H[subscript 2]-Optimal Decentralized Control Over Posets: A State-Space Solution for State-Feedback

$\mathcal{H}_2$-Optimal Decentralized Control over Posets: A State-Space Solution for State-Feedback

Parikshit Shah and Pablo A. Parrilo

Abstract—We develop a complete state-space solution to $\mathcal{H}_2$-optimal decentralized control of poset-causal systems with state-feedback. Our solution is based on the exploitation of a key separability property of the problem, that enables an efficient computation of the optimal controller by solving a small number of uncoupled standard Riccati equations. Our approach gives important insight into the structure of optimal controllers, such as controller degree bounds that depend on the structure of the poset. A novel element in our state-space characterization of the controller is an intuitive description of the controller as an aggregation of local control laws.

I. Introduction

Finding computationally efficient algorithms to design decentralized controllers is a challenging area of research (see e.g. [19], [5] and the references therein). Current research suggests that while the problem is hard in general, certain classes with special information structures are tractable via convex optimization techniques. In past work, the authors have argued that communication structures modeled by partially ordered sets (or *posets*) provide a rich class of decentralized control systems (which we call *poset-causal* systems) that are amenable to such an approach [19]. Posets have appeared in the control theory literature earlier in the context of team theory [11], and specific posets (chains) have been studied in the context of decentralized control [27]. Poset-causal systems are also related to the class of systems studied more classically in the context of *hierarchical systems* [13], [10], where abstract notions of hierarchical organization of large-scale systems were introduced and their merits were argued for.

While it is possible to design optimal decentralized controllers for a fairly large class of systems known as quadratically invariant systems in the frequency domain via the Youla parametrization [16], there are some important drawbacks with such an approach. Typically Youla domain techniques are not computationally efficient, and the degree of optimal controllers synthesized with such techniques is not always well-behaved. In addition to computational efficiency, issues related to numerical stability also arise. Typically, operations at the transfer function level are inherently less stable numerically. Moreover, such approaches do not provide insight into the structure of the optimal controller. These drawbacks emphasize the need for *state-space techniques* to synthesize optimal decentralized controllers. State-space techniques are usually computationally efficient, numerically stable, and provide degree bounds for optimal controllers. In our case we will also show that the solution provides important insight into the structure of the controller.

In this paper we consider the problem of designing $\mathcal{H}_2$ optimal decentralized controllers for poset-causal systems. The control objective is the design of optimal feedback laws that have access to local state information. We emphasize here that different subsystems *do not* have access to the global state, but only the local states of the systems in a sense that will be made precise in the next section. The main contributions in the paper are as follows:

- We show a certain crucial separability property of the problem under consideration. This result is outlined in Theorem 2. This makes it possible to decompose the decentralized control problem over posets into a collection of standard centralized control problems.
- We give an explicit state-space solution procedure in Theorem 3. To construct the solution, one needs to solve standard Riccati equations (corresponding to the different sub-problems). Using the solutions of these Riccati equations, one constructs certain block matrices and provides a state-space realization of the controller.
- We provide bounds on the degree of the optimal controller in terms of a parameter $\sigma_{\mathcal{P}}$ that depends only on the order-theoretic structure of the poset (Corollary 2).
- In Theorem 4 we briefly describe the structural

form of the optimal controller. We introduce a transfer function Θ, which we call the *differential filter*. The transfer function Θ computes a notion of generalized differentials of state predictions at different subsystems. The discussion related to structural aspects of the controller is brief and informal in this paper and have been formalized in the paper [22].

- We state a new and intuitive decomposition of the structure of the optimal controller into local control laws. Each local control law computes constitutes a local gain operating on the generalized differential of the local state predictions. The overall control law is a superposition of these local control laws.

A. Related Work

It is well-known that in general decentralized control is a hard problem, and significant research efforts have been directed towards its different aspects; see for instance the classical survey [17] for some of the earlier results. More recently, Blondel and Tsitsiklis [4] have shown that in certain instances, decentralized control problems are computationally intractable, in particular they show that the problem of finding bounded-norm, block-diagonal stabilizing controllers in the presence of output-feedback is NP-hard. Nevertheless the computation of distributed controllers for different classes of problems remains an active and important research area e.g. [2]. An important development due to Voulgaris [27], [28] was the observation that in several cases, decentralized control problems were amenable to exact convex reparametrization and therefore computationally tractable. Rotkowitz and Lall generalized these ideas in terms of a property called quadratic invariance [16], we discuss connections to their work later. In past work [19], we have shown that posets provide a unifying umbrella to describe these tractable examples under an appealing theoretical framework.

Partially ordered sets (posets) are very well studied objects in combinatorics. The associated notions of incidence algebras and Galois connections were first studied by Rota [15] in a combinatorics setting. Since then, order-theoretic concepts have been used in engineering and computer science; we mention a few specific works below. In control theory, ideas from order theory have been used in different ways.

Ho and Chu used posets to study team theory problems [11]. They were interested in sequential decision making problems where agents must make decisions at different time steps. They study computational and structural properties of optimal decision policies when the problems have poset structure. Mullans and Elliot [14] use posets to model the notions of time and causality, and study evolution of systems on locally finite posets. Wyman [29] has studied time-varying linear-systems evolving on locally finite posets in an algebraic framework, including aspects related to realization theory and duality. In computer science, Cousot and Cousot used these ideas to develop tools for formal verification of computer programs in their seminal paper [8]. Del Vecchio and Murray [26] have used ideas from lattice and order theory to construct estimators for continuous states in hybrid systems.

More recently, the authors of this paper have initiated a systematic study of decentralized control problems from the point of view of partial order theory. In [19], we introduce the partial order framework and show how several well-known classes of problems such as nested systems [27] fit into the partial order framework. In [20], we extend this poset framework to spatio-temporal systems and generalize certain results related to the so called "funnel causal systems" of Bamieh et. al [3]. In [21], we show that a class of time-delayed systems known to be amenable to convex reparametrization [16] also has an underlying poset structure. In that paper, we also study the close connections between posets and another class of decentralized control problems known as quadratically invariant problems. While this poset framework provides a lens to view all these examples in a common intuitive framework, a systematic study of state-space approaches has been lacking.

In an interesting paper by Swigart and Lall [24], the authors consider a state-space approach to the $\mathcal{H}_2$ optimal controller synthesis problem over a particular poset with two nodes. Their approach is restricted to the finite time horizon setting (although in a subsequent paper [25], they extend this to the infinite time horizon setting), and uses a particular decomposition of certain optimality conditions. In this setting, they synthesize optimal controllers and provide insight into the structure of the optimal controller. These results are also summarized in the thesis [23]. By using our new separability condition

(which is related to their decomposition property, but which we believe to be more fundamental) we significantly generalize those results in this paper. We provide a solution for *all posets* and for the *infinite time horizon*. In recent work [16], Rotkowitz and Lall proposed a state-space technique to solve $\mathcal{H}_2$ optimal control problems for quadratically invariant systems (which could be used for poset-causal systems). An important drawback of their reformulation is that one would need to solve larger Riccati equations. Our approach for poset-causal systems is more efficient computationally. Moreover, our approach also provides insight into the structure of the optimal controllers.

The rest of this paper is organized as follows. In Section II we introduce the necessary preliminaries regarding posets, the control theoretic framework and notation. In Section III we describe our solution strategy. In Section IV we present the main results. We devote Section V to a discussion of the main results, and their illustration via examples.

II. Preliminaries

In this section we introduce some concepts from order theory. Most of these concepts are well studied and fairly standard, we refer the reader to [1], [9] for details.

A. Posets

Definition 1: A partially ordered set (or *poset*) $\mathcal{P} = (P, \preceq)$ consists of a set P along with a binary relation $\preceq$ which has the following properties:

1) $a \preceq a$ (reflexivity),
2) $a \preceq b$ and $b \preceq a$ implies $a = b$ (antisymmetry),
3) $a \preceq b$ and $b \preceq c$ implies $a \preceq c$ (transitivity).

We will sometimes use the notation $a \prec b$ to denote the strict order relation $a \preceq b$ but $a \neq b$.

In this paper we will deal with finite posets (i.e. $|P|$ is finite). It is possible to represent a poset graphically via a *Hasse diagram* by representing the transitive reduction of the poset as a graph (i.e. by drawing only the minimal order relations graphically, a downward arrow representing the relation $\preceq$, with the remaining order relations being implied by transitivity).

Example 1: An example of a poset with three elements (i.e., $P = \{1, 2, 3\}$) with order relations $1 \preceq 2$ and $1 \preceq 3$ is shown in Figure 1(b).
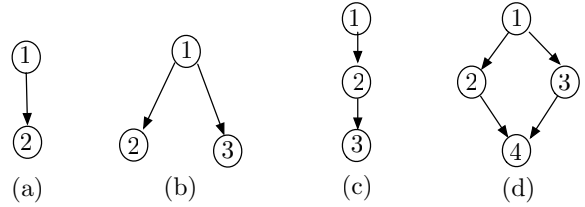

Fig. 1.   Hasse diagrams of some posets.

Let $\mathcal{P} = (P, \preceq)$ be a poset and let $p \in P$. We define $\downarrow p = \{q \in P \mid p \preceq q\}$ (we call this the *downstream set*). [‡] Let $\downarrow\downarrow p = \{q \in P \mid p \preceq q, q \neq p\}$. Similarly, let $\uparrow p = \{q \in P \mid q \preceq p\}$ (called the *upstream set*), and $\uparrow\uparrow p = \{q \in P \mid q \preceq p, q \neq p\}$. We define $\downarrow\uparrow p = \{q \in P \mid q \not\preceq p, q \not\succeq p\}$ (called the *off-stream set*); this is the set of *uncomparable* elements that have no order relation with respect to p. Define an *interval* $[i, j] = \{p \in P \mid i \preceq p \preceq j\}$. A *minimal element* of the poset is an element $p \in P$ such that if $q \preceq p$ for some $q \in P$ then $q = p$. (A *maximal element* is defined analogously).

In the poset shown in Figure 1(d), $\downarrow 1 = \{1, 2, 3, 4\}$, whereas $\downarrow\downarrow 1 = \{2, 3, 4\}$. Similarly $\uparrow\uparrow 1 = \emptyset$, $\uparrow 4 = \{1, 2, 3, 4\}$, and $\uparrow\uparrow 4 = \{1, 2, 3\}$. The set $\downarrow\uparrow 2 = \{3\}$.

Definition 2: Let $\mathcal{P} = (P, \preceq)$ be a poset. Let $\mathcal{R}$ be a ring. The set of all functions $f : P \times P \to \mathcal{R}$ with the property that $f(x, y) = 0$ if $y \not\preceq x$ is called the *incidence algebra* of $\mathcal{P}$ over $\mathcal{R}$. It is denoted by $\mathcal{I}(\mathcal{P})$. [*]

When the poset $\mathcal{P}$ is finite, the elements in the incidence algebra may be thought of as matrices with a specific sparsity pattern given by the order relations of the poset in the following way. One indexes the rows and columns of the matrices by the elements of P. Indeed, if $f \in \mathcal{I}(\mathcal{P})$ then f has a

---

[‡]We have reversed conventions with respect to some of our conference papers, wherein the Hasse diagrams are drawn with upward arrows and the set $\uparrow p$ corresponds to the set $\{q \in P \mid p \preceq q\}$. The present convention has been adopted to make the presentation more intuitive. For example the downstream set at p corresponds to elements drawn lower in the Hasse diagram. It also corresponds to the elements that are "in the future" with respect to p, in keeping with the intuition that information in a river propagates "downstream".

