Information Processing Letters 114 (2014) 728–733

Contents lists available at ScienceDirect

Information Processing Letters

www.elsevier.com/locate/ipl

Testing the irreducibility of nonsquare Perron–Frobenius systems

C. Avin [a], M. Borokhovich [a], Y. Haddad [b], E. Kantor [c,1], Z. Lotker [a,2], M. Parter [d,*,3,4], D. Peleg [d,3]

[a] *Ben Gurion University, Beer-Sheva, Israel*
[b] *Jerusalem College of Technology, Jerusalem, Israel*
[c] *The Technion, Haifa, Israel*
[d] *The Weizmann Institute of Science, Rehovot, Israel*

ARTICLE INFO

Article history:
Received 17 September 2013
Received in revised form 14 May 2014
Accepted 14 May 2014
Available online 23 June 2014
Communicated by M. Chrobak

Keywords:
Algorithms
Irreducibility
Perron–Frobenius theorem

ABSTRACT

The Perron–Frobenius (PF) theorem provides a simple characterization of the eigenvectors and eigenvalues of irreducible nonnegative square matrices. A generalization of the PF theorem to nonsquare matrices, which can be interpreted as representing systems with additional degrees of freedom, was recently presented in [1]. This generalized theorem requires a notion of irreducibility for nonsquare systems. A suitable definition, based on the property that every maximal square (legal) subsystem is irreducible, is provided in [1], and is shown to be necessary and sufficient for the generalized theorem to hold. This note shows that irreducibility of a nonsquare system can be tested in polynomial time. The analysis uses a graphic representation of the nonsquare system, termed the *constraint graph*, representing the flow of influence between the constraints of the system.

1. Introduction

The Perron–Frobenius (PF) theorem is stated for irreducible nonnegative square matrices, and provides a simple characterization of their eigenvectors and eigenvalues. This characterization is applicable in many fields of science and engineering, including dynamical systems theory, economics, statistics and optimization. However, many real-life scenarios give rise to nonsquare matrices. For example, the *Power control* problem in Multiple Input Single Output (MISO) systems [5] has a natural algebraic formulation using *nonsquare* matrices. In such systems, a set of multiple synchronized transmitters, located at different places, can transmit at the same time to the same receiver and hence the total number of transmitters m might be strictly larger than the total number of receivers n, leading to $n \times m$ nonsquare systems [2].

The question of whether the PF theorem (along with its applications) can be generalized to a nonsquare setting calls for extending central notions of the spectral theory, such as eigenvectors, eigenvalues and spectral ratio, to a nonsquare setting. This question has been recently

* Corresponding author.
E-mail addresses: avin@cse.bgu.ac.il (C. Avin), borokhom@cse.bgu.ac.il (M. Borokhovich), yoram.haddad@gmail.com (Y. Haddad), erez.kantor@gmail.com (E. Kantor), zvilo@cse.bgu.ac.il (Z. Lotker), merav.parter@weizmann.ac.il (M. Parter), david.peleg@weizmann.ac.il (D. Peleg).

[1] Supported by Eshkol fellowship, the Israel Ministry of Science and Technology.
[2] Supported by a grant of the Israel Science Foundation.
[3] Supported in part by the Israel Science Foundation (grant 894/09), the I-CORE program of the Israel PBC and ISF (grant 4/11), the United States–Israel Binational Science Foundation (grant 2008348), the Israel Ministry of Science and Technology (infrastructures grant 6478), and the Citigroup Foundation.
[4] Recipient of the Google Europe Fellowship in distributed computing; research supported in part by this Google Fellowship.

http://dx.doi.org/10.1016/j.ipl.2014.05.004
0020-0190/

addressed by [1], where such a generalization was presented, giving rise to the following optimization problem, in which the matrices $\mathcal{M}^+, \mathcal{M}^- \in \mathbb{R}^{n\times m}$ are nonsquare matrices for $m \geq n$:

$$\text{maximize } \beta$$
$$\text{subject to:} \quad \mathcal{M}^- \cdot \overline{X} \leq 1/\beta \cdot \mathcal{M}^+ \cdot \overline{X}, \qquad \|\overline{X}\|_1 = 1, \ \overline{X} \geq \bar{0}. \tag{1}$$

The nonsquare PF system $\mathcal{L} = \langle \mathcal{M}^+, \mathcal{M}^- \rangle$ of nonsquare matrices $\mathcal{M}^+, \mathcal{M}^-$ is interpreted as representing some additional freedom given to the system designer. In this setting, each row *entity* has several columns *affectors*, referred to as its *supporters*, which can cooperate in serving it while potentially *repressing* the others. The task is to find the best way to organize the cooperation between the supporters of each entity.

