

Generalized Perron–Frobenius Theorem for Multiple Choice Matrices, and Applications

Chen Avin ^{*} Michael Borokhovich ^{*} Yoram Haddad ^{†*} Erez Kantor [‡]
Zvi Lotker ^{*} Merav Parter ^{§¶} David Peleg ^{‡||}

October 4, 2012

Abstract

The celebrated Perron–Frobenius (PF) theorem is stated for irreducible nonnegative square matrices, and provides a simple characterization of their eigenvectors and eigenvalues. The importance of this theorem stems from the fact that eigenvalue problems on such matrices arise 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. Despite the extensive development of spectral theories for nonnegative matrices, the applicability of such theories to non-convex optimization problems is not clear. In particular, a natural question is whether the PF Theorem (along with its applications) can be generalized to a nonsquare setting. Our paper provides a generalization of the PF Theorem to nonsquare multiple choice matrices. The extension can be interpreted as representing systems with additional degrees of freedom, where each client entity may choose between multiple servers that can cooperate in serving it (while potentially interfering with other clients). This formulation is motivated by applications to power control in wireless networks, economics and others, all of which extend known examples for the use of the original PF Theorem.

We show that the option of cooperation does not

improve the situation, in the sense that in the optimum solution, no cooperation is needed, and only one server per client entity needs to work. Hence, the additional power of having several potential servers per client translates into choosing the “best” single server and not into sharing the load between the servers in some way, as one might have expected.

The two main contributions of the paper are (i) a generalized PF Theorem that characterizes the optimal solution for a non-convex problem, and (ii) an algorithm for finding the optimal solution in polynomial time.

In addition, we extend the definitions of irreducibility and largest eigenvalue of square matrices to nonsquare ones in a novel and non-trivial way, which turns out to be necessary and sufficient for our generalized theorem to hold.

To characterize the optimal solution, we use techniques from a wide range of areas. In particular, the analysis exploits combinatorial properties of polytopes, graph-theoretic techniques and analytic tools such as spectral properties of nonnegative matrices and root characterization of integer polynomials.

1 Introduction

Motivation and main results. This paper presents a generalization of the well known Perron–Frobenius (PF) Theorem [25, 14], and some of its applications. To begin, let us consider the following motivating example. Power control is one of the most fundamental problems in wireless networks. We are given n receiver-transmitter pairs and their physical locations. All transmitters are set to transmit at the same time with the same frequency, thus causing interference to the other receivers. Therefore, receiving and decoding a message at each receiver depends on the transmitting power of its paired transmitter as well as the power of the rest of the transmitters. If the signal strength received by a device divided by the interfering strength of other simultaneous transmissions is above some *reception threshold*

^{*}Ben Gurion University, Beer-Sheva, Israel. Email: {avin, borokhom, zvilo}@cse.bgu.ac.il, yoram.haddad@gmail.com.

[†]Jerusalem College of Technology, Jerusalem, Israel.

[‡]The Technion, Haifa, Israel. Email: eriz.kantor@gmail.com. Supported by Eshkol fellowship, the Israel Ministry of Science and Technology.

[§]The Weizmann Institute of Science, Rehovot, Israel. Email: {merav.parter, david.peleg}@weizmann.ac.il.

[¶]Recipient of the Google Europe Fellowship in distributed computing; research supported in part by this Google Fellowship.

^{||}Supported in part by the Israel Science Foundation (grant 894/09), the United States-Israel Binational Science Foundation (grant 2008348), the Israel Ministry of Science and Technology (infrastructures grant), and the Citi Foundation.

β , then the receiver successfully receives the message, otherwise it does not [28]. The question of power control is then to find an optimal power assignment for the transmitters, so as to make the reception threshold β as high as possible and ease the decoding process.

As it turns out, this power control problem can be solved elegantly using the Perron–Frobenius (PF) Theorem [37]. The theorem holds for square matrices and can be formulated as dealing with the following optimization problem (where $A \in \mathbb{R}^{n \times n}$):

$$(1.1) \quad \begin{aligned} &\text{maximize } \beta \text{ subject to:} \\ &A \cdot \bar{X} \leq 1/\beta \cdot \bar{X}, \quad \|\bar{X}\|_1 = 1, \quad \bar{X} \geq \bar{0}. \end{aligned}$$

Let β^* denote the optimal solution for Program (1.1). The Perron–Frobenius (PF) Theorem characterizes the solution to this optimization problem:

THEOREM 1.1. (PF THEOREM, SHORT VERSION, [14, 25]) *Let A be an irreducible nonnegative matrix. Then $\beta^* = 1/r$, where $r \in \mathbb{R}_{>0}$ is the largest eigenvalue of A , called the Perron–Frobenius (PF) root of A . There exists a unique (eigen-)vector $\bar{\mathbf{P}} > 0$, $\|\bar{\mathbf{P}}\|_1 = 1$, such that $A \cdot \bar{\mathbf{P}} = r \cdot \bar{\mathbf{P}}$, called the Perron vector of A . (The pair $(r, \bar{\mathbf{P}})$ is hereafter referred to as an eigenpair of A .)*

Returning to our motivating example, let us consider a more complicated variant of the power control problem, where each receiver has several transmitters that can transmit to it (and only to it) synchronously. Since these transmitters are located at different places, it may conceivably be better to divide the power (i.e., work) among them, to increase the reception threshold at their common receiver. Again, the question concerns finding the best power assignment among all transmitters.

In this paper we extend Program (1.1) to nonsquare matrices and consider the following extended optimization problem, which in particular captures the multiple transmitters scenario. (Here $A, B \in \mathbb{R}^{n \times m}$, $n \leq m$.)

$$(1.2) \quad \begin{aligned} &\text{maximize } \beta \text{ subject to:} \\ &A \cdot \bar{X} \leq 1/\beta \cdot B \cdot \bar{X}, \quad \|\bar{X}\|_1 = 1, \quad \bar{X} \geq \bar{0}. \end{aligned}$$

We interpret the nonsquare matrices A, B as representing some additional freedom given to the system designer. In this setting, each *entity* (receiver, in the power control example) has several *effectors* (transmitters, in the example), referred to as its *supporters*, which can cooperate in serving it and share the workload. In such a general setting, we would like to find the best way to organize the cooperation between the “supporters” of each entity.

The original problem was defined for a square matrix, so the rise of eigenvalues seems natural. In contrast, in the generalized setting the situation seems more complex. Our main result is an extension of the PF Theorem to nonsquare matrices and systems.

