Systems & Control Letters 136 (2020) 104575

Contents lists available at ScienceDirect

Systems & Control Letters

journal homepage: www.elsevier.com/locate/sysconle

Controllability-Gramian submatrices for a network consensus model

Sandip Roy [*,1], Mengran Xue

School of Electrical Engineering and Computer Science at Washington State University, United States of America

ARTICLE INFO

Article history:
Received 21 March 2019
Received in revised form 25 September 2019
Accepted 4 November 2019
Available online 7 January 2020

Keywords:
Network consensus
Controllability and target controllability
Security
Gramian

ABSTRACT

Principal submatrices of the controllability Gramian and their inverses are examined, for a network-consensus model with inputs at a subset of network nodes. Several properties of the Gramian submatrices and their inverses – including dominant eigenvalues and eigenvectors, diagonal entries, and sign patterns – are characterized by exploiting the special doubly-nonnegative structure of the matrices. In addition, majorizations for these properties are obtained in terms of cutsets in the network's graph, based on the diffusive form of the model. The asymptotic (long time horizon) structure of the controllability Gramian is also analyzed. The results on the Gramian are used to study metrics for target control of the network-consensus model.

1. Introduction

Dynamical models for consensus or synchronization in networks have been exhaustively studied [1–3]. One focus of this effort has been on open-loop control of the dynamics using inputs at a subset of the network's nodes [4–9]. In particular, graph-theoretic necessary or sufficient conditions for controllability have been obtained, and some characterizations of the required control energy have also been obtained using analyses of the controllability Gramian. Recently, researchers have begun to study *target control* of network models, wherein inputs are designed to manipulate a group of target nodes rather than the whole network [10–14]. The target-control problem is of practical interest in several application domains, in which stakeholders need to use limited actuation capabilities to guide a few key nodes' states. In parallel with the general controllability analysis for networks, the effort on target controllability has also yielded graph-theoretic conditions and analyses of metrics. These studies demonstrate that limited-energy target control of network processes is sometimes possible even when full-state control is prohibitively costly or impossible.

Target control for network models can be analyzed in terms of principal submatrices of the controllability Gramian [10,11]. Precisely, target controllability resolves to invertibility of a principal submatrix of the Gramian, while the minimum energy required to achieve a desired target state and/or guide certain state projections can be found in terms of quadratic forms of the Gramian-submatrix inverses. Thus, the study of target controllability motivates analysis of the Gramian matrix and its principal submatrices for network consensus/synchronization models.

Because of the relevance of Gramian matrices to network controllability as well as dual observability/estimation problems, some structural and graph-theoretic results on the full Gramian have been developed for canonical network-consensus models, as well as for other dynamical network processes [9,11,15–17]. In addition, explicit formulae for the Gramian inverse in terms of the network model's spectrum have been developed, using Cauchy matrix properties [9]. These explicit computations give insight into the relationship between the network model's spectrum and the required control energy.

The purpose of this study is to develop new characterizations of the Gramian and its principal submatrices in the context of a canonical discrete-time network consensus model, with the goal of assessing target control metrics. Relative to the earlier studies on the Gramian of network models, the main contributions of this work are to: (1) characterize principal submatrices of the Gramian and their inverse, in addition to the full Gramian matrix; (2) give new insights into the sign patterns and eigenvalues of Gramian- and Gramian-submatrix inverses; (3) develop graph-theoretic majorizations on the Gramian's entries; and (4) assess target control metrics and optimal inputs using the results on Gramian submatrices. The analyses primarily draw on the diffusive structure of the network model, which imposes a spatial pattern on the input response of the network. A main result is that Gramian submatrix inverses exhibit a special sign pattern as well as a dependence on cutsets of the network graph, which allows majorization of the control energy and analysis of minimum-energy inputs.

Although our focus here is on target control for a consensus model, the analyses of Gramian submatrices are germane to

* Corresponding author.
E-mail addresses: sandip@wsu.edu (S. Roy), morashu@wsu.edu (M. Xue).

1 The first author acknowledges support from United States National Science Foundation Grants 1545104 and 1635184.

https://doi.org/10.1016/j.sysconle.2019.104575

the control and estimation of various network dynamical processes with a diffusive or nonnegative structure. In particular, the analyses are relevant to the controllability analysis of other discrete-time models with nonnegative state matrices and continuous-time models with Metzler state matrices (e.g., models for infection spread, economic systems, etc.) [18,19]. The results also inform other problems in network estimation and control which require consideration of Gramians and Markov parameters, including observability analysis and model reduction [20,21].

The article is organized as follows. The network consensus model is presented in Section 2, and analysis of target control metrics in terms of the Gramian is reviewed in Section 3. The main results on the Gramian, and their implications on target control, are developed in Section 4. Finally, the graph-theoretic characterization of target control is illustrated in an example in Section 5.

2. Model

A network with n nodes, labeled $1, \dots, n$, is considered. Each node i has a scalar state $x_i[k]$ which evolves in discrete time ($k \in Z^+$). A set $\mathcal{S}$ containing m nodes, which we call the *source nodes*, are amenable to external actuation. The nodes' states evolve according to:

$$\mathbf{x}[k+1] = A\mathbf{x}[k] + B\mathbf{u}[k], \tag{1}$$

where the state vector is $\mathbf{x}[k] = \begin{bmatrix} x_1[k] \\ \vdots \\ x_n[k] \end{bmatrix}$, $A = [a_{ij}]$ is a row-stochastic matrix ($a_{ij} \geq 0$, $\sum_{j=1}^{n} a_{ij} = 1$ for each i), $\mathbf{u}[k] = \begin{bmatrix} u_1[k] \\ \vdots \\ u_m[k] \end{bmatrix}$ specifies the input (actuation) signals at the m source nodes, and B is an $n \times m$ matrix whose columns are 0–1 indicators of the source nodes in $\mathcal{S}$.

The manipulation of the states of p *target nodes* specified in a set $\mathcal{T}$ is of interest. The *target state* vector $\mathbf{y}[k]$, defined as containing the states of the p target nodes at time k, can be expressed as:

$$\mathbf{y}[k] = C\mathbf{x}[k], \tag{2}$$

where C is a $p \times m$ matrix whose rows are 0–1 indicators of the target nodes in $\mathcal{T}$. We refer to the model as a whole as the *input–output network consensus model* or simply the *network model*.

A weighted digraph Γ is defined to represent the topology of nodal interactions in the network model. Specifically, Γ is defined to have n nodes labeled $1, \dots, n$, which correspond to the n vertices. A directed edge is drawn from vertex i to vertex j if and only if $a_{ji} > 0$, and the weight of the edge is set to a_{ji}. We note that the graph may include self loops (i.e., edges from vertices to themselves). The sum of the weights of the incoming edges to each vertex is 1. An edge in Γ from node i to node j indicates that node j's state at time $k+1$ is directly influenced by the node i's state at time k. The vertices corresponding to the source and target nodes are referred to as source and target vertices, respectively.

Throughout the article, we assume that the graph Γ is ergodic, which means that it is strongly connected (i.e. has a directed path between each ordered pair of vertices) and also aperiodic (i.e. the greatest common factor among path lengths between any two vertices is 1). Equivalently, this assumption can be expressed as the matrix A being irreducible (having no permutation which yields an upper triangular matrix) and aperiodic (in the sense that there is a power of the matrix that is strictly positive), please see [22] for further details on these definitions. Under this assumption, the unactuated model reaches consensus, i.e. the manifold where all nodes' states are identical is globally asymptotically stable.

3. Preliminaries: Target control and the Gramian

Target control of the network model is primarily concerned with two questions [10,11]: (1) deciding whether the input $\mathbf{u}[k]$ can be designed to guide the target state vector $\mathbf{y}[k]$ to a desired goal (i.e. analyzing *target controllability*); and (2) determining how much actuation energy or effort is needed to do so (i.e. assessing *target control metrics*). Our primary focus here is on assessing target control metrics.