[*]Standard definitions of the incidence algebra use an opposite convention, namely $f(x, y) = 0$ if $x \not\preceq y$. Thus, the matrix representation of the incidence algebra is typically a transposal of the matrix representations that appear here. For example, while the incidence algebra of a chain is the set of lower-triangular matrices in this paper, in standard treatments it would appear as upper-triangular matrices. We reverse the convention so that in a control theoretic setting one may interpret such matrices as representing *poset-causal* maps.

matrix representation M such that $f(i,j) = M_{ij}$. An example of an element of $\mathcal{I}(\mathcal{P})$ for the poset from Example 1 (Fig. 1(b)) is:

$$\zeta_{\mathcal{P}} = \begin{bmatrix} 1 & 0 & 0 \\ 1 & 1 & 0 \\ 1 & 0 & 1 \end{bmatrix}.$$

Given two functions $f, g \in \mathcal{I}(\mathcal{P})$, their sum $f + g$ and scalar multiplication cf are defined as usual. The product $h = f \cdot g$ is defined by $h(x,y) = \sum_{z \in P} f(x,z) g(z,y)$. Note that the above definition of function multiplication is made so that it is consistent with standard matrix multiplication.

Lemma 1: Let $\mathcal{P}$ be a poset. Under the usual definition of addition and multiplication as defined in (1) the incidence algebra is an associative algebra (i.e. it is closed under addition, scalar multiplication and function multiplication).

Proof: The proof is standard, see for example [19]. ∎

Given $i \le j$, let $[i \to j]$ denote the set of all chains from i to j of the form $\{i, i_1\}, \ldots, \{i_k, j\}$ such that $i \le i_1 \le \cdots \le i_k \le j$. For example, in the poset in Fig. 1(c), $[1 \to 3] = \{\{\{1,2\}, \{2,3\}\}, \{1,3\}\}$. A standard corollary of Lemma 1 is the following.

Corollary 1: Suppose $A \in \mathcal{I}(\mathcal{P})$. Then A is invertible if and only if A_{ii} is invertible for all $i \in P$. Furthermore $A^{-1} \in \mathcal{I}(\mathcal{P})$, and the inverse is given by:

$$[A^{-1}]_{ij} = \begin{cases} A_{ii}^{-1} \sum_{p_{ij} \in [j \to i]} \prod_{\{l,k\} \in p_{ij}} (-A_{lk} A_{kk}^{-1}) & \text{if } i \ne j \\ A_{ii}^{-1} & \text{if } i = j. \end{cases}$$

B. Control Theoretic Preliminaries

We consider the following state-space system in continuous time:

$$\begin{aligned} \dot{x}(t) &= Ax(t) + Fw(t) + Bu(t) \\ z(t) &= Cx(t) + Du(t). \end{aligned} \quad (1)$$

In this paper we present the continuous time case only, however, we wish to emphasize that analogous results hold in discrete time in a straightforward manner. In this paper we consider what we will call *poset-causal systems*. We think of the system matrices (A, B, C, D, F) to be partitioned into blocks in the following natural way. Let $\mathcal{P} = (P, \le)$ be a poset with $P = \{1, \ldots, p\}$. We think of this system as being divided into p sub-systems, with sub-system i having some states $x_i(t) \in \mathbb{R}^{n_i}$, and

control inputs $u_i(t) \in \mathbb{R}^{m_i}$ for $i \in \{1, \ldots, p\}$. The external output is $z(t) \in \mathbb{R}^l$. The signal $w(t)$ is a disturbance signal. (To use certain standard state-space factorization results, we assume that $C^T D = 0$ and $D^T D > 0$, these assumptions can be relaxed in a straightforward way). The states and inputs are partitioned in the natural way such that the sub-systems correspond to elements of the poset $\mathcal{P}$ with $x(t) = \begin{bmatrix} x_1(t) | x_2(t) | \ldots | x_p(t) \end{bmatrix}^T$, and $u(t) = \begin{bmatrix} u_1(t) | u_2(t) | \ldots | u_p(t) \end{bmatrix}^T$. This naturally partitions the matrices A, B, C, D, F into appropriate blocks so that $A = \begin{bmatrix} A_{ij} \end{bmatrix}_{i,j \in P}$, $B = \begin{bmatrix} B_{ij} \end{bmatrix}_{i,j \in P}$, $C = \begin{bmatrix} C_j \end{bmatrix}_{j \in P}$ (partitioned into columns), $D = \begin{bmatrix} D_j \end{bmatrix}_{j \in P}$, $F = \begin{bmatrix} F_{ij} \end{bmatrix}_{i,j \in P}$. (We will throughout deal with matrices at this block-matrix level, so that A_{ij} will unambiguously mean the (i,j) block of the matrix A.) Using these block partitions, one can define the incidence algebra at the block matrix level in the natural way. We denote by $\mathcal{I}_A(\mathcal{P}), \mathcal{I}_B(\mathcal{P})$ the block incidence algebras corresponding to the block partitions of A and B.

We will further assume that F is block diagonal and full column rank, so that it is left invertible with a left inverse that is also block diagonal. [†] Often, matrices will have different (but compatible) dimensions and the block structure will be clear from the context. In these cases, we will abuse notation and will drop the subscript and simply write $\mathcal{I}(\mathcal{P})$.

We say that a system is $\mathcal{P}$-*poset-causal* (or simply poset-causal) if its plant $P_{22} := (sI - A)^{-1} B \in \mathcal{I}(\mathcal{P})$. In this paper we will in fact make a stronger realizability assumption, namely that $A \in \mathcal{I}_A(\mathcal{P})$ and $B \in \mathcal{I}_B(\mathcal{P})$. It is known [12] that not every poset-causal system necessarily has a structured realization.

Example 2: We use this example to illustrate ideas and concepts throughout this paper. Consider the system

$$\begin{aligned} \dot{x}(t) &= Ax(t) + Fw(t) + Bu(t) \\ z(t) &= Cx(t) + Du(t) \\ y(t) &= x(t), \end{aligned}$$

[†]More generally we can assume that for the system under consideration (1), $F \in \mathcal{I}(\mathcal{P})$ and the diagonal blocks are full column rank. Operating under this assumption, one can perform an invertible coordinate transformation $T \in \mathcal{I}(\mathcal{P})$ on the states so that $T^{-1}F$ is block diagonal. Since T can be chosen to be in the incidence algebra, $T^{-1}AT, T^{-1}B \in \mathcal{I}(\mathcal{P})$. Hence, without loss of generality, we assume that F is block diagonal.

with matrices

$$A = \begin{bmatrix} -0.5 & 0 & 0 & 0 \\ -1 & -0.25 & 0 & 0 \\ -1 & 0 & -0.2 & 0 \\ -1 & -1 & -1 & -0.1 \end{bmatrix} \quad B = \begin{bmatrix} 1 & 0 & 0 & 0 \\ 1 & 1 & 0 & 0 \\ 1 & 0 & 1 & 0 \\ 1 & 1 & 1 & 1 \end{bmatrix}$$
(2)

$$C = \begin{bmatrix} I^{4\times4} \\ 0^{4\times4} \end{bmatrix} \qquad F = I \quad D = \begin{bmatrix} 0^{4\times4} \\ I^{4\times4} \end{bmatrix}.$$
(3)

This system is poset-causal with the underlying poset described in Fig. 1(d). Note that in this system, each subsystem has a single input, a single output and a single state. The matrices A and B are in the incidence algebra of the poset. Furthermore, $F = I$.

Recall that the standard notion of causality in systems theory is based crucially on an underlying totally ordered index set (time). Systems (in LTI theory these are described by impulse responses) are said to be *causal* if the support of the impulse response is consistent with the ordering of the index set: an impulse at time zero is only allowed to propagate in the increasing direction with respect to the ordering. This notion of causality can be readily extended to situations where the underlying index set is only *partially ordered*. Indeed this abstract setup has been studied by Mullans and Elliott [14], and an interesting algebraic theory of systems has been developed.

Our notion of poset-causality is very much in the same spirit. We call such systems poset-causal due to the following analogous property among the sub-systems. If an input is applied to sub-system i via u_i at some time t, the effect of the input is seen by the states x_j for all sub-systems $j \in \downarrow i$ (at or after time t). Thus $\downarrow i$ may be seen as the cone of influence of input i. We refer to this causality-like property as *poset-causality*. This notion of causality enforces (in addition to causality with respect to time), causality with respect to the subsystems via a poset. For most of this paper we will deal with systems that are poset-causal (with respect to some arbitrary but fixed finite poset $\mathcal{P}$). Before we turn to the problem of optimal control we state an important result regarding stabilizability of poset-causal systems by poset-causal controllers.

Theorem 1: The poset-causal system (1) is stabilizable by a poset-causal controller $K \in \mathcal{I}(\mathcal{P})$ if and only if the (A_{ii}, B_{ii}) are stabilizable for all $i \in P$.

Proof: See Appendix. ∎

In this paper, we make the following important assumption about the stabilizability of the sub-systems. By the preceding theorem, this assumption is necessary and sufficient to ensure that the systems under consideration have feasible controllers.

Assumption 1: Given the poset-causal system of the form (1), we assume that the sub-systems (A_{ii}, B_{ii}) are stabilizable for all $i \in \{1, \dots, p\}$.
In the absence of this assumption, there is no poset-causal stabilizing controller, and hence the problem of finding an optimal one becomes vacuous. This assumption is thus necessary and sufficient for the problem to be well-posed. Moreover, in what follows, we will need the solution of certain standard Riccati equations. Assumption 1 ensures that all of these Riccati equations have well-defined stabilizing solutions. This stabilizing property of the Riccati solutions will be useful for proving internal stability of the closed loop system.

Assumption 2: We assume that $(C(\downarrow j), A(\downarrow j, \downarrow j))$ have no unobservable modes on the imaginary axis for all $j \in \{1, \dots, p\}$.

Assumption 2 is a technical assumption that we use later in (14) to ensure that the solutions of certain Riccati equations exist and are unique. The optimal controller (16) in Theorem 3 will require these Riccati equations to have well-defined solutions.

The system (1) may be viewed as a map from the inputs w, u to outputs z, x via

$$z = P_{11}w + P_{12}u$$
$$x = P_{21}w + P_{22}u$$

where

$$\begin{bmatrix} P_{11} & P_{12} \\ P_{21} & P_{22} \end{bmatrix} = \begin{bmatrix} C(sI - A)^{-1}F & C(sI - A)^{-1}B + D \\ (sI - A)^{-1}F & (sI - A)^{-1}B \end{bmatrix}$$
$$= \begin{bmatrix} A & F & B \\ \hline C & 0 & D \\ I & 0 & 0 \end{bmatrix}.$$
(4)

A controller $u = Kx$ induces a map T_{zw} from the disturbance input w to the exogenous output z via $T_{zw} = P_{11} + P_{12}K(I - P_{22}K)^{-1}P_{21}$. Thus, after the controller is interconnected with the system, the closed-loop map is T_{zw}. The objective function of interest is to minimize the $\mathcal{H}_2$ norm [32] of T_{zw} which we denote by $\|T_{zw}\|$.

C. Information Constraints on the Controller

Given the system (1), we are interested in designing a controller K that meets certain specifications.

In traditional control problems, one requires K to be proper, causal and stabilizing. One can impose additional constraints on the controller, for example require it to belong to some subspace. Such seemingly mild requirements can actually make the problem significantly more challenging. This paper focuses on addressing the challenge posed by subspace constraints arising from particular decentralization structures. The decentralization constraint of interest in this paper is one where the controller mirrors the structure of the plant, and is therefore also in the block incidence algebra $\mathcal{I}_K(\mathcal{P})$ (we will henceforth drop the subscripts and simply refer to the incidence algebra $\mathcal{I}(\mathcal{P})$). This translates into the requirement that input u_i only has access to x_j for $j \in \uparrow i$ thereby enforcing poset-causality constraints also on the controller. In this sense the controller has access to *local* states, and we thus refer to it as *decentralized*.