THEOREM 1.2. (MULTIPLE CHOICE PF THEOREM, SHORT VERSION) *Let $\langle A, B \rangle$ be an irreducible nonnegative system (to be made formal later). Then $\beta^* = 1/r$, where $r \in \mathbb{R}_{>0}$ is the smallest Perron–Frobenius (PF) root of all $n \times n$ square sub-systems (defined formally later). There exists a vector $\bar{\mathbf{P}} \geq 0$ such that $A \cdot \bar{\mathbf{P}} = r \cdot B \cdot \bar{\mathbf{P}}$ and $\bar{\mathbf{P}}$ has n entries greater than 0 and $m - n$ zero entries (referred to as a $\mathbf{0}^*$ solution).*

In other words, we show that the option of cooperation does not improve the situation, in the sense that in the optimum solution, no cooperation is needed and only one supporter per entity needs to work. Hence, the additional power of having several potential supporters per entity translates into choosing the “best” single supporter and not into sharing the load between the supporters in some way, as one might have expected. Actually, the lion’s share of our analysis involves such a characterization of the optimal solution for (the non-convex) problem of Program (1.2). The challenge was to show that at the optimum, there exists a solution in which only one supporter per entity is required to work; we call such a solution a $\mathbf{0}^*$ solution. Namely, the structure that we aim to establish is that the optimal solution for our multiple choices system is in fact the optimal solution of an *embedded* PF system. Indeed, to enjoy the benefits of an equivalent square system, one should show that there exists a solution in which only one supporter per entity is required to work. Interestingly, it is relatively easy to show that there exists an optimal “almost $\mathbf{0}^*$ ” solution, in which each entity *except at most one* has a single active supporter and the remaining entity has at most *two* active supporters. Despite the presumably large “improvement” of decreasing the number of servers from m to $n + 1$, this still leaves us in the frustrating situation of a nonsquare $n \times (n + 1)$ system, where no spectral characterization for optimal solutions exists. In order to allow us to characterize the optimal solution using the eigenpair of the best square matrix embedded within the nonsquare system, one must overcome this last hurdle, and reach the critical (or “phase transition”) point of n servers, in which the system is *square*. Our main efforts went into showing that the remaining entity, too, can select just one supporter while maintaining optimality, ending with a *square* $n \times n$ irreducible system where the traditional PF Theorem can be applied. Proving the existence of an optimal $\mathbf{0}^*$ solution requires techniques from a wide range of areas to come into play and provide a rich understanding of the system on various levels. In particular, the analysis exploits combinatorial properties of polytopes, graph-theoretic techniques and analytic tools such as spectral properties of nonnegative matrices and root characterization of integer polynomials.

For the example of power control in wireless network with multiple transmitters per receiver, a $\mathbf{0}^*$ solution means that the best reception threshold is achieved when only a single transmitter transmits to each receiver. This is illustrated by the SIR (Signal to Interference Ratio) diagram (cf. [2]) in Fig. 3, depicting a system of three receivers and two transmitters per receiver using the optimal value β^* of the system. The figure illustrates that in the optimal $\mathbf{0}^*$ solution for the system, each receiver is covered by one of its transmitters, but other solutions, including the one where all transmitters act simultaneously, may fail to cover the receivers.

Other known applications of PF Theorem can also be extended in a similar manner. Examples for such applications are the *input-output economic model* [26]. In the economy model, each industry produces a commodity and buys commodities from other industries. The percentage profit margin of an industry is the ratio of its total income and total expenses (for buying its raw materials). It is required to find a pricing maximizing the ratio of the total income and total expenses of all industries. The extended PF variant concerns the case where an industry can produce multiple commodities instead of just one. In all of these examples, the same general phenomenon holds: only a single supporter needs to “work” for each entity in the optimum solution, i.e., the optimum is a $\mathbf{0}^*$ solution.

While the original PF root and PF vector can be computed in polynomial time, this is not clear in the extended case, since the problem is not convex [5] (and also not log-convex) and there are exponentially many choices in the system even if we know that the optimal solution is $\mathbf{0}^*$ and each entity (e.g., receiver) has only two supporters (e.g., transmitters) to choose from. Our second main contribution is providing a polynomial time algorithm to find β^* and $\bar{\mathbf{P}}$. The algorithm uses the fact that for a given β we get a relaxed problem which is convex (actually it becomes linear). This allows us to employ the well known interior point method [5], for testing a specific β for feasibility. Hence, the problem reduces to finding the maximum feasible β , and the algorithm does so by applying binary search on β . Clearly, the search results in an approximate solution (in fact yielding an FPTAS for program (1.2)). This, however, leaves open the intriguing question of whether program (1.2) is polynomial. Obtaining an exact optimal β^* , along with an appropriate vector $\bar{\mathbf{P}}$, is thus another challenging aspect of the algorithm, which is successfully solved via an original approach based on the extended PF Theorem, which states that there is an optimal $\mathbf{0}^*$ solution, and proving that the proposed algorithm is polynomial.

A central notion in the generalized PF theorem is the *irreducibility* of the system. While irreducibility is well-defined and known for square systems, it is not clear how to define irreducibility for a nonsquare matrix or system as in Program (1.2). We provide a suitable definition based on the property that every maximal square (legal) subsystem is irreducible, and show that our definition is necessary and sufficient for the theorem to hold. But, since there could be exponentially many such square subsystems, it is not a priori clear if one can check that a nonsquare system is irreducible in polynomial time. We address this issue using what we call the *constraint graph* of the system, whose vertex set is the set on 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. Essentially, although the matrices are not square, the notion of constraint graph is well defined and provide in a way a valuable *square* representation of the non square system (i.e., the adjacency matrix of the graph). Interestingly, we find a polynomial time algorithm for testing irreducibility of the system, which exploits the properties of the constraint graph.

Related work. The PF Theorem establishes the following two important “PF properties” for a nonnegative irreducible square matrix $A \in \mathbb{R}^{n \times n}$: (1) the *Perron–Frobenius property*: A has a maximal nonnegative eigenpair. If in addition, the maximal eigenvalue is strictly positive, dominant and with a strictly positive eigenvector, then the matrix A is said to enjoy the *strong Perron–Frobenius property*. (2) the *Collatz–Wielandt property* (a.k.a. min-max characterization): the maximal eigenpair is the optimal solution of Program (1.1). Matrices with these properties have played an important role in a wide variety of applications. The wide applicability of the PF Theorem, as well as the fact that it is still not fully understood what are the necessary and sufficient properties of a matrix A for the PF properties to hold, have led to the emergence of many generalizations. We note that whereas all generalizations concern the Perron–Frobenius property, the Collatz–Wielandt property is not always established. The long series of existing PF extensions includes [22, 13, 29, 18, 31, 19, 27, 21]. Section 2 discusses these extensions in comparison to the current work. In addition, in Section 8.1 we discuss the existing literature for the wireless power control problem with multiple transmitters.

2 Existing Extensions for the PF Theorem

Existing PF extensions can be broadly classified into four classes. The first concerns matrices that do not follow the irreducibility and nonnegativity requirement. For example, [22, 13] establish the Perron-Frobenius property for *almost* nonnegative matrices or *eventually* nonnegative matrices. A second class of generalizations concerns square matrices over different domains. For example, in [29], the PF Theorem was established for complex matrices $A \in \mathbb{C}^{n \times n}$. In the third type of generalization, the linear transformation obtained by applying the nonnegative irreducible matrix A is generalized to a nonlinear mapping [18, 31], a concave mapping [19] or a matrix polynomial mapping [27].