As a starting point, we consider the following definition of target controllability:

Definition 1. The input–output network-consensus model is said to be target controllable over $[0, k_f]$ if, for any goal $\bar{\mathbf{y}} \in \mathcal{R}^p$, an input signal $\mathbf{u}[0], \dots, \mathbf{u}[k_f - 1]$ can be designed to drive the network model from a relaxed initial state $\mathbf{x}[0] = 0$ to the target state $\mathbf{y}[k_f] = \bar{\mathbf{y}}$.

In the case that the network model is target controllable over an interval $[0, k_f]$, a number of metrics for the energy required to guide the target state may be interest. Broadly, these various metrics give an indication of the ease or difficulty with which the external actuation can be used to guide the target, and hence capture the manipulability of the dynamics through actuation. These manipulability metrics may be interpreted as measures of network *flexibility*, in the sense that they capture the ease with which an operator can guide target states to desirable values. Conversely, however, the metrics may be interpreted as measures of *security*, in situations where they capture the ease/difficulty with which an adversary can hijack the target state via actuation. Our work in this area has largely been motivated by threat assessment problems for infrastructural and social networks, hence our definitions below reflect a security perspective. At the end of this section, we present a couple of conceptual examples that further illustrate the relevance of the definitions to security and threat assessment problems.

The first metric of interest is minimum input energy required to achieve a particular goal state. This notion is formalized in the following definition:

Definition 2. The *target-control energy* for an interval $[0, k_f]$ and goal state $\bar{\mathbf{y}}$ is defined as $E(k_f, \bar{\mathbf{y}}) = \min_{\mathbf{u}[0],\dots,\mathbf{u}[k_f-1]} \sum_{i=0}^{k_f-1} \mathbf{u}^T[i]\mathbf{u}[i]$, subject to the constraint that the input sequence drives the system from a relaxed state to $\mathbf{y}[k_f] = \bar{\mathbf{y}}$. We refer to an argument (input sequence) that achieves the minimum as an *optimal target-control input*.

Additionally, the energy required to drive a scalar projection of the target state to a unit value may be important to characterize, as an indication of the manipulability of key output statistics. This notion is formalized as follows:

Definition 3. The *projection-manipulation energy* for an interval $[0, k_f]$ and projection vector $\boldsymbol{\alpha}$ is defined as $F(k_f, \overline{\boldsymbol{\alpha}}) = \min_{\mathbf{u}[0],\dots,\mathbf{u}[k_f-1]} \sum_{i=0}^{k_f-1} \mathbf{u}^T[i]\mathbf{u}[i]$, subject to the constraint that the input sequence drives the system from a relaxed state to $\boldsymbol{\alpha}^T\mathbf{y}[k_f] = 1$. We refer to an input sequence that achieves the minimum as an *optimal projection-manipulation input*.

The projection-manipulation energy captures the effort required to alter output statistics, such as the average of the target nodes' states, or the difference between two target nodes' states.

Often, it is of interest to characterize extremal values of the target-control energy across goal states with a particular norm, or of the projection-manipulation energy across projection vectors of a certain norm. In particular, the minimum value of the target-control energy across unit-norm goal states indicates the lowest effort using which the target state can be driven away from its equilibrium by a unit amount; thus, it is a holistic measure of the difficulty that an adversary would have in manipulating the target state, i.e. a security measure. This notion is formalized in the following definition:

Definition 4. The minimum of the target-control energy over unit-two-norm goal states, i.e. $E_{min}(k_f) = \min_{\overline{\mathbf{y}}\, s.t.\, \|\overline{\mathbf{y}}\|_2=1} E(k_f, \overline{\mathbf{y}})$, is referred to as the *target security* of the network model. A goal state $\overline{\mathbf{y}}$ that achieves the minimum is denoted as $\overline{\mathbf{y}}_{min}$, and is termed a *minimally-secure goal.*

Meanwhile, the minimum of the projection-manipulation energy across unit-norm projection vectors is an indication of the easiest-to-manipulate statistic of the target state. Again, the minimum indicates in a holistic sense the difficulty that an adversary would have in manipulating target statistics or projections, and hence also can be considered as a security measure. This notion is formalized next:

Two metrics across goal states and projection vectors, respectively, are indications of the security of the network model to manipulation. These global security notions are formalized in the following definitions:

Definition 5. The minimum of the projection-manipulation energy over projection vectors with unit one-norm, i.e. $F_{min}(k_f) = \min_{\overline{\boldsymbol{\alpha}}\, s.t.\, \|\overline{\boldsymbol{\alpha}}\|_1=1} F(k_f, \overline{\boldsymbol{\alpha}})$, is referred to as the *projection security* of the network model. A projection vector $\overline{\boldsymbol{\alpha}}$ that achieves the minimum is denoted as $\overline{\boldsymbol{\alpha}}_{min}$, and is termed a *minimally-secure projection.*

Remark. Other norms may be used in the security definitions. We have assumed a 1-norm constraint on the projection vector in the projection-security definition, because the weighted sum of nodal quantities is often of interest for diffusive processes. Meanwhile, we have assumed a two-norm constrain on the goal state in the target security definition, in keeping with standard assessments of controllability/security of linear models.

Target controllability and the target-control metrics can readily be characterized in terms of principal submatrices of the controllability Gramian of the network model. The controllability Gramian for the network model over the interval $[0, k_f]$ is given by:

$$W(k_f) = \sum_{i=0}^{k_f-1} (A^i B)(A^i B)^T. \tag{3}$$

We define principal submatrices of the controllability Gramian using a set $\mathcal{B}$ which lists a subset of the nodes $1, \ldots, n$ in the network. The $\mathcal{B}$-controllability Gramian $W(\mathcal{B}, k_f)$ is defined as the principal submatrix of $W(k_f)$ in which the rows and columns indicated in $\mathcal{B}$ are maintained.

The following lemma provides characterizations of target controllability and the target control metrics in terms of the $\mathcal{T}$-controllability Gramian (i.e. the principal submatrix of controllability Gramian associated with the target nodes $\mathcal{T}$). These results follow directly from standard analyses of output controllability [23,24], hence the proof is omitted.

Lemma 1. *The input–output network-consensus model is target controllable if and only if the $\mathcal{T}$-controllability Gramian $W(\mathcal{T}, k_f)$ is invertible.*

If the network model is target controllable, then the target-control energy is given by:

$$E(k_f, \overline{\mathbf{y}}) = \overline{\mathbf{y}}^T W(\mathcal{T}, k_f)^{-1} \overline{\mathbf{y}}, \tag{4}$$

and an optimal target-control input is

$$\widehat{\mathbf{u}}[i] = (CA^{k_f-i-1}B)^T W(\mathcal{T}, k_f)^{-1} \overline{\mathbf{y}} \tag{5}$$

for $i = 0, \ldots, k_f - 1$. Further, the target security is given by:

$$E_{min}(k_f) = \frac{1}{\lambda_{max}(W(\mathcal{T}, k_f))}, \tag{6}$$

where $\lambda_{max}(W(\mathcal{T}, k_f))$ refers to the largest eigenvalue of matrix $W(\mathcal{T}, k_f)$. The minimally secure goal is given by

$$\overline{\mathbf{y}}_{min} = \mathbf{v}_{max}(W(\mathcal{T}, k_f)), \tag{7}$$

where $\mathbf{v}_{max}(W(\mathcal{T}, k_f))$ is the right eigenvector of $W(\mathcal{T}, k_f)$ associated with the largest eigenvalue.