D. Problem Statement

Given the poset-causal system (4) with poset $\mathcal{P} = (P, \preceq)$, $|P| = p$, solve the optimization problem:

$$
\begin{aligned}
\underset{K}{\text{minimize}} \quad & \|P_{11} + P_{12}K(I - P_{22}K)^{-1}P_{21}\|^2 \\
\text{subject to} \quad & K \in \mathcal{I}(\mathcal{P}) \\
& K \text{ stabilizing.}
\end{aligned}
\tag{5}
$$

The main problem under consideration is to solve the above stated optimal control problem in the controller variable K. The feasible set is the set of all rational proper transfer function matrices that internally stabilize the system (1). In the absence of the decentralization constraints $K \in \mathcal{I}(\mathcal{P})$ this is a standard, well-studied control problem that has an efficient finite-dimensional *state-space* solution [32]. The main objective of this paper is to construct such a solution for the poset-causal case.

E. Notation

Given a matrix Q, let $Q(j)$ denote the j^{th} column of Q. We denote the i^{th} component of the vector $Q(j)$ to be $Q(j)_i$. For a poset $\mathcal{P}$ with incidence algebra $\mathcal{I}(\mathcal{P})$, if $M \in \mathcal{I}(\mathcal{P})$ then recall that M is sparse, i.e. has a zero pattern given by $M_{ij} = 0$ if $j \not\preceq i$. We denote the sparsity pattern of the j^{th} column of the matrices in $\mathcal{I}(\mathcal{P})$ by $\mathcal{I}(\mathcal{P})^j$. Let $v \in \mathbb{R}^p$ and v_i denote its i^{th} component.

$$
\mathcal{I}(\mathcal{P})^j := \{v \in \mathbb{R}^p | v_i = 0 \text{ for } j \not\preceq i\}.
$$

In the above definition v is understood to be a vector composed of $|P|$ blocks, with sparsity being enforced at the block level.

Given the data (A, B, C, D), we will often need to consider sub-matrices or embed a sub-matrix into a full dimensional matrix by zero padding. Some notation for that purpose we will use is the following:

1) Define $Q^{\downarrow j} = [Q_{ij}]_{i \in \downarrow j}$ (so that it is the j^{th} column shortened to include only the nonzero entries).
2) Also define $A(\downarrow j) = [A(i)]_{i \in \downarrow j}$ so that it is the sub-matrix of A containing exactly those columns corresponding to the set $\downarrow j$.
3) Define $A(\downarrow j, \downarrow j) = [A_{kl}]_{k,l \in \downarrow j}$ so that it is the $(\downarrow j, \downarrow j)$ sub-matrix of A (containing exactly those rows and columns corresponding to the set $\downarrow j$).
4) Given a block $|\downarrow j| \times |\downarrow j|$ matrix we will need to embed it into a block matrix indexed by the original poset (i.e. a $p \times p$ matrix) by padding it with zeroes. Given K (a block $|\downarrow j| \times |\downarrow j|$ matrix) we define the embedded matrix $\hat{K}$ with elements given by:

$$
[\hat{K}]_{l,m} = \begin{cases} K_{lm} & \text{if } l, m \in \downarrow j \\ 0 & \text{otherwise.} \end{cases}
$$

5) $E_i = [\, 0 \ \ldots \ I \ \ldots \ 0 \,]^T$ be the tall block matrix (indexed with the elements of the poset) with an identity in the i^{th} block row.
6) Let $S \subseteq P$. Define E_S to be the matrix whose columns are E_i for $i \in S$. Note that given a block $p \times p$ matrix M, $ME_{\downarrow j} = M(\downarrow j)$ is a matrix containing the columns indexed by $\downarrow j$.
7) Given matrices $A_i, i \in P$, we define the block diagonal matrix:

$$
\text{diag}(A_i) = \begin{bmatrix} A_1 & & \\ & \ddots & \\ & & A_p \end{bmatrix}.
$$

Recall that every poset $\mathcal{P}$ has a linear extension (i.e. a total order on P that is consistent with the partial order $\preceq$). For convenience, we fix such a linear extension of $\mathcal{P}$, and all indexing of our matrices throughout the paper will be consistent with this linear extension (so that elements of the incidence algebra are lower triangular).

Example 3: Let $\mathcal{P}$ be the poset shown in Fig. 1(d). We continue with Example 2 to illustrate

notation. (Note that $\downarrow 2 = \{2, 4\}$). As per the notation defined above,

$$A^{\downarrow 2} = \begin{bmatrix} -.25 \\ -1 \end{bmatrix} \qquad A(\downarrow 2) = \begin{bmatrix} 0 & 0 \\ -0.25 & 0 \\ 0 & 0 \\ -1 & -0.1 \end{bmatrix}$$

$$A(\downarrow 2, \downarrow 2) = \begin{bmatrix} -0.25 & 0 \\ -1 & -0.1 \end{bmatrix}.$$

Also, if $K(\downarrow 2, \downarrow 2) = \begin{bmatrix} 1 & 2 \\ 3 & 4 \end{bmatrix}$, then

$$K(\widehat{\downarrow 2, \downarrow 2}) = \begin{bmatrix} 0 & 0 & 0 & 0 \\ 0 & 1 & 0 & 2 \\ 0 & 0 & 0 & 0 \\ 0 & 3 & 0 & 4 \end{bmatrix}.$$

III. SOLUTION STRATEGY

In this section we first remind the reader of a standard reparametrization of the problem known as the Youla parametrization. Using this reparametrization, we illustrate the main technical idea of this paper using an example.

A. Reparametrization

Problem (5) as stated has a nonconvex objective function. Typically [16], [19], this is convexified by a bijective change of parameters given by $R := K(I - P_{22}K)^{-1}$ (though one typically needs to make a stability or prestabilization assumption). When the sparsity constraints are poset-causal (or quadratically invariant, more generally), this change of parameters preserves the sparsity constraints, and R inherits the sparsity constraints of K. The resulting infinite-dimensional problem is convex in R.

For poset-causal systems with state-feedback we will use a slightly different parametrization. We note that for poset-causal systems, the matrices A and B are both in the block incidence algebra. As a consequence of (4), P_{21} and P_{22} are also in the incidence algebra. This structure, which follows from the closure properties of an incidence algebra, will be extensively used. Since $P_{21}, P_{22} \in \mathcal{I}(\mathcal{P})$ the optimization problem (5) maybe be reparametrized as follows. Set

$$Q := K(I - P_{22}K)^{-1}P_{21}. \qquad (6)$$

Note that P_{21} is left invertible, and a left inverse is given by

$$P_{21}^\dagger = F^\dagger(sI - A),$$

where $F^\dagger$ is the pseudoinverse of F, so that $F^\dagger F = I$ (note also that the pseudoinverse is block diagonal and hence in $\mathcal{I}(\mathcal{P})$). As a consequence, given Q, K can be recovered using

$$K = QP_{21}^\dagger(I + P_{22}QP_{21}^\dagger)^{-1}. \qquad (7)$$

Since $I, P_{21}, P_{21}^\dagger, P_{22}$ all lie in the incidence algebra, $K \in \mathcal{I}(\mathcal{P})$ if and only if $Q \in \mathcal{I}(\mathcal{P})$. Using this reparametrization the optimization problem (5) can be relaxed to:

$$\begin{aligned} \underset{Q}{\text{minimize}} \quad & \|P_{11} + P_{12}Q\|^2 \\ \text{subject to} \quad & Q \in \mathcal{I}(\mathcal{P}). \end{aligned} \qquad (8)$$

Remarks 1) We note that $P_{21}^\dagger$, and hence (7) may potentially be improper. However, we will *prove* that for the optimal Q in (8), this expression is proper and corresponds to a rational controller $K^* \in \mathcal{I}(\mathcal{P})$.

2) For the objective function to be bounded, the optimal Q would have to render $P_{11} + P_{12}Q$ stable. However, one also requires that the overall system is internally stable. We relax this requirement on Q and later show that K^* is nevertheless internally stabilizing. Thus (8) is in fact a *relaxation* of (5). We show that the solution of the relaxation actually corresponds to a feasible controller.

We would like to emphasize the very important role played by the availability of *full state-feedback*. As a consequence of state-feedback, we have that $P_{21} = (sI - A)^{-1}F$. Thus P_{21} is left invertible (though the inverse is improper), and in the (block) incidence algebra. It is this very important feature of P_{21} that allows us to use this modified parametrization mentioned (6) in the preceding paragraph. This parametrization enables us to rewrite the problem in the form (8). This form will turn out to be crucial to our main separability result (Theorem 2), which enables us to separate the decentralized problem into a set of decoupled centralized problems.

A main step in our solution strategy will be to reduce the optimal control problem to a set of standard centralized control problems, whose solutions may be obtained by solving standard Riccati equations. The key result about centralized $\mathcal{H}_2$ optimal control is as follows.

Lemma 2: Consider a system H given by

$$H = \begin{bmatrix} H_{11} & H_{12} \end{bmatrix} = \left[\begin{array}{c|cc} A_H & F_H & B_H \\ \hline C_H & 0 & D_H \end{array} \right]$$

along with the following optimal control problem:

$$\begin{array}{cl} \underset{Q}{\text{minimize}} & \|H_{11} + H_{12}Q\|^2 \\ \text{subject to} & Q \text{ stable.} \end{array} \quad (9)$$

Suppose the pair (A_H, B_H) is stabilizable, (C_H, A_H) has no unobervable modes on the imaginary axis, and $C_H^T D_H = 0$, and $D_H^T D_H > 0$. Then the Hamiltonian matrix associated to this problem (denoted by $\mathcal{H}(H)$) is such that $\mathcal{H}(H) \in \text{dom(Ric)}$ [31, Chap. 13] and its associated Riccati equation:

$$A_H^T X + X A_H - X B_H (D_H^T D_H)^{-1} B_H^T X + C_H^T C_H = 0 \quad (10)$$

has a stabilizing, symmetric and positive semidefinite solution $X = \text{Ric}(\mathcal{H}(H))$ determined by the invariant subspaces of the Hamiltonian.

Let L be obtained from this solution via:

$$L = (D_H^T D_H)^{-1} B_H^T X. \quad (11)$$

Then the optimal solution to (9) is given by:

$$Q = \left[\begin{array}{c|c} A_H - B_H L & F_H \\ \hline -L & 0 \end{array} \right]. \quad (12)$$

(We will often refer to the trio of equations (10), (11), (12) by $(L, Q) = \mathcal{H}_2^{\text{opt}}(H)$.)

Remark Since the above Riccati equation is in dom(Ric) [31, Chap. 13.2], it has a unique canonical solution determined by the invariant subspaces of the Hamiltonian [31, Chap. 13.2]. In this paper, we will always use this unique canonical solution.

Proof: The proof is based on standard techniques and can be argued via a completion-of-squares argument. In particular, it follows from the solution to the standard $\mathcal{H}_2$ optimal control problem [31, Theorem 14.7]. Using Theorem 14.7, the solution to the $\mathcal{H}_2$ optimal control problem for the standard problem with the data

$$G = \left[\begin{array}{cc} G_{11} & G_{12} \\ I & 0 \end{array} \right] = \left[\begin{array}{c|cc} A_H & F_H & B_H \\ \hline C_H & 0 & D_H \\ 0 & I & 0 \end{array} \right]$$

gives the required formula.

