2011 50th IEEE Conference on Decision and Control and
European Control Conference (CDC-ECC)
Orlando, FL, USA, December 12-15, 2011

An Optimal Controller Architecture for Poset-Causal Systems

Parikshit Shah and Pablo A. Parrilo

Abstract— **We propose a novel and natural architecture for decentralized control, that is applicable whenever the underlying system has the structure of a partially ordered set (poset). This controller architecture is based on the Möbius transform of the poset, and enjoys simple and appealing separation properties, since the closed-loop dynamics can be analyzed in terms of decoupled subsystems. The controller structure provides rich and interesting connections between concepts from order theory such as Möbius inversion and control-theoretic concepts such as state prediction, correction, and separability. In addition, using our earlier results on $\mathcal{H}_2$-optimal decentralized control for arbitrary posets, we prove that the $\mathcal{H}_2$-optimal controller in fact possesses the proposed structure, thereby establishing the optimality of the new controller architecture.**

I. Introduction

Motivated by the intuition that *acyclic* structures within the context of decentralized control should be tractable, the authors began a systematic study of a class of systems known as poset-causal systems in [8]. In follow-up work [7], [9] we showed that the problem of computing $\mathcal{H}_2$-optimal controllers using state-space techniques over this class of systems was tractable, with efficient solutions in terms of uncoupled Riccati equations. We also provided several intuitive explanations of the controller structure, though a detailed analysis of the same was not presented.

In this paper we are concerned with the following questions: "What is a sensible architecture of controllers for poset-causal systems? What should be the role of controller states, and what computations should be involved in the controller?" This paper focuses on answering this *architectural* question. Our main contributions in this paper are:

- We propose a controller architecture that involves natural concepts from order theory and control theory as building blocks.
- We show that a natural coordinate transformation of the state variables yields a novel *separation principle*
- We show that the optimal $\mathcal{H}_2$ controller (with state-feedback) studied in [9] has precisely the proposed controller structure.

The controller structure that we propose in this paper is as follows. At each subsystem of the overall system, the partial ordering of the information structure allows one to decompose the global state into "upstream" states (i.e. states that are available), "downstream" (these are unavailable) and "off-stream" states (corresponding to uncomparable elements of the poset). The downstream and off-stream states are (partially) predicted using available upstream information, this prediction is the role of the controller states. The best available information of the global state at each subsystem is then described using a matrix X; each column of X corresponds to the best local guess or estimate of the overall state at a particular subsystem.

Having computed these local partial estimates, the controller then performs certain natural local operations on X that preserve the structure of the poset. These local operations are the well-known ζ and μ operations in Möbius inversion. These operations, which are intimately related to the inclusion-exclusion formula and its generalizations, have a rich and interesting theory, and appear in a variety of mathematical contexts [5]. The control inputs are of the form $U = \zeta(\mathbf{G} \circ \mu(X))$. As we will see later, the operators μ and ζ can be interpreted as generalized notions of differentiation and integration on the poset so that $\mu(X)$ may be interpreted as the differential improvement in the prediction of the local state. The quantity $\mathbf{G} \circ \mu(X)$ may therefore be interpreted as a local "differential contribution" to the overall control signal. The overall control law then aggregates all these local contributions by "integration" along the poset using ζ.

Computational and architectural issues in decentralized control have been important areas of study; we mention some related works below. From a computational standpoint, the problem of computing $\mathcal{H}_2$-optimal controllers for quadratically invariant systems was studied in [6], however that approach does not provide much insight into the structure of the optimal controller. In the context of decentralized control, the computational and architectural issues for the "Two-Player Case" were studied in [11]. This work was extended to arbitrary posets in [9] (similar results were obtained in [10]), and some hints regarding the structure of the optimal controller were provided in our previous work. Another important related work is the simpler but related *team-theory* problem over posets studied in [4] which provides us with an interesting starting point in this paper. We mention the work of Witsenhausen [12], [13] who provided important insight regarding different types of information constraints in control problems. Finally, connections between information structures, team theory and decentralized control have also been studied in [3].

The rest of this paper is organized as follows: In Section II we introduce the necessary order-theoretic and control-theoretic preliminaries. In Section III we present the basic building blocks involved in the controller architecture. In Section IV we describe in detail the proposed architecture, establish the separability principle and explain its optimality property with respect to the $\mathcal{H}_2$ norm.

The authors are with the Laboratory for Information and Decision Systems, Department of Electrical Engineering and Computer Science, Massachusetts Institute of Technology, Cambridge, MA 02139. Email: pari@mit.edu, parrilo@mit.edu.

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], [2] 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 is reflexive, anti-symmetric and transitive [1].
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 only 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 [1].

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).

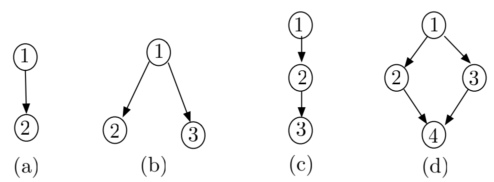

Fig. 1. Hasse diagrams of some posets.

