Pattern Identification in Distributed Systems

Melkior Ornik[1], Adam C. Sniderman[2], Mireille E. Broucke[1], and Gabriele M. T. D'Eleuterio[2]

***Abstract*— A distributed system's interconnection structure emerges as a pattern in the system matrices. This pattern must be preserved through system analysis and control synthesis, and much has been written on these topics. A problem which has not received any attention to date is how to identify a pattern, given the linear system model. This paper proposes a method for identifying a pattern that is mathematically encoded through a commuting relationship with a base matrix. Our method generates the commuting relationship, when it exists. When it does not exist, our method produces the closest approximation to the commuting relationship. Further, it indicates which additional subsystem interconnections would render it achievable. We provide both an exact solution and an almost sure polynomial-time solution in the probabilistic sense. Finally, we give several examples to demonstrate the utility of this method for finding patterns in distributed systems.**

I. Introduction

A distributed linear system can be modelled by system matrices partitioned into blocks representing subsystem dynamics and their interactions. Typically, the blocks are not completely arbitrary: some blocks might be equal, some might be algebraically related, others might be fixed at zero. We have developed a class of distributed systems called *patterned systems* to capture such special structure [1]. In order for these patterns to be useful for analysis and synthesis, we require a mathematical characterization of them. A number of suitable characterization methods have been introduced in the literature; in particular, patterns can be described by polynomials of a base matrix [1] and by commuting relationships with a base matrix [2]. In this paper we follow the latter approach. As shown in prior work [2], [3], a number of common patterns in distributed systems can be encoded by commuting relationships.

This paper addresses the problem of identification of a distributed system's pattern. More specifically, given a block matrix in which certain blocks are equal and others are zero, we aim to find a commuting relationship that uniquely determines the pattern. This question of pattern identification has not been treated in the literature either in terms of commuting relationships or in terms of other encoding methods [1], [4]. Nevertheless, pattern identification remains one of the major barriers to a complete theory of patterned linear systems [1, §9.3]. This paper breaks down that barrier for patterns encoded by commuting relationships: whenever such an encoding is possible, our method will find it.

* This work is supported by the Natural Sciences and Engineering Research Council of Canada (NSERC).

[1] Department of Electrical & Computer Engineering, University of Toronto, Ontario, Canada.

[2] Institute for Aerospace Studies, University of Toronto, Ontario, Canada.

Now we outline the plan for the paper and introduce key concepts in plain terms. Given a distributed system, we first partition the system into subsystems and then catalogue the distinct subsystem matrices. Next we characterize the pattern through a so-called *pattern-generating matrix* whose role is to capture which subsystem blocks are equal. An *instantiation* of a pattern-generating matrix is a matrix whose blocks adhere to the structure of a pattern-generating matrix. The next step is to encode the resultant patterns in terms of commuting relationships, following prior work on patterned systems [2]. A matrix that commutes with all instantiations of a pattern-generating matrix is called a *base matrix*. Thus, we arrive at the fundamental problem considered in this paper. The *Pattern Identification Problem* is to find a base matrix that commutes with all instantiations of a given pattern-generating matrix, but no others. The Pattern Identification Problem is not always solvable, but we can nevertheless always produce a matrix that preserves the given pattern, and as few other patterns as possible. We provide two solution approaches: the first is exact while the second is an almost sure solution in the probabilistic sense.

The paper is organized as follows. In Section II, we introduce patterned matrices. In Section III, we propose an algorithm to solve the Pattern Identification Problem. In Section IV, we provide a measure-theoretic result that enables us to solve the Pattern Identification Problem almost surely, in polynomial time. In Section V, we give several examples, showing how patterns emerge and can be encoded in commuting relationships.

II. Generating and Encoding Patterned Matrices

The idea of patterns emerging as repeated blocks in matrices was considered in [2], through examples that showed the patterns in ring and chain systems. Similar block matrices have arisen in many analyses of symmetrically interconnected systems (e.g., [5], [6], [7], [8]). In these papers, it was not necessary to expound the idea behind these repeated blocks, but in the current work, we need to be precise. Thus, to begin our study, we will formally define these patterns; specifically, we give a way to easily characterize every system with a certain pattern, irrespective of block size.

Definition 1: Let $r, k \in \mathbb{N}$. Matrix $\tilde{A} \in \mathbb{R}^{r\times r}$ is a *pattern-generating matrix* with k *pattern components* if $[\tilde{A}]_{ij} \in \{0, 1, \ldots, k\}$ for all $1 \leq i, j \leq r$.

We denote the set of all $r \times r$ pattern-generating matrices, with any number of pattern components, by $\mathcal{P}_r$.

Definition 2: A pattern-generating matrix $\tilde{B} \in \mathcal{P}_r$ is *sparser* than a pattern-generating matrix $\tilde{A} \in \mathcal{P}_r$ if

(i) $[\tilde{A}]_{ij} = [\tilde{A}]_{i'j'} \Rightarrow [\tilde{B}]_{ij} = [\tilde{B}]_{i'j'}$, $1 \le i,j,i',j' \le r$,
(ii) $[\tilde{A}]_{ij} = 0 \Rightarrow [\tilde{B}]_{ij} = 0$, $\quad 1 \le i,j \le r$.

We denote this by $\tilde{B} \preceq \tilde{A}$, i.e., $\tilde{A} \succeq \tilde{B}$.