Note that to use this theorem, the four assumptions stated in [31, pp. 376] need to be verified. The first part of assumption (i) on pp. 376 is clearly satisfied since (A_H, B_H) is assumed to be stabilizable. The second part of (i) on pp. 376 can be ignored. This assumption is required to ensure that the "observer" Riccati equation (defined by J_2 on pp. 376) has a solution. This Riccati equation does not play a role in our analysis since its solution Y_2 does not enter into the formula of the optimal controller in our case. Condition (ii) on pp. 376 is satisfied since $D_H^T D_H > 0$ (we can relax the unitariness assumption by modifying the Riccati equation). Condition (iii) is satisfied because (C_H, A_H) is assumed to have no unobservable imaginary modes. By [31, Lemma 13.9] and the following remark (Remark 13.3, pp.332), it follows that (C_H, A_H) having no unobservable imaginary modes in addition to $C_H^T D_H = 0$ implies that condition (iii) is satisfied. Finally condition (iv) is also irrelevant, since it is needed only to ensure that the Riccati equation corresponding to J_2 has a well-defined solution. ∎

B. Separability of Optimal Control Problem

We next illustrate the main solution strategy via a simple example. Consider the decentralized control problem (5) for the poset in Fig. 1(b). Using the reformulation (8) the optimal control problem (5) may be recast as:

$$\underset{Q}{\text{minimize}} \quad \left\| P_{11} + P_{12} \begin{bmatrix} Q_{11} & 0 & 0 \\ Q_{21} & Q_{22} & 0 \\ Q_{31} & 0 & Q_{33} \end{bmatrix} \right\|^2$$

Note that $P_{12}(\downarrow 1) = P_{12}, P_{12}(\downarrow 2) = P_{12}(2)$ (second column of P_{12}), and $P_{12}(\downarrow 3) = P_{12}(3)$. Similarly $Q^{\downarrow 1} = \begin{bmatrix} Q_{11}^T & Q_{21}^T & Q_{31}^T \end{bmatrix}^T$, $Q^{\downarrow 2} = Q_{22}$, and $Q^{\downarrow 3} = Q_{33}$. Due to the column-wise separability of the $\mathcal{H}_2$ norm, the problem can be recast as:

$$\underset{Q}{\text{minimize}} \quad \left\| P_{11}(1) + P_{12}(\downarrow 1) \begin{bmatrix} Q_{11} \\ Q_{21} \\ Q_{31} \end{bmatrix} \right\|^2 +$$
$$\|P_{11}(2) + P_{12}(\downarrow 2)Q_{22}\|^2 + \|P_{11}(3) + P_{12}(\downarrow 3)Q_{33}\|^2.$$

Since the sets of variables appearing in each of the three quadratic terms are disjoint, the problem now may be decoupled into three separate sub-problems, each of which is a standard centralized control problem. For instance, the solution to the second sub-problem can be obtained by noting the

realizations of $P_{11}(2)$ and $P_{12}(\downarrow 2)$ and then using (12). In this instance,

$$
\begin{aligned}
(L_{22}, Q_{22}) &= \mathcal{H}_2^{\text{opt}}\left(\begin{bmatrix} P_{11}(2) & P_{12}(\downarrow 2) \end{bmatrix}\right) \\
&= \mathcal{H}_2^{\text{opt}}\left(\left[\begin{array}{c|cc} A_{22} & F_{22} & B_{22} \\ \hline C(2) & 0 & D(2) \end{array}\right]\right).
\end{aligned}
$$

In a similar way, the entire optimal solution matrix Q^* can be obtained, and by design $Q^* \in \mathcal{I}(\mathcal{P})$ (and is stabilizing). To obtain the optimal K^*, one can use (7). In fact, it is possible to give an explicit state-space formula for K^*, this is the main content of Theorem 3 in the next section.

IV. Main Results

In this section, we present the main results of the paper. The proofs are available in Section VI.

A. Problem Decomposition and Computational Procedure

Theorem 2 (Decomposition Theorem): Let $\mathcal{P}$ be a poset and $\mathcal{I}(\mathcal{P})$ be its incidence algebra. Consider a poset-causal system given by (4). The problem (8) is equivalent to the following set of $|P|$ independent decoupled problems:

$$
\underset{Q^{\downarrow j}}{\text{minimize}} \quad \|P_{11}(j) + P_{12}(\downarrow j)Q^{\downarrow j}\|^2 \quad \forall j \in P. \quad (13)
$$

Theorem 2 is essentially the first step towards a state-space solution. The advantage of this equivalent reformulation of the problem is that we now have $p = |P|$ sub-problems, each over a different set of variables (thus the problem is *decomposed*). Moreover, each sub-problem corresponds to a particular *standard centralized* control problem, and thus the optimal Q in (5) can be computed by simply solving each of these sub-problems.

The subproblems described in (13) have the following interpretation. Once a controller K, or equivalently Q is chosen, a map T_{zw} from the exogenous inputs w to the outputs z is induced. Let us denote by $T_{zw}(1)$ to be the map from the first input w_1 to all the outputs z (this corresponds to the first column of T_{zw}). Similarly, the map from w_i to z for $i \in P$ is given by $T_{zw}(i)$. These subproblems correspond to the computation of the *optimal* maps $T_{zw}^*(i)$ for all $i \in P$ from the i^{th} input w_i to the output z. The decomposability of the $\mathcal{H}_2$ norm implies that these maps may be computed separately, and the performance of the overall system is simply the aggregation of these individual maps.

Our next theorem provides an efficient computational technique to obtain the required state-space solution. To obtain the solution, one needs to solve Riccati equations corresponding to the sub-problems we saw in Theorem 2. We combine these solutions to form certain simple block matrices, and after simple LFT transformations, one obtains the optimal controller K^*.

Before we state the theorem, we introduce some relevant notation.

Definition 3: We define the operator $\mathcal{H}_2^{\text{opt}}(\downarrow j)$ for $j \in P$ by:

$$
\mathcal{H}_2^{\text{opt}}(\downarrow j) := \mathcal{H}_2^{\text{opt}}\left(\left[\begin{array}{c|cc} A(\downarrow j, \downarrow j) & E_1 F_{jj} & B(\downarrow j, \downarrow j) \\ \hline C(\downarrow j) & 0 & D(\downarrow j) \end{array}\right]\right). \quad (14)
$$

(We remind the reader that in the above E_1 is the block $|\downarrow j| \times 1$ matrix which picks out the first column corresponding of the block $|\downarrow j| \times |\downarrow j|$ matrix before it.) We define $K(\downarrow j, \downarrow j)$ via $(K(\downarrow j, \downarrow j), Q(j)) = \mathcal{H}_2^{\text{opt}}(\downarrow j)$ for $j \in P$. Note that these quantities are well-defined by Assumptions 1 and 2 and Lemma 2. We introduce two matrices related to the above solution, namely:

$$
\begin{aligned}
\mathbf{A} &= \text{diag}(A(\downarrow j, \downarrow j) - B(\downarrow j, \downarrow j)K(\downarrow j, \downarrow j)) \\
\mathbf{K} &= \text{diag}(K(\downarrow j, \downarrow j)).
\end{aligned}
$$

We will see later on that $\mathbf{A}$ is the closed-loop state transition matrix under a particular indexing of the states.

It will be convenient to introduce this particular indexing of the states now. At the j^{th} subsystem, denote the local plant state by x_j. Recall that n_i denotes the degree of the i^{th} sub-system in (1). Let $n_{\max} = \max_i n_i$ be the largest degree of the subsystems, and $N_p = \sum_{i \in P} n_i$ be the total degree of the plant. Let $n(\downarrow\downarrow i) = \sum_{j \in \downarrow\downarrow i} n_j$. The controller states associated to the j^{th} subsystem will be denoted by $q(j) \in \mathbb{R}^{n(\downarrow\downarrow j)}$. (The subsystems downstream of j is precisely the set $\downarrow\downarrow j$, and for each $i \in \downarrow\downarrow j$ subsystem j has a controller state $q_i(j) \in q(j)$ to track the plant states of each of these subsystems, see Fig. 2.) We further define $N_q = \sum_{i \in P} n(\downarrow\downarrow i)$, this is the total degree of the controller. Let $N = N_p + N_q$ be the total degree of the closed-loop. Let $\sigma_{\mathcal{P}} = \sum_{j \in P} |\downarrow\downarrow j|$ (note that this is a purely combinatorial quantity, dependent only on the poset).
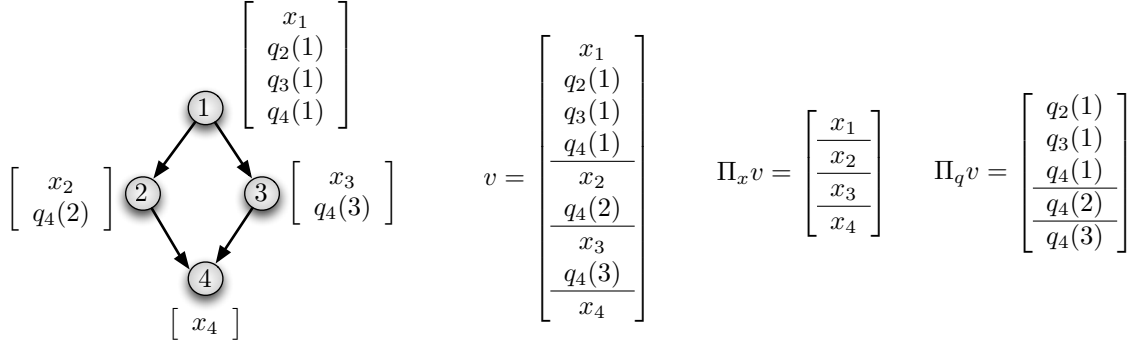
Fig. 2. Example illustrating the matrices Π_x and Π_q.

Using this notation we introduce a vector $v \in \mathbb{R}^N$ (to be thought of as an indexing of the states for the closed loop system) and a pair of linear maps $\Pi_x : \mathbb{R}^N \to \mathbb{R}^{N_p}$ and $\Pi_q : \mathbb{R}^N \to \mathbb{R}^{N_q}$ which act as projection operators as described below:

$$v = \begin{bmatrix} x_1 \\ \hline q(1) \\ \hline x_2 \\ \hline q(2) \\ \hline \vdots \\ \hline x_p \\ \hline q(p) \end{bmatrix} \qquad \Pi_x v = \begin{bmatrix} x_1 \\ \hline x_2 \\ \hline \vdots \\ \hline x_p \end{bmatrix} = x \qquad \Pi_q v = \begin{bmatrix} q(1) \\ \hline q(2) \\ \hline \vdots \\ \hline q(p) \end{bmatrix} = q.$$

It will be convenient to think of Π_x and Π_q via their natural matrix representations in which they are $0-1$ matrices. The action of these projection operators onto the x (plant states) and q (controller states) components of v are illustrated in Fig. 2. The linear maps $x = \Pi_x v$ and $q = \Pi_q v$ have natural adjoints, with the adjoint operators $\Pi_x^T : \mathbb{R}^{N_p} \to \mathbb{R}^N$ and $\Pi_q^T : \mathbb{R}^{N_q} \to \mathbb{R}^N$ representing embeddings as follows:

$$\Pi_x^T x = \begin{bmatrix} x_1 \\ \hline 0 \\ \hline x_2 \\ \hline 0 \\ \hline \vdots \\ \hline x_p \\ \hline 0 \end{bmatrix} \qquad \Pi_q^T q = \begin{bmatrix} 0 \\ \hline q(1) \\ \hline 0 \\ \hline q(2) \\ \hline \vdots \\ \hline 0 \\ \hline q(p) \end{bmatrix}.$$

In addition, one can also define $\Pi_{x_i} : \mathbb{R}^N_p \to \mathbb{R}^{n_i}$ such that $\Pi_{x_i} v = x_i$, and similarly $\Pi_{q(i)} : \mathbb{R}^N \to \mathbb{R}^{n(\downarrow\downarrow i)}$ such that $\Pi_{q(i)} v = q(i)$. Their adjoints are embedding operators in the usual way. We now introduce one more operator $\Sigma : \mathbb{R}^N \to \mathbb{R}^{N_p}$ which acts on v by taking partial sums along the poset:

$$\Sigma v = \begin{bmatrix} x_1 + \sum_{k \in \downarrow\downarrow 1} q_1(k) \\ \hline x_2 + \sum_{k \in \downarrow\downarrow 2} q_2(k) \\ \hline \vdots \\ \hline x_p + \sum_{k \in \downarrow\downarrow p} q_p(k) \end{bmatrix}.$$

As a consequence of this relation we have:

$$\Sigma \Pi_q^T q = \Sigma \Pi_q^T \Pi_q v = \begin{bmatrix} \sum_{k \in \downarrow\downarrow 1} q_1(k) \\ \hline \sum_{k \in \downarrow\downarrow 2} q_2(k) \\ \hline \vdots \\ \hline \sum_{k \in \downarrow\downarrow p} q_p(k) \end{bmatrix}.$$

Note that from the above two relations it is easy to deduce that $\Sigma \Pi_q^T \Pi_q + \Pi_x = \Sigma$. The optimal controller and other related objects can be expressed in terms of the following matrices:

$$A_\Phi = \Pi_q \mathbf{A} \Pi_q^T, \qquad B_\Phi = \Pi_q \mathbf{A} \Pi_x^T. \qquad (15)$$

Theorem 3 (Computation of Optimal Controller): Consider the poset-causal system of the form (4), such that Assumptions 1 and 2 are satisfied. Consider the following Riccati equations:

$$(K(\downarrow j, \downarrow j), Q(j)) = \mathcal{H}_2^{\text{opt}}(\downarrow j) \quad \forall j \in P.$$

Then the optimal solution to the problem (5) is given by the controller:

$$K^* = \left[\begin{array}{c|c} A_\Phi - B_\Phi \Sigma \Pi_q^T & B_\Phi \\ \hline -\Sigma \mathbf{K}(\Pi_q^T - \Pi_x^T \Sigma \Pi_q^T) & -\Sigma \mathbf{K} \Pi_x^T \end{array} \right]. \qquad (16)$$

Moreover, the controller $K^* \in \mathcal{I}(\mathcal{P})$ and is internally stabilizing.