The projection manipulation energy is given by:

$$F_{min}(k_f, \boldsymbol{\alpha}) = \frac{1}{\boldsymbol{\alpha}^T W(\mathcal{T}, k_f) \boldsymbol{\alpha}}, \tag{8}$$

The projection security is given by

$$F_{min}(k_f) = \frac{1}{\max_i [W(\mathcal{T}, k_f)]_{i,i}}, \tag{9}$$

The minimally-secure projection is given by $\overline{\boldsymbol{\alpha}}_{min} = \mathbf{e}_j$, where $j = \arg\max_i [W(\mathcal{T}, k_f)]_{i,i}$.

Lemma 1 is the starting point for the main graph-theoretic and structural analyses developed in the paper.

Conceptual examples

We conclude the formulation and preliminary analysis of the target-control metrics with two conceptual examples, which illustrate the relevance of the definitions to threat-assessment and security problems.

Example 1: Industrial-process security. Industrial systems often involve diffusive or flow processes, which take the form of the consensus dynamics. For instance, industrial systems may include mass-balance processes among reservoirs, which can be represented using a network synchronization or consensus model [25]. Likewise, chemical processes in industrial systems admit linearizations which are consensus processes, to within a scaling of the state variables (i.e. they are scaled consensus processes). Controllers for these processes increasingly use networked cyber-systems, leading to a growing concern that cyber attackers can manipulate the physical states of the industrial systems; indeed, several recent events have highlighted such vulnerabilities. A cyber-attacker which has the ability to manipulate independent components in these industrial processes can be modeled as setting actuations at one or a small set of network components. The attacker may be focused on modifying certain critical target states (e.g. certain critical reservoir levels or chemical concentrations), with the intent of causing systemic failures or damage to hardware. Furthermore, the ease with which the attacker can manipulate the target can be abstractly measured by the required size (energy) of the input. Thus, the analysis of cyber attacks can be abstracted to a target control metric analysis. We note that several of the defined metrics may be of interest, including the target-control energy for certain sensitive goal states, the projection-manipulation energy for certain key statistics, and the holistic security metrics. The highly abstracted formulation is appealing because it can potentially give simple graph-theoretic insights into the manipulability of industrial processes, in addition to allowing numerical computation using the Gramian-based analysis presented above.

Example 2: Manipulation of information-diffusion processes. Consensus models have been abstractly used to capture information diffusion and decision-making in social networks [26]. While these models are a significant abstraction of reality, they are representative of some information-flow processes that occur in social networks. They also provide a framework for characterizing manipulation of decision-making processes by external adversaries, using a linear network consensus model with imposed actuation [27]. The energy metrics defined above abstractly capture the effort needed by a stakeholder to manipulate remote target states or state projections (e.g., the average opinion of a group of target individuals). Thus, the Gramian-based analyses of the consensus model are applicable in understanding manipulation of information-diffusion processes from a graph-theoretic perspective.

4. Main results

Per Lemma 1, target controllability and the target control metrics are tied to properties of the controllability Gramian of the network model. Our focus here is to develop structural and graph-theoretic results on the Gramian, with the aim of giving insights into target control. Because a number of graph-theoretic results have already been developed for the binary question of target controllability [10,11,28], we will primarily focus on the target-control metrics.

First, we identify some matrix-theoretic properties of Gramian submatrices and their inverses for the network model. These properties depend on the diffusive structure of the network-consensus model, but not on the specifics of the network's topology.

Theorem 1. *Consider any principal submatrix of the Gramian for the network model, say $Q = W(\mathcal{B}, k_f)$. Also, for invertible Q, consider $R = Q^{-1}$. The matrices Q and R have the following properties:*

(1) Q is doubly nonnegative, i.e. it is symmetric, positive semidefinite, and entry-wise nonnegative. For sufficiently large k_f, Q is entry-wise strictly positive
(2) The eigenvalues of Q are real, nonnegative, and non-defective. For sufficiently large k_f, the largest eigenvalue $\lambda_{max}(Q)$ has algebraic multiplicity of 1, and its associated eigenvector $\mathbf{v}_{max}(Q)$ is strictly positive (to within a scale factor).
(3) $\lambda_{max}(Q) \leq \lambda_{max}(W[k_f])$. Furthermore, for sufficiently large k_f, the inequality is strict.
(4) The matrix R is symmetric and positive definite. For sufficiently large k_f, R is irreducible.
(5) Consider any permutation $T = PRP^{-1}$ of the matrix R, and consider any block-partition of T as $T = \begin{bmatrix} T_{11} & T_{12} \\ T_{12}^T & T_{22} \end{bmatrix}$ where T_{11} is square. Assuming that k_f is sufficiently large, the matrix T_{12} has at least one negative entry.
(6) For the special case that the cardinality of the set $\mathcal{B}$ (denoted $|\mathcal{B}|$) is 2, the matrix R is a nonsingular M matrix.

Proof. *Proof of Item 1:* The Gramian is symmetric and positive semidefinite, hence it is immediate that the principal submatrix Q is also symmetric and positive semidefinite. We characterize the signs of the entries in $Q = [q_{ij}]$ as follows. First, these entries are expressed in terms of the impulse responses of the network model at network nodes due to inputs at each source node. To simplify indexing in this analysis, we assume without loss of generality that the target nodes are the nodes $1, \ldots, p$, and hence the Gramian submatrix of interest is a leading principal submatrix. In this case, the entry q_{ij} can be written as $q_{ij} = \sum_{z \in \mathcal{S}} h_{zi,k_f}^T h_{zj,k_f}$, where $h_{zl,k_f} = \begin{bmatrix} \mathbf{e}_l^T \mathbf{e}_z \\ \mathbf{e}_l^T A \mathbf{e}_z \\ \vdots \\ \mathbf{e}_l^T A^{k_f - 1} \mathbf{e}_z \end{bmatrix}$, where the notation $\mathbf{e}_w$ is used for a 0–1 indicator vector with wth entry equal to 1. We notice that h_{zl,k_f} encodes the impulse response at node l due to an input at node z. Since the matrix A is nonnegative, it is immediate that h_{zl,k_f} is nonnegative, and hence $q_{ij} \geq 0$. From the fact that A is irreducible and aperiodic, it follows that the final entry in the vector h_{zl,k_f} is strictly positive for all z and l, for all sufficiently large k_f (see [22]). Thus, q_{ij} is necessarily strictly positive for sufficiently large k_f.

Proof of Item 2: Since Q is symmetric and positive semidefinite, it is immediate that its eigenvalues are real, nonnegative, and nondefective (i.e. each eigenvalue's algebraic and geometric multiplicities are identical). For sufficiently large k_f, we have shown above that Q is strictly positive. It thus follows from the Frobenius–Perron theory that Q has a dominant eigenvalue (an eigenvalue with magnitude larger than any other eigenvalue) with algebraic multiplicity 1, whose eigenvector is strictly positive (upon appropriate scaling) [22].

Proof of Item 3: The inequality $\lambda_{max}(Q) \leq \lambda_{max}(W[k_f])$ is an immediate consequence of the fact that Q is a principal submatrix of the positive semidefinite matrix $W[k_f]$. The strictness of the inequality for sufficiently large k_f can be proved by contradiction. If $\lambda_{max}(Q)$ was equal to $\lambda_{max}(W[k_f])$, then from the Courant–Fisher theorem [29], $\max_{\mathbf{v}} \frac{\mathbf{v}^T W[k_f] \mathbf{v}}{\mathbf{v}^T \mathbf{v}}$ would equal $\lambda_{max}(Q)$, and further the argument maximizing the quadratic form would be the dominant eigenvector of $W[k_f]$. However, notice that substituting $\mathbf{v} = \begin{bmatrix} \mathbf{v}_{max}(Q) \\ \mathbf{0} \end{bmatrix}$ into the quadratic form yields $\lambda_{max}(Q)$, but this $\mathbf{v}$ cannot be a dominant eigenvector since it is not strictly positive. Hence, a contradiction is reached.