Definition 3: Let $r, n \in \mathbb{N}$ such that r divides n ($r|n$). Let $A \in \mathbb{R}^{n\times n}$ be a block matrix given by

$$A = \begin{bmatrix} A_{11} & A_{12} & \cdots & A_{1r} \\ A_{21} & A_{22} & \cdots & A_{2r} \\ \vdots & \vdots & \ddots & \vdots \\ A_{r1} & A_{r2} & \cdots & A_{rr} \end{bmatrix},$$

with blocks of the same size. A is an *instantiation* of the pattern-generating matrix $\tilde{A} \in \mathcal{P}_r$ if the following conditions are satisfied:

(i) $A_{ij} = 0$ when $[\tilde{A}]_{ij} = 0$,
(ii) $A_{ij} = A_{i'j'}$ when $[\tilde{A}]_{ij} = [\tilde{A}]_{i'j'}$.

Denote by $\mathcal{I}_n(\tilde{A})$ the set of all $n \times n$ instantiations of a pattern-generating matrix $\tilde{A}$. Throughout this paper, we always assume $r|n$, as otherwise Definition 3 does not make sense.

In the sense of Definitions 1 and 3, a pattern denotes certain matrix entries that are fixed at zero, and certain others that must be equal but can have any value. While this paper is restricted to patterns of that form, we note that the following results can be extended to cases where, for example, some blocks are linear combinations of others (rather than only identical or different).

An illustration of the above definitions to a standard example is provided in Example 4.

Example 4: Let $\dot{x} = Ax + Bu$ be a block circulant system with r agents, where each agent has n/r states and n/r inputs, following the definition of [3], [9]. In other words,

$$\begin{aligned} A &= \begin{bmatrix} A_1 & A_2 & \cdots & A_r \\ A_r & A_1 & \cdots & A_{r-1} \\ \vdots & \vdots & \ddots & \vdots \\ A_2 & A_3 & \cdots & A_1 \end{bmatrix}, \\ B &= \begin{bmatrix} B_1 & B_2 & \cdots & B_r \\ B_r & B_1 & \cdots & B_{r-1} \\ \vdots & \vdots & \ddots & \vdots \\ B_2 & B_3 & \cdots & B_1 \end{bmatrix}. \end{aligned} \tag{1}$$

Let $\tilde{A} \in \mathcal{P}_r$ be given by

$$\tilde{A} = \begin{bmatrix} 1 & 2 & \cdots & r \\ r & 1 & \cdots & r-1 \\ \vdots & \vdots & \ddots & \vdots \\ 2 & 3 & \cdots & 1 \end{bmatrix}. \tag{2}$$

By Definition 3, $\mathcal{I}_n(\tilde{A})$ consists exactly of matrices of the form (1). Thus, a control system $\dot{x} = Ax + Bu$ with r agents is block circulant if and only if $A, B \in \mathcal{I}_n(\tilde{A})$ for $\tilde{A}$ as in (2), where $r|n$. ◁

We note that the results of Example 4 can easily be extended to agents with n/r states and m/r inputs (where $r|n$ and $r|m$).

In order to effectively use methods of linear algebra on $\mathcal{I}_n(\tilde{A})$, we first show that this set is a vector space. The proof of this lemma is both intuitive and entirely computational, and we hence omit it.

Lemma 5: Assume $\tilde{A} \in \mathcal{P}_r$ is a pattern-generating matrix. Then, $\mathcal{I}_n(\tilde{A})$ is a vector subspace of $\mathbb{R}^{n\times n}$.

Remark 6: The basis vectors for $\mathcal{I}_n(\tilde{A})$ can easily be found explicitly. Assume that $\tilde{A}$ has k pattern components. Let $\ell = n/r$, and let $1 \le i,j \le \ell$, and $1 \le m \le k$. Let $e_{ij} \in \mathbb{R}^{\ell\times\ell}$ be a matrix with all zeros, except for 1 at position (i,j).

Now, let $e_{mij} \in \mathbb{R}^{n\times n}$ be a block matrix

$$e_{mij} = \begin{bmatrix} e_{mij}^{11} & e_{mij}^{12} & \cdots & e_{mij}^{1r} \\ e_{mij}^{21} & e_{mij}^{22} & \cdots & e_{mij}^{2r} \\ \vdots & \vdots & \ddots & \vdots \\ e_{mij}^{r1} & e_{mij}^{r2} & \cdots & e_{mij}^{rr} \end{bmatrix}$$

where blocks $e_{mij}^{i'j'}$ are as follows:

$$e_{mij}^{i'j'} = \begin{cases} 0 & \text{if } [\tilde{A}]_{i'j'} = 0\,, \\ e_{ij} & \text{if } [\tilde{A}]_{i'j'} = m\,. \end{cases}$$

It can be shown that

$$\mathcal{E} = \{e_{mij} \mid 1 \le i,j \le \ell, 1 \le m \le k, e_{mij} \ne 0\} \tag{3}$$

is a basis for $\mathcal{I}_n(\tilde{A})$.

An example of the basis described in Remark 6 is given in Example 7.

