) $\xi_k = \mathbb{E}\big(x^{0k} \,|\, y^{0k}\big)$ *for* $k = 1, \dots, N$

Proof. Items (i) and (iii) are immediate. For example, the Kalman filter $\mathbb{E}(x^{0k} \,|\, y^{0k})$ can be computed by extracting the relevant state equations from the plant:

$$\begin{aligned} \dot{x}^{0k} &= A^{0k} x^{0k} + B^{0k} u^{0k} + w^{0k} \\ y^{0k} &= C^{0k} x^{0k} + v^{0k} \end{aligned}$$

Because the k^{th} leaf node has access to y^{0k}, it may also compute u^{0k}. Thus, this is a centralized Kalman filtering problem, and the optimal estimator is:

$$\dot{\xi}_k = A^{0k} \xi_k + B^{0k} u^{0k} - L_{0k}\big(y^{0k} - C^{0k} \xi_k\big)$$

which matches the state equation in Theorem 1. Item (ii) is more involved because x_k depends on u_k, which is not known to the root node. The equations can nevertheless be verified by closing the loop on the u_k using the optimal controller from Theorem 1 and solving for the corresponding estimator, which has an augmented state. The details are omitted due to space constraints. ∎

Another useful interpretation of the optimal controller is as the sum of a baseline policy that uses the information common to all controllers, plus correction terms that account for additional private information. Structural results of this type were explored in [8] for more general Markov decision processes. In the broadcast case, y_0 is the common information, and indeed, the optimal control input is of the form:

$$u = K\zeta + (\text{correction terms})$$

where $\zeta = \mathbb{E}(x \,|\, y_0)$. The leaf nodes benefit from the additional knowledge of y_k, and so the control policy for u_k contains correction terms that depend on the estimation error $(\xi^{0k} - \zeta^{0k}) = \big(\mathbb{E}(x^{0k} \,|\, y^{0k}) - \mathbb{E}(x^{0k} \,|\, y_0)\big)$. In the case where the y_k are uninformative for $k = 1, \dots, N$, the correction terms are zero and we recover the optimal centralized controller based on the measurement y_0.

6 Discussion

The assumption $A6$ turns out to be very important. As an illustrative example, consider the case $N = 2$. The sparsity patterns are:

$$A = \begin{bmatrix} A_{00} & 0 & 0 \\ A_{10} & A_{11} & 0 \\ A_{20} & 0 & A_{22} \end{bmatrix}, \qquad Q = \begin{bmatrix} Q_{00} & Q_{01} & Q_{02} \\ Q_{10} & Q_{11} & 0 \\ Q_{20} & 0 & Q_{22} \end{bmatrix}$$

where B and C have patterns similar to A, and R and S have patterns similar to Q. The physical interpretation is that the cost function is not coupled amongst the leaf nodes of the graph. No such constraints are required for the covariance terms W, V, and U.

A key simplification happens in the proof of Theorem 1, namely when $\mathcal{Q}_{kk}$ is computed as a function of the other controllers. In the absence of Assumption $A6$, the expression does not simplify, and our approach fails to obtain the optimal controller.

Numerical computations were carried out using the vectorization method of [12], and it was found that Assumption $A6$ is necessary in Theorem 1; if Q, R, S are general, the optimal controller has a larger state dimension than the one reported in Theorem 1. A possible explanation is that having a coupled cost function between leaf nodes would require each leaf node to also estimate the states of the other leaf nodes. Simply computing $\mathbb{E}(x_0 \,|\, y_0)$ and $\mathbb{E}\big(x^{0k} \,|\, y^{0k}\big)$ as indicated by Theorem 2 might be insufficient.

This evidence suggests that a new approach might be required to solve the case with a general cost function, and that the structure of the optimal controller would be different from the one reported herein. This an avenue of potential future research.

Nevertheless, the result of Theorem 1 is a generalization of the previous result on two-player architectures [6]. Indeed, when we set $N = 1$, we recover the two-player case. Note that Assumption $A6$ is nonrestrictive in the case $N = 1$, so it may be removed.