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 a *upstream set*), and $\uparrow\uparrow p = \{q \in P \mid q \preceq p, q \neq p\}$. We define $\downarrow\uparrow p = \{q \in P \mid q \npreceq p, q \nsucceq p\}$ (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 $\mathbb{Q}$ be a ring. The set of all functions $f : P \times P \to \mathbb{Q}$ with the property that $f(x, y) = 0$ if $y \npreceq x$ is called the *incidence algebra* of $\mathcal{P}$ over $\mathbb{Q}$. It is denoted by $I(\mathcal{P})$. *
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. 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 $f, g \in I(\mathcal{P})$, 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. It is well-known that the incidence algebra is an associative algebra [1], [8].

B. Control Theoretic Preliminaries

1) Poset-causal systems: We consider the following state-space system in discrete time:

$$\begin{aligned} x[t+1] &= Ax[t] + w[t] + Bu[t] \\ z[t] &= Cx[t] + Du[t] \\ y[t] &= x[t]. \end{aligned} \tag{1}$$

In this paper we present the discrete time case only, however, we wish to emphasize that analogous results hold in continuous 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) to be partitioned into blocks in the following natural way. Let $\mathcal{P} = (P, \preceq)$ be a poset with $P = \{1, \ldots, s\}$. We think of this system as being divided into s subsystems, with subsystem i having some states $x_i[t] \in \mathbb{R}^{n_i}$, and we let $N = \sum_{i \in P} n_i$ be the total degree of the system. The control inputs at the subsystems are $u_i[t] \in \mathbb{R}^{m_i}$ for $i \in \{1, \ldots, s\}$. The external output is $z[t] \in \mathbb{R}^p$. The signal $w[t]$ is a disturbance signal. The states and inputs are partitioned in the natural way such that the subsystems correspond to elements of the poset $\mathcal{P}$ with $x[t] = [x_1[t] | x_2[t] | \ldots | x_s[t]\ ]^T$, and $u[t] = [u_1[t] | u_2[t] | \ldots | u_s[t]\ ]^T$. This naturally partitions the matrices A, B, C, D into appropriate blocks so that $A = \left[A_{ij}\right]_{i,j \in P}$, $B = \left[B_{ij}\right]_{i,j \in P}$, $C = \left[C_j\right]_{j \in P}$ (partitioned into columns), $D = \left[D_j\right]_{j \in P}$. (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. The block sizes will be obvious from the context and we denote by $\mathcal{I}(\mathcal{P})$ the block incidence algebra.

Remark In this paper, for notational simplicity we will assume $n_i = 1$, and $m_i = 1$. We emphasize that this is only done to simplify the presentation; the results hold for arbitrary block sizes n_i and m_i by interpreting the formulas "block-wise" in the obvious way.

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

$$\begin{aligned} z &= P_{11} w + P_{12} u \\ x &= P_{21} w + P_{22} u \end{aligned}$$

*Standard definitions of the incidence algebra use the opposite convention, namely $f(x, y) = 0$ if $x \npreceq y$ so their matrix representation typically has upper triangular structure. We reverse the convention so that they are lower-triangular, and thus in a control-theoretic setting one may interpret them as representing *poset-causal* maps. This reversal of convention entails transposing other standard objects like the zeta and the Möbius operators. For the same reason, we also reverse the convention of drawing Hasse diagrams so that minimal elements appear at the top of the poset.

where

$$\left[\begin{array}{cc} P_{11} & P_{12} \\ P_{21} & P_{22} \end{array}\right] = \left[\begin{array}{c|cc} A & I & B \\ \hline C & 0 & D \\ I & 0 & 0 \end{array}\right]. \quad (2)$$

(We refer the reader to [14] as a reminder of standard LFT notation used above). In this paper we will assume that $A \in \mathcal{I}(\mathcal{P})$ and $B \in \mathcal{I}(\mathcal{P})$. Indeed, this assumption ensures that the plant $P_{22}(z) = (zI - A)^{-1}B \in \mathcal{I}(\mathcal{P})$.

We call such systems *poset-causal* due to the following causality-like property among the subsystems. If an input is applied to subsystem i via u_i at some time t, the effect of the input is seen by the downstream states x_j for all subsystems $j \in\downarrow i$ (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), a causality relation between the subsystems with respect to a poset.

2) Information Constraints on Controller: In this paper, we will be interested in the design of poset-causal controllers of the form:

$$K = \left[\begin{array}{c|c} A_K & B_K \\ \hline C_K & D_K \end{array}\right]. \quad (3)$$

We will require that the controller also be poset-causal, i.e. that $K \in \mathcal{I}(\mathcal{P})$. In later sections we will present a general architecture for controllers with this structure with some elegant properties.

A control law (3) with $K \in \mathcal{I}(\mathcal{P})$ is said to be *poset-causal* since u_i depends only on x_j for $j \in\uparrow i$ (i.e. upstream information) thereby enforcing poset-causality constraints also on the controller.