Example 7: Let $\tilde{A}$ be the pattern-generating matrix from Example 4, and let $r = 3$, $n = 6$. The basis of $\mathcal{I}_n(\tilde{A})$ is given by matrices

$$\begin{bmatrix} e_{ij} & 0 & 0 \\ 0 & e_{ij} & 0 \\ 0 & 0 & e_{ij} \end{bmatrix}, \begin{bmatrix} 0 & 0 & e_{ij} \\ e_{ij} & 0 & 0 \\ 0 & e_{ij} & 0 \end{bmatrix}, \begin{bmatrix} 0 & e_{ij} & 0 \\ 0 & 0 & e_{ij} \\ e_{ij} & 0 & 0 \end{bmatrix}$$

for $1 \le i,j \le 2 = n/r$. ◁

Having characterized the block structure that arises through a pattern, the next step is to encode it algebraically. We have chosen to encode patterns through commuting relationships, and the examples in Section V show that this encoding works for a number of canonical patterns. The above definitions and following results also hold for other possible encodings of patterns, namely using polynomials of matrices [1], [4].

III. Pattern Identification Problem

Now, we turn to the main problem of taking a specific pattern, and trying to encode it in a commuting relationship. Once that encoding is achieved, it can be used to analyze and control patterned systems while preserving their patterns (as in [2], [3], [1]). To that end, we define commuting relationships for pattern-generating matrices:

Definition 8: Let $\tilde{A} \in \mathcal{P}_r$ be a pattern-generating matrix. Matrix $V \in \mathbb{R}^{n\times n}$ is an *$\tilde{A}$-base matrix* if

$$AV = VA \text{ for all } A \in \mathcal{I}_n(\tilde{A})\,. \tag{4}$$

The set of all matrices V that satisfy (4) is denoted $\Pi(\tilde{A})$. Our goal is to relate the difficult combinatorial structure of patterns from Section II with the notion of base matrices, in order to show that the corresponding commuting relationship adequately characterizes the underlying pattern. Ideally, we would like to identify each pattern $\tilde{A}$ with a matrix in $\Pi(\tilde{A})$. However, this is not possible. For instance, every $\tilde{A}$-base matrix $V \in \Pi(\tilde{A})$ is also contained in $\Pi(\tilde{B})$ for all matrices $\tilde{B} \preceq \tilde{A}$. Additionally, all matrices $V \in \mathbb{R}^{n\times n}$ are automatically base matrices for the pattern of the identity matrix $I_n \in \mathbb{R}^{n\times n}$. Hence, a more attainable goal is to find $V \in \Pi(\tilde{A})$ that do not preserve any patterns that they do not automatically *need* to satisfy as a consequence of (4).

In other words, we want to solve the following problem:

Problem 9 (Pattern Identification Problem): Let $r, n \in \mathbb{N}$, and $\tilde{A} \in \mathcal{P}_r$. Find a matrix $V \in \mathbb{R}^{n\times n}$ such that

$$\begin{aligned} &V \in \Pi(\tilde{A})\,, \\ &V \notin \Pi(\tilde{B}) \text{ for all } \tilde{B} \in \mathcal{P}_r \text{ such that } \Pi(\tilde{A})\backslash\Pi(\tilde{B}) \neq \emptyset\,. \end{aligned}$$

Solving Problem 9 will enable us to distinguish different patterns through base matrices inasmuch as is possible. In order to approach Problem 9, we will be using the structure of spaces $\Pi(\tilde{A})$, $\tilde{A} \in \mathcal{P}_r$. The following lemma shows that these sets are vector spaces.

Lemma 10: Let $r, n \in \mathbb{N}$, and $\tilde{B} \in \mathcal{P}_r$. Then, $\Pi(\tilde{A})$ is a vector space.

Proof: By definition, $0 \in \Pi(\tilde{A}) \subseteq \mathbb{R}^{n\times n}$. Suppose that $V_1, V_2 \in \Pi(\tilde{A})$ and $\alpha \in \mathbb{R}$. Now, for any $A \in \mathcal{I}_n(\tilde{A})$, we have $V_1A = AV_1$, $V_2A = AV_2$. Thus, $(V_1 + V_2)A = A(V_1 + V_2)$ and $(\alpha V_1)A = A(\alpha V_1)$, confirming that $\Pi(\tilde{A})$ is a subspace of $\mathbb{R}^{n\times n}$. ■

Remark 11: Problem 9 can be split into two parts: first, we seek a matrix V for which $AV = VA$ for any instantiation A of a pattern-generating matrix; and second, we ensure that V does not satisfy $BV = VB$ for as many other patterns as possible. The first part is the common algebraic problem of finding the centralizer of $\mathcal{I}_n(\tilde{A})$ (e.g., [**?**, §1.4]), which can be solved in this case as follows. First, the dimension and a basis of $\Pi(\tilde{A})$ can be found using Remark 6. In particular, by Lemma 5, $VA = AV$ for all $A \in \mathcal{I}_n(\tilde{A})$ if and only if $Ve_{mij} = e_{mij}V$ for the basis matrices $e_{mij} \in \mathcal{I}_n(\tilde{A})$ from Remark 6. Then, this commuting relationship can be rearranged into $e'_{mij}\mathrm{vec}(V) = 0$, a standard system of linear equations. Solutions to this equation form a basis for $\Pi(\tilde{A})$, from which all its elements can easily be determined. Thus, only the second part of Problem 9 remains — we need a way to exclude matrices $V \in \Pi(\tilde{A}) \cap \Pi(\tilde{B})$ for $\tilde{B} \neq \tilde{A}$. ◁

By Definition 1, $\mathcal{P}_r$ is a finite set, as we can always set $k \leq r^2$. Let us order the elements of $\mathcal{P}_r = \{\tilde{B}_1, \ldots, \tilde{B}_g\}$. We will do the following: first, pick $V \in \Pi(\tilde{A})$; second, for each $\tilde{B}_i$ such that $V \in \Pi(\tilde{B}_i)$, perturb V in a way which ensures that the updated matrix is not contained in $\Pi(\tilde{B}_i)$. Since the $\Pi(\tilde{B}_i)$ are closed sets, such a small perturbation will not cause the updated V to belong to any $\Pi(\tilde{B}_i)$ that it did not belong to before. In this way, step by step, we eliminate the portions of V that correspond to undesired patterns. This method is shown formally in Algorithm 12.