Last, a much less well studied generalization deals with nonsquare matrices, i.e., matrices in $\mathbb{R}^{n \times m}$ for $m \neq n$. Note that when considering a nonsquare system, the notion of eigenvalues requires definition. There are several possible definitions for eigenvalues in nonsquare matrices. One possible setting for this type of generalizations considers 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 $\lambda \in \mathbb{C}$. Of special interest here are the values that reduce the pencil rank, namely, the λ values satisfying $(A - \lambda B) \cdot \bar{X} = \bar{0}$ for some nonzero \bar{X} . This problem is known as the *generalized eigenvalue problem* [21, 10, 4, 20] which can be stated as follows: Given matrices $A, B \in \mathbb{R}^{n \times m}$, find a vector $\bar{X} \neq \bar{0}$, $\lambda \in \mathbb{C}$, so that $A \cdot \bar{X} = \lambda B \cdot \bar{X}$. The complex number λ is said to be an *eigenvalue of A relative to B* iff $A\bar{X} = \lambda \cdot B \cdot \bar{X}$ for some nonzero \bar{X} and \bar{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, [21] considered pairs of nonsquare matrices A, B and was the first to characterize the relation between A and B required to establish their PF property, i.e., guarantee that the generalized eigenpair is nonnegative. Essentially, this is done by generalizing the notion of positivity and nonnegativity in the following manner. A matrix A is said to be *positive* (respectively, *nonnegative*) with respect to B , if $B^T \cdot \bar{Y} \geq 0$ implies that $A^T \cdot \bar{Y} > 0$ (resp., $A^T \cdot \bar{Y} \geq 0$). Note that for $B = I$, these definitions reduce to the classical definitions of a positive (resp., nonnegative) matrix. Let $A, B \in \mathbb{R}^{n \times m}$, for $n \geq m$, be such that the rank of A or the rank of B is n . It is shown in [21] that if A is positive (resp., nonnegative) with respect to B , then the generalized eigenvalue problem $A \cdot \bar{X} = \lambda \cdot B \cdot \bar{X}$ has a discrete and finite spectrum, the eigenvalue with the largest absolute value is real and positive (resp., nonnegative), and the corresponding eigenvector is positive (resp., nonnegative). Observe that under the defi-

nition used therein, the cases where $m > n$ (which is the setting studied here) is uninteresting, as the columns of $A - \lambda \cdot B$ are linearly dependent for any real λ , and hence the spectrum $sp(A_B)$ is unbounded.

Despite the significance of [21] and its pioneering generalization of the PF Theorem to nonsquare systems, it is not clear what are the applications of such a generalization, and no specific implications are known for the traditional applications of the PF theorem, such as the power-control problem or the economy model. Moreover, although [21] established the PF property for a class of pairs of nonsquare matrices, the Collatz–Wielandt property, which provides the algorithmic power for the PF Theorem, does not necessarily hold with the spectral definition of [21].

In addition, in [21], since no notion of irreducibility was defined, the spectral radius of a nonnegative system (in the sense of the definition of [21]) might be zero, and the corresponding eigenvector might be nonnegative in the strong sense (with some zero coordinates). These degenerations can be handled only by considering irreducible nonnegative matrices, as was done in [14].

The goal of the current work is to develop the spectral theory for a pair of nonnegative matrices in a way that is both necessary and sufficient for both the PF property and the Collatz–Wielandt property to hold (necessary and sufficient in the sense for the nonsquare system to be of the “same power” as the square systems considered by Perron and Frobenius). We consider nonsquare matrices of dimension $n \times m$ for $n \leq m$, which can be interpreted as describing a system with multiple choices (of columns) per entity (row). We define the spectrum of pairs of matrices A and B in a novel manner. We note that although according to [21] the spectrum $sp(A_B)$ is not bounded if $n < m$, with our definition the spectrum is bounded. Interestingly, the maximum eigenvalue of the spectrum we define, is also the maximum of spectrum according to the definition of [21] and therefore we can show that the Collatz–Wielandt property is extended as well. It is important to note that although the generalized eigenvalue problem has been studied for many years, and multiple approaches for nonsquare spectral theory in general have been developed, the algorithmic aspects of such theories with respect to the the Collatz–Wielandt property have been neglected when concerning nonsquare matrices (and also in other extensions). This paper is the first, to the best of our knowledge, to provide spectral definitions for nonsquare systems that have the same algorithmic power as those made for square systems (in the context of PF Theorem). The extended optimization problem that corresponds to this nonsquare setting, is a nonconvex problem (which is also not log-convex),

therefore its polynomial solution and characterization are of interest.

Another way to extend the notion of eigenvalues and eigenvectors of a square matrix to a nonsquare matrix is via *singular value decomposition (SVD)* [23]. Formally, the singular value decomposition of an $n \times m$ real matrix M is a factorization of the form $M = U\Sigma V^*$, where U is an $m \times m$ real or complex unitary matrix, Σ is an $m \times n$ diagonal matrix with nonnegative reals on the diagonal, and V^* (the conjugate transpose of V) is an $n \times n$ real or complex unitary matrix. The diagonal entries $\Sigma_{i,i}$ of Σ are known as the singular values of M . After performing the product $U\Sigma V^*$, it is clear that the dependence of the singular values of M is linear. In case all the inputs of M are positive, we can add the absolute value, and thus the SVD has a flavor of L^1 dependence. In contrast to the SVD definition, here we are interested in finding the maximum, so our interpretation has the flavor of L^∞ .

In a recent paper [34], Vazirani defined the notion of *rational convex programs* as problems that have a rational number as a solution. Our paper can be considered as an example for *algebraic programming*, since we show that a solution to our problem is an algebraic number.

3 Algebraic Preliminaries

Definitions and Terminology. Let $A \in \mathbb{R}^{n \times n}$ be a square matrix. Let $\text{EigVal}(A) = \{\lambda_1, \dots, \lambda_k\}$, $k \leq n$, be the set of (at most n) real eigenvalues of A . The *characteristic polynomial* of A , denoted by $P(A, t)$, is a polynomial whose roots are precisely the eigenvalues of A , $\text{EigVal}(A)$, and it is given by

$$(3.3) \quad P(A, t) = \det(t \cdot I - A)$$

where I is the $n \times n$ identity matrix. Note that $P(A, t) = 0$ iff $t \in \text{EigVal}(A)$. The *spectral radius* of A is defined as $\rho(A) = \max_{\lambda \in \text{EigVal}(A)} |\lambda|$. The i^{th} element of a vector \bar{X} is given by $X(i)$, and the i, j entry of a matrix A is $A(i, j)$. Let $A_{i,0}$ (respectively, $A_{0,i}$) denote the i -th row (resp., column) of A . Vector and matrices inequalities are interpreted in the component-wise sense. A is *positive* (respectively, *nonnegative*) if all its entries are positive. A is *primitive* if there exists a natural number k such that $A^k > 0$. A is *irreducible* if for every i, j , there exists a natural $k_{i,j}$ such that $(A^{k_{i,j}})_{i,j} > 0$. An *irreducible* matrix A is *periodic* with period h if $(A^t)_{ii} = 0$ for $t \neq k \cdot h$.

PF Theorem for nonnegative irreducible matrices. The PF Theorem states the following.