C. Notation

Since we are dealing with poset-causal systems (with respect to the poset $\mathcal{P} = (P, \preceq)$), most vectors and matrices will be naturally indexed with respect to the set P (at the block level). Recall that every poset $\mathcal{P}$ has a linear extension (i.e. a total order on P which 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).

Given a matrix M, M_{ij} will as usual denote the $(i, j)^{th}$ entry. The i^{th} column will be denoted by M^i. If M is a block $|P|\times|P|$ matrix, we will denote $M(\downarrow i, \downarrow i)$ to be the sub-matrix of M whose rows and columns are in $\downarrow i$. We will also need to deal with the inverse operation: we will be given an $|S| \times |S|$ matrix K (indexed by some subset $S \subseteq P$) and we will wish to embed it into a $|P| \times |P|$ matrix by zero-padding the locations corresponding to row and column locations in $P\backslash S$. We will denote this embedded matrix by $\hat{K}$.

III. Ingredients of the Architecture

The controller architecture that we propose is composed of three main ingredients:

- The notion of *local variables*,
- A notion of a local product, denoted by "$\circ$",
- A pair of operators ζ, μ that operate on the local variables in a way that is consistent with the order-theoretic structure of the poset. These operators, called the *zeta* operator and the *Möbius* operator respectively, are classical objects and play a central role in much of order theory, number theory and combinatorics [5].

A. Local Variables and Local Products

We begin with the notion of global variables.

Definition 3: A *global variable* is a function $z : P \to \mathbb{R}$

Remark Typical global variables that we encounter will be the overall state x and the input u.

Note that the overall system is composed of $s = |P|$ subsystems. Subsystem i has access to components of the global variable corresponding to $\uparrow i$, and components corresponding to $\downarrow\downarrow i$ are unavailable. One can imagine each subsystem maintaining a local prediction of the global variable. This notion is captured by the following.

Definition 4: Let z be a global variable. A matrix $Z \in \mathbb{R}^{s\times s}$ such that $Z_{ii} = z_i$ is a *local variable* associated to z.

Remark The i^{th} column of Z, denoted by Z^i is to be thought of as the local variable at subsystem i. The components corresponding to $\downarrow\downarrow i$ correspond to the predictions of the unknown (downstream) components of z. Note that $Z_{ii} = z_i$ so that at subsystem i the component z_i of the global variable is available.

We will use the indexing $Z^i = [Z^i_j]_{j\in P}$, so that Z^i_j denotes the local prediction of z_j at subsystem i. We will sometimes also denote Z^i_j by $z_j(i)$. While local variables in general are full matrices, an important class of local variables that we will encounter will have the property that they are in $\mathcal{I}(\mathcal{P})$.

The two important local variables we will encounter are X (local state variables) and U (local input variables).

Example 2: We illustrate the concepts of global variables and local variables with an example. Consider the poset shown in Fig. 1(d). Then we can define the global variable x and a corresponding local variable X as follows:

$$x = \begin{bmatrix} x_1 \\ x_2 \\ x_3 \\ x_4 \end{bmatrix} \qquad X = \begin{bmatrix} x_1 & x_1 & x_1 & x_1 \\ x_2(1) & x_2 & x_2(1) & x_2 \\ x_3(1) & x_2(1) & x_3 & x_3 \\ x_4(1) & x_4(2) & x_4(3) & x_4 \end{bmatrix}.$$

We define the following important product:

Definition 5: Let $\mathbf{G} = \{G(1), \ldots, G(s)\}$ be a collection of maps $G(i) : \downarrow i \times \downarrow i \to \mathbb{R}$ (viewed as matrices). Let X be a local variable. We define the *local product* $\mathbf{G} \circ X$ columnwise via

$$(\mathbf{G} \circ X)^i \triangleq \hat{G}(i) X^i \quad \text{for all } i \in P. \quad (4)$$

Remark Note that if $X \in \mathcal{I}(\mathcal{P})$ and $Y = \mathbf{G} \circ X$, then it is easy to verify that $Y \in \mathcal{I}(\mathcal{P})$. We call the matrices $G(i)$ the *local gains*. Local products give rise to decoupled local relationships in the following natural way. Let X, Y be local variables. If they are related via $Y = \mathbf{G} \circ X$ then the relationship between X and Y is said to be *decoupled*. This is because, by definition,

$$Y^k = \hat{G}(k)X^k \quad \text{for all } k \in P.$$

Thus the maps relating the pairs (X^k, Y^k) are *decoupled* across all $k \in P$ (i.e. Y^k depends only on X^k and not on X^j for any other $j \neq k$).