An important component in the generalized PF theorem of [1] is the extension of spectral theory concepts to a nonsquare system $\mathcal{L} = \langle \mathcal{M}^+, \mathcal{M}^- \rangle$. This extension defines the spectrum of $\mathcal{L}$ based on the spectra of all maximal *legal* square subsystems "hidden" within $\mathcal{L}$. (A "hidden" square subsystem is *legal* if each row entity i selects exactly one column entity j such that $\mathcal{M}^+(i, j) > 0$ and in addition the selections of all row entities are distinct, as will be explained formally later.) For example, the *spectral ratio* of the nonsquare system $\mathcal{L}$ is defined as the minimal spectral ratio of all legal square systems "hidden" in $\mathcal{L}$.

Another central notion in the generalized PF theorem is the *irreducibility* of a nonsquare system. A suitable definition is provided in [1], based on the property that every maximal square and legal subsystem hidden in $\mathcal{L}$ is irreducible. This approach has been shown in [1] to yield nonsquare systems with properties similar to those of a square system with respect to the Collatz–Wielandt property, which provides the algorithmic power for the PF theorem. Moreover, it is shown that this irreducibility requirement is both necessary and sufficient for the generalized theorem to hold.

Note, however, that since there could be exponentially many legal square subsystems in a given nonsquare system $\mathcal{L}$, it is not a priori clear if one can check that $\mathcal{L}$ is irreducible in polynomial time. In this note we address this issue using a representation called the *constraint graph* of the system, whose vertices are the n constraints (one per entity) and whose edges represent direct influence between the constraints. For a square system, irreducibility is equivalent to the constraint graph being strongly connected, but for nonsquare systems the situation is more delicate. Although the matrices are not square, the constraint graph is well-defined and provides a valuable *square* representation of the nonsquare system (i.e., the adjacency matrix of the graph). We present a polynomial time algorithm for testing irreducibility of the system, which exploits the properties of the constraint graph. In other words, we show that one can verify in polynomial time that every square system hidden in the nonsquare system is irreducible. Since irreducibility of nonsquare systems is a sufficient and necessary condition for the generalized PF theorem of [1], our algorithm provides an efficient way to test whether this theorem applies for a particular nonsquare system.

Despite the extensive development of spectral theories for square matrices, the generalization of key spectral concepts such as eigenvector or spectral ratio to the nonsquare setting has been less well-studied. There are several possible alternative definitions for eigenvalues in nonsquare matrices. A pioneering generalization of the PF theorem to nonsquare systems, given in [7], applies to a setting involving a pair of nonsquare "pencil" matrices $A, B \in \mathbb{R}^{n\times m}$, where the term "pencil" refers to the expression $A - \lambda \cdot B$, for complex $\lambda \in \mathbb{C}$. Of special interest here are the values that reduce the pencil rank, namely, the λ values satisfying $(A - \lambda B) \cdot \overline{X} = \bar{0}$ for some nonzero $\overline{X}$. This problem is known as the *generalized eigenvalue problem* [7,4,3,6], which can be stated as follows: Given matrices $A, B \in \mathbb{R}^{n\times m}$, find a vector $\overline{X} \neq \bar{0}$, $\lambda \in \mathbb{C}$, so that $A \cdot \overline{X} = \lambda B \cdot \overline{X}$. The complex number λ is said to be an *eigenvalue of A relative to B* iff $A\overline{X} = \lambda \cdot B \cdot \overline{X}$ for some nonzero $\overline{X}$ and $\overline{X}$ is called the *eigenvector of A relative to B*. The set of all eigenvalues of A relative to B is called the *spectrum of A relative to B*, denoted by $sp(A_B)$. Using the above definition, [7] characterized the relation between A and B required to establish their PF property, i.e., guarantee that the generalized eigenpair is nonnegative. As explained before, the approach taken in [1] is different. The spectral notions of eigenvalues, eigenvectors as well as the notion of irreducibility are adapted to the nonsquare setting by treating the nonsquare system as an ensemble of *hidden square* systems on which the standard measures applies.

Another example for defining a property π of *nonsquare* matrices based on demanding that every *square* hidden matrix satisfies property π' involves the notion of totally unimodular matrices. Since the determinant function is defined only for square matrices, the "extension" to the nonsquare setting involves testing the subsquare matrices of the nonsquare matrix. Hence a (nonsquare) matrix A is *totally unimodular* if each square submatrix of A has determinant equal to $0, +1$, or -1. Indeed, total unimodularity of matrices has been turned out to form an important tool in studying integer vectors in polyhedron due to the fact that the vertex set of the polyhedron $P = \{x \mid Ax \leq b\}$ is an *integer* vector. A polynomial time tester for totally unimodularity is given in [8]. Here, too, since there could be exponentially many such square subsystems, it is not a priori clear if one can check it in polynomial time. In this note such polynomial time tester for irreducibility is provided.