THEOREM 3.1. (PF THEOREM) *Let $A \in \mathbb{R}_{\geq 0}^{n \times n}$ be a nonnegative irreducible matrix with spectral ratio $\rho(A)$.*

Then $\max \text{EigVal}(A) > 0$. There exists an eigenvalue $\lambda \in \text{EigVal}(A)$ such that $\lambda = \rho(A)$. λ is called the Perron–Frobenius (PF) root of A (denoted here by r). The algebraic multiplicity of r is one. There exists an eigenvector $\bar{X} > 0$ such that $A \cdot \bar{X} = r \cdot \bar{X}$. The unique vector $\bar{\mathbf{P}}$ defined by $A \cdot \bar{\mathbf{P}} = r \cdot \bar{\mathbf{P}}$ and $\|\bar{\mathbf{P}}\|_1 = 1$ is called the Perron–Frobenius (PF) vector. There are no nonnegative eigenvectors for A with r except for positive multiples of $\bar{\mathbf{P}}$. If A is a nonnegative irreducible periodic matrix with period h , then A has exactly h eigenvalues equal to $\lambda_j = \rho(A) \cdot \exp^{2\pi i \cdot j/h}$, $j = 1, 2, \dots, h$, and all other eigenvalues of A are of strictly smaller magnitude than $\rho(A)$.

Collatz–Wielandt characterization (the min-max ratio). Collatz and Wielandt [11, 35] established the following formula for the PF root, also known as the min-max ratio characterization.

LEMMA 3.1. [11, 35] $r = \min_{\bar{X} \in \mathcal{N}} f(\bar{X})$ where $f(\bar{X}) = \max_{1 \leq i \leq n, X(i) \neq 0} \left\{ \frac{(A \cdot \bar{X})_i}{X(i)} \right\}$ and $\mathcal{N} = \{\bar{X} \geq \bar{0}, \|\bar{X}\|_1 = 1\}$.

Alternatively, this can be written as the following optimization problem.

$$(3.4) \quad \max \beta \text{ s.t. } A \cdot \bar{X} \leq 1/\beta \cdot \bar{X}, \|\bar{X}\|_1 = 1, \bar{X} \geq \bar{0}.$$

Let β^* be the optimal solution of Program (3.4) and let \bar{X}^* be the corresponding optimal vector. Using the representation of Program (3.4), Lemma 3.1 translates into the following.

THEOREM 3.2. $\beta^* = 1/r$ where $r \in \mathbb{R}_{>0}$ is the maximal eigenvalue of A and \bar{X}^* is given by eigenvector $\bar{\mathbf{P}}$ corresponding for r . Hence at the optimum value β^* , the set of n constraints given by $A \cdot \bar{X}^* \leq 1/\beta^* \cdot \bar{X}^*$ of Program (3.4) holds with equality.

This can be interpreted as follows. Consider the ratio $Y(i) = (A \cdot \bar{X})_i / X(i)$, viewed as the ‘repression factor’ for entity i . The task is to find the input vector \bar{X} that minimizes the maximum repression factor over all i , thus achieving balanced growth. In the same manner, one can characterize the max-min ratio. Again, the optimal value (resp., point) corresponds to the PF eigenvalue (resp., eigenvector) of A . In summary, when taking \bar{X} to be the PF eigenvector, $\bar{\mathbf{P}}$, and $\beta^* = 1/r$, all repression factors are equal, and optimize the max-min and min-max ratios.

4 A Generalized PF Theorem for Nonsquare Systems

System definitions. Our 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 $g(i, j)$. We describe the affector-entity relation by the *supporters gain matrix* $\mathcal{M}^+ \in \mathbb{R}^{n \times m}$

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

and the *repressors gain matrix* $\mathcal{M}^- \in \mathbb{R}^{n \times m}$, given by

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

Again, for clarity we may write $\mathcal{M}^-(v_i, \mathcal{A}_j)$ for $\mathcal{M}^-(i, j)$ (and similarly for \mathcal{M}^+).

We can now formally define a *system* as $\mathcal{L} = \langle \mathcal{M}^+, \mathcal{M}^- \rangle$, where $\mathcal{M}^+, \mathcal{M}^- \in \mathbb{R}_{\geq 0}^{m \times n}$, $n = |\mathcal{V}|$ and $m = |\mathcal{A}|$. We denote the supporter (resp., repressor) set of v_i by

$$\begin{aligned} \mathcal{S}_i(\mathcal{L}) &= \{\mathcal{A}_j \mid \mathcal{M}^+(v_i, \mathcal{A}_j) > 0\} \\ \mathcal{R}_i(\mathcal{L}) &= \{\mathcal{A}_j \mid \mathcal{M}^-(v_i, \mathcal{A}_j) > 0\}. \end{aligned}$$

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 $v_i \in \mathcal{V}$. We classify the systems into three types:

- (a) $\mathfrak{L}^s = \{\mathcal{L} \mid m \leq n, |\mathcal{S}_i| = 1, \text{ for every } v_i \in \mathcal{V}\}$ is the family of *Square Systems (SS)*.
- (b) $\mathfrak{L}^w = \{\mathcal{L} \mid m = n + 1, \exists j \text{ s.t. } |\mathcal{S}_j| = 2 \text{ and } |\mathcal{S}_i| = 1 \text{ for every } v_i \in \mathcal{V} \setminus \{v_j\}\}$ is the family of *Weak Systems (WS)*, and
- (c) $\mathfrak{L}^{MS} = \{\mathcal{L} \mid m > n + 1\}$ is the family of *Multiple Systems (MS)*.

The generalized PF optimization problem. Consider a set of n entities and gain matrices $\mathcal{M}^+, \mathcal{M}^- \in \mathbb{R}^{n \times m}$, for $m \geq n$. The main application of the generalized PF Theorem is the following optimization problem, which is an extension of Program (3.4).

$$(4.5) \quad \max \beta$$

$$(4.6) \quad \text{s.t. } \mathcal{M}^- \cdot \bar{X} \leq 1/\beta \cdot \mathcal{M}^+ \cdot \bar{X}$$

$$(4.7) \quad \bar{X} \geq \bar{0}$$

$$\|\bar{X}\|_1 = 1.$$

We begin with a simple observation. An affector \mathcal{A}_j is *redundant* if $\mathcal{M}^+(v_i, \mathcal{A}_j) = 0$ for every i .

OBSERVATION 4.1. *If \mathcal{A}_j is redundant, then $X(j) = 0$ in any optimal solution \bar{X} .*

In view of Obs. 4.1, we may hereafter restrict attention to the case where there are no redundant affectors in the system, as any redundant affector \mathcal{A}_j can be removed and simply assigned $X(j) = 0$.