Definition 6: Let $M \in \mathbb{R}^{s \times s}$ be a matrix. Define

$$\Pi(M) = \begin{cases} M_{ij} & \text{for } i \le j \\ 0 & \text{otherwise.} \end{cases} \qquad \Pi_\perp(M) = \begin{cases} M_{ij} & \text{for } j < i \\ 0 & \text{otherwise.} \end{cases}$$

Thus $\Pi(M)$ simply corresponds to the projection of the matrix M onto the incidence algebra $\mathcal{I}(\mathcal{P})$ viewed as a subspace of matrices, and $\Pi_\perp(M)$ onto its orthogonal complement.

B. The Möbius and zeta operators

We first remind the reader of two important order-theoretic notions, namely the *zeta and Möbius operators*. These are well-known concepts in order theory that generalize discrete integration and finite differences (i.e. discrete differentiation) to posets.

Definition 7: Let $\mathcal{P} = (P, \preceq)$. The *zeta matrix* ζ is defined to be the matrix $\zeta : P \times P \to \mathbb{R}$ such that $\zeta(i, j) = 1$ whenever $j \preceq i$ and zeroes elsewhere. The *Möbius matrix* is its inverse, $\mu := \zeta^{-1}$.

These matrices may be viewed as operators acting on functions on the poset $f : P \to \mathbb{R}$ (the functions being expressed as row vectors). The matrices ζ, μ, which are members of the incidence algebra, act as linear transformations on f in the following way:

$$\begin{aligned} \zeta : \mathbb{R}^{|P|} &\to \mathbb{R}^{|P|} & \qquad \mu : \mathbb{R}^{|P|} &\to \mathbb{R}^{|P|} \\ f &\mapsto f\zeta^T & f &\mapsto f\mu^T. \end{aligned}$$

Note that $\zeta(f)$ is also a function on the poset given by

$$(\zeta(f))_i = \sum_{j \preceq i} f_j. \tag{5}$$

This may be naturally interpreted as a discrete integral of the function f over the poset.

The role of the Möbius operator is the opposite: it is a generalized finite difference (i.e. a discrete form of differentiation over the poset). If $f : P \to \mathbb{R}$ is a local variable then the function $\mu(f) : P \to \mathbb{R}$ may be computed recursively by:

$$(\mu(f))_i = \begin{cases} f_i & \text{for } i \text{ a minimal element,} \\ f_i - \sum_{j \prec i} (\mu(f))_j & \text{otherwise.} \end{cases} \tag{6}$$

Example 3: Consider the poset in Figure 1(c). The zeta and the Möbius matrices are given by:

$$\zeta = \begin{bmatrix} 1 & 0 & 0 \\ 1 & 1 & 0 \\ 1 & 1 & 1 \end{bmatrix} \qquad \mu = \begin{bmatrix} 1 & 0 & 0 \\ -1 & 1 & 0 \\ 0 & -1 & 1 \end{bmatrix}.$$

If $f = \begin{bmatrix} f_1 & f_2 & f_3 \end{bmatrix}$, then

$$\begin{aligned} \zeta(f) &= \begin{bmatrix} f_1 & f_1 + f_2 & f_1 + f_2 + f_3 \end{bmatrix} \\ \mu(f) &= \begin{bmatrix} f_1 & f_2 - f_1 & f_3 - f_2 \end{bmatrix}. \end{aligned}$$

We now define modified versions of the zeta and Möbius operators that extend the actions of μ and ζ from global variables x to local variables X. Let ζ and μ be matrices as defined in Definition 7.

Definition 8: Let X be a local variable. Define the operators $\mu : \mathbb{R}^{s \times s} \to \mathcal{I}(\mathcal{P})$ and $\zeta : \mathbb{R}^{s \times s} \to \mathcal{I}(\mathcal{P})$ acting via

$$\zeta(X) = \Pi(X\zeta^T) \qquad \mu(X) = \Pi(X\mu^T). \tag{7}$$

Lemma 1: The operators ζ and μ may be written more explicitly as

$$\zeta(X)^i_j \triangleq \sum_{k \preceq i} X^k_j \qquad \mu(X)^i_j \triangleq X^i_j - \sum_{k \prec i} \mu(X)^k_j \tag{8}$$

for $i \preceq j$ and 0 otherwise.

Proof: The proofs follow in a straightforward fashion from (5) and (6). ∎

Note that if $Y = \mu(X)$ then Y is a local variable in $\mathcal{I}(\mathcal{P})$. The operator ζ has the natural interpretation of *aggregating or integrating* the local variables X^k for $k \in P$, whereas μ performs the inverse operation of differentiation of the local variables.

Example 4: We illustrate the action of μ acting on a local variable. Consider the local variable X from Example 2. It is easy to verify that

$$\mu(X) = \begin{bmatrix} x_1 & 0 & 0 & 0 \\ x_2(1) & x_2 - x_2(1) & 0 & 0 \\ x_3(1) & 0 & x_3 - x_3(1) & 0 \\ x_4(1) & x_4(2) - x_4(1) & x_4(3) - x_4(1) & x_4 - x_4(3) - x_4(2) + x_4(1) \end{bmatrix}.$$

Lemma 2: The operators (μ, ζ) satisfy the following properties:

1) (μ, ζ) are invertible restricted to $\mathcal{I}(\mathcal{P})$ and are inverses of each other so that for all local variables $X \in \mathcal{I}(\mathcal{P})$,
$$\zeta(\mu(X)) = \mu(\zeta(X)) = X.$$
2) $\mu(X) = \mu(\Pi(X))$ and $\zeta(X) = \zeta(\Pi(X))$.
3) Let $A, X \in \mathcal{I}(\mathcal{P})$. Then $\mu(AX) = A\mu(X)$, and $\zeta(AX) = A\zeta(X)$.