As we mentioned in the introduction, one of the advantages of state-space techniques is that they provide graceful degree bounds for the optimal controller. As a consequence of Theorem 3 we have
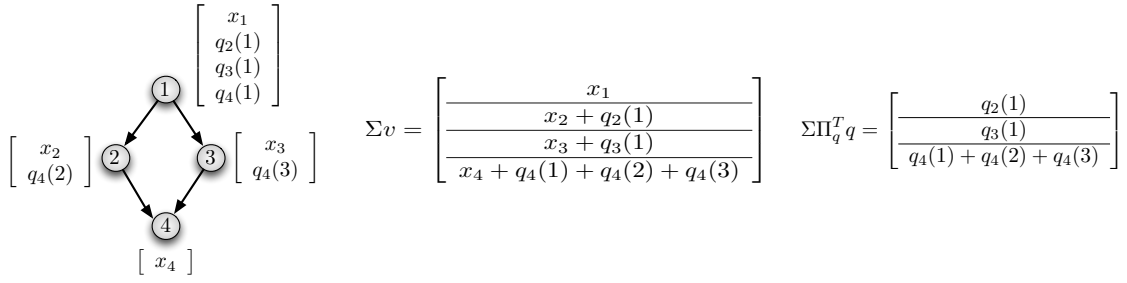
Fig. 3. Example illustrating Σ operating on v and q.

the following:

Corollary 2 (Degree Bounds): The degree d_{K^*} of the overall optimal controller is bounded above by

$$d_{K^*} \leq \sum_{j \in P} n(\downarrow\downarrow j) = N_q.$$

In particular, $d_{K^*} \leq \sigma_{\mathcal{P}} n_{\max}$. Moreover, the degree of the controller implemented by subsystem j is bounded above by $n(\downarrow\downarrow j)$.

B. Ingredients of the Optimal Controller

Having established the computational aspects, we now turn to some structural aspects of the optimal controller. Note that the dynamic equations of the controller (16) are the following:

$$\dot{q} = (A_\Phi - B_\Phi \Sigma \Pi_q^T)q + B_\Phi x$$
$$u = -\Sigma \mathbf{K}(\Pi_q^T - \Pi_x^T \Sigma \Pi_q^T)q - \Sigma \mathbf{K} \Pi_x^T x$$
$$= -\Sigma \mathbf{K}\left(\Pi_q^T q + \Pi_x^T(x - \Sigma \Pi_q^T q)\right).$$

This controller may be more readily understood via a block diagram representation as shown in Fig. 4.

We next examine the role of the controller states. Let us introduce the transfer function Θ with LFT realization:

$$\Theta = \left[\begin{array}{c|c} A_\Phi - B_\Phi \Sigma \Pi_q^T & B_\Phi \\ \hline \Pi_q^T - \Pi_x^T \Sigma \Pi_q^T & \Pi_x^T \end{array} \right]. \quad (17)$$

Note that by (16), $K^* = -\Sigma \mathbf{K}\Theta$, so that the states of Θ are precisely $q(i)$, the controller states. In particular, Θ is a $N \times N_p$ transfer function and thus has a natural partition as:

$$\Theta = \left[\begin{array}{c} \Theta(1) \\ \vdots \\ \Theta(p) \end{array} \right],$$

where each

$$\Theta(i) = \left[\begin{array}{c} \Pi_{x_i} \\ \Pi_{q(i)} \end{array} \right] \Theta.$$

is a $(n(i) + n(\downarrow\downarrow i)) \times N_p$ sized transfer function. By zero-padding, one can embed $\Theta(i)$ into a transfer function matrix of size $N_p \times N_p$, this is done by defining

$$\hat{\Theta}(i) = E_{\downarrow i}\Theta(i).$$

We also introduce

$$\hat{\Theta} = \left[\begin{array}{c} \hat{\Theta}(1) \\ \vdots \\ \hat{\Theta}(p) \end{array} \right].$$

One can show that Θ (and equivalently $\hat{\Theta}$, which is just a zero-padded version of Θ), in fact, corresponds to a *specific filter* called the *differential filter*. At a high level, the controller at subsystem i predicts the unknown states downstream in the poset. For example, in Fig. 1(a), if $x_1(t)$ is the state at subsystem 1, subsystem 1 constructs a prediction of state $x_2(t)$ locally. As one proceeds "downstream" through the poset, more information is available, and consequently the prediction of the global state becomes more accurate.

Remark We remark that the term "prediction" in common usage is employed to describe outcomes likely to happen in the "temporal future". In the context of this paper, when the term prediction is used, it is with reference to "spatial future" as defined by the poset $\mathcal{P}$ that captures causality among the subsystems, rather than a temporal notion of causality.

The transfer function Θ plays the role of computing the generalized differential in the prediction of the global state. This is precisely the role of the differential filter: to compute these generalized
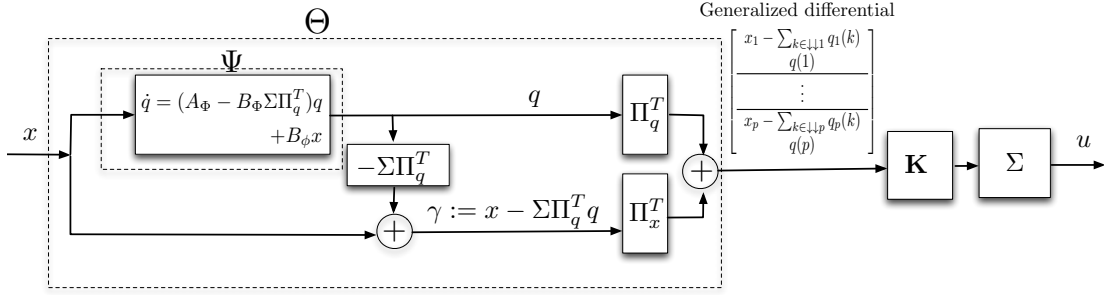
Fig. 4. Block diagram representation of the controller. The block marked Θ performs the role of computing the generalized differentials of the predicted state. The transfer function from x to γ is called Φ^{-1}, and plays a role in controller factorization.

finite differences of predictions that capture "local improvements" in the local predictions. We remark that we use the nomenclature "generalized differential" since it is intimately related to the notion of *Möbius inversion* on a poset, a generalization of differentiation to posets. We briefly discuss these ideas in the ensuing discussion.

Theorem 4 (Structure of Optimal Controller): The optimal controller (16) is of the form:

$$u(t) = -\Sigma \mathbf{K} \Theta x(t)$$
$$= -\sum_{j \in P} K(\widehat{\downarrow j, \downarrow j}) \hat{\Theta}(j) x(t).$$

Remark Let us denote the vector $e(j) = \hat{\Theta}(j)x$. We will interpret $e(j)$ as the generalized differential in the prediction of the global state x at subsystem j. Denoting $K(\widehat{\downarrow j, \downarrow j})$ by K_j, note that the control law takes the form $u(t) = \sum_{j \in P} K_j e(j)$. This structural form suggests that the controller uses the generalized differentials at the different subsystems as the atoms of local control laws, and that the overall control law is an aggregation of these local control laws.

C. State Prediction, Differentiation, and Integration

Due to the information constraints in the problem, at subsystem j only states in $\uparrow j$ are available, states of other subsystems are unavailable. A reasonable architecture for the controller would involve *predicting* the unknown states at subsystem j from the available information. This is illustrated by the following example.

Example 4: Consider the system shown in Fig. 5

with dynamics

$$\begin{bmatrix} \dot{x}_1 \\ \dot{x}_2 \\ \dot{x}_3 \end{bmatrix} = \begin{bmatrix} A_{11} & 0 & 0 \\ 0 & A_{22} & 0 \\ A_{31} & A_{32} & A_{33} \end{bmatrix} \begin{bmatrix} x_1 \\ x_2 \\ x_3 \end{bmatrix} + \begin{bmatrix} B_{11} & 0 & 0 \\ 0 & B_{22} & 0 \\ B_{31} & B_{32} & B_{33} \end{bmatrix} \begin{bmatrix} u_1 \\ u_2 \\ u_3 \end{bmatrix}.$$

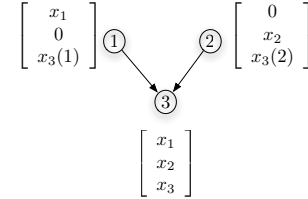Note that subsystem 1 has no information about

Fig. 5. Local state information at the different subsystems. The quantities $x_3(1)$ and $x_3(2)$ are partial state predictions of x_3.

the state of subsystem 2. Moreover, the state x_1 or input u_1 do not affect the dynamics of 2 (their respective dynamics are uncoupled). Hence the only sensible prediction of x_2 at subsystem 1 (which we denote by $x_2(1)$) is $x_2(1) = 0$. Subsystem 1 also does not have access to the state x_3. However, it can predict x_3 based on the influence that the state x_1 has on x_3. (Note that both the states x_1, x_2 and inputs u_1, u_2 affect x_3 and u_3.) Let us denote $x_3(1)$ to be the prediction of state x_3 at subsystem 1. Since x_2 and u_2 are unknown, the state $x_3(1)$ is a *partial* prediction of x_3 (i.e. $x_3(1)$ is the prediction of the *component of* x_3 that is affected by subsystem 1). Similarly, subsystem 2 maintains a prediction of x_3 denoted by $x_3(2)$, which is also a partial prediction of x_3. Each subsystem thus maintains (possibly partial) predictions of unknown downstream states, as shown in Fig. 5.

In this paper we will not discuss how the state predictions are computed, a detailed discussion of the same is available in [18, Chapter 5], [22].

The next notion that is intimately related to the structure of the controller is that of generalized integration and differentiation with respect to posets. These concepts can be formalized via the notions of μ and ζ functions of posets. We will not explore these concepts formally here, but rather explain them in the context of Example 4. Suppose we have a function $z : P \to \mathcal{R}$ which is expressed as a vector $z = \begin{bmatrix} z_1 & z_2 & z_3 \end{bmatrix}$. One can define a generalized integral along the poset in Example 4 as $\zeta(z) = \begin{bmatrix} z_1 & z_2 & z_1 + z_2 + z_3 \end{bmatrix}$. One can similary define a notion of a generalized differential on the poset as $\mu(z) = \begin{bmatrix} z_1 & z_2 & z_3 - z_1 - z_2 \end{bmatrix}$ as the inverse operation of ζ.