Theorem 13: Let $r, n \in \mathbb{N}$, and $\tilde{A} \in \mathcal{P}_r$. Algorithm 12 returns a matrix V that solves Problem 9.

Proof: We claim the following: The sequence $V^{(0)}, \ldots, V^{(g)}$ satisfies

Algorithm 12:

Let $\mathcal{E}$ be the basis of $\Pi(\tilde{A})$, $\mathcal{E}_i$ bases of $\Pi(\tilde{B}_i)$;
$V^{(0)} := 0$;
for $i = 1, \ldots, g$
  if $\Pi(\tilde{A})\backslash\Pi(\tilde{B}_i) = \emptyset$ or $V^{(i-1)} \notin \Pi(\tilde{B}_i)$
    $V^{(i)} := V^{(i-1)}$;
  else
    Find $V' \in \mathcal{E}$ such that $V \notin \Pi(\tilde{B}_g)$;
    for $i' = 1, \ldots, i$, $\Pi(\tilde{A})\backslash\Pi(\tilde{B}_{i'}) \neq \emptyset$
      Find $E_{i'} \in \mathcal{E}_{i'}$ such that $\|E_{i'}V - VE_{i'}\| > 0$;
    end for
    if $\max_{i'} \|E_{i'}V' - V'E_{i'}\| = 0$
      $\varepsilon := 1$;
    else
      $\varepsilon := \dfrac{\min_{i'} \|E_{i'}V - VE_{i'}\|}{2\max_{i'} \|E_{i'}V' - V'E_{i'}\|}$;
    end if
    $V^{(i)} = V^{(i-1)} + \varepsilon V'$;
  end if
end for
$V := V^{(g)}$;

(i) $V^{(i)} \in \Pi(\tilde{A})$, $i = 0, \ldots, g$,
(ii) $V^{(i)} \notin \Pi(\tilde{B}_{i'})$ for all $i = 0, \ldots, g$, $i' = 1, \ldots, i$, $\Pi(\tilde{A})\backslash\Pi(\tilde{B}_{i'}) \neq \emptyset$.

We will prove this claim inductively by i. For $i = 0$, $V^{(0)} = 0$. So, (i) trivially holds, and (ii) is vacuous. Now, assume that the claim holds for all $0, \ldots, i-1$. Let us observe $V^{(i)}$.

If $\Pi(\tilde{A})\backslash\Pi(\tilde{B}_i) = \emptyset$, $V^{(i)} = V^{(i-1)}$. Hence, by the inductive assumption, (i) holds, and we have $V^{(i-1)} \notin \Pi(\tilde{B}_{i'})$ for all $i' = 1, \ldots, i-1$ such that $\Pi(\tilde{A})\backslash\Pi(\tilde{B}_{i'}) \neq \emptyset$. Since $\Pi(\tilde{A})\backslash\Pi(\tilde{B}_i) = \emptyset$, it is also true that $V^{(i)} = V^{(i-1)} \notin \Pi(\tilde{B}_{i'})$ for all $i' = 1, \ldots, i$, $\Pi(\tilde{A})\backslash\Pi(\tilde{B}_{i'}) \neq \emptyset$.

If $\Pi(\tilde{A})\backslash\Pi(\tilde{B}_i) \neq \emptyset$ and $V^{(i-1)} \notin \Pi(\tilde{B}_i)$, again $V^{(i)} = V^{(i-1)}$. Hence, (i) still holds, and as (ii) held for i and $V^{(i-1)}$, and $V^{(i)} = V^{(i-1)} \notin \Pi(\tilde{B}_i)$, (ii) continues to hold for $V^{(i)}$ as well.

Assume now $\Pi(\tilde{A})\backslash\Pi(\tilde{B}_i) \neq \emptyset$ and $V^{(i-1)} \in \Pi(\tilde{B}_i)$. As $V^{(i-1)}, V' \in \Pi(\tilde{A})$, by Lemma 10,

$$V^{(i)} = V^{(i-1)} + \varepsilon V' \in \Pi(\tilde{A})\,.$$

This verifies (i).

Let now $\mathcal{E}_{i'}$ be a basis for $\mathcal{I}_n(\tilde{B}_{i'})$, $i' = 1, \ldots, i-1$. By (ii), $E_{i'}V^{(i-1)} \neq V^{(i-1)}E_{i'}$ for some $E_{i'} \in \mathcal{E}_{i'}$. We get

$$\begin{aligned} &\|E_{i'}V^{(i)} - V^{(i)}E_{i'}\| \\ &\quad \geq \|E_{i'}V^{(i-1)} - V^{(i-1)}E_{i'}\| - \varepsilon\|E_{i'}V' - V'E_{i'}\| \\ &\quad \geq \min_{i'=1,\ldots,i-1} \|E_{i'}V^{(i-1)} - V^{(i-1)}E_{i'}\|/2 \ > 0\,. \end{aligned}$$

It remains to verify that $V^{(i)} \notin \Pi(\tilde{B}_{i'})$ holds for $i' = i$. As $V^{(i-1)} \in \Pi(\tilde{B}_i)$ and $V' \notin \Pi(\tilde{B}_i)$, $V^{(i)} = V^{(i-1)} + \varepsilon V' \notin \Pi(\tilde{B}_i)$. This verifies that (ii) holds for $V^{(i)}$.