7 Broadcast-in Structures

In Sections 4–6, we only considered the broadcast-out structure of Figure 1. It turns out that the broadcast-in case can be solved in a similar manner, and the result provides a dual interpretation of Theorem 1. The model-matching formulation (12) is given by:

$$\begin{aligned} \text{minimize} \quad & \big\| \mathcal{P}_{11} + \mathcal{P}_{12} \mathcal{Q} \mathcal{P}_{21} \big\| \\ \text{subject to} \quad & \mathcal{Q} \in \mathcal{RH}_\infty \cap \Lambda \end{aligned}$$

Suppose $\mathcal{Q}_{\text{opt}}$ is the solution to this problem. Since the $\mathcal{H}_2$ norm does not change if we transpose the operand, the solution to the dual optimization problem

$$\begin{aligned} \text{minimize} \quad & \big\| \mathcal{P}_{11}^{\mathsf{T}} + \mathcal{P}_{21}^{\mathsf{T}} \overline{\mathcal{Q}} \mathcal{P}_{12}^{\mathsf{T}} \big\| \\ \text{subject to} \quad & \overline{\mathcal{Q}} \in \mathcal{RH}_\infty \cap \Lambda^{\mathsf{T}} \end{aligned}$$

is given by $\overline{\mathcal{Q}}_{\text{opt}} = \mathcal{Q}_{\text{opt}}^{\mathsf{T}}$. Note that Λ^{T} is the sparsity pattern of the broadcast-in architecture. Thus, we may solve broadcast-in problems by transposing the appropriate system parameters and solving the dual broadcast-out problem via Theorem 1.

A point worth noticing is that under a broadcast-in structure, Assumption $A6$ becomes a constraint on the noise covariance matrices W, V, and U rather than a constraint on the cost matrices. So in the broadcast-in cast, the noise processes driving each of the leaf nodes must be mutually independent.

8 Proof Outline of the Main Result

In this section, we provide an outline of the proof of Theorem 1 in the case where A is Hurwitz. The result still holds for unstable plants, but the proof is more involved and omitted due to space constraints. Roughly speaking, a Youla-like parameterization of stabilizing controllers is carried out similar to the one found in [9]. This converts the problem to a stable structured model-matching problem similar to (12).

The problem of Section 2 can be converted to a sparsity-constrained $\mathcal{H}_2$ model-matching problem. For a thorough introduction on $\mathcal{H}_2$ optimal control and associated notation, see for example [2, 20]. In the case when A is Hurwitz, we may pose the problem as:

$$\begin{array}{ll} \text{minimize} & \left\| \mathcal{P}_{11} + \mathcal{P}_{12}\mathcal{K}\left(I - \mathcal{P}_{22}\mathcal{K}\right)^{-1}\mathcal{P}_{21} \right\| \\ \text{subject to} & \mathcal{K} \in \mathcal{RL}_\infty \cap \Lambda \\ & \mathcal{K} \text{ is stabilizing} \end{array} \tag{11}$$

where:

$$\begin{bmatrix} \mathcal{P}_{11} & \mathcal{P}_{12} \\ \mathcal{P}_{21} & \mathcal{P}_{22} \end{bmatrix} = \left[\begin{array}{c|cc} A & B_1 & B \\ \hline C_1 & 0 & D_{12} \\ C & D_{21} & 0 \end{array}\right]$$

and B_1, C_1, D_{12}, and D_{21} satisfy:

$$\begin{bmatrix} Q & S \\ S^\mathsf{T} & R \end{bmatrix} = \begin{bmatrix} C_1 & D_{12} \end{bmatrix}^\mathsf{T} \begin{bmatrix} C_1 & D_{12} \end{bmatrix}$$

$$\begin{bmatrix} W & U \\ U^\mathsf{T} & V \end{bmatrix} = \begin{bmatrix} B_1 \\ D_{21} \end{bmatrix} \begin{bmatrix} B_1 \\ D_{21} \end{bmatrix}^\mathsf{T}$$

Now use the change of variables $\mathcal{Q} = \mathcal{K}\left(I - \mathcal{P}_{22}\mathcal{K}\right)^{-1}$. Since $\mathcal{P}_{22} \in \Lambda$ and $\mathcal{K} \in \Lambda$, the sparsity constraint is preserved: $\mathcal{Q} \in \Lambda$. Equation (11) becomes:

$$\begin{array}{ll} \text{minimize} & \left\| \mathcal{P}_{11} + \mathcal{P}_{12}\mathcal{Q}\mathcal{P}_{21} \right\| \\ \text{subject to} & \mathcal{Q} \in \mathcal{RH}_\infty \cap \Lambda \end{array} \tag{12}$$

The following lemma shows how (12) can be split up into $N+1$ coupled optimization problems.