Proof: The proof is straightforward, we omit it due to space constraints. ∎

Note that if $X \notin \mathcal{I}(\mathcal{P})$ then $\zeta(\mu(X)) = \Pi(X)$. The second part of the preceding lemma says that $\mu(X)$ and $\zeta(X)$ *depend only on the components of X that lie in $\mathcal{I}(\mathcal{P})$*, i.e. on $\Pi(X)$.

Since ζ and μ may be interpreted as integration and differentiation operators, the first part of the above lemma may be viewed as a "poset" version of the fundamental theorem of calculus.

IV. Proposed Architecture

A. Local States and Local Inputs

Having defined local and global variables, we now specialize these concepts to our state-space system (1). We will denote x_j to be the true state at subsystem j. We denote $x_j(i)$ to be a prediction of state x_j at subsystem i. Recall the information constraints at subsystem i:

- Information about $\downarrow\downarrow i$: This state information is unavailable, so a (possibly partial) prediction of x_j for $j \in \downarrow\downarrow i$ is formed. We denote this prediction by $x_j(i)$. *Computing these partial predictions is the role of the controller states.*

- Information about $\uparrow i$: Complete state information about x_j for $j \in \uparrow i$ is available, so that $x_j(i) = x_j$. Moreover, the predictions from upstream $x_k(j)$ for all $k \in P$ and $j \leq i$ are also available.
- At subsystem i, state information about x_j for j not comparable to i is unavailable. The prediction of x_j is computed using $x_j(k)$ for $k \prec i$.

Analogous information constraints hold also for the inputs. At a particular subsystem, information about downstream inputs is not available. Consequently, we introduce the notion of prediction of unknown inputs, with similar notation as that for the states. These ideas can be formalized by defining local variables that capture the best available information at the subsystems. We introduce two local variables:

1) The local state X associated with the system state x,
2) The local input U associated with the controller input u.

The local state (as also the input) satisfies the following properties:

1) $X^i_j = x_j$ for $j \leq i$ (true states available for upstream subsystems)
2) X^i_j is a prediction of x_j for $j \not\leq i$.

Example 5: Consider the poset shown in Fig. 1(d). The matrix X shown in Example 2 is a local state variable. The predicted partial states are $x_2(1), x_3(1), x_4(1), x_4(2), x_4(3)$. The plant states are x_1, x_2, x_3, x_4. Note that since subsystems 1 and 2 have the same information about subsystem 3 (2 and 3 are unrelated in the poset), the best estimate of x_3 at subsystem 2 is $x_3(1)$.

We now clarify the notion of a partial prediction with an example.

Example 6: Consider the system composed of three subsystems with $P = \{1, 2, 3\}$ with $1 \leq 3$ and $2 \leq 3$:

$$\begin{bmatrix} x_1 \\ x_2 \\ x_3 \end{bmatrix}[t+1] = \begin{bmatrix} A_{11} & 0 & 0 \\ 0 & A_{22} & 0 \\ A_{31} & A_{32} & A_{33} \end{bmatrix}\begin{bmatrix} x_1 \\ x_2 \\ x_3 \end{bmatrix}[t] + \begin{bmatrix} B_{11} & 0 & 0 \\ 0 & B_{22} & 0 \\ B_{31} & B_{32} & B_{33} \end{bmatrix}\begin{bmatrix} u_1 \\ u_2 \\ u_3 \end{bmatrix}[t].$$

Note that subsystem 1 has no information about the state of subsystem 2. Moreover, the state x_1 or input u_1 does not affect the dynamics of 2 (their respective dynamics are uncoupled). Hence the only sensible prediction of x_2 at subsystem 1 is $x_2(1) = 0$ (the situation for $u_2(1)$ is identical). However, both the states x_1, x_2 and inputs u_1, u_2 affect x_3 and u_3. Since x_2 and u_2 are unknown, the state $x_3(1)$ can at best be 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 $x_3(2)$ is only a partial prediction of x_3. Indeed, one can show that $x_3(1) + x_3(2)$ is a more accurate prediction of the state x_3, and when suitably designed, their sum converges to the true state x_3.