Properties (i) and (ii) for $V^{(g)}$ correspond to the conditions of Problem 9. As $V = V^{(g)}$, we are done. ∎

IV. Measure-Theoretic Solution

Algorithm 12 and Theorem 13 provide an exact solution to Problem 9. However, this solution requires that one passes through all possible pattern-generating matrices $\tilde{B}$. As $|\mathcal{P}_r| \geq 2^{r^2}$, this presents problems for larger values of r. Hence, it is of interest to find a computationally less difficult alternative to Algorithm 12. In this section, we will use a measure-theoretic argument to obtain a solution to Problem 9 almost surely (a.s.), in polynomial time. The argument is based on the fact that proper vector subspaces of a vector space have Lebesgue measure $\lambda = 0$ in that space.

Lemma 14: Suppose $\tilde{A}, \tilde{B}_1, \ldots, \tilde{B}_f \in \mathcal{P}_r$ are pattern-generating matrices such that $\Pi(\tilde{A}) \backslash \Pi(\tilde{B}_i) \neq \emptyset$ for all $i = 1, \ldots, f$. Let λ be the Lebesgue measure on $\Pi(\tilde{A})$. Then

$$\lambda\left(\Pi(\tilde{A}) \cap \left(\bigcup_{i=1}^{f} \Pi(\tilde{B}_i)\right)\right) = 0\,.$$

Proof: Let $i = 1, \ldots, f$. By Lemma 10, $\Pi(\tilde{A}), \Pi(\tilde{B}_i)$ are vector subspaces of $\mathbb{R}^{n \times n}$. Hence, $\Pi(\tilde{A}) \cap \Pi(\tilde{B}_i)$ is a vector subspace of $\Pi(\tilde{A})$. By the assumption $\Pi(\tilde{A}) \backslash \Pi(\tilde{B}_i) \neq \emptyset$, we get

$$\dim\left(\Pi(\tilde{A}) \cap \Pi(\tilde{B}_i)\right) < \dim(\Pi(\tilde{A}))\,. \tag{5}$$

From (5), it is well-known (see, e.g., [10]) that

$$\lambda\left(\Pi(\tilde{A}) \cap \Pi(\tilde{B}_i)\right) = 0\,. \tag{6}$$

Using the finite subadditivity of λ, it follows from (6) that

$$\lambda\left(\Pi(\tilde{A}) \cap \left(\bigcup_{i=1}^{f} \Pi(\tilde{B}_i)\right)\right) \leq \sum_{i=1}^{f} \lambda\left(\Pi(\tilde{A}) \cap \Pi(\tilde{B}_i)\right) = 0\,.$$

∎

The above lemma shows that almost all elements of $\Pi(\tilde{A})$ will satisfy the conditions of Problem 9. In other words, choosing elements of $\Pi(\tilde{A})$ in any way (with respect to the Lebesgue measure) will produce a solution of Problem 9. Thus, by Lemma 14 and the definition of a probability density function, we immediately obtain the following result:

Theorem 15: Let $\{V_1, \ldots, V_\kappa\}$ be a basis for $\Pi(\tilde{A})$. Let $(\alpha_1, \ldots, \alpha_\kappa) \in \mathbb{R}^p$ be a point selected according to any probability density function on $\mathbb{R}^p$, and define

$$V := \sum_{i=1}^{\kappa} \alpha_i V_i\,.$$

V is a solution to Problem 9 (a.s.).

The computational complexity of the procedure in Theorem 15 is significantly lower than the complexity of Algorithm 12 from Section III. The only computationally heavy step is choosing a basis for $\Pi(\tilde{A})$. However, as in Remark 11, $\Pi(\tilde{A})$ is obtained as a kernel of matrix D, where $D \in \mathbb{R}^{\left(n^2 \dim\left(\mathcal{I}_n(\tilde{A})\right)\right) \times n^2}$. By Remark 6,

$$\dim\left(\mathcal{I}_n(\tilde{A})\right) \leq k(n/r)^2 \leq n^2\,.$$

Hence, D has at most n^4 rows, and $\mathrm{Ker}(D)$ can be found in polynomial time using, for example, QR or SVD decompositions [11]. The price paid for this computationally feasible solution is that V is no longer a definite solution, but only a solution almost surely (in the probabilistic sense). Nevertheless, it will be shown in the examples in Section V that this approach works in practice.

V. Examples

In this section, we discuss the patterns that emerge in common distributed systems, and show how solutions to Problem 9 can provide algebraic relationships to characterize those patterns. In each case, we focus on the matrix structure in the linearized system dynamics $\dot{x} = Ax$. This structure carries over to system inputs, measurements, etc. — all aspects used in linear control.

We begin by showing that patterns and commuting relationships are equivalent for two canonical distributed systems, the ring and the unidirectional chain.

Example 16: Consider a ring system, as shown in Figure 1 for three agents, each of which has ℓ states.