2. Preliminaries

Consider a directed graph $G = (V, E)$. A subset of the vertices $W \subseteq V$ is called a *strongly connected component* if G contains a directed path from v to u for every $v, u \in W$. G is said to be *strongly connected* if V is a strongly connected component.

Throughout, vector and matrix inequalities are interpreted in the component-wise sense. A matrix A is *positive* (respectively, *nonnegative*) if all its entries are positive. Let $A \in \mathbb{R}^{n\times n}$ be a square matrix. A is *irreducible* if for every i and j, there exists a natural $k_{i,j}$ such that $(A^{k_{i,j}})_{i,j} > 0$. An alternative definition to irreducibility can be given by considering the directed graph $G_A = (V, E)$ where $V = \{1, \ldots, n\}$ and $(i, j) \in E$ iff $A_{i,j} \neq 0$. (I.e., A

with all nonzero entries changed to 1 is the adjacency matrix of G_A.) Then, A is irreducible iff G_A is *strongly connected*.

The nonsquare PF framework consists of a set $\mathcal{V} = \{v_1, \dots, v_n\}$ of entities whose growth is regulated by a set of *affectors* $\mathcal{A} = \{\mathcal{A}_1, \mathcal{A}_2, \dots, \mathcal{A}_m\}$, for some $m \geq n$. As part of the solution, we set each affector to be either *passive* or *active*. If an affector $\mathcal{A}_j$ is set to be active, then it affects each entity v_i, by either increasing or decreasing it by a certain amount, denoted $g(i, j)$ (which is specified as part of the input). If $g(i, j) > 0$ (resp., $g(i, j) < 0$), then $\mathcal{A}_j$ is referred to as a *supporter* (resp., *repressor*) of v_i. For clarity we may write $g(v_i, \mathcal{A}_j)$ for $g(i, j)$. The affector-entity relation is described by two matrices, the *supporters gain* matrix $\mathcal{M}^+ \in \mathbb{R}^{n \times m}$ and the *repressors gain* matrix $\mathcal{M}^- \in \mathbb{R}^{n \times m}$, given by

$$\mathcal{M}^+(i, j) = \begin{cases} g(v_i, \mathcal{A}_j), & \text{if } g(v_i, \mathcal{A}_j) > 0; \\ 0, & \text{otherwise}, \end{cases}$$

$$\mathcal{M}^-(i, j) = \begin{cases} -g(v_i, \mathcal{A}_j), & \text{if } g(v_i, \mathcal{A}_j) < 0; \\ 0, & \text{otherwise}. \end{cases}$$

A *system* is given by $\mathcal{L} = \langle \mathcal{M}^+, \mathcal{M}^- \rangle$, where $\mathcal{M}^+, \mathcal{M}^- \in \mathbb{R}_{\geq 0}^{m \times n}$, $n = |\mathcal{V}|$ and $m = |\mathcal{A}|$. The supporter (resp., repressor) set of v_i is denoted by

$$\mathcal{S}_i(\mathcal{L}) = \{\mathcal{A}_j \mid \mathcal{M}^+(v_i, \mathcal{A}_j) > 0\} \quad \text{and}$$
$$\mathcal{R}_i(\mathcal{L}) = \{\mathcal{A}_j \mid \mathcal{M}^-(v_i, \mathcal{A}_j) > 0\}.$$

When $\mathcal{L}$ is clear from context, we may omit it and simply write $\mathcal{R}_i$ and $\mathcal{S}_i$. Throughout, we restrict attention to systems in which $|\mathcal{S}_i| \geq 1$ for every $v_i \in \mathcal{V}$. A system is *square* if $m = n$. If $m > n$ it is *nonsquare*.

To define *irreducibility* for a nonsquare PF system $\mathcal{L}$, we first present the notion of a *selection matrix*. A selection matrix $F \in \{0, 1\}^{m \times n}$ is *legal* for $\mathcal{L}$ iff for every entity $v_i \in \mathcal{V}$ there exists exactly one supporter $\mathcal{A}_j \in \mathcal{S}_i$ such that $F(j, i) = 1$. In other words, the sum of each column is 1. Such a matrix F can be thought of as representing a selection performed on $\mathcal{S}_i$ by each entity v_i, picking exactly one of its supporters. Since every supporter is also a repressor of some other entity it holds that $\mathcal{A} = \bigcup_i \mathcal{S}_i = \bigcup_i \mathcal{R}_i$. It follows that $|\mathcal{A}| = |\mathcal{V}| = n$, the number of active affectors becomes equal to the number of entities, resulting in a square system. Denote the family of legal selection matrices, capturing the ensemble of all square systems hidden in $\mathcal{L}$, by

$$\mathcal{F}(\mathcal{L}) = \{F \mid F \text{ is legal for } \mathcal{L}\}. \tag{2}$$

Let $\mathcal{L}(F)$ be the square system corresponding to the legal selection matrix F, namely, $\mathcal{L}(F) = \langle \mathcal{M}^+ \cdot F, \mathcal{M}^- \cdot F \rangle$.