We now proceed with some definitions. Let $X(\mathcal{A}_j)$ denote the value of \mathcal{A}_j in \bar{X} . Denote the set of affectors taken to be active in a solution \bar{X} by $NZ(\bar{X}) = \{\mathcal{A}_j \mid X(\mathcal{A}_j) > 0\}$. Let $\beta^*(\mathcal{L})$ denote the optimal value of Program (4.5), i.e., the maximal positive value for which there exists a nonnegative, nonzero vector \bar{X} satisfying the constraints of Program (4.5). When the system \mathcal{L} is clear from the context we may omit it and simply write β^* . A vector $\bar{X}_{\tilde{\beta}}$ is *feasible* for $\tilde{\beta} \in (0, \beta^*]$ if it satisfies all the constraints of Program (4.5) with $\beta = \tilde{\beta}$. A vector \bar{X}^* is *optimal* for \mathcal{L} if it is feasible for $\beta^*(\mathcal{L})$, i.e., $\bar{X}^* = \bar{X}_{\beta^*}$. The system \mathcal{L} is *feasible* for β if $\beta \leq \beta^*(\mathcal{L})$, i.e., there exists a feasible \bar{X}_{β} solution for Program (4.5). For vector \bar{X} , the *total repression* on v_i in \mathcal{L} for a given \bar{X} is $T^-(\bar{X}, \mathcal{L})_i = (\mathcal{M}^- \cdot \bar{X})_i$. Analogously, the *total support* for v_i is $T^+(\bar{X}, \mathcal{L})_i = (\mathcal{M}^+ \cdot \bar{X})_i$. It now follows that \bar{X} is feasible with β iff

$$(4.8) \quad T^-(\bar{X}, \mathcal{L})_i \leq 1/\beta \cdot T^+(\bar{X}, \mathcal{L})_i \text{ for every } i.$$

When \mathcal{L} is clear from context, we may omit it and simply write $T^-(\bar{X})_i$ and $T^+(\bar{X})_i$. As a direct application of the generalized PF theorem, there is an exact polynomial time algorithm for solving Program (4.5) for irreducible systems, as defined next.

Irreducibility of square systems. A square system $\mathcal{L} \in \mathfrak{L}^s$ is *irreducible* iff (a) \mathcal{M}^+ is nonsingular and (b) \mathcal{M}^- is irreducible. Given an irreducible square \mathcal{L} , let

$$Z(\mathcal{L}) = (\mathcal{M}^+)^{-1} \cdot \mathcal{M}^-.$$

Note the following two observations.

OBSERVATION 4.2. (a) *If \mathcal{M}^+ is nonsingular, then $\mathcal{S}_i \cap \mathcal{S}_j = \emptyset$.* (b) *If \mathcal{L} is an irreducible system, then $Z(\mathcal{L})$ is an irreducible matrix as well.*

(For lack of space, some proofs are deferred to the full version.) Throughout, when considering square systems, it is convenient to assume that the entities and affectors are ordered in such a way that \mathcal{M}^+ is a diagonal matrix, i.e., in \mathcal{M}^+ (and \mathcal{M}^-) the i^{th} column corresponds to $\mathcal{A}_k \in \mathcal{S}_i$, the unique supporter of v_i .

Selection matrices and irreducibility of non-square systems. To define a notion of irreducibility for a nonsquare system $\mathcal{L} \notin \mathfrak{L}^s$, 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 $F(j, i) = 1$. 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 there are no redundant affectors, 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

$$(4.9) \quad \mathcal{F}(\mathcal{L}) = \{F \mid F \text{ is legal for } \mathcal{L}\}.$$

When \mathcal{L} is clear from context, we simply write \mathcal{F} . Let $\mathcal{L}(F)$ be the square system corresponding to the legal selection matrix F , namely, $\mathcal{L}(F) = \langle \mathcal{M}^+ \cdot F, \mathcal{M}^- \cdot F \rangle$.

OBSERVATION 4.3. (a) $\mathcal{L}(F) \in \mathfrak{L}^s$ for every $F \in \mathcal{F}$. (b) $\beta^*(\mathcal{L}) \geq \beta^*(\mathcal{L}(F))$ for every selection $F \in \mathcal{F}$.

We are now ready to define the notion of irreducibility for nonsquare systems: A nonsquare system \mathcal{L} is *irreducible* iff $\mathcal{L}(F)$ is irreducible for every selection matrix $F \in \mathcal{F}$. Note that this condition is the “minimal” *necessary* condition for our theorem to hold, as explained next. Our theorem states that the optimum solution for the nonsquare system is the optimum solution for the best *embedded* square system. It is easy to see that for any nonsquare system $\mathcal{L} = \langle \mathcal{M}^+, \mathcal{M}^- \rangle$, one can increase or decrease any entry $g(i, j)$ in the matrices, while maintaining the sign of each entry in the matrices, such that a particular selection matrix $F^* \in \mathcal{F}$ would correspond to the optimal square system. With an optimal embedded square system at hand, which is also guaranteed to be irreducible (by the definition of irreducible nonsquare systems), our theorem can then apply the traditional PF Theorem, where a spectral characterization for the solution of Program (3.4) exists. Note that irreducibility is a *structural* property of the system, in the sense that it does not depend on the exact gain values, but rather on the sign of the gains, i.e., to determine irreducibility, it is sufficient to observe the binary matrices $\mathcal{M}_B^+, \mathcal{M}_B^-$, treating $g(i, j) > 0$ (resp. $g(i, j) < 0$) as 1 (resp., -1). On the other hand, deciding which of the embedded square system has the maximal eigenvalue (and hence is optimal), depends on the *precise* values of the entries of these matrices. It is therefore necessary that the structural property of irreducibility would hold for any specification of gain values (while maintaining the binary representation of $\mathcal{M}_B^+, \mathcal{M}_B^-$). Indeed, consider a reducible nonsquare system, for which there exists an embedded square system $\mathcal{L}(F)$ that is reducible. It is not hard to see that there exists a specification of gain values that would render this square system $\mathcal{L}(F)$ optimal (i.e., with the maximal eigenvalue among all other embedded square systems). But since $\mathcal{L}(F)$ is reducible, the PF Theorem cannot be applied, and in particular, the corresponding eigenvector is no longer guaranteed to be *positive*.

COROLLARY 4.1. In an irreducible system \mathcal{L} , $\mathcal{S}_i \cap \mathcal{S}_j = \emptyset$ for every v_i, v_j .

PF Theorem for nonnegative irreducible systems. Recall that the root of a square system $\mathcal{L} \in \mathfrak{L}^s$ is $r(\mathcal{L}) = \max \{EigVal(Z(\mathcal{L}))\}$. $\bar{\mathbf{P}}(\mathcal{L})$ is eigenvector of $Z(\mathcal{L})$ corresponding to $r(\mathcal{L})$. We now turn to define the *generalized Perron–Frobenius (PF) root* of a nonsquare system $\mathcal{L} \notin \mathfrak{L}^s$, which is given by

$$(4.10) \quad r(\mathcal{L}) = \min_{F \in \mathcal{F}} \{r(\mathcal{L}(F))\}.$$

Let F^* be the selection matrix that achieves the minimum in Eq. (4.10). We now describe the corresponding eigenvector $\bar{\mathbf{P}}(\mathcal{L})$. Note that $\bar{\mathbf{P}}(\mathcal{L}) \in \mathbb{R}^m$, whereas $\bar{\mathbf{P}}(\mathcal{L}(F^*)) \in \mathbb{R}^n$.