Lemma 3. *$\mathcal{Q}$ is optimal for (12) if and only if:*

(i) $\mathcal{Q}_{00}, \mathcal{Q}_{10}, \ldots, \mathcal{Q}_{N0}$ are optimal for:

$$\begin{array}{ll} \textit{minimize} & \left\| \mathcal{P}_{11} + \mathcal{P}_{12}\mathcal{Q}\mathcal{P}_{21} \right\| \\ \textit{subject to} & \mathcal{Q}_{k0} \in \mathcal{RH}_\infty,\ k = 0, \ldots, N \end{array} \tag{13}$$

(ii) for $k = 1, \ldots, N$, $\mathcal{Q}_{k0}$ and $\mathcal{Q}_{kk}$ are optimal for:

$$\begin{array}{ll} \textit{minimize} & \left\| \mathcal{P}_{11} + \mathcal{P}_{12}\mathcal{Q}\mathcal{P}_{21} \right\| \\ \textit{subject to} & \mathcal{Q}_{k0}, \mathcal{Q}_{kk} \in \mathcal{RH}_\infty \end{array} \tag{14}$$

Proof. The result follows from comparing the optimality conditions of (12)–(14). ∎

Lemma 3 shows that solving the structured model-matching problem (12) is equivalent to *simultaneously* solving $N+1$ optimization problems. The key property of (13)–(14) is that they are each centralized. For example, for any $\mathcal{Q}_{11}, \ldots, \mathcal{Q}_{NN}$, (13) can be written as:

$$\begin{array}{ll} \text{minimize} & \left\| \overline{\mathcal{P}}_{11} + \mathcal{P}_{12} \begin{bmatrix} \mathcal{Q}_{00} \\ \vdots \\ \mathcal{Q}_{N0} \end{bmatrix} E_0^\mathsf{T} \mathcal{P}_{21} \right\| \\ \text{subject to} & \mathcal{Q}_{k0} \in \mathcal{RH}_\infty,\ k = 0, \ldots, N \end{array}$$

where

$$\overline{\mathcal{P}}_{11} = \mathcal{P}_{11} + \mathcal{P}_{12} \begin{bmatrix} 0 & 0 & \cdots & 0 \\ 0 & \mathcal{Q}_{11} & & \\ \vdots & & \ddots & \\ 0 & & & \mathcal{Q}_{NN} \end{bmatrix} \mathcal{P}_{21}$$

Similarly, (14) can be written as a centralized problem in the variable $\begin{bmatrix} \mathcal{Q}_{k0} & \mathcal{Q}_{kk} \end{bmatrix}$. The main approach is to solve each of these problems separately as if the remaining variables were fixed. Then, compare the solutions and enforce compatibility. First consider (13), and suppose $\mathcal{Q}_{kk}$ has a state-space realization $\left[\begin{array}{c|c} A_{Qkk} & B_{Qkk} \\ \hline C_{Qkk} & 0 \end{array}\right]$ for $k = 1, \ldots, N$. If we solve for $\mathcal{Q}_{00}$ in terms of the $\mathcal{Q}_{kk}$, we obtain:

$$\mathcal{Q}_{00} = \left[\begin{array}{cc|c} A + BK & \begin{bmatrix} M_0 \\ \hat{L} \end{bmatrix} C_{00} & \begin{bmatrix} M_0 \\ \hat{L} \end{bmatrix} \\ 0 & A_{00} + M_0 C_{00} & M_0 \\ \hline -E_0^\mathsf{T} K & 0 & 0 \end{array}\right]$$

where M_0 and K come from (3)–(4), the block-rows of $\hat{L}$ are given by (10), and the Ψ_k satisfy the following independent equations for $k = 1, \ldots, N$:

$$A_{kk}\Theta_Y^k + \Theta_Y^k (A_{00} + M_0 C_{00})^\mathsf{T} + (A_{k0}\tilde{Y}_0 + U_{k0} M_0^\mathsf{T} + W_{k0}) = 0 \tag{15}$$

$$A_{Qkk}\Gamma_Y^k + \Gamma_Y^k (A_{00} + M_0 C_{00})^\mathsf{T} + B_{Qkk}(V_{k0} M_0^\mathsf{T} + U_{0k}^\mathsf{T} + C_{k0} + C_{kk}\Theta_Y^k) = 0 \tag{16}$$

$$A_{kk}\Psi_k + \Psi_k (A_{00} + M_0 C_{00})^\mathsf{T} + (A_{k0}\tilde{Y}_0 + U_{k0} M_0^\mathsf{T} + W_{k0} + B_{kk} C_{Qkk} \Gamma_Y^k) = 0 \tag{17}$$

A similar procedure is explained in greater detail in [7] for the case $N = 1$. Note that the poles of $\mathcal{Q}_{00}$ do not depend on the unknown $\mathcal{Q}_{kk}$.