Proof of 4: Since Q is invertible, it is in fact positive definite (in addition to being symmetric and positive semidefinite). It is immediate that $R = Q^{-1}$ is symmetric and positive definite. Since Q is elementwise strictly positive, it follows that $R = Q^{-1}$ is irreducible.

Proof of 5: This result on the sign pattern of the inverse was proved for the class of doubly-positive matrices (positive-definite matrices with strictly-positive entries) by Fiedler in [30], and generalized to the class of irreducible doubly nonnegative matrices in our recent work (see Theorem 1 of [31]).

Proof of 6: R is positive definite matrix, and hence its diagonal entries are positive as is its determinant. The off-diagonal entries are seen to be negative from the matrix inversion formula for 2×2 matrices. Thus, the matrix is an M matrix. ■

Remark 1. The characterizations in Theorem 1 crucially depend on the doubly-nonnegative structure of the Gramian, which is a consequence of the nonnegative structure of the network-consensus model (as defined by nonnegative state, input, and output matrices). Doubly-nonnegative matrices also arise in other contexts, such as semi-definite programming and covariance-matrix analysis [32,33].

Item 5 of Theorem 1 indicates that inverses of Gramian submatrices have a sophisticated sign pattern. While the diagonal entries of the inverse are positive, the off-diagonal entries may be of either sign. Item 5 indicates, however, that some of the off-diagonal entries must be negative. The pattern of nonnegative off-diagonal entries can be given a graph-theoretic interpretation, which helps to give insight into the target-control metrics. To formalize this interpretation, it is helpful to define a graph that represents the sign pattern of the inverse of a Gramian submatrix. Specifically, we define the *negative-inverse graph* for the invertible Gramian submatrix $W(\mathcal{B}, k_f)$ as an (unweighted, undirected) graph on $|\mathcal{B}|$ vertices labeled $1, \ldots, |\mathcal{B}|$, where the notation $|\mathcal{B}|$ indicates the cardinality of the set. An edge is drawn between vertex i and j if $[W^{-1}]_{ij}$ is negative. The following corollary is an immediate consequence of Item 5 of Theorem 1:

Corollary 1. *The negative-inverse graph for any invertible Gramian submatrix $W(\mathcal{B}, k_f)$ is connected.*

The matrix-theoretic properties of the Gramian's principal submatrices developed in Theorem 1 allow characterization of the defined target-control metrics and associated optimal input signals. Several results on the target-control metrics are listed in the following theorem, and then interpreted in the following discussion. The theorem and ensuing proof uses the notation $|\mathbf{x}|$ to indicate the entry-wise absolute value of a vector $\mathbf{x}$.

Theorem 2. *Assume that the input–output network-consensus model is target controllable. The target-control metrics and associated optimal input signals have the following properties, for all sufficiently large k_f:*

(1) The network model has a minimally-secure goal $\bar{\mathbf{y}}_{min}$ which is strictly positive. Additionally, the optimal input signal for the minimally-secure goal is nonnegative for $k = 0, \ldots .k_f - 1$.
(2) When the network has two target nodes, the target control energy satisfies $E(k_f, |\bar{\mathbf{y}}|) \leq E(k_f, \bar{\mathbf{y}})$ for any goal state $\mathbf{y}$.
(3) The projection-manipulation energy satisfies $F_{min}(k_f, |\boldsymbol{\alpha}|) \leq F_{min}(k_f, \boldsymbol{\alpha})$ for any projection vector $\boldsymbol{\alpha}$. Further, the optimal input signal for manipulation of any projection is nonnegative.
(4) The global security metrics satisfy the inequality: $E_{min,full}(k_f) < E_{min}(k_f) < F_{min}(k_f)$, where $E_{min,full}(k_f)$ refers to the target-security metric when the set of target nodes $\mathcal{T}$ contains all nodes in the network.

Proof. *Proof of Item 1:* The minimally-secure goal $\bar{\mathbf{y}}_{min}$ is the dominant eigenvector of $W(\mathcal{T}, k_f)$. From Item 2 of Theorem 1, this dominant eigenvector is strictly positive. From Eq. (5), the optimal target-control input for this goal is $\widehat{\mathbf{u}}[i] = (CA^{k_f-i-1}B)^T W(\mathcal{T}, k_f)^{-1} \bar{\mathbf{y}}_{min}$. Since $\bar{\mathbf{y}}_{min}$ is an eigenvector of $W(\mathcal{T}, k_f)$, it is also an eigenvector of $W(\mathcal{T}, k_f)^{-1}$. Thus, $W(\mathcal{T}, k_f)^{-1}\mathbf{y}_{min}$ is positive, and it follows that $\widehat{\mathbf{u}}[i]$ is nonnegative.

Proof of Item 2: The result follows immediately from the fact that $W(\mathcal{T}, k_f)^{-1}$ is an M-matrix in this case.

Proof of 3: Since $W(\mathcal{T}, k_f)$ is a nonnegative matrix, it follows that $|\boldsymbol{\alpha}^T|W(\mathcal{T}, k_f)|\boldsymbol{\alpha}| \geq \boldsymbol{\alpha}^T W(\mathcal{T}, k_f)\boldsymbol{\alpha}$. The inequality on the projection-manipulation energy follows immediately. The nonnegativity of the input sequence then follows from a direct computation of the optimal input sequence.

Proof of 4: The inequality relating $E_{min,full}(k_f)$ and $E_{min}(k_f)$ follows from Item 3 of Theorem 1. The inequality relating $E_{min}(k_f)$ and $F_{min}(k_f)$ can be derived by noting that $\lambda_{max}(W(\mathcal{T}, k_f))$ majorizes the diagonal entries of $W(\mathcal{T}, k_f)$; this follows using the same argument as used to derive Item 3 of Theorem 1. ■

Item 1 of Theorem 2 indicates that the network model always has a minimally-secure goal (the unit-norm goal that takes the minimum energy to reach) in the positive orthant. Further, the lowest-energy input needed to reach this goal is nonnegative. It is worth noting, however, that the minimum-energy input needed to reach other positively-valued goals need not be positive.

Because of the diffusive structure of the network dynamics, one might postulate that goal states in the positive orthant (i.e., goals with identically-signed entries) require less energy to achieve. Item 2 in the theorem demonstrates that goal states in the positive orthant indeed require less energy to reach compared to their sign-reversed versions, when there are only two target nodes. In other words, it is easier to move any pair of nodes' states to a more synchronized goal (with both nodes' goal states having the same sign), than to a comparable differentially-signed goal. The low-energy characteristic of the positive orthant results specifically form the M-matrix structure of the inverse Gramian submatrix in the two-target case. However, the result does not generalize to models with more than two target nodes: as Theorem 1 indicates, the inverse Gramian submatrix has a complicated sign pattern when the target set has more than two nodes, which means that mixed-sign goals may sometimes require less energy to reach than their positive orthant counterparts. However, the connectedness of the negative-sign graph (Corollary 1) does indicate that low-energy control is possible for many goal states in the nonnegative orthant.

Per Item 3, projections defined by vectors in the first quadrant are easier to manipulate than their sign-reversed counterparts, regardless of the number of target states. This characterization is a direct consequence of the nonnegative form of the Gramian submatrix.

Finally, Item 4 provides a comparison of different global security metrics. The inequalities follow immediately from the positive definiteness of the Gramian and its submatrices, however the fact that they are strict is a consequence of the doubly-nonnegative structure of the Gramian.

Remark. All goal states in the positive orthant require less energy to reach than their sign-reversed counterparts if the corresponding inverse Gramian submatrix is an M-matrix. Per the discussion above, this is guaranteed when the network has two target nodes. When the network has three or more target nodes, the inverse Gramian submatrix may or may not be an M-matrix. The class of nonnegative matrices whose inverses are M-matrices has been characterized algebraically in the linear-algebra literature (see e.g. [34], and these results can be brought to bear to check whether positive-orthant goal states necessarily can be reached with low energy.