Consider $\bar{X}' = \bar{\mathbf{P}}(\mathcal{L}(F^*))$ and let $\bar{\mathbf{P}}(\mathcal{L}) = \bar{X}$, where

$$(4.11) \quad X(\mathcal{A}_j) = \begin{cases} X'(\mathcal{A}_j), & \text{if } \sum_{i=1}^n F^*(j, i) > 0; \\ 0, & \text{otherwise.} \end{cases}$$

We next state our main result, which is a generalized variant of the PF Theorem for every nonnegative irreducible system.

THEOREM 4.1. Let \mathcal{L} be an irreducible and nonnegative system. Then

$$(Q1) \quad r(\mathcal{L}) > 0,$$

$$(Q2) \quad \bar{\mathbf{P}}(\mathcal{L}) \geq 0,$$

$$(Q3) \quad |NZ(\bar{\mathbf{P}}(\mathcal{L}))| = n,$$

$$(Q4) \quad \bar{\mathbf{P}}(\mathcal{L}) \text{ is not unique.}$$

$$(Q5) \quad \text{The generalized Perron root of } \mathcal{L} \text{ satisfies } r = \min_{\bar{X} \in \mathcal{N}} \{f(\bar{X})\}, \text{ where}$$

$$f(\bar{X}) = \max_{1 \leq i \leq n, (\mathcal{M}^+ \cdot \bar{X})_i \neq 0} \left\{ \frac{(\mathcal{M}^- \cdot \bar{X})_i}{(\mathcal{M}^+ \cdot \bar{X})_i} \right\} \text{ and}$$

$\mathcal{N} = \{\bar{X} \geq 0, \|\bar{X}\|_1 = 1, \mathcal{M}^+ \cdot \bar{X} \neq 0\}$. I.e., the Perron–Frobenius (PF) eigenvalue is $1/\beta^*$ where β^* is the optimal value of Program (4.5), and the PF eigenvalue is the corresponding optimal point. Hence, at the optimum value β^* , the set of n constraints of Eq. (4.6) hold with equality.

5 Proof of the Generalized PF Theorem

We first discuss a natural approach one may consider for proving Thm. 4.1 in general and solving Program (4.5) in particular, and explain this approach fails in this case.

The difficulty. A common approach is to turn a non-convex program into an equivalent convex one by performing a standard variable exchange. An example for a program that’s amenable to this technique is Program (3.4) which is *log-convex* (see Lemma 5.1a),

namely, it becomes convex after replacing terms $X(i)$ with new variables $\widehat{X}(i)$. Unfortunately, in contrast to Program (3.4), its generalization, namely, Program (4.5), is not log-convex (see Lemma 5.1b) and hence cannot be transformed into a convex one in this manner.

Our main efforts in the paper went into showing that at the optimum point, the system loses one degree of freedom, hence guaranteeing the existence of an optimal solution.

5.1 Proof overview Our main challenge is to show that the optimal value of Program (4.5) is related to an *eigenvalue* of some hidden square system \mathcal{L}^* in \mathcal{L} (where “hidden” implies that there is a selection on \mathcal{L} that yields \mathcal{L}^*). The flow of the analysis is as follows. We first provide a graph theoretic characterization of irreducible systems. In particular, we introduce the notion of constraint graph and discuss its properties. We then consider a convex relaxation of Program (4.5) and show that the set of feasible solutions of Program (4.5), for every $\beta \in (0, \beta^*]$, corresponds to a bounded polytope. Moreover, we show that for irreducible systems, the vertex set of such a polytope corresponds to a hidden *weak* system $\mathcal{L}^* \in \mathcal{L}^w$. That is, there exists an hidden weak system in \mathcal{L} that achieves β^* . Note that a solution for such a hidden system can be extended to a solution \overline{X}^* for \mathcal{L} simply by setting the entries of the non-selected effectors to zero in \overline{X}^* .

Next, we exploit a generalization of Cramer’s rule for homogeneous linear systems as well as a separation theorem for nonnegative matrices to show that there is a hidden optimal *square* system in \mathcal{L} that achieves β^* , which establishes the lion’s share of the theorem.

A surprising conclusion of our generalized theorem is that although the given system of matrices is not square, and eigenvalues cannot be straightforwardly defined for it, the nonsquare system contains a *hidden optimal* square system, optimal in the sense that a solution for this system can be translated into a solution to the original system (simply by putting zeros for non-selected effectors) and to satisfy Program (4.5) with the optimal value β^* . The power of nonsquare systems is thus not in the ability to create a solution better than *any* hidden square system it possesses, but rather in the *option* to select the best possible hidden square system out of the optionally exponential many ones.

5.2 Tools

The constraint graph. We begin by providing a graph theoretic characterization of irreducible systems. We define two versions of a (directed) *constraint graph* for system \mathcal{L} . Let $\mathcal{CG}_{\mathcal{L}}(V, E)$ be the *constraint graph* for system \mathcal{L} , defined as follows: $V = \mathcal{V}$, and the rule for a

directed edge $e_{i,j}$ from v_i to v_j is

$$(5.12) \quad e_{i,j} \in E \quad \text{iff} \quad \mathcal{S}_i \cap \mathcal{R}_j \neq \emptyset.$$

Let $\mathcal{SCG}_{\mathcal{L}}(V, \widehat{E})$ be the *strong constraint graph* for system \mathcal{L} defined as follows: $V = \mathcal{V}$, and the directed edge $\widehat{e}_{i,j}$ from v_i to v_j is defined by

$$(5.13) \quad \widehat{e}_{i,j} \in \widehat{E} \quad \text{iff} \quad \mathcal{S}_i \subseteq \mathcal{R}_j.$$

Note that the main difference between $\mathcal{SCG}_{\mathcal{L}}(V, E)$ and $\mathcal{CG}_{\mathcal{L}}(V, E)$ is that $\mathcal{SCG}_{\mathcal{L}} = \mathcal{SCG}_{\mathcal{L}(F)}$ for every selection $F \in \mathcal{F}$. However, for the constraint graph $\mathcal{CG}_{\mathcal{L}}(V, E)$, it is possible that $\mathcal{CG}_{\mathcal{L}} \not\subseteq \mathcal{CG}_{\mathcal{L}(F)}$ for some $F \in \mathcal{F}$. A graph $\mathcal{CG}_{\mathcal{L}}(V, E)$ is *robustly strongly connected* if $\mathcal{CG}_{\mathcal{L}(F)}(V, E)$ is strongly connected for every $F \in \mathcal{F}$.

OBSERVATION 5.1. *Let \mathcal{L} be an irreducible system.*

- (a) *If \mathcal{L} is square, then $\mathcal{SCG}_{\mathcal{L}}(V, E) = \mathcal{CG}_{\mathcal{L}}(V, E)$ and $\mathcal{CG}_{\mathcal{L}}(V, E)$ is strongly connected.*
- (b) *If \mathcal{L} is nonsquare, then $\mathcal{CG}_{\mathcal{L}}(V, E)$ is robustly strongly connected.*

Strongly irreducible systems. The strong constraint graph $\mathcal{SCG}_{\mathcal{L}}(V, E)$ can be used to define a stronger notion of irreducibility. A system \mathcal{L} is *strongly irreducible* iff $\mathcal{S}_i \cap \mathcal{S}_j = \emptyset$, for every $i, j \in [1, n]$, and $\mathcal{SCG}_{\mathcal{L}}(V, E)$ is strongly connected. The following properties are satisfied by strongly irreducible system.