A square system $\mathcal{L} = \langle \mathcal{M}^+, \mathcal{M}^- \rangle$ is *irreducible* iff (a) $\mathcal{M}^+$ is nonsingular and (b) $\mathcal{M}^-$ is irreducible. A nonsquare system $\mathcal{L}'$ is *irreducible* iff $\mathcal{L}'(F)$ is irreducible for every selection matrix $F \in \mathcal{F}$. The following corollary is from [1].

Corollary 2.1. *In an irreducible system $\mathcal{L}$, $\mathcal{S}_i \cap \mathcal{S}_j = \emptyset$ for every v_i, v_j.*

The constraint graph provides a graph-theoretic characterization of irreducible systems. Let $\mathcal{CG}_{\mathcal{L}}(\mathcal{V}, E)$ be the *constraint graph* for system $\mathcal{L}$, defined by including in E a directed edge $e_{i,j}$ from v_i to v_j iff $\mathcal{S}_i \cap \mathcal{R}_j \neq \emptyset$. Note that for a legal selection matrix F, the constraint graph $\mathcal{CG}_{\mathcal{L}(F)}$ is a subgraph of $\mathcal{CG}_{\mathcal{L}}$, and moreover, $\mathcal{CG}_{\mathcal{L}} = \bigcup_{F \in \mathcal{F}} \mathcal{CG}_{\mathcal{L}(F)}$.

3. Algorithm for testing irreducibility

In this section, we provide a polynomial-time algorithm for testing the irreducibility of a given nonnegative system $\mathcal{L}$. Note that if $\mathcal{L}$ is a square system, then irreducibility can be tested in a straightforward manner by checking that $\mathcal{M}^-$ is irreducible and that $\mathcal{M}^+$ is nonsingular.

However, recall that a nonsquare system $\mathcal{L}$ is irreducible iff every hidden square system $\mathcal{L}(F)$, $F \in \mathcal{F}$, is irreducible. Since $\mathcal{F}$ might be exponentially large, a brute-force testing of $\mathcal{L}(F)$ for every F is too costly, hence another approach is needed. Before presenting the algorithm, we provide some notation. For a digraph D, denote the set of incoming neighbors of a node v_k by $\Gamma^{in}(v_k, D) = \{v_j \mid e_{j,i} \in E(D)\}$. The incoming neighbors of a set of nodes $V' \in \mathcal{V}$ is denoted $\Gamma^{in}(V', D) = \bigcup_{v_k \in V'} \Gamma^{in}(v_k, D)$.

Algorithm description To test irreducibility, Algorithm Irr_Test (see Fig. 1) must verify that the constraint graph $\mathcal{CG}_{\mathcal{L}(F)}$ of every $F \in \mathcal{F}$ is strongly connected. The algorithm consists of at most $n - 1$ rounds. In round t, it is given as input a partition $\mathcal{C}^t = \{C_1^t, \dots, C_{k_t}^t\}$ of $\mathcal{V}$ into k_t disjoint clusters such that $\bigcup_i C_i^t = \mathcal{V}$. For round $t = 0$, the input is a partition $\mathcal{C}^0 = \{C_1^0, \dots, C_n^0\}$ of the entity set $\mathcal{V}$ into n singleton clusters $C_i^0 = \{v_i\}$. The output at round t is a coarser partition $\mathcal{C}^{t+1}$, in which at least two clusters of $\mathcal{C}^t$ were merged into a single cluster in $\mathcal{C}^{t+1}$. The partition $\mathcal{C}^{t+1}$ is formed as follows. The algorithm first forms a graph $D_t = (\mathcal{C}^t, E_t)$ on the clusters of the input partition $\mathcal{C}^t$, treating each cluster $C_i^t \in \mathcal{C}^t$ as a node, and including in E_t a directed edge (i, j) from C_i^t to C_j^t if and only if there exists an entity node $v_k \in C_i^t$ such that *each* of its supporters $\mathcal{A}_i \in \mathcal{S}_k$ is a repressor of *some* entity $v_{k'} \in C_j^t$, i.e., $\mathcal{S}_k \subseteq \bigcup_{v_{k'} \in C_j^t} \mathcal{R}_{k'}$.