Note that at subsystem i one can naturally decompose the local state into components belonging to $\downarrow i$ (downstream elements), and $(\downarrow i)^c$ (upstream and off-stream elements). The downstream components correspond to $X_d \triangleq \Pi(X) \in \mathcal{I}(\mathcal{P})$ and the other components to $X_u \triangleq \Pi_\perp(X)$. Thus we can decompose the state into

$$X = \Pi(X) + \Pi_\perp(X) = X_d + X_u.$$

One can similarly decompose $U = U_d + U_u$. We will see subsequently that the diagonal components of X_d are the plant states, the other elements in $\mathcal{I}(\mathcal{P})$ are the controller states. Moreover, the components in X_u will be completely determined by the elements in X_d. (Analogous properties for U hold).

B. Role of μ

We now give a natural interpretation of the operator $\mu(X)$ in terms of the differential improvement in predicted states with the help of an example.

Example 7: Consider the poset shown in Fig. 2, and let us inspect the predictions of the state x_6 at the various subsystems. The prediction of x_6 at subsystem 1 is $x_6(1)$ and

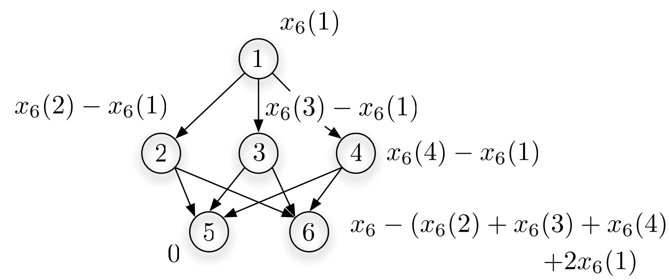

Fig. 2. Poset showing the differential improvement of the prediction of state x_6 at various subsystems.

the prediction of x_6 at subsystem 2 is $x_6(2)$. The differential improvement in the prediction at subsystem 2 regarding the state x_6 is $x_6(2) - x_6(1)$. At subsystems 3 and 4, the formulae for the differential improvements are similar. The differential improvement in x_6 at subsystem 5 is zero. These are depicted in Fig. 2.

C. Control Law

We now formally propose the following control law:

$$U_d = \zeta(\mathbf{G} \circ \mu(X)). \tag{9}$$

We make the following remarks about this control law.

Remarks

1) We note that (9) specifies U_d which amounts to specifying the input $(U_d)^i_i = u_i$ for all $i \in P$. It also specifies $(U_d)^i_j = u_j(i)$ for $i \prec j$ which is the prediction of the input u_j at an downstream subsystem i.
2) Since $\hat{G}(i)$ is non-zero only on rows and columns in $\downarrow i$, the controller respects the information constraints. Thus for any choice of gains $G(i)$, the resulting controller respects the information constraints. In this sense (9) may be viewed as a parameterization of controllers.
3) The control law (9) may be alternatively written as $U^i_d = \sum_{k \leq i} G(k)\mu(X)^k$. The control law has the following interpretation. If i is a minimal element of the poset $\mathcal{P}$, then $\mu(X)^i = X^i_d$, the vector of partial predictions of the state at i. The local control law uses these partial predictions with the gain $G(i)$. If i is a non-minimal element it aggregates all the control laws from $\uparrow\uparrow i$

and adds a *correction term* based on the differential improvement in the global state-prediction $\mu(X)^i$. This correction term is precisely $G(i)\mu(X)^i$.

Example 8: Consider a poset causal system where the underlying poset is shown in Fig 1(d). The controller architecture described above is of the form $U_d^i = \sum_{k \prec i} G(k)\mu(X)^k$ (where U^i is a vector containing the predictions of the global input at subsystem i). Noting that $(U_d)_i^i = u_i$, we write out the control law explicitly to obtain:

$$\begin{bmatrix} u_1 \\ u_2 \\ u_3 \\ u_4 \end{bmatrix} = G(1)\begin{bmatrix} x_1 \\ x_2(1) \\ x_3(1) \\ x_4(1) \end{bmatrix} + G(2)\begin{bmatrix} 0 \\ x_2 - x_2(1) \\ 0 \\ x_4(2) - x_4(1) \end{bmatrix} +$$

$$G(3)\begin{bmatrix} 0 \\ 0 \\ x_3 - x_3(1) \\ x_4(3) - x_4(1) \end{bmatrix} + G(4)\begin{bmatrix} 0 \\ 0 \\ 0 \\ x_4 - x_4(2) - x_4(3) + x_4(1) \end{bmatrix}.$$

D. State Prediction

Recall that at subsystem i the states x_j for $j \in \downarrow\downarrow i$ are unavailable and must be predicted. Typically, one would predict those states via an observer. However, those states are *unobservable*; only the state x_k for $k \in \uparrow i$ are observable, and are in fact directly available. In this situation, rather than using an observer one constructs a *predictor* to predict the unobservable states. These predictions are computed by the controller via prediction dynamics, which we now specify. Recall the decomposition $X = X_d + X_u$ where X_d contains the plant states and the predicted downstream states. Given X_d, we propose that X_u be computed via:

$$X_u = \Pi_\perp(\mu(X_d)\zeta^T). \tag{10}$$

(Thus specifying X_d completely specifies X_u and hence X). This ensures that $(X_u)_j^i = x_i$ for $j \prec i$ (easily verified), i.e. that subsystem i uses the true states in $\uparrow\uparrow i$. Furthermore, the predictions for the off-stream components are computed via $(X_u)_j^i = \sum_{k \prec i} \mu(X)_j^k$.