OBSERVATION 5.2. *Let \mathcal{L} be strongly irreducible. Then (a) \mathcal{L} is irreducible; (b) The matrix $\mathcal{M}^+ \cdot (\mathcal{M}^-)^T$ is irreducible.*

It is important to note that the irreducibility of \mathcal{L} does not imply that $\mathcal{SCG}_{\mathcal{L}}(V, E)$ is strongly connected. We now describe an example of a system that is irreducible but not strongly irreducible. Consider a system $\mathcal{L} = \langle \mathcal{M}^+, \mathcal{M}^- \rangle$ of three entities and four effectors, where

$$\mathcal{M}^- = \begin{pmatrix} 0 & 1 & 1 & 0 \\ 1 & 0 & 0 & 1 \\ 1 & 0 & 0 & 0 \end{pmatrix} \quad \text{and} \quad \mathcal{M}^+ = \begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & 1 \end{pmatrix}.$$

We have $\mathcal{S}_1 = \{\mathcal{A}_1\}$, $\mathcal{S}_2 = \{\mathcal{A}_2\}$ and $\mathcal{S}_3 = \{\mathcal{A}_3, \mathcal{A}_4\}$; and $\mathcal{R}_1 = \{\mathcal{A}_2, \mathcal{A}_3\}$, $\mathcal{R}_2 = \{\mathcal{A}_1, \mathcal{A}_4\}$ and $\mathcal{R}_3 = \{\mathcal{A}_1\}$. There are two complete selections, $\mathbf{S}_1 = \{\mathcal{A}_1, \mathcal{A}_2, \mathcal{A}_3\}$ and $\mathbf{S}_2 = \{\mathcal{A}_1, \mathcal{A}_2, \mathcal{A}_4\}$. Both systems, $\mathcal{L}(\mathbf{S}_1)$ and $\mathcal{L}(\mathbf{S}_2)$, are irreducible (see the constraint graphs in Figures 1(b) and 1(c), respectively). However, the system is not strongly irreducible (see the constraint graph in Figure 1(a)).

Hence, our definition for irreducibility is *less stringent* than the requirement that the strong constraint graph is strongly connected. For example, if the matrix $\mathcal{M}^+ + \mathcal{M}^-$ is *positive* and the supporter sets




Figure 1: (a) The constraint graph $\mathcal{CG}_{\mathcal{L}}$, which is *not* strongly connected, hence also *not* strongly irreducible; (b) the constraint graph of $\mathcal{L}(\mathbf{S}_1)$, which is strongly connected, hence irreducible. (c) the constraint graph of $\mathcal{L}(\mathbf{S}_2)$, which is also strongly connected and irreducible.

$\mathcal{S}_1(\mathcal{L}), \dots, \mathcal{S}_n(\mathcal{L})$ are disjoint, then the system is irreducible. But in fact, much less is required to establish irreducibility.

Finally, we provide a poly-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 straightforwardly, e.g., by checking that the directed graph corresponding the matrix $Z(\mathcal{L})$ is strongly connected. However, recall that a nonsquare system \mathcal{L} is irreducible iff every hidden square system $\mathcal{L}(F)$, $F \in \mathcal{F}$, is irreducible, i.e., the matrix $Z(\mathcal{L}(F))$ is irreducible for every $F \in \mathcal{F}$. Since \mathcal{F} might be exponentially large, a brute-force testing of $\mathcal{L}(F)$ for every F is too costly.

LEMMA 5.1. *There exists a polynomial time algorithm for deciding irreducibility on nonnegative systems.*

Partial Selection for Irreducible System. We consider an irreducible system, where $\mathcal{S}_i \cap \mathcal{S}_j = \emptyset$, for every $v_i, v_j \in \mathcal{V}$. Let $\mathbf{S}' \subseteq \mathcal{A}$. We say that \mathbf{S}' is a *partial selection*, if there exists a subset of entities $V' \subseteq \mathcal{V}$ such that

- (a) $|\mathbf{S}'| = |V'|$, and
- (b) for every $v_i \in V'$, $|\mathcal{S}_i \cap V'| = 1$.

That is, every entity in V' has a single representative supporter in \mathbf{S}' . In the system $\mathcal{L}(\mathbf{S}')$ the supporters \mathcal{A}_k of any $v_i \in V'$ that were not selected by v_i , i.e., $\mathcal{A}_k \notin \mathbf{S}' \cap \mathcal{S}_i$, are discarded. In other words, system's affectors set consists of the selected supporters \mathbf{S}' , and the supporters of entities that have not made up their selection in \mathbf{S}' . Formally, the set of the affectors in $\mathcal{L}(\mathbf{S}')$ is given by $\mathcal{A}(\mathcal{L}(\mathbf{S}')) = \mathbf{S}' \cup \bigcup_{v_i \mid \mathcal{S}_i \cap \mathbf{S}' = \emptyset} \mathcal{S}_i$. The number of affectors in $\mathcal{L}(\mathbf{S}')$ is denoted by $m(\mathbf{S}') = |\mathcal{A}(\mathcal{L}(\mathbf{S}'))|$. We now turn to describe $\mathcal{L}(\mathbf{S}')$ formally. Without loss of generality, we consider an ordering on the affectors $\mathcal{A}_1, \dots, \mathcal{A}_m$ such that the i^{th} column of the matrices $\mathcal{M}^+, \mathcal{M}^-$ corresponds to \mathcal{A}_i . Let $\text{ind}(\mathcal{A}_i) = i - |\{\mathcal{A}_j \notin \mathcal{A}(\mathcal{L}(\mathbf{S}')), j \leq i - 1\}|$ be the index of the affector \mathcal{A}_i in the new system, $\mathcal{L}(\mathbf{S}')$ (i.e., the $\text{ind}(\mathcal{A}_i)^{\text{th}}$ column in $\mathcal{M}^+, \mathcal{M}^-$ corresponds to \mathcal{A}_i). Define $F(\mathbf{S}') \in$

$\{0, 1\}^{m \times m(\mathbf{S}')}$ such that $F(\mathbf{S}')_{i, \text{ind}(\mathcal{A}_i)} = 1$ for every $\mathcal{A}_i \in \mathcal{A}(\mathcal{L}(\mathbf{S}'))$, and $F(\mathbf{S}')_{i, j} = 0$ otherwise. Finally, let $\mathcal{L}(\mathbf{S}') = \langle \mathcal{M}^+(\mathbf{S}'), \mathcal{M}^-(\mathbf{S}') \rangle$, where $\mathcal{M}^+(\mathbf{S}') = \mathcal{M}^+ \cdot F(\mathbf{S}')$ and $\mathcal{M}^-(\mathbf{S}') = \mathcal{M}^- \cdot F(\mathbf{S}')$. Note that $\mathcal{M}^+(\mathbf{S}'), \mathcal{M}^-(\mathbf{S}') \in \mathbb{R}^{n \times m(\mathbf{S}')}$. Observe that if the selection \mathbf{S}' is a complete legal selection, then $|\mathbf{S}'| = n$ and the system $\mathcal{L}(\mathbf{S}')$ is a square system. In sum, we have two equivalent representations for square systems in the nonsquare system \mathcal{L} ;

- (a) by specifying a complete selection \mathbf{S} , $|\mathbf{S}| = n$, and
- (b) by specifying the selection matrix, $F \in \mathcal{F}$.