An interesting aspect of the controller is that Θ plays the role of Möbius inversion (i.e. the computation of μ) with respect to the partial state predictions. On the other hand, the operator Σ defined earlier plays the role of generalized integration. We define variables $q_k(i)$ which correspond to the generalized differentials of predicted states for $k \in \downarrow\downarrow i$. At subsystem j the true state x_j becomes available for the first time (with respect to the subposet $\uparrow j$). The quantity $\Theta_j(j)x = x_j - \sum_{i<j} q_j(i)$ measures the *generalized differential* in the knowledge of state x_j, i.e. the difference between the true state x_j and its best prediction from upstream information. We let $q(i) = [q_j(i)]_{j\in\downarrow\downarrow i}$, so that $q(i)$ corresponds to the generalized differential in state predictions at the i^{th} subsystem. This $q(i)$ is a vector of length $|\downarrow\downarrow i|$.

D. Structure of the Optimal Controller

Using Theorem 4, the optimal control law can be expressed as:

$$u = -\sum_{i \in P} K(\widehat{\downarrow i}, \downarrow i)\hat{\Theta}(i)x. \qquad (18)$$

As explained above, $\Theta(i)x$ is a vector containing the generalized differential in the prediction of the global state at subsystem i. Each term $K(\widehat{\downarrow i}, \downarrow i)\hat{\Theta}(i)x$ may be viewed as a local control law acting on the local generalized differential in the predicted state. The overall control law has the elegant interpretation of being an aggregation of these local control laws.

Example 5: Let us consider the poset from Fig. 1(d), and examine the structure of the controller. (For simplicity, we let $K_j = -K(\widehat{\downarrow j}, \downarrow j)$, the gains

obtained by solving the Riccati equations). The control law may be decomposed into local controllers as:

$$u = K_1\hat{\Theta}_1 x + K_2\hat{\Theta}_2 x + K_3\hat{\Theta}_3 x + K_4\hat{\Theta}_4 x$$

$$= K_1 \begin{bmatrix} x_1 \\ q_2(1) \\ q_3(1) \\ q_4(1) \end{bmatrix} + K_2 \begin{bmatrix} 0 \\ x_2 - q_2(1) \\ 0 \\ q_4(2) \end{bmatrix} + K_3 \begin{bmatrix} 0 \\ 0 \\ x_3 - q_3(1) \\ q_4(3) \end{bmatrix} +$$

$$K_4 \begin{bmatrix} 0 \\ 0 \\ 0 \\ (x_4 - q_4(1)) - q_4(2) - q_4(3) \end{bmatrix}.$$

Each term in the above expression has the natural interpretation of being a local control signal corresponding to generalized differential in predicted states, and the final controller can be viewed as an aggregation of these.

Note that zeros in the above expression imply no improvement on the local state. For example, at subsystem 2 there is no improvement in the predicted value of x_3 because the state x_2 does not affect subsystem 3 due to the poset-causal structure. There is no improvement in the predicted value of state x_3 at subsystem 4 either, because the best available prediction of x_3 from downstream information $\uparrow\uparrow 4$ is x_3 itself. While this interpretation has been stated informally here, it has been made precise in [18, Chapter 5], [22].

V. DISCUSSION AND EXAMPLES

A. The Nested Case

Consider the poset on two elements $\mathcal{P} = (\{1, 2\}, \le)$ with the only order relation being $1 \le 2$ (Fig. 1(a)). This is the poset corresponding to the communication structure in the "Two-Player Problem" considered in [24]. We show that their results are a specialization of our general results in Section IV restricted to this particular poset.

We begin by noting that from the problem of designing a nested controller (again we assume $F = I$ for simplicity) can be recast as:

$$\underset{Q}{\text{minimize}} \left\| P_{11} + P_{12} \begin{bmatrix} Q_{11} & 0 \\ Q_{21} & Q_{22} \end{bmatrix} \right\|^2.$$

13

By Theorem 2 this problem can be recast as:

$$\underset{Q}{\text{minimize}} \quad \left\| P_{11}^1 + P_{12} \begin{bmatrix} Q_{11} \\ Q_{21} \end{bmatrix} \right\|^2 + \left\| P_{11}^2 + P_{12}^2 Q_{22} \right\|^2.$$

We wish to compare this to the results obtained in [24]. It is possible to obtain precisely this same decomposition in the finite time horizon where the $\mathcal{H}_2$ norm can be replaced by the Frobenius norm and separability can be used to decompose the problem. For each of the sub-problems, the corresponding optimality conditions may be written (since they correspond to simple constrained-least squares problems). These optimality conditions correspond exactly to the decomposition of optimality conditions they obtain (the crucial Lemma 3 in their paper). We point out that the decomposition is a simple consequence of the separability of the Frobenius norm.

Let us now examine the structure of the optimal controller via Theorem 4. Note that $\downarrow 1 = \{1, 2\}$ and $\downarrow 2 = \{2\}$. Based on Theorem 3, we are required to solve $(K, Q(1)) = \mathcal{H}_2^{\text{opt}}(\downarrow 1)$, and $(J, Q(2)) = \mathcal{H}_2^{\text{opt}}(\downarrow 2)$. A straightforward application of Theorem 4 yields the following:

$$u_1(t) = -(K_{11} + K_{12}\Theta_2(1))x_1(t)$$
$$u_2(t) = -(K_{21} + K_{22}\Theta_2(1))x_1(t) - J(x_2(t) - \Theta_2(1)x_1(t)),$$

which is precisely the structure of the optimal controller given in [24], [25]. It is possible to show (as Swigart et. al indeed do in [24]) that $\Theta_2(1)$ is an predictor of x_2 based on x_1. Thus the controller for u_1 predicts the state of x_2 from x_1, uses the estimate as a surrogate for the actual state, and uses the gain K_{21} in the feedback loop. The controller for u_2 (perhaps somewhat surprisingly) also estimates the state x_2 based on x_1 using $\hat{x}_2 = \Theta_2(1)x_1$ (this can be viewed as a "simulation" of the controller for u_1). The prediction error for state 2 is then given by $e_2 := x_2 - \hat{x}_2 = x_2 - \Theta_2(1)x_1$. The control law for u_2 may be rewritten as

$$u_2 = -(K_{21}x_1 + K_{22}\hat{x}_2 + Je_2).$$

Thus this controller uses predictions of x_2 based on x_1 along with prediction errors in the feedback loop. We will see in a later example, that this prediction of states higher up in the poset is prevalent in such poset-causal systems, which results in somewhat larger order controllers.

Analogous to the results in [24], it is possible to derive the results in this paper for the finite time horizon case (this is a special case corresponding to FIR plants in our setup). We do not devote attention to the finite time horizon case in this paper, but just mention that similar results follow in a straightforward manner.

B. Discussion Regarding Computational Complexity

Note that the main computational step in the procedure presented in Theorem 3 is the solution of the p sub-problems. The j^{th} sub-problem requires the solution of a Riccati equation of size at most $|\downarrow j|n_{\max} = O(p)$ (when the degree $n_{\max}$ is fixed). Assuming the complexity of solving a Riccati equation using linear algebraic techniques is $O(p^4)$ [7] the complexity of solving p of them is at most $O(p^5)$. We wish to compare this with the only other known state-space technique that works on all poset-causal systems, namely the results of Rotkowitz and Lall [16]. In this paper, they transform the problem to a standard centralized problem using Kronecker products. In the final computational step, one would be required to solve a single large Riccati equation of size $O(p^2)$, resulting in a computational complexity of $O(p^8)$.

C. Discussion Regarding Degree Bounds

It is insightful to study the asymptotics of the degree bounds in the setting where the sub-systems have fixed degree and the number of sub-systems p grows. As an immediate consequence of the corollary, the degree of the optimal controller (assuming that the degree of the sub-systems $n_{\max}$ is fixed) is at most $O(p^2)$ (since $n(\downarrow j) \leq p$). In fact, the asymptotic behaviour of the degree can be sub-quadratic. Consider a poset $(\{1, \ldots, p\}, \preceq)$ with the only order relations being $1 \preceq i$ for all i. Here $|\downarrow 1| = p$, and $|\downarrow i| = 1$ for all $i \neq 1$. Hence, $\sum_j |\downarrow j| - p \leq p$, and thus $d^* \leq sn_{\max}$. In this sense, *the degree of the optimal controller is governed by the poset parameter $\sigma_{\mathcal{P}}$.*

VI. Proofs of the Main Results

Proof of Theorem 1: Note that one direction is trivial. Indeed if the (A_{ii}, B_{ii}) are stabilizable, one can pick a diagonal controller with diagonal elements K_{ii} such that $A_{ii} + B_{ii}K_{ii}$ is stable for all $i \in P$. This constitutes a stabilizing controller.

For the other direction let

$$K = \left[\begin{array}{c|c} A_K & B_K \\ \hline C_K & D_K \end{array} \right]$$

be a poset-causal controller for the system. We will first show that without loss of generality, we can assume that A_K, B_K, C_K, D_K are block lower triangular (so that K has a realization where all matrices are block lower triangular).

First, note that since $K \in \mathcal{I}(\mathcal{P})$, $D_K \in \mathcal{I}(\mathcal{P})$. Recall, that we assumed throughout that the indices of the matrices in the incidence algebra are labeled so that they are consistent with a linear extension of the poset, so that D_K is lower triangular. Note that the controller K is a block $p \times p$ transfer function matrix which has a realization of the form:

$$K = \left[\begin{array}{c|ccc} A_K & B_K(1) & \ldots & B_K(p) \\ \hline C_K(1) & D_K(1,1) & \ldots & D_K(1,p) \\ \vdots & \vdots & \ddots & \\ C_K(p) & D_K(p,1) & \ldots & D_K(p,p) \end{array} \right]$$

Since the controller $K \in \mathcal{I}(\mathcal{P})$, we have that $K_{jp} = 0$ for all $j \neq p$ (recall that p is the cardinality of the poset). This vector of transfer functions (given by the last column of K with the (p,p) entry deleted) is given by the realization:

$$\bar{K}_p := \left[\begin{array}{c} C_K(1) \\ \vdots \\ C_K(p-1) \end{array} \right] (sI - A_K)^{-1} B_K(p) + \left[\begin{array}{c} D_K(1,p) \\ \vdots \\ D_K(p-1,p) \end{array} \right] = 0.$$

Since this transfer function is zero, in addition to $D_K(j,p) = 0$ for all $j = 1, \ldots, p-1$, it must also be the case that the controllable subspace of $(A_K, B_K(p))$ is contained within the unobservable subspace of $\left(\left[\begin{array}{ccc} C_K(1)^T & \ldots & C_K(p-1)^T \end{array} \right]^T, A_K \right)$. By the Kalman decomposition theorem [6, pp. 247], there is a realization of this system of the form: $\bar{K}_p = \left[\begin{array}{c|c} \bar{A} & \bar{B} \\ \hline \bar{C} & \bar{D} \end{array} \right]$, where $(\bar{A}, \bar{B}, \bar{C}, \bar{D})$ are of the form:

$$\bar{A} = \left[\begin{array}{ccc} A_{11} & 0 & 0 \\ A_{21} & A_{22} & 0 \\ A_{21} & A_{32} & A_{33} \end{array} \right] \quad \bar{B} = \left[\begin{array}{c} 0 \\ 0 \\ B_3 \end{array} \right] \quad (19)$$

$$\bar{C} = \left[\begin{array}{ccc} C_1 & 0 & 0 \end{array} \right] \quad \bar{D} = 0.$$

As an aside, we remind the reader that this decomposition has a natural interpretation. For example, the subsystem $(A_{11}, 0, C_1)$ corresponds to the observable subspace, where the system is uncontrollable, etc. (The usual Kalman decomposition as stated in

standard control texts is a block 4×4 decomposition of the state-transition matrix. Here we have a smaller block 3×3 decomposition because of the collapse of the subspace where the system is required to be both controllable and observable). Thus this decomposition allows us to infer the specific block structure (19) on the matrices $(\bar{A}, \bar{B}, \bar{C}, \bar{D})$. Consequently there is a realization of the overall controller $(\tilde{A}_K, B_K, C_K, D_K)$, where all the matrices have the block structure

$$\left[\begin{array}{cccc} M_{1,1} & \ldots & M_{1,p-1} & 0 \\ \vdots & \ddots & & \vdots \\ M_{p-1,1} & \ldots & M_{p-1,p-1} & 0 \\ M_{p,1} & \ldots & M_{p,p-1} & M_{p,p} \end{array} \right].$$

One can now repeat this argument for the upper $(p-1) \times (p-1)$ sub-matrix of K. By repeating this argument for first $p-1, p-2, \ldots, 1$ we obtain a realization of K where all four matrices are block lower triangular.