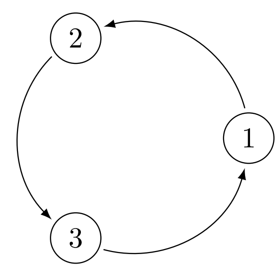

Fig. 1. Ring System

It is well known [12] that the ring structure manifests as block circulant matrices, with the pattern

$$\tilde{A}_{\text{ring}} = \begin{bmatrix} 1 & 2 & 3 \\ 3 & 1 & 2 \\ 2 & 3 & 1 \end{bmatrix} \quad \Rightarrow \; A = \begin{bmatrix} A_1 & A_2 & A_3 \\ A_3 & A_1 & A_2 \\ A_2 & A_3 & A_1 \end{bmatrix},\; A_i \in \mathbb{R}^{\ell \times \ell}$$

so $A \in \mathcal{I}_\ell(\tilde{A}_{\text{ring}})$, as in Example 4. Now, take $\ell = 2$. Running $\tilde{A}_{\text{ring}}$ through Algorithm 12 gives solutions to the equation $VA = AV$ of the form

$$V = \left[\begin{array}{cc|cc|cc} v_1 & 0 & v_2 & 0 & v_3 & 0 \\ 0 & v_1 & 0 & v_2 & 0 & v_3 \\ \hline v_3 & 0 & v_1 & 0 & v_2 & 0 \\ 0 & v_3 & 0 & v_1 & 0 & v_2 \\ \hline v_2 & 0 & v_3 & 0 & v_1 & 0 \\ 0 & v_2 & 0 & v_3 & 0 & v_1 \end{array}\right].$$

A particular basis of solutions $\{V_1, V_2, V_3\}$ can be found by taking $v_i = 1$ and all other $v_j = 0$ for each $i = 1, 2, 3$. By Theorem 15, a random linear combination of these solutions will almost surely solve the Pattern Identification Problem for the ring system. Here, we focus on the particular solution V_2, which is the block fundamental permutation matrix. It is well known that a matrix is block circulant if and only if it commutes with the block fundamental permutation matrix [9]. Thus, imposing that system matrices commute with V_2 ensures that they preserve the ring structure, and so our algorithm correctly produces a way to encode the pattern in a ring system. ◁

Example 17: Consider a chain system, as shown in Figure 2 for three agents, each of which has ℓ states.

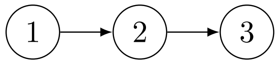

Fig. 2. Unidirectional Chain System

The structure of this chain, in which each agent can only get information from those to its left, manifests as *lower triangular Toeplitz* matrices, with the pattern

$$\tilde{A}_{\text{chain}} = \begin{bmatrix} 1 & 0 & 0 \\ 2 & 1 & 0 \\ 3 & 2 & 1 \end{bmatrix} \quad \Rightarrow \quad A = \begin{bmatrix} A_1 & 0 & 0 \\ A_2 & A_1 & 0 \\ A_3 & A_2 & A_1 \end{bmatrix}, \; A_i \in \mathbb{R}^{\ell \times \ell}$$

so $A \in \mathcal{I}_\ell(\tilde{A}_{\text{chain}})$. Now, take $\ell = 2$. Running $\tilde{A}_{\text{chain}}$ through Algorithm 12 gives solutions to the equation $VA = AV$ of the form

$$V = \left[\begin{array}{cc|cc|cc} v_1 & 0 & 0 & 0 & 0 & 0 \\ 0 & v_1 & 0 & 0 & 0 & 0 \\ \hline v_2 & 0 & v_1 & 0 & 0 & 0 \\ 0 & v_2 & 0 & v_1 & 0 & 0 \\ \hline v_3 & 0 & v_2 & 0 & v_1 & 0 \\ 0 & v_3 & 0 & v_2 & 0 & v_1 \end{array}\right].$$

As before, a basis of solutions $\{V_1, V_2, V_3\}$ can be found by taking $v_i = 1$ and all other $v_j = 0$ for each $i = 1, 2, 3$. By Theorem 15, almost all linear combinations of these solutions solve the Pattern Identification Problem for the chain system. Here, we focus on the particular solution V_2 — the block fundamental nilpotent matrix — which is well known to commute with all block lower triangular Toeplitz matrices [1]. It can also easily be shown that a matrix is block lower triangular Toeplitz if *and only if* it commutes with the block fundamental nilpotent matrix. Thus, imposing that system matrices commute with V_2 ensures that they preserve the chain structure, and so our algorithm again correctly produces a way to encode the system's pattern. ◁

Next, we find commuting relationships for some other common patterns. In these cases, we find that the specific form of the pattern matters for recovering a commuting relationship.

Example 18: Consider a bidirectional chain system, as shown in Figure 1 for three agents, each of which has ℓ states. Unlike the unidirectional chain of Example 17, every agent can send information to and receive information from every other agent; unlike the ring of Example 16, all information shared must pass through agent 2, rather than moving around the cycle.