Next, we consider (14), and solve for $\mathcal{Q}_{kk}$ Again, we find that the poles are independent of the unknown $\mathcal{Q}$'s:

$$\mathcal{Q}_{kk} = \left[\begin{array}{cc|c} A_{kk} + B_{kk} J_k & B_{kk} \begin{bmatrix} \hat{K}_k & J_k \end{bmatrix} & 0 \\ 0 & A^{0k} + L_{0k} C^{0k} & -L_{0k} \begin{bmatrix} 0 \\ I \end{bmatrix} \\ \hline J_k & \begin{bmatrix} \hat{K}_k & J_k \end{bmatrix} & 0 \end{array}\right]$$

This is the point where we rely on Assumption *A*6. Indeed, if Assumption *A*6 fails, we cannot simplify $\mathcal{Q}_{kk}$, and we find that in general, its poles depend on the poles of the other leaf nodes.

Compatibility is enforced by substituting the expression for $\mathcal{Q}_{kk}$ into (15)–(17), since we now know A_{Qkk}, B_{Qkk}, and C_{Qkk}. We then show that (15)–(17) hold when

$$\Gamma_Y^k = \begin{bmatrix} \Psi_k - \Theta_Y^k \\ \tilde{Y}_0 - Y_{00} \\ \Theta_Y^k - Y_{k0} \end{bmatrix}$$

The final result is (8), which depends only on Φ_k and Ψ_k. Now perform a similar procedure by substituting the expression for $\mathcal{Q}_{00}$ into the analogous equations to (15)–(17), derived this time from (14). The final result is (7). Now that we have found the diagonal entries $\mathcal{Q}_{kk}$, it remains to return to (13), and solve for the missing $\mathcal{Q}_{10}, \ldots, \mathcal{Q}_{N0}$. Once we have assembled the optimal $\mathcal{Q}$, we can recover $\mathcal{K}$ via $\mathcal{K} = \mathcal{Q}(I + \mathcal{P}_{22}\mathcal{Q})^{-1}$, and the main result of Theorem 1 follows after simplifications. ∎

9 Conclusion

In this paper, we presented the solution to decentralized output-feedback control problems with a broadcast architecture. The main result, Theorem 1, provides an explicit state-space realization of the optimal controller, which was previously not known. Computing the optimal controller requires solving some standard algebraic Riccati equations as well as a set of coupled linear equations. The computational complexity of the decentralized case is comparable to that of the centralized case. An interpretation of the controller states as estimators which compute conditional expectations is also provided.

References

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

[2] G. E. Dullerud and F. Paganini. *A Course in Robust Control Theory: A Convex Approach*. Springer Verlag, 2000.

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

[4] J.-H. Kim and S. Lall. Separable optimal cooperative control problems. In *American Control Conference*, 2012. (To Appear).

[5] A. Lamperski and J. C. Doyle. On the structure of state-feedback LQG controllers for distributed systems with communication delays. In *IEEE Conference on Decision and Control*, 2011. (To Appear).

[6] 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, 2011.

[7] L. Lessard and S. Lall. Optimal controller synthesis for the decentralized two-player problem with output feedback. In *American Control Conference*, 2012. (To Appear).

[8] A. Mahajan, A. Nayyar, and D. Teneketzis. Identifying tractable decentralized problems on the basis of information structures. In *Allerton Conference on Communication, Control, and Computing*, pages 1440–1449, 2008.

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

[10] A. Rantzer. Linear quadratic team theory revisited. In *American Control Conference*, pages 1637–1641, 2006.

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

[13] N. Sandell, Jr. and M. Athans. Solution of some nonclassical LQG stochastic decision problems. *IEEE Transactions on Automatic Control*, 19(2):108–116, 1974.

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

[15] P. Shah and P. A. Parrilo. $\mathcal{H}_2$-optimal decentralized control over posets: A state space solution for state-feedback. In *IEEE Conference on Decision and Control*, pages 6722–6727, 2010.

[16] J. Swigart. *Optimal Controller Synthesis for Decentralized Systems*. PhD thesis, Stanford University, 2010.

[17] J. Swigart and S. Lall. An explicit state-space solution for a decentralized two-player optimal linear-quadratic regulator. In *American Control Conference*, pages 6385–6390, 2010.

[18] J. Swigart and S. Lall. Optimal controller synthesis for a decentralized two-player system with partial output feedback. In *American Control Conference*, pages 317–323, 2011.

[19] H. S. Witsenhausen. A counterexample in stochastic optimum control. *SIAM Journal on Control*, 6:131, 1968.

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