Next, we study how the graph topology of the input–output network consensus model constrains the associated Gramian submatrix and its inverse. The main outcome of this analysis is that the magnitudes of the entries in the Gramian submatrix are small if the target nodes are far from the source. In fact, the entries decrease monotonically as the target nodes are moved further away from the source nodes, in a certain sense (related to cutsets of the network graph). Conversely, metrics related to the inverse Gramian are large if the target nodes are far from the source nodes. To formalize these notions, we find it convenient to consider vertex-cutsets in the network graph that separate the source and target vertices. Formally, a set of vertices $\mathcal{C}$ in the graph Γ (equivalently, nodes in the network) is referred to as a *separating cutset*, if all directed paths between source and target nodes in Γ pass through a vertex in $\mathcal{C}$.

The following theorem characterizes the principal submatrix of the Gramian associated with the network model (i.e. the $\mathcal{T}$-Gramian submatrix), and its inverse, in terms of a separating cutset:

Theorem 3. *Consider the Gramian submatrix $W(\mathcal{T}, k_f)$ for the input–output network-consensus model. Also, let $\mathcal{C}$ be a separating cutset of the network graph Γ. Then the following inequalities hold:*

(1) $[W(\mathcal{T}, k_f)]_{ij} \leq max_l[W(\mathcal{C}, k_f)]_{ll}$. That is, all entries in the Gramian submatrix associated with the target nodes are smaller than at least one of the diagonal entries of the Gramian matrix corresponding to the separating cutset nodes.

(2) Consider any vector $\boldsymbol{\alpha}$ such that $|\boldsymbol{\alpha}|^T\mathbf{1} = 1$. Then $\boldsymbol{\alpha}^T W(\mathcal{T}, k_f)\boldsymbol{\alpha} \leq \max_i[W(\mathcal{C}, k_f)]_{ii}$.

(3) $\lambda_{max}(W(\mathcal{T}, k_f)) \leq p \max_i[W(\mathcal{C}, k_f)]_{ii}$, where p is the number of target nodes.

Proof. From the proof of Theorem 1, the diagonal entries of the Gramian $Q = W(k_f)$ can be written in terms of the impulse responses of the network model at the corresponding nodes. Specifically, the lth diagonal entry can be written as $q_{ll} = \sum_{z\in\mathcal{S}} h^T_{zl,k_f} h_{zl,k_f}$.

To prove the theorem, we first compare q_{ll} for $l \in \mathcal{C}$ with q_{ll} for $l \in \mathcal{T}$. The crux of the proof lies in recognizing that the impulse responses h_{zl,k_f} for $l \in \mathcal{T}$ can be expressed in terms of the impulse responses h_{zl,k_f} for $l \in \mathcal{C}$. To express the impulse response in this way, let us first define a set $\mathcal{V}$ which contains all nodes that are isolated from the source nodes by $\mathcal{C}$. (We note that the set $\mathcal{V}$ contains $\mathcal{T}$ as well as all other nodes separated from the source nodes by the cutset). The vector $\widehat{\mathbf{x}}[k]$ is defined to contain the states of the nodes in $\mathcal{V}$. Then notice that the response at any node l within $\mathcal{V}$ due to an impulse input at node z can be found by solving:

$$\widehat{\mathbf{x}}[k+1] = \widehat{A}\widehat{\mathbf{x}}[k] + \widehat{B}\mathbf{h}_{z,k_f}[k] \tag{10}$$

$$h_{zl,k_f}[k] = \mathbf{e}_p^T\widehat{\mathbf{x}}[k], \tag{11}$$

where $\widehat{A}$ is the principal submatrix of A formed by maintaining the rows/columns identified in $\mathcal{V}$, $\widehat{B}$ is the submatrix of A with rows specified by $\mathcal{V}$ and columns specified by $\mathcal{C}$, $\mathbf{h}_{z,k_f}[k]$ concatenates the impulse responses h_{zl,k_f} for $l \in \mathcal{C}$ at time k, and the 0–1 indicator vector $\mathbf{e}_p$ is a 0–1 indicator vector which selects the response at node l from the vector $\widehat{\mathbf{x}}[k]$, We note that the expression holds for $k = 0, 1, \ldots, k_f$. We stress that this expression allows computation of the impulse response at any node l in $\mathcal{T}$ without requiring tracking of the states of nodes outside $\mathcal{V}$, provided that the impulse responses at nodes in $\mathcal{C}$ are known.

From Eq. (10), it follows that

$$h_{zl,k_f}[k] = \mathbf{e}_p^T\begin{bmatrix}\widehat{A}^{k-1}\widehat{B} & \ldots & \widehat{A}\widehat{B} & \widehat{B}\end{bmatrix}\mathbf{h}_{z,k_f}, \tag{12}$$

where $\mathbf{h}_{z,k_f} = \begin{bmatrix}\mathbf{h}_{z,k_f}[0] \\ \vdots \\ \mathbf{h}_{z,k_f}[k-1]\end{bmatrix}$. Now consider the sum of the entries in each row of the matrix $\begin{bmatrix}\widehat{A}^{k-1}\widehat{B} & \ldots & \widehat{A}\widehat{B} & \widehat{B}\end{bmatrix}$. The sum of each row in this matrix is less than or equal to the sum of the corresponding row in the top-right block of $\begin{bmatrix}\widehat{A} & \widehat{B} \\ 0 & I\end{bmatrix}^{k-1}$, where the dimension of the identity matrix has been chosen so that exponentiated matrix is square. However, as the matrix $\begin{bmatrix}\widehat{A} & \widehat{B} \\ 0 & I\end{bmatrix}$ has unity row sums, so does its powers. Thus, the vector $\mathbf{e}_p^T\begin{bmatrix}\widehat{A} & \widehat{B} \\ 0 & I\end{bmatrix}^{k-1}$ has entries that are nonnegative and sum to less than 1. Considering Eq. (12) for $k = 0, 1, \ldots, k_f$, we thus find that h_{zl,k_f} can be found by convolving the impulse responses h_{zi,k_f} for $i \in C$ with nonnegative signals, whose total sum is less than 1, and then summing. Further, the same convolution can be applied to find the impulse response for each source node $z \in \mathcal{S}$. However, it is known that convolution by a nonnegative signal whose entries sum to less than 1 serves to decrease the energy (two-norm) of a signal. We thus recover that $q_{ll} = \sum_{z\in\mathcal{S}} h^T_{zl,k_f} h_{zl,k_f}$ must be smaller for each $l \in \mathcal{V}$ as compared to at least one $l \in \mathcal{C}$. Since $\mathcal{T} \in \mathcal{V}$, we have thus shown that the diagonal entries of $W(\mathcal{T}, k_f)$ are less than or equal to $max_l[W(\mathcal{C}, k_f)]_{ll}$. Finally, from the fact that $W(\mathcal{T}, k_f)$ is doubly nonnegative, it is immediate that the off-diagonal entries are less than or equal to the largest diagonal entry. Thus, Item 1 of the theorem statement is verified. Items 2 and 3 then follow immediately from standard properties of positive definite matrices. ■

The graph-theoretic analyses of Gramian submatrix properties in Theorem 3 immediately yield graph-theoretic bounds on the defined target control metrics. In particular, the target control metrics can be majorized in terms of the energy required to manipulate the nodes on any separating cutset of the network graph. To present these comparisons, it is helpful to explicitly define control energy metrics for the nodes on a separating cutset. Specifically, first consider any vertex c contained in a separating cutset $\mathcal{C}$ of the network graph. We refer to the minimum input energy required to move the state of the corresponding network node c to a unity value over the interval $[0, k_f]$ (assuming that the network is initially relaxed) as $E_c(k_f)$. In analogy with Definition 2, $E_c(k_f)$ can be formally defined as $E_c(k_f) = \min_{\mathbf{u}[0],\ldots,\mathbf{u}[k_f-1]} \sum_{i=0}^{k_f-1} \mathbf{u}^T[i]\mathbf{u}[i]$, subject to the constraint that the input sequence drives the system from a relaxed state to $x_c[k_f] = 1$. We then define the *cutset-control energy* as $E_{\mathcal{C}}(k_f) = \min_{c\in\mathcal{C}} E_c(k_f)$. The target-control metrics can be majorized in terms of the cutset-control energy, as follows:

Theorem 4. *Consider the input–output network consensus model. The following inequalities hold for the target-control metrics, for any separating cutset $\mathcal{C}$ of the network graph:*

(1) The projection-manipulation energy satisfies $F(k_f, \boldsymbol{\alpha}) \geq E_{\mathcal{C}}(k_f)$ for any $\boldsymbol{\alpha}$ such that $|\boldsymbol{\alpha}|^T\mathbf{1} = 1$.