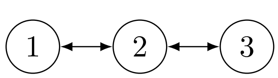

Fig. 3. Bidirectional Chain System

There are a few possible ways to characterize the bidirectional chain. Here, we look at the pattern

$$\tilde{A}_{\text{chain2}} = \begin{bmatrix} 1 & 2 & 3 \\ 2 & 1 & 2 \\ 3 & 2 & 1 \end{bmatrix} \quad \Rightarrow \quad A = \begin{bmatrix} A_1 & A_2 & A_3 \\ A_2 & A_1 & A_2 \\ A_3 & A_2 & A_1 \end{bmatrix}, \; A_i \in \mathbb{R}^{\ell \times \ell} \tag{7}$$

so $A \in \mathcal{I}_\ell(\tilde{A}_{\text{chain2}})$, signifying that each agent's internal dynamics are the same, all first-level interconnections (1→2, 2→1, 2→3, 3→2) operate in the same way, and all second-level interconnections (1→3, 3→1) also operate in the same way. Now, take $\ell = 1$. Running $\tilde{A}_{\text{chain2}}$ through Algorithm 12 gives solutions to the equation $VA = AV$ of the form

$$V = \begin{bmatrix} v_1 & 0 & v_2 \\ 0 & v_1 + v_2 & 0 \\ v_2 & 0 & v_1 \end{bmatrix}.$$

Again, a basis of solutions $\{V_1, V_2\}$ can be found by taking $v_i = 1$ and all other $v_j = 0$ for each $i = 1, 2$, and by Theorem 15, almost all linear combinations of these solutions solve the Pattern Identification Problem for the bidirectional chain system. In this case, though, the recovered commutating relationship has a more general form than that specified by $\tilde{A}_{\text{chain2}}$: for any choice of v_1 and $v_2 \neq 0$, V commutes with any matrix of the form

$$A_{\text{recovered}} = \begin{bmatrix} a_1 & a_2 & a_3 \\ \color{red}{a_4} & \color{red}{a_5} & \color{red}{a_4} \\ a_3 & a_2 & a_1 \end{bmatrix}$$

where the differences from (7) are in red, and it can clearly be seen that our algorithm recovers a denser pattern — in other words, any found V *must* commute with more matrices than originally asked. In terms of the underlying system, the pattern in $A_{\text{recovered}}$ means that the middle agent of the chain in Figure 3 can have different dynamics than the outer two agents, and its outgoing interconnections can be different than those incoming. Thus, unlike in the previous two examples, a single commuting relationship cannot uniquely recover the pattern in (7), but it can still recover a valid pattern for a bidirectional chain. ◁

Example 19: Consider the three-agent tree system on the left side of Figure 4, where each agent has ℓ states.

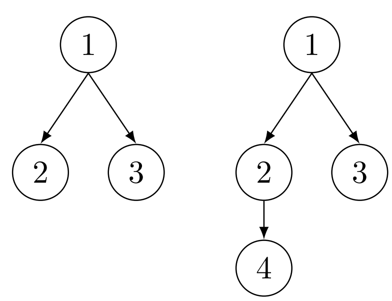

Fig. 4. Two Tree Systems

A pattern for this tree system is given by

$$\tilde{A}_{\text{tree}} = \begin{bmatrix} 1 & 0 & 0 \\ 2 & 1 & 0 \\ 2 & 0 & 1 \end{bmatrix} \quad \Rightarrow \quad A = \begin{bmatrix} A_1 & 0 & 0 \\ A_2 & A_1 & 0 \\ A_2 & 0 & A_1 \end{bmatrix}, \; A_i \in \mathbb{R}^{\ell \times \ell} \tag{8}$$

so $A \in \mathcal{I}_\ell(\tilde{A}_{\text{tree}})$, and the interconnections $1 \to 2$ and $1 \to 3$ are identical. Now, take $\ell = 1$. Running $\tilde{A}_{\text{tree}}$ through Algorithm 12 gives solutions to the equation $VA = AV$ of the form

$$V = \begin{bmatrix} v_1 & 0 & 0 \\ v_2 & v_4 & v_1 - v_4 \\ v_3 & v_1 - v_5 & v_5 \end{bmatrix}.$$

A particular basis of solutions $\{V_1, \ldots, V_5\}$ can be found by taking $v_i = 1$ and all other $v_j = 0$ for each $i = 1, \ldots, 5$, and by Theorem 15, a random linear combination of these solutions will almost surely solve the Pattern Identification Problem for the tree system. However, once again, any choice of $v_1, \ldots, v_5$ allows V to commute with matrices instantiated by a denser pattern than that of $\tilde{A}_{\text{chain}}$. In particular, taking $v_1 = v_2 = v_3 = 1$ and $v_4 = v_5 = 0$ allows V to commute with any matrix of the form

$$A_{\text{recovered}} = \begin{bmatrix} a_3 & 0 & 0 \\ a_2 & a_1 & a_3 - a_1 \\ a_2 & a_3 - a_1 & a_1 \end{bmatrix}$$

where the differences from (8) are in red. (Different choices for the v_i will, in this case, give different forms for the matrix $A_{\text{recovered}}$.) In terms of the underlying system, the pattern in $A_{\text{recovered}}$ means that the root agent (1) can have different internal dynamics than those of the second-level agents (2 and 3), and the difference between these dynamics appears in the coupling terms between the second-level agents. It is important to note that a system in the form (8) still fits the form of $A_{\text{recovered}}$; however, imposing only commutation with the chosen V does not *uniquely* translate back to the pattern $\tilde{A}_{\text{tree}}$. Again, the commuting relationship requires a denser pattern than that initially asked, while still fitting the overall tree structure. ◁