The partition $\mathcal{C}^{t+1}$ is now formed by merging clusters C_j^t that belong to the same strongly connected component in D_t into a single cluster $C_{k'}^{t+1}$ in $\mathcal{C}^{t+1}$. Each cluster of $\mathcal{C}^{t+1}$ corresponds to a unique strongly connected component in D_t. If D_t contains no strongly connected component except for singletons, which implies that no two cluster nodes of D_t can be merged, then the algorithm declares the system $\mathcal{L}$ as reducible and halts. Otherwise, it proceeds with the new partition $\mathcal{C}^{t+1}$. Importantly, in $\mathcal{C}^{t+1}$ there are at least two entity subsets that belong to distinct clusters in $\mathcal{C}^t$ but to the same cluster node in $\mathcal{C}^{t+1}$. If none of the rounds ends with the algorithm declaring the system reducible (due to clusters "merging" failure), then the procedure proceeds with the cluster merging until at some round $t^* \leq n - 1$ the remaining partition $\mathcal{C}^{t^*} = \{\{\mathcal{V}\}\}$ consists of a single cluster node that encompasses the entire entity set.

Algorithm Irr_Test($\mathcal{L}$)

1. $t \leftarrow 0$;
2. $k_t \leftarrow n$;
3. $C_i^0 \leftarrow \{v_i\}$ for every $i \in [1, k_t]$;
4. $\mathcal{C}^0 \leftarrow \{C_1^0, \ldots, C_{k_t}^0\}$;
5. While $|\mathcal{C}^t| > 1$ do:
   (a) $\mathcal{R}(C_i^t) \leftarrow \bigcup_{v_k \in C_i^t} \mathcal{R}_k$, for every $i \in [1, k_t]$;
   (b) $E_t \leftarrow \{e(i,j) \mid \exists v_k \in C_i^t,$ such that $\mathcal{S}_k \subseteq \mathcal{R}(C_j^t)\}$;
   (c) Let $D_t = (\mathcal{C}^t, E_t)$;
   (d) $k_{t+1} \leftarrow$ number of strongly connected components in D_t;
   (e) If $k_{t+1} = k_t$ and $|\mathcal{C}^t| \geq 2$, then return "no";
   (f) Decompose $D_t(\mathcal{C}^t, E_t)$ into strongly connected components $\widehat{\mathcal{C}}^1, \ldots, \widehat{\mathcal{C}}^{k_{t+1}}$;
   (g) $C_i^{t+1} \leftarrow \bigcup_{C_j \in \widehat{\mathcal{C}}^i} C_j$ for every $i \in [1, k_{t+1}]$;
   (h) $\mathcal{C}^{t+1} \leftarrow \{C_1^{t+1}, \ldots, C_{k_{t+1}}^{t+1}\}$;
   (i) $t \leftarrow t+1$;
6. Return "yes".

Fig. 1. The pseudocode of Algorithm Irr_Test.

Analysis We first provide some high level intuition for the correctness of the algorithm. Recall, that the goal of the algorithm is to test whether the entire entity set $\mathcal{V}$ resides in a single strongly connected component in the constraint graph $\mathcal{CG}_{\mathcal{L}(F)}$ for every selection matrix $F \in \mathcal{F}$. This test is performed by the algorithm in a gradual manner by monotonically increasing the subsets of nodes that belong to the same strongly connected component in every $\mathcal{CG}_{\mathcal{L}(F)}$. In the beginning of the execution, the most one can claim is that every entity v_k is in its own strongly connected component. Over time, clusters are merged while maintaining the invariant that all entities of the same cluster belong to the same strongly connected component in every $\mathcal{CG}_{\mathcal{L}(F)}$. More formally, the following invariant is maintained in every round t: the entities of each cluster $C_i^t \subseteq \mathcal{V}$ of the graph D_t are guaranteed to be in the same strongly connected component in the constraint graph $\mathcal{CG}_{\mathcal{L}(F)}$ for every selection matrix $F \in \mathcal{F}$. We later show that if the system $\mathcal{L}$ is irreducible, then the merging process never fails and therefore the last partition $\mathcal{C}^{t^*} = \{\{\mathcal{V}\}\}$ consists of a single cluster node that contains all entities, and by the invariant, all entities are guaranteed to be in the same strongly connected component in the constraint graph of any hidden square subsystem.