(2) The projection security satisfies $F_{min}(k_f) \geq E_{\mathcal{C}}(k_f)$.

(3) The target security satisfies $E_{min(k_f)} \geq \frac{1}{p}E_{\mathcal{C}}(k_f)$.

Proof. Item 1 in the theorem follows from Item 2 of Theorem 3, together with the expression for the projection-manipulation energy in Lemma 1. Item 2 is verified by noticing that the inequality in Item 1 holds for all $\boldsymbol{\alpha}$ with unit one-norm, and hence holds for the vector $\boldsymbol{\alpha}$ that minimizes $F(k_f, \boldsymbol{\alpha})$. Item 3 follows from Item 3 of Theorem 3, together with the expression for the target security in Lemma 1. ■

Theorem 4 demonstrates that the target control metrics follow a spatial majorization, with respect to cutsets in the network graph away from the source nodes. Specifically, the energy required to manipulate any projection of the target state is larger than the energy required to manipulate at least one of the nodes on a separating cutset. Thus, state projections become more secure (harder to manipulate) away from the source nodes. A similar result also holds for the target security metric, but with a scale factor related to the number of nodes being manipulated.

Remark. The results in Theorems 3 and 4 based on separating cutsets can be specialized for classes of graphs. One particularly simple case is that of tree graphs (i.e. graphs where there is a single directed path between any two vertices). In this case, consider any connected subgraph of the network graph which does not have any actuation vertices; since the graph is a tree, any such subgraph can be separated from the actuation vertices via a single-node cutset. It thus follows that the entries of the inverse submatrix Gramian corresponding to the subgraph are majorized by the inverse of the Gramian's diagonal entry corresponding to the single node cut. An immediate consequence is that the target-control metrics are monotonically non-decreasing along paths away from the actuation vertices. The results in the Theorem can be also specialized to characterize manipulation of nodes that are at different distances from the actuation vertices, since the set of nodes at a given distance form a separating cutset.

The values of the target-control metrics for long time horizons (i.e., in the limit of large k_f) are of interest, since they serve as lower bounds on energy requirements for arbitrary horizons. Because the network-consensus process naturally asymptotes to a synchronized state, one might expect that manipulation of the dynamics to a desired synchronized state can be achieved with a vanishingly-small energy requirement, given a long time horizon. This notion can be formalized by characterizing Gramian submatrices and their spectra for large k_f. This characterization of Gramian submatrices, and consequent analyses of the target-control metrics, are formalized in the following theorem:

Theorem 5. *Consider any principal submatrix of the Gramian for the network model, say $Q = W(\mathcal{B}, k_f)$. The matrices Q has the following properties:*

(1) Q can be written as $Q = (k_f \sum_{i \in \mathcal{S}} w_i^2)\mathbf{1}\mathbf{1}^T + H(k_f)$, where the absolute values of the entries in the matrix $H(k_f)$ have an upper bound that is independent of k_f. In the expression, w_i is ith entry in the left eigenvector of A associated with its unity eigenvalue, where the eigenvector has been normalized so that its entries sum to 1. Also, $\mathbf{1}$ represents a vector with all entries equal to 1, of appropriate dimension.

(2) The dominant eigenvalue of the matrix Q is given by $\lambda_{max}(Q) = |\mathcal{B}| k_f \sum_{i \in \mathcal{S}} w_i^2 + \widehat{\lambda}(k_f)$, where $|\widehat{\lambda}(k_f)|$ has an upper-bound that is independent of k_f. The corresponding dominant eigenvector is given by $\mathbf{v} = \mathbf{1} + \mathbf{c}(k_f)$, where each entry in $|\mathbf{c}(k_f)|$ is upper bounded by an asymptotically-vanishing function of k_f.

Now consider target control for the input–output network-consensus model. Provided that the model is target controllable, the target security metric is given by: $E_{min}(k_f) = \frac{1}{pk_f \sum_{i \in \mathcal{S}} w_i^2} + \widehat{E}(k_f)$, where $|\widehat{E}(k_f)|$ is upper-bounded by an asymptotically-vanishing function of k_f. The minimally-secure goal is given by $\overline{\mathbf{y}}_{min} = \mathbf{1} + \mathbf{c}(k_f)$, where each entry in $|\mathbf{c}(k_f)|$ is upper bounded by an asymptotically-vanishing function of k_f.

Proof. The state matrix A for the network consensus model has a single strictly dominant eigenvalue at 1, with a corresponding right eigenvector of $\mathbf{1}$ and a left eigenvector which is entrywise strictly positive. Thus, from the Jordan form of A, it follows immediately that powers of the matrix can be expressed as $A^n = \mathbf{1}\mathbf{w}^T + Q(n)$, where $\mathbf{w}^T$ is the left eigenvector of A associated with the 0 eigenvalue (whose entries have been normalized to sum to 1), and $Q(n)$ is a matrix whose entries are each upper-bounded by a decaying exponential function of n. Substituting the expression for A^n into the Gramian formula, we find that $W(k_f) = \sum_{n=0}^{k_f-1} (\mathbf{1}\mathbf{w}^T + R(n)) BB^T (\mathbf{1}\mathbf{w}^T + R(n))$. With some algebra, we find that $W(k_f) = k_f \sum_{i \in \mathcal{S}} w_i^2 \mathbf{1}\mathbf{1}^T + J(k_f)$, where the entries in $J(k_f)$ each have an upper bound that is independent of k_f. The form of the principal submatrix Q of the Gramian given in the theorem statement (Item 1) follows immediately.

To characterize the dominant eigenvalue/eigenvector for Q, notice that the matrix can be written as $k_f((\sum_{i \in \mathcal{S}} w_i^2)\mathbf{1}\mathbf{1}^T + \frac{1}{k_f} H(k_f))$. Next, notice that the matrix $(\sum_{i \in \mathcal{S}} w_i^2)\mathbf{1}\mathbf{1}^T$ has a non-repeated eigenvalue equal to $|\mathcal{B}| \sum_{i \in \mathcal{S}} w_i^2$, with corresponding eigenvector equal to $\mathbf{1}$; the remaining eigenvalues of the matrix are equal to 0. From standard eigenvalue and eigenvector perturbation results for non-repeated eigenvalues [35], it follows that the matrix $(\sum_{i \in \mathcal{S}} w_i^2)\mathbf{1}\mathbf{1}^T + \frac{1}{k_f} H(k_f)$ has a dominant eigenvalue equal to $|\mathcal{B}| \sum_{i \in \mathcal{S}} w_i^2 + g(k_f)$, where $|g(k_f)|$ is upper-bounded by a function of the form $\frac{g}{k_f}$ for some constant g. The corresponding dominant eigenvector is equal to $\mathbf{v} = \mathbf{1} + \mathbf{c}(k_f)$, where each entry in $|\mathbf{c}(k_f)|$ is upper bounded by an asymptotically-vanishing function of k_f. The characterization of the dominant eigenvalue and eigenvector in Item 2 of the theorem statement follows immediately.