Note that given the controller K (henceforth assumed to have a lower triangular realization), the closed loop matrix A_{cl} is given by $A_{cl} = \left[\begin{array}{cc} A + BD_K & BC_K \\ B_K & A_K \end{array} \right]$. By assumption the (open loop) system is poset-causal, hence A and B are block lower triangular. As a result, each of the blocks $A + BD_K$, BC_K, B_K, A_K are block lower triangular. A straightforward permutation of the rows and columns enables us to put A_{cl} into block lower triangular form where the diagonal blocks of the matrix are given by

$$\left[\begin{array}{cc} A_{jj} + B_{jj}D_{K_{jj}} & B_{jj}C_{K_{jj}} \\ B_{K_{jj}} & A_{K_{jj}} \end{array} \right]. \quad (20)$$

Note that the eigenvalues of this lower triangular matrix (and thus of A_{cl}, since simultaneous permutations of rows and columns are spectrum-preserving) are given by the eigenvalues of the diagonal blocks. The matrix A_{cl} is stable if and only if all its eigenvalues are in the left half plane, i.e. the above blocks are stable for each $j \in P$. Note that (20) is obtained as the closed-loop matrix precisely by the interconnection of

$$\left[\begin{array}{c|c} A_{jj} & B_{jj} \\ \hline I & 0 \end{array} \right] \quad \text{with the controller} \quad \left[\begin{array}{c|c} A_{K_{jj}} & B_{K_{jj}} \\ \hline C_{K_{jj}} & D_{K_{jj}} \end{array} \right].$$

Hence, (20) (and thus the overall closed loop) is stable if and only if (A_{jj}, B_{jj}) are stabilizable for all $j \in P$, and $(A_{K_{jj}}, B_{K_{jj}}, C_{K_{jj}}, D_{K_{jj}})$ are chosen to stabilize the pair. $\blacksquare$

Proof of Theorem 2: If $G = [G_1, \ldots, G_k]$ is a transfer function matrix with G_i as its columns, then $\|G\|^2 = \sum_{i=1}^{k} \|G_i\|^2$. The separability property of the $\mathcal{H}_2$ norm can be used to simplify (8). Recall that $P_{11}(j), Q(j)$ denote the j^{th} columns of P_{11} and Q respectively. Using the separability we can rewrite (8) as

$$
\begin{aligned}
&\underset{Q}{\text{minimize}} \quad \sum_{j \in P} \|P_{11}(j) + P_{12}Q(j)\|^2 \\
&\text{subject to} \quad Q(j) \in \mathcal{I}(\mathcal{P})^j
\end{aligned}
\tag{21}
$$

The formulation in (21) can be further simplified by noting that for $Q^j \in \mathcal{I}(\mathcal{P})^j$,

$$
P_{12}Q(j) = P_{12}(\downarrow j)Q^{\downarrow j}.
\tag{22}
$$

The advantage of the representation (22) is that, in the right hand side the variable $Q^{\downarrow j}$ is *unconstrained*. Using this we may reformulate (21) as:

$$
\underset{Q^{\downarrow 1}, \ldots, Q^{\downarrow p}}{\text{minimize}} \quad \sum_{j \in P} \|P_{11}(j) + P_{12}(\downarrow j)Q^{\downarrow j}\|^2
\tag{23}
$$

Since the variables in the $Q^{\downarrow j}$ are distinct for different j, this problem can be separated into p standard centralized sub-problems as follows:

$$
\begin{aligned}
&\underset{Q^{\downarrow j}}{\text{minimize}} \quad \|P_{11}(j) + P_{12}(\downarrow j)Q^{\downarrow j}\|^2 \\
&\qquad \text{for all } j \in P.
\end{aligned}
\tag{24}
$$

The sub-problems can be solved using canonical procedures as described in the next lemma.

Lemma 3: Let (A, B, C, D) be as given in (1) with A, B in the block incidence algebra $\mathcal{I}(\mathcal{P})$. Let $(K(\downarrow j, \downarrow j), Q(j)) = \mathcal{H}_2^{\text{opt}}(\downarrow j)$. Then the optimal solution of each sub-problem (13) is given by:

$$
(Q^{\downarrow j}) = \left[\begin{array}{c|c} A(\downarrow j, \downarrow j) - B(\downarrow j, \downarrow j)K(\downarrow j, \downarrow j) & E_1 F_{jj} \\ \hline -K(\downarrow j, \downarrow j) & 0 \end{array} \right].
\tag{25}
$$

(We remind the reader that in the above E_1 is the block $|\downarrow j| \times 1$ matrix which picks out the first column corresponding of the block $|\downarrow j| \times |\downarrow j|$ matrix before it.)

Proof: The proof follows directly from Lemma 2 by choosing

$$
\begin{aligned}
H &= \left[\begin{array}{cc} P_{11}(j) & P_{12}(\downarrow j) \end{array} \right] \\
&= \left[\begin{array}{c|cc} A(\downarrow j, \downarrow j) & E_1 F_{jj} & B(\downarrow j, \downarrow j) \\ \hline C(\downarrow j) & 0 & D(\downarrow j) \end{array} \right].
\end{aligned}
$$

■

Lemma 4: The optimal solution to (8) is given

by

$$
Q^* = \left[\begin{array}{c|c} \mathbf{A} & \Pi_x^T F \\ \hline -\Sigma \mathbf{K} & 0 \end{array} \right].
\tag{26}
$$

Proof: We note that Lemma 3 gives an expression for the individual columns of Q^*. Using Lemma 3 and the LFT formula for column concatenation:

$$
\left[\begin{array}{cc} G_1 & G_2 \end{array} \right] = \left[\begin{array}{cc|cc} A_1 & 0 & B_1 & 0 \\ 0 & A_2 & 0 & B_2 \\ \hline C_1 & C_2 & D_1 & D_2 \end{array} \right],
$$

we obtain the required expression. ■

Lemma 5: The matrix $\mathbf{A}$ is stable.

Proof: Recall that $\mathbf{A} = \text{diag}(A(\downarrow j, \downarrow j) - B(\downarrow j, \downarrow j)K(\downarrow j, \downarrow j))$. Since $A(\downarrow j, \downarrow j)$ and $B(\downarrow j, \downarrow j)$ are lower triangular with A_{kk}, B_{kk}, $k \in \downarrow j$ along the diagonals respectively, we see that the pair $(A(\downarrow j, \downarrow j), B(\downarrow j, \downarrow j))$ is stabilizable by Assumption 1 (simply picking a diagonal K which stabilizes the diagonal terms would suffice to stabilize $(A(\downarrow j, \downarrow j), B(\downarrow j, \downarrow j)))$. Hence, there exists a stabilizing solution to $\mathcal{H}_2^{\text{opt}}(\downarrow j)$ and the corresponding controller $K(\downarrow j, \downarrow j)$ is stabilizing. Thus $A(\downarrow j, \downarrow j) - B(\downarrow j, \downarrow j)K(\downarrow j, \downarrow j))$ is stable, and thus so is $\mathbf{A}$. ■

Lemma 6: Given transfer function matrices M and K with realizations

$$
M = \left[\begin{array}{c|cc} A & B_1 & B_2 \\ \hline C_1 & D_{11} & D_{12} \\ C_2 & D_{21} & 0 \end{array} \right], \quad K = \left[\begin{array}{c|c} A_K & B_K \\ \hline C_K & D_K \end{array} \right],
$$

the Linear Fractional Transformation (LFT) $f(M, K) = M_{11} + M_{12}K(I - M_{22}K)^{-1}M_{21}$ is given by:

$$
f(M, K) = \left[\begin{array}{cc|c} A + B_2 D_K C_2 & B_2 C_K & B_1 + B_2 D_K D_{21} \\ B_K C_2 & A_K & B_K D_{21} \\ \hline C_1 + D_{12}D_K C_2 & D_{12}C_K & D_{11} + D_{12}D_K D_{21} \end{array} \right].
\tag{27}
$$

Proof: The proof is standard, see for example [32, pp. 179] and the references therein. ■

Proof of Theorem 3: Consider again the optimal control problem (5): Let v_1^* be the optimal value of (5). Consider, on the other hand the optimization problem:

$$
\begin{aligned}
&\underset{Q}{\text{minimize}} \quad \|P_{11} + P_{12}Q\|^2 \\
&\text{subject to} \quad Q \in \mathcal{I}(\mathcal{P}).
\end{aligned}
\tag{28}
$$

Let v_2^* be the optimal value of (28). Recall that the optimal solution Q^* of (28) was obtained in Lemma 4 as (26). We note that if K^* is an optimal solution to (5) then the corresponding $\bar{Q} := K^*(I - P_{22}K^*)^{-1}P_{21}$ is feasible for (28). Hence $v_2^* \leq v_1^*$. We will show

that the controller in (16) is optimal by showing that $\bar{Q} = Q^*$ (so that $v_1^* = v_2^*$). We will also show that $K^* \in \mathcal{I}(\mathcal{P})$ and is internally stabilizing. Since it achieves the lower bound v_2^* and is internally stabilizing, it *must* be optimal.

Given K^*, one can evaluate $\bar{Q} := K^*(I - P_{22}K^*)^{-1}P_{21}$. To do so we use K^* as per (16) and

$$M = \begin{bmatrix} 0 & I \\ P_{21} & P_{22} \end{bmatrix} = \left[\begin{array}{c|cc} A & F & B \\ \hline 0 & 0 & I \\ I & 0 & 0 \end{array} \right]$$ and use (27) to

obtain

$$\bar{Q} = \left[\begin{array}{cc|c} A - B\Sigma\mathbf{K}\Pi_x^T & -B\Sigma\mathbf{K}(\Pi_q^T - \Pi_x^T\Sigma\Pi_q^T) & F \\ B_\Phi & A_\Phi - B_\Phi\Sigma\Pi_q^T & 0 \\ \hline -\Sigma\mathbf{K}\Pi_x^T & -\Sigma\mathbf{K}(\Pi_q^T - \Pi_x^T\Sigma\Pi_q^T) & 0 \end{array} \right].$$

Recall that Q^* given by (26) is the optimal solution to (8) (which constitutes a lower bound to the problem we are trying to solve). We are trying to show that it is achievable by explicitly producing K^* such that $\bar{Q} := K^*(I - P_{22}K^*)^{-1}P_{21}$ and $\bar{Q} = Q^*$, thereby proving optimality of K^*.

While Q^* in (26) and $\bar{Q}$ obtained above appear different at first glance, their state-space realizations are actually equivalent modulo a coordinate transformation. Recall that Π_q and Π_x are coordinate projection operators on complementary subspaces. As a result the matrix $\begin{bmatrix} \Pi_x^T & \Pi_q^T \end{bmatrix}$ is a permutation matrix. Define the matrices

$$\Lambda := \begin{bmatrix} \Pi_x^T & \Pi_q^T \end{bmatrix} \begin{bmatrix} I & -\Sigma\Pi_q^T \\ 0 & I \end{bmatrix},$$

$$\Lambda^{-1} = \begin{bmatrix} I & \Sigma\Pi_q^T \\ 0 & I \end{bmatrix} \begin{bmatrix} \Pi_x \\ \Pi_q \end{bmatrix}.$$

Note that Λ is a square, invertible matrix. Changing state coordinates on Q^* using Λ via:

$$\mathbf{A} \mapsto \Lambda^{-1}\mathbf{A}\Lambda \qquad \Pi_x^T F \mapsto \Lambda^{-1}\Pi_x^T F \qquad -\Sigma\mathbf{K} \mapsto -\Sigma\mathbf{K}\Lambda$$

along with the relations $\Sigma\Pi_q^T\Pi_q + \Pi_x = \Sigma$, $\Sigma\mathbf{A} + B\Sigma\mathbf{K} = A\Sigma$, and $A\Sigma\Pi_x^T = A$, we see that the transformed realization of Q^* is equal to the realization of $\bar{Q}$, and hence $Q^* = \bar{Q}$.