We now provide some high level explanation for the validity of this invariant. Starting with round $t = 0$, each cluster node $C_i^0 = \{v_i\}$ is a singleton and every singleton entity is trivially in its own strongly connected component in any constraint graph $\mathcal{CG}_{\mathcal{L}(F)}$. Assume the invariant holds up to round t, and consider round $t+1$. The key observation in this context is that the new partition $\mathcal{C}^{t+1}$ is defined based on the graph $D_t = (\mathcal{C}^t, E_t)$, whose edges are independent of the specific supporter selection that is made by the entities (and that determines the resulting hidden square subsystem). This holds due to the fact that a directed edge $(i,j) \in E_t$ between the clusters $C_i^t, C_j^t \in \mathcal{C}^t$ exists if and only if there exists an entity node $v_k \in C_i^t$ such that *each* of its supporter $\mathcal{A}_i \in \mathcal{S}_k$ is a repressor of *some* entity $v_{k'} \in C_j^t$. Therefore, if the edge (i,j) exists in the D_t, then it exists also in the cluster graph corresponding to the constraint graph $\mathcal{CG}_{\mathcal{L}(F)}$ (i.e., the graph formed by representing every strongly connected component of $\mathcal{CG}_{\mathcal{L}(F)}$ by a single node) for *every* hidden square subsystem $\mathcal{L}(F)$, no matter which supporter $\mathcal{A}_i \in \mathcal{S}_k$ was selected by F for v_k. Hence, under the assumption that the invariant holds for $\mathcal{C}^t$, the coarse-grained representation of the clusters of $\mathcal{C}^t$ in $\mathcal{C}^{t+1}$ is based on their membership in the same strongly connected component in the "selection invariant" graph D_t, thus the invariant holds also for $t+1$.

We next formalize this argumentation. We say that round t is *successful* if D_t contains a strongly connected component of size greater than 1. We begin by proving the following.

Claim 3.1. *For every successful round t, the partition $\mathcal{C}^{t+1}$ satisfies the following properties.*

(A1) *$\mathcal{C}^{t+1}$ is a partition of $\mathcal{V}$, i.e., $C_i^{t+1} \subseteq \mathcal{V}$, $C_j^{t+1} \cap C_i^{t+1} = \emptyset$ for every $i, j \in [1, k_{t+1}]$, and $\bigcup_{j \leq k_{t+1}} C_j^{t+1} = \mathcal{V}$.*

(A2) *Every $C_j^{t+1} \in \mathcal{C}^{t+1}$ is a strongly connected component in the constraint graph $\mathcal{CG}_{\mathcal{L}(F)}$ for every selection matrix $F \in \mathcal{F}$.*

Proof. By induction on t. Clearly, since $C_i^0 = \{v_i\}$ for every i, Properties (A1) and (A2) trivially hold for $\mathcal{C}^0$. We now show that if round $t = 0$ is successful, then (A1) and (A2) hold for $\mathcal{C}^1$. Since the edges of D_0 exist also in the corresponding cluster graph of $\mathcal{CG}_{\mathcal{L}(F)}$ under any selection F of the entities, the clusters of $\mathcal{C}^0$ that are merged into a single strongly connected component in $\mathcal{C}^1$, belong also to the same strongly connected component in the constraint graph $\mathcal{CG}_{\mathcal{L}(F)}$ of every $F \in \mathcal{F}$. Next, assume these properties to hold for every round up to $t-1$ and consider round t. Since round t is successful, any prior round $t' < t$ was successful as well, and thus the induction assumption can be applied on round $t-1$. In particular, since $\mathcal{C}^{t+1}$ corresponds to strongly connected components of D_t, it represents a partition of the clusters of $\mathcal{C}^t$. By the induction assumption for round $t-1$, Property (A1) holds for $\mathcal{C}^t$

and therefore $\mathcal{C}^t$ is a partition of the entity set $\mathcal{V}$. Since $\mathcal{C}^{t+1}$ corresponds to a partition of $\mathcal{C}^t$, it is a partition of $\mathcal{V}$ as well so (A1) is established. Property (A2) holds for $\mathcal{C}^{t+1}$ by the same argument provided for the induction base. The claim follows. □

We next show that the algorithm return "yes" for every irreducible system. Specifically, we show that for an irreducible system, if $|\mathcal{C}^t| > 1$ then round t is *successful*, i.e., the merging operation of the cluster graph D_t succeeds. Once $\mathcal{C}^t$ contains a single cluster (containing all entities), the algorithm terminates and returns "yes". We first provide an auxiliary claim.