The characterization of the target security metric then follows from Lemma 1. ■

Remark. The graph-theoretic analysis of Gramian submatrices in Theorem 3 is relevant to myriad techniques which use the Gramian, beyond assessment of the open-loop control energy. For instance, the Gramian is used in several model reduction methods such as the balanced truncation algorithm [21,36]. The graph-theoretic analysis suggestion that, if the system being reduced has a diffusive-network structure, these model reduction methods can be adapted to also maintain contiguous portions of the network's graph. We leave a careful study to future work.

Remark. The analyses of controllability-Gramian submatrices developed in this letter (Theorems 1 and 3) can be immediately adapted for observability-Gramian submatrices for consensus processes that are measured at a subset of nodes. However, the results admit a different interpretation in the context of target-state observability. As clarified in our previous work [11], estimability of target states is related to submatrices of the inverse Gramian rather than inverses of Gramian submatrices. Thus, the results developed here do not directly give insight into the estimability of target states. Instead, inverses of observability Gramian submatrices measure estimability of target states, when the remaining states are known *a priori*.

5. Examples

We present two examples to illustrate the formal results on Gramian submatrix inverses and target controllability.

Example 1. We illustrate matrix-theoretic properties of the Gramian submatrix inverses developed in Theorem 1 and Lemma 1, and discuss their implications on target control as developed in Theorem 2. For this example, a network model with 5 nodes is considered, which has state matrix

$$A = \begin{bmatrix} 1/6 & 1/3 & 1/6 & 1/3 & 0 \\ 1/4 & 1/8 & 1/8 & 1/4 & 1/4 \\ 1/41/4 & 1/4 & 1/4 & 0 & \\ 1/3 & 1/3 & 1/6 & 1/6 & 0 \\ 0 & 2/3 & 0 & 0 & 1/3 \end{bmatrix}.$$

The network is assumed to have a single source node, namely node 1, and control over the time horizon $[0, 50]$ is considered. The model dynamics are representative of a time-sampled flow-balance process, such as may occur in fluid flow among reservoirs or chemical reaction processes.

First, we find the inverse of the 2×2 controllability Gramian submatrix for the target set $\mathcal{T} = \{3, 4\}$. The inverse is given by $\begin{bmatrix} 143.5 & -142.4 \\ -142.4 & 141.6 \end{bmatrix}$. As expected per Theorem 1, the inverse matrix is seen to have negative off-diagonal entries. An immediate consequence, as indicated in Theorem 2, is that goal states in the first orthant require less energy than goal states in other quadrants with the same-magnitude entries. Also of note, the minimally-secure goal state is found to be $\begin{bmatrix} 0.70 \\ 0.71 \end{bmatrix}$, which lies close to the model's equilibrium manifold wherein all nodes' states are identical, which is in accordance the result on the minimally-secure goal given in Theorem 5.

Second, we also find the sign pattern of the inverse of the full Gramian matrix (i.e., the submatrix inverse when the target set includes all nodes). The sign pattern is as follows:

$$\begin{bmatrix} + & + & - & + & - \\ + & + & + & - & - \\ - & + & + & - & - \\ + & - & - & + & + \\ - & - & - & + & + \end{bmatrix}. \tag{13}$$

We see that the inverse is in accordance with Lemma 1, in that the negative-inverse graph is connected.

Example 2. The spatial majorization of target-control metrics, as developed in Theorems 3 and 4, is illustrated in an example 50-node network. The example is inspired by problems related to manipulation of opinion-dispersion processes in groups, such as targeted advertisement of products in communities or social networks [27]. Dynamic opinion-dispersion processes have been represented using consensus-type models, and further targeted manipulation of opinions has been studied within this context as a target control problem [27]. These models are significant abstractions of real-world opinion-dispersion processes. However, they may be interesting as a means for obtaining graph-theoretic insights into opinion manipulation – e.g., identifying vulnerable locations and possible extent of impact of manipulators.

The example was constructed as follows. The network's graph Γ was formed by placing vertices randomly in the unit square, and connecting vertices within a certain radius. The edge weights of the incoming edges into each vertex were selected to be identical, and were scaled to sum to 1. The consensus dynamics considered in the study was instantiated on the graph (see Fig. 1). We assume that the manipulative adversary can access one node in the network, and can subject the node to actuation, as indicated in red in Fig. 1.

We study the effort required by the adversary to manipulate each individual node's state in the network over a time period of 200 steps. We do so by determining and comparing the target security when each individual node is the target. Per Theorem 4, the target security should improve along cutsets away from the source node: specifically, the target security should be larger for all nodes beyond a cutset as compared to the minimum target security along the cutset. Computation of the target security metrics for the example network show that this is the case. Specifically, the target security metric when each node is the target is displayed on the network graph in Fig. 1. Specifically, the target security is indicated by the radius of the disk at each node. It is seen that the nodes that are close to the source (actuated) node in the graph can be manipulated with limited energy, while distant nodes require significant energy to manipulate, as formalized in Theorem 4. Interestingly, the target security grows rapidly with distance in a neighborhood close to the source, but grows only gradually at further distances. The example suggests that manipulative actors have a sphere of influence where they can alter opinions with limited effort, but also can gradually manipulate opinions at remote locations with larger effort.

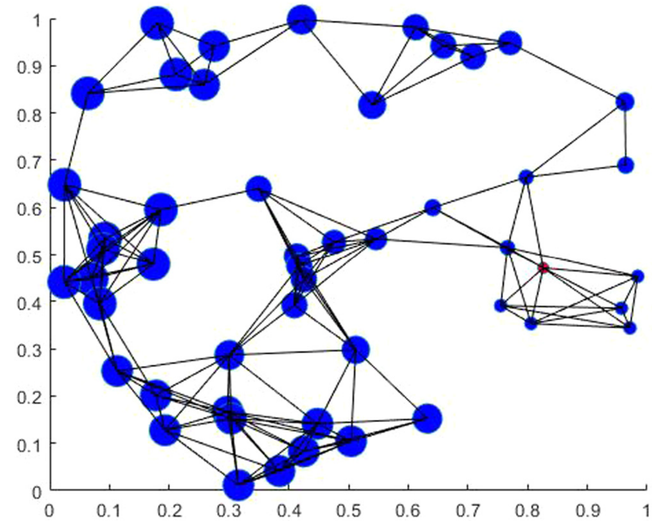

Fig. 1. The energy required to manipulate each individual node in a network-consensus process from a source node is shown, for a 50-node network. The source node is colored red (see bottom right part of the plot). For each node, the size of the disk indicates the energy requirement. Nodes near the source can be manipulated with less energy. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

6. Conclusions and future work

A number of properties of controllability-Gramian submatrices and their inverses have been determined, for a network consensus model with actuation applied at a subset of network nodes. Of particular note, we have shown that the submatrix inverses display a sophisticated sign structure, which results from their doubly positive (positive definite and also entry-wise nonnegative) structure. We have also demonstrated that the entries in the inverse matrices display a majorization with respect to cutsets of the network graph. We have also demonstrated that these properties of controllability Gramian submatrices give insight into target controllability metrics for network consensus processes. The results are appealing from the perspective of assessing threats and evaluating manipulation of network processes, because they allow generic insights into the control energy without requiring exhaustive numerical computation of Gramian inverses.