Example 20: In Example 19, we saw that while the pattern in a simple tree system is not recovered, a similar but denser pattern can be found. Now, consider the slightly more complicated four-agent tree system on the right side of Figure 4. Notably, in addition to having an extra agent, this system is not symmetric about its root agent. A pattern for this tree system is given by

$$\tilde{A}_{\text{tree2}} = \begin{bmatrix} 1 & 0 & 0 & 0 \\ 2 & 1 & 0 & 0 \\ 2 & 0 & 1 & 0 \\ 3 & 2 & 0 & 1 \end{bmatrix} \quad \Rightarrow \quad A = \begin{bmatrix} A_1 & 0 & 0 & 0 \\ A_2 & A_1 & 0 & 0 \\ A_2 & 0 & A_1 & 0 \\ A_3 & A_2 & 0 & A_1 \end{bmatrix} \tag{9}$$

again with $A_i \in \mathbb{R}^{\ell \times \ell}$, so $A \in \mathcal{I}_\ell(\tilde{A}_{\text{tree2}})$ and all first-level interconnections are identical. Now, take $\ell = 1$. Running $\tilde{A}_{\text{tree2}}$ through Algorithm 12 gives solutions to the equation $VA = AV$ of the form

$$V = \begin{bmatrix} v_1 & 0 & 0 & 0 \\ v_2 & v_1 & 0 & 0 \\ v_3 & v_1 - v_5 & v_5 & 0 \\ v_4 & v_6 & v_2 - v_6 & v_1 \end{bmatrix}.$$

As in all previous examples, a random linear combination of the solutions $V_1, \ldots, V_6$ (with, in turn, $v_i = 1$ and all other $v_j = 0$) will almost surely solve the Pattern Identification Problem for this tree system by Theorem 15. However, as in the simpler tree in Example 19, any choice of $v_1, \ldots, v_6$ allows V to commute with matrices instantiated by a denser pattern than that of $\tilde{A}_{\text{chain2}}$. In particular, taking $v_2 = v_3 = v_5 = v_6 = 1$ and $v_1 = v_4 = 0$ gives commutation with any matrix of the form

$$A_{\text{recovered}} = \begin{bmatrix} a_1 & 0 & 0 & 0 \\ a_2 & a_1 & 0 & 0 \\ a_2 & a_1 - a_4 & a_4 & 0 \\ a_3 & a_2 & 0 & a_1 \end{bmatrix}$$

where the differences from (9) are in red. (Again, different choices for the v_i will give different forms for $A_{\text{recovered}}$.) The asymmetry in the underlying tree system is reflected in the found $A_{\text{recovered}}$: in this case, the dynamics and interconnections in the $1 \to 2 \to 4$ chain must be identical (as in the unidirectional chain in Example 17), whereas the dynamics in the $1 \to 3$ chain can be different. As with the simpler tree, a system in the form (9) still fits the form of $A_{\text{recovered}}$, and the found pattern (while different from the initial pattern) still translates to the desired tree structure. ◁

VI. Concluding Remarks

This paper addresses an open problem of pattern identification in distributed systems. The pattern is encoded algebraically via commuting relationships with a base matrix. Heretofore, there has been no algorithmic method to determine these commuting relationships. We propose a solution that recovers the best choice of base matrix, in the sense that the found base matrix commutes with all system matrices that have the desired pattern, and with as few other patterns as possible. Through this procedure, we provide a solid mathematical foundation for patterned control design.

References

[1] S. C. Hamilton and M. E. Broucke, *Geometric Control of Patterned Linear Systems*. Lecture Notes in Control and Information Sciences: Springer, 2012.

[2] A. C. Sniderman, M. E. Broucke, and G. M. T. D'Eleuterio, "Controllability Is Not Sufficient for Pole Placement in Patterned Systems," in *54th IEEE Conference on Decision and Control*. Osaka: IEEE, 2015.

[3] ——, "Block Circulant Control: A Geometric Approach," in *52nd IEEE Conference on Decision and Control*. Florence: IEEE, Dec. 2013, pp. 2097–2102.

[4] P. Massioni and M. Verhaegen, "Distributed Control for Identical Dynamically Coupled Systems: A Decomposition Approach," *IEEE Transactions on Automatic Control*, vol. 54, no. 1, pp. 124–135, Jan. 2009.

[5] M. Hazewinkel and C. F. Martin, "Symmetric Linear Systems: An Application of Algebraic Systems Theory," *International Journal of Control*, vol. 37, no. 6, pp. 1371–1384, 1983.

[6] M. Hovd and S. Skogestad, "Control of Symmetrically Interconnected Plants," *Automatica*, vol. 30, no. 6, pp. 957–973, Jun. 1994.

[7] B. Bamieh, F. Paganini, and M. A. Dahleh, "Distributed Control of Spatially Invariant Systems," *IEEE Transactions on Automatic Control*, vol. 47, no. 7, pp. 1091–1107, Jul. 2002.

[8] C. Holmes and M. E. Broucke, "Pattern preserving pole placement and stabilization for linear systems," in *2016 American Control Conference*, 2016, accepted.

[9] P. J. Davis, *Circulant Matrices*. New York, NY: Wiley, 1979.

[10] A. J. Weir, *Lebesgue Integration and Measure*. New York, NY: Cambridge University Press, 1973.

[11] G. H. Golub and C. F. van Loan, *Matrix Computations*. JHU Press, 2012, vol. 3.

[12] R. W. Brockett and J. L. Willems, "Discretized Partial Differential Equations: Examples of Control Systems Defined on Modules," *Automatica*, vol. 10, no. 5, pp. 507–515, Sep. 1974.