Claim 3.2. *If $\mathcal{L}$ is irreducible and $|\mathcal{C}^t| > 1$, then $|\Gamma^{in}(C_j^t, D_t)| \geq 1$ for every $C_j^t \in \mathcal{C}^t$.*

Proof. First note that if $\mathcal{C}^t$ is defined, then round $t-1$ was successful. Therefore, by Property (A1) of Claim 3.1, $\mathcal{C}^t$ is a partition of the entity set $\mathcal{V}$. Assume, towards contradiction, that the claim does not hold, and let $C_j^t \in \mathcal{C}^t$ be such that $\Gamma^{in}(C_j^t, D_t) = \emptyset$. Denote the set of incoming neighbors of component C_j^t in the constraint graph $\mathcal{CG}_{\mathcal{L}}$ by $W = \Gamma^{in}(C_j^t, \mathcal{CG}_{\mathcal{L}}) \setminus C_j^t$. Since $\mathcal{CG}_{\mathcal{L}}$ is irreducible, the vertices of C_j^t are reachable from the outside, so $W \neq \emptyset$. Let the repressors set of C_j^t be $\mathcal{R}(C_j^t) = \bigcup_{v_k \in C_j^t} \mathcal{R}_k$. We now construct a square hidden system $\mathcal{L}(F^*)$ which is reducible, in contradiction to the irreducibility of $\mathcal{L}$. Specifically, we look for a selection matrix F^* satisfying that for every entity $v_k \in W$, its selected supporter $\mathcal{A}_k$ in $\mathcal{L}(F^*)$ (i.e., the one for which $F^*(\mathcal{A}_k, v_k) = 1$) is not a repressor of any of the entities in C_j^t, i.e., $\mathcal{A}_k \in \mathcal{S}_k \setminus \mathcal{R}(C_j^t)$. Recall, that since $\mathcal{L}$ is irreducible, the supporter sets $\mathcal{S}_i, \mathcal{S}_j$ are pairwise disjoint (see Corollary 2.1). Note that since $\Gamma^{in}(C_j^t, D_t) = \emptyset$, such a selection matrix F^* exists. To see this, assume, towards contradiction, that F^* does not exist. This implies that there exists an entity $v_k \in W$ such that $\mathcal{S}_k \subseteq \mathcal{R}(C_j^t)$, and therefore an affector in $\mathcal{S}_k \setminus \mathcal{R}(C_j^t)$ could not be selected for F^*. Let $C_i^t \in \mathcal{C}^t$ be the cluster such that $v_k \in C_i^t$. Since $\mathcal{C}^t$ is a partition of the entity set $\mathcal{V}$, such C_i^t exists. Since $\mathcal{S}_k \subseteq \mathcal{R}(C_j^t)$, the edge $e_{i,j}$ exists in D_t, in contradiction to the fact that C_j^t has no incoming neighbors in D_t. We therefore conclude that F^* exists.

We now show that $\mathcal{L}(F^*)$ is reducible. In particular, we show that the incoming degree of the component C_j^t (from entities in other components) in the constraint graph $\mathcal{L}(F^*)$ of the square system $\mathcal{L}(F^*)$, is zero, i.e., $\Gamma^{in}(C_j^t, \mathcal{CG}_{\mathcal{L}(F^*)}) = \emptyset$. Assume, towards contradiction, that there exists a directed edge $e_{x,y}$ from entity $v_x \in \mathcal{V} \setminus C_j^t$ to some $v_y \in C_j^t$ in $\mathcal{CG}_{\mathcal{L}(F^*)}$. This implies that $e_{x,y} \in \mathcal{CG}_{\mathcal{L}}$ exists in the constraint graph of the original (nonsquare) system $\mathcal{L}$ and thus v_x is in W. Let $\mathcal{A}_{x'} \in \mathcal{S}_x$ be the selected supporter of v_x in F^*. By construction of F^*, $\mathcal{A}_{x'} \notin \mathcal{R}(C_j^t)$, in contradiction to the fact that the edge $e_{x,y} \in \mathcal{CG}_{\mathcal{L}(F^*)}$ exists.

Since there exists a node in $\mathcal{CG}_{\mathcal{L}(F^*)}$ with no incoming neighbors, this graph is not strongly connected, implying that $\mathcal{L}(F^*)$ is reducible.

Finally, as $\mathcal{L}$ is irreducible, it holds that every hidden square system is irreducible, in particular $\mathcal{L}(F^*)$, hence, contradiction. The claim follows. □

Lemma 3.3. *If $\mathcal{L}$ is irreducible then Algorithm Irr_Test($\mathcal{L}$) returns "yes".*

Proof. By Claim 3.2, we have that if $\mathcal{L}$ is irreducible and $|\mathcal{C}^t| > 1$, then every node in D_t has an incoming edge, which necessitates that there exists a (directed) cycle $C = (C_{i_1}, \ldots, C_{i_k})$, for $k \geq 2$ in D^t. Since the nodes in such cycle C are strongly connected, they can be merged in $\mathcal{C}^{t+1}$, and therefore round t is successful. Moreover, since at least two clusters of $\mathcal{C}^t$ are merged into a single cluster in $\mathcal{C}^{t+1}$, we have that $|\mathcal{C}^{t+1}| < |\mathcal{C}^t|$. This means that the merging never fails as long as $|\mathcal{C}^t| > 1$, so $k_t = |\mathcal{C}^t|$ is monotonically decreasing. It follows that the algorithm terminates within at most $n-1$ rounds with a "yes". The lemma follows. □

We now consider a reducible system $\mathcal{L}$ and show that Irr_Test($\mathcal{L}$) returns "no".