The results developed here apply to a narrow class of linear dynamical network models, which are significant abstractions of real-world network processes. For instance, the models do not capture the nonlinearities and stochastics that are typical of many network processes, including most industrial processes and social-network dynamics. Thus, the results developed here are at best a starting point for assessing manipulation of many real-world network processes. However, we believe that our analysis approach can be extended toward a much broader class of models with diffusive structures. In particular, several key results in our development only depend on the nonnegative structure of the state matrix, and hence may naturally generalize to the much broader class of linear and nonlinear positive systems. Likewise, several graph-theoretic analyses primarily depend on the diffusive or flow-like structure of the model, and hence may be extensible to a broader class of stochastic and nonlinear diffusive systems (e.g., queueing-network or stochastic automaton models). We expect to pursue these generalizations as future work.

Another interesting direction of future work is to study whether actuation can be used to not only drive the target state to a desired set, but to maintain the state in this set for a period of time, as may be required in a number of applications. We note that maintaining the state in a set may be infeasible even when target control is feasible, and in general will take more effort than simply driving the state to the target. The special structure of the network consensus model may, however, allow maintenance of the state with minimum additional effort in some cases, for instance when the goal states are close to the consensus manifold. We leave it to future work to delineate these cases.

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

References

[1] Chai Wah Wu, Leon O. Chua, Application of Kronecker products to the analysis of systems with uniform linear coupling, IEEE Trans. Circuits Syst. I 42 (10) (1995) 775–778.

[2] Lin Xiao, Stephen Boyd, Seung-Jean Kim, Distributed average consensus with least-mean-square deviation, J. Parallel Distrib. Comput. 67 (1) (2007) 33–46.

[3] Wei Ren, Randal W. Beard, Ella M. Atkins, Information consensus in multivehicle cooperative control, IEEE Control Syst. Mag. 27 (2) (2007) 71–82.

[4] Amirreza Rahmani, Meng Ji, Mehran Mesbahi, Magnus Egerstedt, Controllability of multi-agent systems from a graph-theoretic perspective, SIAM J. Control Optim. 48 (1) (2009) 162–186.

[5] Zhijian Ji, Hai Lin, Haisheng Yu, Protocols design and uncontrollable topologies construction for multi-agent networks, IEEE Trans. Automat. Control 60 (3) (2015) 781–786.

[6] Lin Wang, Guanrong Chen, Xiaofan Wang, Wallace K.S. Tang, Controllability of networked MIMO systems, Automatica 69 (2016) 405–409.

[7] Mengran Xue, Sandip Roy, Input–output properties of linearly-coupled dynamical systems: interplay between local dynamics and network interactions, in: 2017 IEEE 56th Annual Conference on Decision and Control, CDC, IEEE, 2017, pp. 487–492.

[8] Fabio Pasqualetti, Sandro Zampieri, Francesco Bullo, Controllability metrics, limitations and algorithms for complex networks, IEEE Trans. Control Netw. Syst. 1 (1) (2014) 40–52.

[9] Rahul Dhal, Sandip Roy, Vulnerability of network synchronization processes: a minimum energy perspective, IEEE Trans. Automat. Control 61 (9) (2016) 2525–2530.

[10] Henk J. Van Waarde, M. Kanat Camlibel, Harry L. Trentelman, A distance-based approach to strong target control of dynamical networks, IEEE Trans. Automat. Control 62 (12) (2017) 6266–6277.

[11] Amirkhosro Vosughi, Charles Johnson, Mengran Xue, Sandip Roy, Sean Warnick, Target control and source estimation metrics for dynamical networks, Automatica 100 (2019) 412–416.

[12] Yongqiang Guan, Long Wang, Yongqiang Guan, Long Wang, Target controllability of multiagent systems under fixed and switching topologies, Internat. J. Robust Nonlinear Control 29 (9) (2019) 2725–2741.

[13] Jianxi Gao, Yang-Yu Liu, Raissa M. D'souza, Albert-László Barabási, Target control of complex networks, Nature Commun. 5 (2014) 5415.

[14] Isaac Klickstein, Afroza Shirin, Francesco Sorrentino, Energy scaling of targeted optimal control of complex networks, Nature Commun. 8 (2017) 15145.

[15] Tyler H. Summers, Fabrizio L. Cortesi, John Lygeros, On submodularity and controllability in complex dynamical networks, IEEE Trans. Control Netw. Syst. 3 (1) (2016) 91–101.

[16] Dinuka Sahabandu, Jackeline Abad Torres, Rahul Dhal, Sandip Roy, Local open-and closed-loop manipulation of multiagent networks, Internat. J. Robust Nonlinear Control 29 (5) (2019) 1339–1360.

[17] Shiyu Zhao, Fabio Pasqualetti, Networks with diagonal controllability Gramian: Analysis, graphical conditions, and design algorithms, Automatica 102 (2019) 10–18.

[18] Abraham Berman, Robert J. Plemmons, Nonnegative Matrices in the Mathematical Sciences, Vol. 9, SIAM, 1994.

[19] Sandip Roy, Mengran Xue, Sajal K. Das, Security and discoverability of spread dynamics in cyber-physical networks, IEEE Trans. Parallel Distrib. Syst. 23 (9) (2012) 1694–1707.

[20] Mengran Xue, Wei Wang, Sandip Roy, Security concepts for the dynamics of autonomous vehicle networks, Automatica 50 (3) (2014) 852–857.

[21] Hanna Hardin, Jan van Schuppen, System reduction of nonlinear positive systems by linearization and truncation, Posit. Syst. (2006) 431–438.

[22] Robert G. Gallager, Discrete Stochastic Processes, Vol. 321, Springer Science and Business Media, 2012.

[23] E. Kreindler, P. Sarachik, On the concepts of controllability and observability of linear systems, IEEE Trans. Automat. Control 9 (2) (1964) 129–136.

[24] Wilson J. Rugh, Linear System Theory, Vol. 2, Prentice Hall, Upper Saddle River, NJ, 1996.

[25] Arjan van der Schaft, Modeling of physical network systems, Systems Control Lett. 101 (2017) 21–27.

[26] Morris H. DeGroot, Reaching a consensus, J. Amer. Statist. Assoc. 69 (345) (1974) 118–121.

[27] Kasra Koorehdavoudi, Sandip Roy, Mengran Xue, Jackeline Abad Torres, Distributed decision-making algorithms with multiple manipulative actors, Automatica 107 (2019) 317–326.

[28] Jingqi Li, Ximing Chen, Sergio Pequito, George J. Pappas, Victor M. Preciado, Structural target controllability of undirected networks, in: 2018 IEEE Conference on Decision and Control, CDC, IEEE, 2018, pp. 6656–6661.

[29] Nair Maria Maia De Abreu, Old and new results on algebraic connectivity of graphs, Linear Algebra Appl. 423 (1) (2007) 53–73.

[30] Miroslav Fiedler, Old and new about positive definite matrices, Linear Algebra Appl. 484 (2015) 496–503.

[31] S. Roy, M. Xue, Sign patterns of inverse doubly nonnegative matrices, Linear Algebra Appl. (2019) submitted for publication. Also available at https://arxiv.org/abs/1903.04141.

[32] Akiko Yoshise, Yasuaki Matsukawa, On optimization over the doubly nonnegative cone, in: 2010 IEEE International Symposium on Computer-Aided Control System Design, IEEE, 2010, pp. 13–18.

[33] Charles R. Johnson, Brian Lins, Olivia Walch, The critical exponent for continuous conventional powers of doubly nonnegative matrices, Linear Algebra Appl. 435 (9) (2011) 2175–2182.

[34] Ralph A. Willoughby, The inverse M-matrix problem, Linear Algebra Appl. 18 (1) (1977) 75–94.

[35] James Hardy Wilkinson, The Algebraic Eigenvalue Problem, Vol. 662, Clarendon, Oxford, 1965.

[36] Serkan Gugercin, Athanasios C. Antoulas, A survey of model reduction by balanced truncation and some new results, Internat. J. Control 77 (8) (2004) 748–766.