Example 9: For the poset in Fig. 1(d),

$$X_u = \begin{bmatrix} 0 & x_1 & x_1 & x_1 \\ 0 & 0 & x_2(1) & x_2 \\ 0 & x_3(1) & 0 & x_3 \\ 0 & 0 & 0 & 0 \end{bmatrix}.$$

In an analogous manner to X, the local variable U can be decomposed into $U = U_d + U_u$, and U_u can be computed from U_d analogous to (10).

We now describe the prediction dynamics. Since the dynamics of the true state evolve according to $x[t+1] = Ax[t] + Bu[t]$, each subsystem can simulate these dynamics using the local states and inputs. Locally each subsystem implements the dynamics $X^i[t+1] = AX^i[t] + BU^i[t]$. This can be compactly written as

$$X[t] = AX[t] + BU[t].$$

However, at subsystem i only the states in $\downarrow i$ (corresponding to the components in X_d) need to be predicted (the other components are determined by (10)). Writing $X = X_d + X_u$ and projecting the dynamics onto X_d, we obtain:

$$\begin{aligned} X_d[t+1] &= AX_d[t] + BU_d[t] + R[t] \\ R[t] &= \Pi(AX_u[t] + BU_u[t]). \end{aligned} \tag{11}$$

We think of $R[t]$ as the influence of the upstream components (and also the unrelated components) in predicting X_d.

E. Separation Principle

As a consequence of Lemma 2, we see that $\mu(X) = \mu(X_d)$ and also $\mu(R) = A\mu(X_u) + B\mu(U_u) = 0$. Applying μ to (11) we obtain the following *modified closed-loop dynamics* in the new variables $\mu(X)$:

$$\mu(X)[t+1] = A\mu(X)[t] + B\mu(U)[t]. \tag{12}$$

Let us define $\mathbf{A} + \mathbf{BG} = \left\{A + B\hat{G}(1), \ldots, A + B\hat{G}(s)\right\}$. From (9), and the fact that $\mu(\zeta(\mathbf{G} \circ \mu(X))) = \mathbf{G} \circ \mu(X)$ we will momentarily see that the modified closed-loop dynamics are:

$$\mu(X)[t+1] = (\mathbf{A} + \mathbf{BG}) \circ \mu(X)[t]. \tag{13}$$

These dynamics describe how the differential improvements in the state evolve. If one picks U such that $\mu(U)$ stabilizes $\mu(X)$, the differential improvements are all stabilized. Thus $\mu(X)$ converges to zero, the state predictions become accurate asymptotically and the closed-loop is also stabilized. We show that (9) achieves this with an appropriate choice of local gains.

Theorem 1: Let $G(i)$ be chosen such that $A(\downarrow i, \downarrow i) + B(\downarrow i, \downarrow i)G(i)$ is stable for all $i \in P$. Then the control law (9) with local gains $G(i)$ renders (12) stable.

Proof: Since $U_d = \zeta(\mathbf{G} \circ \mu(X))$ it follows that

$$\begin{aligned} \mu(U_d) = \mu(U) &= \mu\left(\zeta(\mathbf{G} \circ \mu(X))\right) \\ &= \mathbf{G} \circ \mu(X). \end{aligned}$$

The last equality follows from Lemma 2 and the fact that $\mathbf{G} \circ \mu(X) \in \mathcal{I}(\mathcal{P})$. As a consequence, $\mu(U)^i = \hat{G}(i)\mu(X)^i$ for all $i \in P$. Hence the closed-loop dynamics (12) become:

$$\mu(X)^i[t+1] = \left(A + B\hat{G}(i)\right)\mu(X)^i[t].$$

Recalling that $\mu(X)$ is a local variable so that $\mu(X)^i$ (viewed as a vector) is non-zero only on $\downarrow i$ it is easy to see that these dynamics are stabilized exactly when $G(i)$ are picked such that $A(\downarrow i, \downarrow i) + B(\downarrow i, \downarrow i)G(i)$ are stable.

■

The dynamics of the different subsystems $\mu(X)^i$ are *decoupled*, so that the gains $G(i)$ may be picked independent of each other. This may be viewed as a *separation principle*. Henceforth, we will assume that the gains $G(i)$ have been picked in this manner. Since the closed loop dynamics of the states $x_i(j)$ are related by an invertible transformation (i.e. $X_d = \zeta(\mu(X_d))$), if the modified closed-loop dynamics (13) are stable, so are the closed-loop dynamics (11).