Lemma 3.4. *If $\mathcal{L}$ is reducible, then Algorithm Irr_Test($\mathcal{L}$) returns "no".*

Proof. Towards contradiction, assume otherwise, i.e., suppose that the algorithm accepts $\mathcal{L}$. This implies that every round $t \in [1, t^*]$ in which $|\mathcal{C}^t| > 1$ is successful.

The reducibility of $\mathcal{L}$ implies that there exists (at least one) hidden square system $\mathcal{L}(F)$ which is reducible, namely, its constraint graph $\widehat{D} = \mathcal{CG}_{\mathcal{L}(F)}$ is not strongly connected. Thus $\widehat{D}$ contains at least two nodes v_i and v_j that belong to distinct strongly connected components in $\widehat{D}$. Note that v_i and v_j are in distinct clusters in $\mathcal{C}^0$, but belong to the same cluster in the partition of the final $\mathcal{C}^{t^*}$. Therefore, there must exists a round $t' \in (0, t^*)$ in which the cluster $C_{i'}^{t'}$ that contains v_i and the cluster $C_{j'}^{t'}$ that contains v_j appeared in the same strongly connected component in $D_{t'}$ and were merged into a single strongly connected component in $\mathcal{C}^{t'+1}$. (Note that since $t'-1$ is a successful round, $\mathcal{C}^{t'}$ is a partition of the entity set (Property (A1) of Claim 3.1) and therefore $C_{i'}^{t'}$ and $C_{j'}^{t'}$ exist.) Since round t' is successful (otherwise the algorithm would terminates with "no"), by Property (A2) of Claim 3.1, it follows that the entity subset of the unified cluster $C \in \mathcal{C}^{t'+1}$ is in the same connected component in the constraint graph $\mathcal{CG}_{\mathcal{L}(F')}$ for every $F' \in \mathcal{F}$. Since $F \in \mathcal{F}$ as well it holds that v_i and v_j are in the same connected component in $\widehat{D}$. Hence, contradiction. The lemma follows. □

By Lemmas 3.3 and 3.4 it follows that Algorithm Irr_Test($\mathcal{L}$) returns "yes" iff the system $\mathcal{L}$ is irreducible, which establish the correctness of the algorithm.

Claim 3.5. *Algorithm Irr_Test terminates in $O(m \cdot n^2)$ rounds.*

Proof. The algorithm consists of at most $n-1$ rounds. In each round t, it constructs the cluster graph $D_t = (\mathcal{C}^{t-1}, E_t)$ in time $O(n \cdot m)$. The decomposition into strongly connected components can be done in $O(|D_t|) = O(n^2)$. The claim follows. □

Theorem 3.6. *There exists a polynomial time algorithm for deciding irreducibility on nonnegative systems.*

References

[1] C. Avin, M. Borokhovich, Y. Haddad, E. Kantor, Z. Lotker, M. Parter, D. Peleg, Generalized Perron–Frobenius theorem for multiple choice matrices, and applications, in: Proc. SODA, 2013.

[2] C. Avin, M. Borokhovich, Y. Haddad, E. Kantor, Z. Lotker, M. Parter, D. Peleg, Generalized Perron–Frobenius theorem for multiple choice matrices, and applications, http://arxiv.org/pdf/1308.5915.pdf.

[3] G. Boutry, M. Elad, G.H. Golub, P. Milanfar, The generalized eigenvalue problem for nonsquare pencils using a minimal perturbation approach, SIAM J. Matrix Anal. Appl. 27 (2) (2005) 582–601.

[4] D. Chu, G.H. Golub, On a generalized eigenvalue problem for nonsquare pencils, SIAM J. Matrix Anal. Appl. 28 (3) (2006) 770–787.

[5] A. Goldsmith, Wireless Communications, Cambridge University Press, 2005.

[6] D. Kressner, E. Mengi, I. Nakic, N. Truhar, Generalized eigenvalue problems with specified eigenvalues, IMA J. Numer. Anal. (2011), submitted for publication.

[7] O.L. Mangasarian, Perron–Frobenius properties of $Ax = \lambda Bx$, J. Math. Anal. Appl. 36 (1971).

[8] K. Truemper, Matroid Decomposition, Academic Press Inc., Boston, MA, 1992.