Using (4) for the open loop, (16) for the controller and the LFT formula (27) to compute the closed loop map, one obtains that the closed-loop state transition matrix is given by

$$\begin{bmatrix} A - B\Sigma\mathbf{K}\Pi_x^T & -B\Sigma\mathbf{K}(\Pi_q^T - \Pi_x^T\Sigma\Pi_q^T) \\ B_\Phi & A_\Phi - B_\Phi\Sigma\Pi_q^T \end{bmatrix} \cong \mathbf{A}.$$

By Lemma 5, the closed loop is internally stable.

To prove that $K^* \in \mathcal{I}(\mathcal{P})$ we produce an explicit factorization $K^* = K_\Phi\Phi^{-1}$, where both factors are in

$\mathcal{I}(\mathcal{P})$. First, we define the factors Φ and K_Φ via

$$\begin{bmatrix} \Phi \\ K_\Phi \end{bmatrix} = \left[\begin{array}{c|c} A_\Phi & B_\Phi \\ \hline \Sigma\Pi_q^T & I \\ -\Sigma\mathbf{K}\Pi_q^T & -\Sigma\mathbf{K}\Pi_x^T \end{array} \right]. \qquad (29)$$

Using the state-space factorization formula [30, pp. 52] and the formula (16), it is straightforward to verify that this is indeed a valid coprime factorization of K^*. Note that the transfer function Φ^{-1} is shown in Fig. 4. The map from γ to x is in fact invertible, and this inverse map is $x = \Phi\gamma$. Then $q = \Psi\Phi\gamma$, which in turn gives a factorization:

$$\begin{bmatrix} q \\ \gamma \end{bmatrix} = \begin{bmatrix} \Psi \\ I \end{bmatrix} \Phi^{-1}.$$

This in turn allows us to interpret the factorization of the entire controller graphically.

Note that Φ is invertible since its feed-through term is the identity. Lastly, we verify that the factors $K_\Phi, \Phi \in \mathcal{I}(\mathcal{P})$. and that Φ is invertible. Note that A_Φ and B_Φ are block diagonal matrices. Let us introduce the following notation for their diagonal blocks:

$$A_\Phi(j) = E_{\downarrow\downarrow j}^T \left(A - BK(\widehat{\downarrow j}, \downarrow j) \right) E_{\downarrow\downarrow j}$$
$$B_\Phi(j) = E_{\downarrow\downarrow j}^T \left(A - BK(\widehat{\downarrow j}, \downarrow j) \right) E_j.$$

Note that the j^{th} columns of Φ, K_Φ are given by the formula:

$$\Phi(j) = \left[\begin{array}{c|c} A_\Phi(j) & B_\Phi(j) \\ \hline E_{\downarrow\downarrow j} & I \end{array} \right]$$

$$K_\Phi(j) = \left[\begin{array}{c|c} A_\Phi(j) & B_\Phi(j) \\ \hline -K(\widehat{\downarrow j}, \downarrow j)E_{\downarrow\downarrow j} & -K(\widehat{\downarrow j}, \downarrow j)E_j \end{array} \right]. \qquad (30)$$

Furthermore, $K_\Phi(j) = -K(\widehat{\downarrow j}, \downarrow j)\Phi(j)$, and if i is such that $j \not\preceq i$ then the i^{th} entry of $\Phi(j)$ is zero since the corresponding row of $E_{\downarrow\downarrow j}$ is zero. By similar reasoning, $K_\Phi(j)$ also has this property. Thus, when we construct the matrices $\Phi = \begin{bmatrix} \Phi(1) & \ldots & \Phi(p) \end{bmatrix}$, $K_\Phi = \begin{bmatrix} K_\Phi(1) & \ldots & K_\Phi(p) \end{bmatrix}$ by column concatenation, we see that both $\Phi \in \mathcal{I}(\mathcal{P})$ and $K_\Phi \in \mathcal{I}(\mathcal{P})$. ∎

Proof of Theorem 4: Using equations (16) and (17) it follows that $K^* = -\Sigma\mathbf{K}\Theta$, from which the first expression in the statement follows directly. The second expression is a simple manipulation of the first. ∎

VII. Conclusions

In this paper we provided a state-space solution to the problem of computing an $\mathcal{H}_2$-optimal

decentralized controller for a poset-causal system. We introduced a new decomposition technique that enables one to separate the decentralized problem into a set of centralized problems. We gave explicit state-space formulae for the optimal controller and provided degree bounds on the controller. We illustrated our technique with a numerical example. Our approach also enabled us to provide insight into the structure of the optimal controller. We introduced a transfer function Θ that relates the role of the controller to state prediction. In future work, it would be interesting to attempt to apply this decomposition technique to a wider class of decentralization structures. Other interesting directions include the study of control laws over posets in the presence of output feedback, and optimal design with respect to the $\mathcal{H}_\infty$ norm.

REFERENCES

[1] M. Aigner. *Combinatorial theory*. Springer-Verlag, 1979.

[2] S. M. V. Andalam and N. Elia. Design of distributed controllers realizable over arbitrary directed networks. In *Proceedings of the 49th IEEE Conference on Decision and Control*, 2010.

[3] B. Bamieh and P. Voulgaris. A convex characterization of distributed control problems in spatially invariant systems with communication constraints. *Systems and Control Letters*, 54(6):575–583, 2005.

[4] V. Blondel and J. N. Tsitsiklis. NP-hardness of some linear control design problems. *SIAM Journal on Control and Optimization*, 35(6):2118–2127, 1997.

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

[6] F. M. Callier and C. A. Desoer. *Linear System Theory*. Springer-Verlag, 1991.

[7] R. B. Carroll and D. A. Brask. A practical implementation for solutions to the algebraic matrix Riccati equation in a LQCM setting. *Journal of Economic Dynamics and Control*, 23(1):1–7, 1998.

[8] P. Cousot and R. Cousot. Abstract interpretation: a unified lattice model for static analysis of programs by construction or approximation of fixpoints. In *Conference Record of the Fourth Annual ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages*, pages 238–252, Los Angeles, California, 1977. ACM Press, New York, NY.

[9] B. A. Davey and H. A. Priestley. *Introduction to Lattices and Order*. Cambridge University Press, Cambridge, 1990.

[10] W. Findeisen, F. N. Bailey, M. Brdys, K. Malinowski, P. Tatjewoki, and A. Wozniak. *Control and Coordination in Hierarchical Systems*. Wiley, 1980.

[11] 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.

[12] L. Lessard, M. Kristalny, and A. Rantzer. On structured realizability and stabilizability of linear systems. *arXiv:1212.1839*, 2012.

[13] M. D. Mesarovic, D. Macko, and Y. Takahara. *Theory of Hierarchical, Multilevel Systems*. Academic Press, 1970.

[14] R. E. Mullans and D.L. Elliott. Linear systems on partially ordered time sets. In *Proceedings of the IEEE Conference on Decision and Control including the 12th Symposium on Adaptive Processes*, volume 12, pages 334–337, 1973.

[15] G.-C. Rota. On the foundations of combinatorial theory I. Theory of Möbius functions. *Probability theory and related fields*, 2(4):340–368, 1964.

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

[17] N. Sandell, P. Varaiya, M. Athans, and M. Safonov. Survey of decentralized control methods for large scale systems. *IEEE Transactions on Automatic Control*, 23(2):108–128, 1978.

[18] P. Shah. *A Partial Order Approach to Decentralized Control*. PhD thesis, Masschusetts Institute of Technology, (available at http://www.mit.edu/˜pari), 2011.

[19] P. Shah and P. A. Parrilo. A partial order approach to decentralized control. In *Proceedings of the 47th IEEE Conference on Decision and Control*, 2008.

[20] P. Shah and P. A. Parrilo. A partial order approach to decentralized control of spatially invariant systems. In *Forty-Sixth Annual Allerton Conference on Communication, Control, and Computing*, 2008.

[21] P. Shah and P. A. Parrilo. A poset framework to model decentralized control problems. In *Proceedings of the 48th IEEE Conference on Decision and Control*, 2009.

[22] P. Shah and P. A. Parrilo. An optimal architecture of controllers on poset-causal systems. In *Proceedings of the 50th IEEE Conference on Decision and Control*, 2011.

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

[24] J. Swigart and S. Lall. An explicit state-space solution for a decentralized two-player optimal linear-quadratic regulator. In *Proceedings of the American Control Conference*, 2010.

[25] J. Swigart and S. Lall. Optimal synthesis and explicit state-space solution for a decentralized two-player linear-quadratic regulator. In *Proceedings of the 49th IEEE Conference on Decision and Control*, 2010.

[26] D. Del Vecchio and R. M. Murray. Discrete state estimators for a class of hybrid systems on a lattice. In *HSCC*, pages 311–325, 2004.

[27] P. Voulgaris. Control of nested systems. In *Proceedings of the American Control Conference*, 2000.

[28] P. Voulgaris. A convex characterization of classes of problems in control with specific interaction and communication structures. In *Proceedings of the American Control Conference*, 2001.

[29] B. F. Wyman. Time varying linear discrete-time systems: II: duality. *Pacific Journal of Mathematics*, 86(1):361–377, 1980.

[30] Q.-C. Zhong. *Robust Control of Time-Delay Systems*. Springer, 2006.

[31] K. Zhou, J. Doyle, and K. Glover. *Robust and Optimal Control*. Prentice Hall, 1996.

[32] K. Zhou and J. C. Doyle. *Essentials of Robust Control*. Prentice Hall, 1998.

Parikshit Shah Parikshit Shah is currently a Member of Research Staff at Philips Research. He has been a research associate at the University of Wisconsin. He received his PhD from the Massachusetts Institute of Technology in June 2011. Earlier, he received a Masters degree from Stanford and a Bachelors degree from the Indian Institute of Technology, Bombay. His research interests lie in the areas of convex optimization, control theory, system identification, and statistics.

Pablo Parrilo Pablo A. Parrilo is a Professor of Electrical Engineering and Computer Science at the Massachusetts Institute of Technology. He is currently an Associate Director of the Laboratory for Information and Decision Systems (LIDS), and is also affiliated with the Operations Research Center (ORC). Past appointments include Assistant Professor at the Automatic Control Laboratory of the Swiss Federal Institute of Technology (ETH Zurich), Visiting Associate Professor at the California Institute of Technology, as well as short-term research visits at the University of California at Santa Barbara (Physics), Lund Institute of Technology (Automatic Control), and University of California at Berkeley (Mathematics). He received an Electronics Engineering undergraduate degree from the University of Buenos Aires, and a PhD in Control and Dynamical Systems from the California Institute of Technology.

His research interests include optimization methods for engineering applications, control and identification of uncertain complex systems, robustness analysis and synthesis, and the development and application of computational tools based on convex optimization and algorithmic algebra to practically relevant engineering problems.

Prof. Parrilo has received several distinctions, including a Finmeccanica Career Development Chair, the Donald P. Eckman Award of the American Automatic Control Council, the SIAM Activity Group on Control and Systems Theory (SIAG/CST) Prize, and the IEEE Antonio Ruberti Young Researcher Prize. He is currently on the Editorial Board of the MOS/SIAM Book Series on Optimization.