F. Controller Realization

We now describe two explicit controller realizations. The natural controller realization arises from the closed-loop dynamics (11) along with the control law (9) to give:

$$X_d[t+1] = AX_d[t] + BU_d[t] + R[t]$$
$$U_d[t] = \zeta(\mathbf{G} \circ \mu(X_d))[t].$$

While the above corresponds to a natural description of the controller, it is possible to specify an alternate realization. This is motivated from the following observation. The control input U depends only on $\mu(X)$. Hence, rather than implementing controller states that track the state predictions X, it is natural to implement controller states that compute $\mu(X)$ directly. Hence an equivalent realization of the controller is:

$$\begin{aligned} \mu(X)[t+1] &= A\mu(X)[t] + B\mu(U)[t] \\ U_d[t] &= \zeta(\mathbf{G} \circ \mu(X))[t]. \end{aligned} \tag{14}$$

G. Structure of the Optimal Controller

Consider again the poset-causal system considered in (2). Consider the optimal control problem:

$$\begin{aligned} &\underset{K}{\text{minimize}} \quad \|P_{11} + P_{12}K(I - P_{22}K)^{-1}P_{21}\|^2 \\ &\text{subject to} \quad K \ \text{ stabilizes } \ P, \ K \in \mathcal{I}(\mathcal{P}). \end{aligned} \tag{15}$$

The solution K^* is the $\mathcal{H}_2$-optimal controller that obeys the poset-causality information constraints described in Section II. The solution to this optimization problem was presented in [9, Theorem 2]. The main idea behind the solution procedure is as follows. Using the fact that $P_{21}, P_{22} \in \mathcal{I}(\mathcal{P})$ are square and invertible (due to the availability of state feedback) it is possible to reparametrize the above problem via $Q = K(I - P_{22}K)^{-1}P_{21}$. Indeed, this relationship is invertible and the incidence algebra structure ensures that $Q \in \mathcal{I}(\mathcal{P})$ if and only if $K \in \mathcal{I}(\mathcal{P})$. Using this the above optimization problem may be rewritten as:

$$\begin{aligned} &\underset{Q}{\text{minimize}} \quad \|P_{11} + P_{12}Q\|^2 \\ &\text{subject to} \quad Q \in \mathcal{I}(\mathcal{P}). \end{aligned} \tag{16}$$

Using the fact that the $\mathcal{H}_2$ norm is column-separable, it is possible to decouple this optimization problem into a set of s optimization problems. Each optimization problem involves the solution to a standard Riccati equation. The solution to each yields the columns of $Q^* \in \mathcal{I}(\mathcal{P})$, from which the optimal controller $K^* \in \mathcal{I}(\mathcal{P})$ may be recovered. An explicit formula for the optimal controller and other details may be found in [7], [9].

In [9], we obtain matrices $K(\downarrow j, \downarrow j)$ by solving a system of decoupled Riccati equations via $(K(\downarrow j, \downarrow j), Q(j), P(j)) = \mathrm{Ric}(A(\downarrow j, \downarrow j), B(\downarrow j, \downarrow j), C(\downarrow j), D(\downarrow j))$ (we use slightly different notation and reversed conventions in that paper, see [9] for details). The optimal solution K^* to (15) is related to the proposed architecture as follows.

Theorem 2: The controller (14) with gains $G(i) = K(\downarrow i, \downarrow i)$ for all $i \in P$ is the optimal solution to the control problem (15).

Proof: The formula for the optimal controller is provided in [9, Theorem 2]. It is straightforward to verify that the controller in (14) (the gains being $\hat{K}(\downarrow i, \downarrow i)$) is equal to the formula in [9]. We omit the detailed proof here. ■

This theorem establishes that the controller architecture proposed in this paper is *optimal* in the sense of the $\mathcal{H}_2$ norm.

V. Conclusions

In this paper we considered the problem of designing decentralized poset-causal controllers for poset-causal systems. We studied the architectural aspects of controller design, addressing issues such as the role of the controller states, and how the structure of the poset should affect the architecture. We proposed a novel architecture in which the role of the controller states was to locally predict the unknown "downstream" states. Within this architecture the controller itself performs certain natural local operations on the known and predicted states. These natural operations are the well-known zeta and Möbius operations on posets.

Having proposed an architecture, we proved two of its important structural properties. The first was a *separation principle* that enabled a decoupled choice of gains for each of the local subsystems. The second was establishing the *optimality properties* of this architecture with respect to the $\mathcal{H}_2$-optimal decentralized control problem. The proposed Möbius-based architecture is quite natural, has very appealing interpretations, and can be easily extended to more complicated and realistic formulations. These extensions will be the subject of future work.

Acknowledgement

We thank the referees for their insightful comments that enabled an improved presentation of this paper.

References

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

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

[3] A. Gattami. *Optimal Decisions with Limited Information*. PhD thesis, Lund University, 2007.

[4] 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.

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

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

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

[8] 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.

[9] P. Shah and P. A. Parrilo. $\mathcal{H}_2$-optimal decentralized control over posets: A state-space solution for state-feedback. In *Proceedings of the 49th IEEE Conference on Decision and Control*, 2010.

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

[11] 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.

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

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

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