(a) and (b) are equivalent in the sense that the two square systems $\mathcal{L}(F(\mathbf{S}))$ and $\mathcal{L}(\mathbf{S})$ are the same. We now show that if the system \mathcal{L} is irreducible, then so must be any $\mathcal{L}(\mathbf{S}')$, for any partial selection \mathbf{S}' .

OBSERVATION 5.3. *For an irreducible system \mathcal{L} , $\mathcal{L}(\mathbf{S}')$ is also irreducible, for every partial selection \mathbf{S}' .*

Agreement of partial selections. Let $\mathbf{S}_1, \mathbf{S}_2 \subseteq \mathcal{A}$ be partial selections for $V_1, V_2 \subseteq \mathcal{V}$ respectively. Then we denote by $\mathbf{S}_1 \sim \mathbf{S}_2$, the property that the partial selections *agree*, meaning that $\mathbf{S}_1 \cap \mathcal{S}_i = \mathbf{S}_2 \cap \mathcal{S}_i$ for every $v_j \in V_1 \cap V_2$.

OBSERVATION 5.4. *Let V_1, V_2, V_3 with selections $\mathbf{S}_1, \mathbf{S}_2, \mathbf{S}_3$ such that $V_1 \subset V_2$ ($|V_2| > |V_1|$), $\mathbf{S}_1 \sim \mathbf{S}_2$ and $\mathbf{S}_2 \sim \mathbf{S}_3$. Then $\mathbf{S}_3 \sim \mathbf{S}_1$.*

Proof. \mathbf{S}_2 is more restrictive than \mathbf{S}_1 since it defines a selection for a strictly larger set of entities. Therefore no partial selection \mathbf{S}_3 that agrees with \mathbf{S}_2 agrees also with \mathbf{S}_1 . \square

5.3 The Geometry of the Generalized PF Theorem

We now turn to characterize the feasible solutions of Program (4.5). We begin by classifying the $m + n$ linear inequality constraints. The program consists of

- (1) SR (Support-Repression) Constraints: the n constraints of Eq. (4.6).
- (2) Nonnegativity Constraints: the m constraints of Eq. (4.7).

For vector $\bar{X} = (X(1), \dots, X(m))$ and $\alpha \in \mathbb{R}$, let $\bar{X}^\alpha = (X(1)^\alpha, \dots, X(m)^\alpha)$. An optimization program Π is *log convex* if given two feasible solutions \bar{X}_1, \bar{X}_2 for Π , their log convex combination $\bar{X}_\delta = \bar{X}_1^\delta \cdot \bar{X}_2^{(1-\delta)}$ (where “ \cdot ” represents entry-wise multiplication) is also a solution for Π , for every $\delta \in [0, 1]$. In the following we ignore the constraint $\|\bar{X}\|_1 = 1$ since we only validate the feasibility of nonzero nonnegative vectors; the constraint can be established afterwards by normalization.

CLAIM 5.1. (a) Program (3.4) is log-convex (without the $\|\bar{X}\|_1 = 1$ constraint).
 (b) Program (4.5) is not log-convex (even without the $\|\bar{X}\|_1 = 1$ constraint).

Note that log-convexity of Program (3.4) implies that by changing variables it can be solved by convex optimization techniques (see [32] for more information). However, Program (4.5) is not log-convex.

We now turn to consider a *convex relaxation* of Program (4.5). Essentially, the convex relaxation is no longer an optimization problem for β , but rather is given β as input.

$$\begin{aligned} (5.14) \quad & \min 1 \\ (5.15) \quad & \text{s.t. } \mathcal{M}^- \cdot \bar{X} \leq 1/\beta \cdot \mathcal{M}^+ \cdot \bar{X} \\ (5.16) \quad & \bar{X} \geq \bar{0} \\ (5.17) \quad & \|\bar{X}\|_1 = 1. \end{aligned}$$

Note that Program (5.14) has the same set of constraints as those of Program (4.5). However, due to the fact that β is no longer a variable, we get the following.

CLAIM 5.2. Program (5.14) is convex.

It is worth noting at this point, that using the above convex relaxation, one may apply a binary search for finding a *near-optimal* solution for Program (5.14), up to any predefined accuracy. In contrast, our approach, which is based on exploiting the special geometric characteristics of the optimal solution, enjoys the theoretically pleasing (and mathematically interesting) advantage of leading to an efficient algorithm for computing the optimal solution precisely.

Throughout, we restrict attention to values of $\beta \in (0, \beta^*]$. Let $\mathcal{P}(\beta)$ be the polyhedra corresponding to Program (5.14) and denote by $V(\mathcal{P}(\beta))$ the set of vertices of $\mathcal{P}(\beta)$. The following characterization holds even for reducible systems.

CLAIM 5.3. (a) $\mathcal{P}(\beta)$ is bounded (or a polytope). (b) For every $\bar{X} \in V(\mathcal{P}(\beta))$, $|NZ(\bar{X})| \leq n + 1$.

Proof. Part (a) holds by the Equality constraint (5.17) which enforces $\|\bar{X}\|_1 = 1$. We now prove Part (b). Every vertex $\bar{X} \in \mathbb{R}^m$ is defined by a set of m linearly independent equalities. Recall that one equality is imposed by the constraint $\|\bar{X}\|_1 = 1$ (Eq. (5.17)). Therefore it remains to assign $m - 1$ linearly independent equalities out of the $n + m$ (possibly dependent) inequalities of Program (5.14). Hence even if all the (at most n) linearly independent SR constraints (5.15) become equalities, we are still left with at least $m - 1 - n$ unassigned equalities, which must be taken

from the remaining m Nonnegativity Constraints (5.16). Hence, at most $n + 1$ Nonnegativity Inequalities were not fixed to zero, which establishes the proof. \square

Handling the last mile. It remains to handle the last step in the case where, in addition, \mathcal{L} is irreducible. In this case, a more delicate characterization of $V(\mathcal{P}(\beta))$ can be deduced, allowing us to make the last remaining step towards Theorem 4.1.

We begin with some definitions. A solution \bar{X} is called a **$\mathbf{0}^f$ solution** if it is a feasible solution $\bar{X}_{\tilde{\beta}}$, $\tilde{\beta} \in (0, \beta^*]$, in which for each $v_i \in \mathcal{V}$ only one affector has a non-zero assignment, i.e., $NZ(\bar{X}) \cap \mathcal{S}_i = 1$ for every i . A solution \bar{X} is called a **$\mathbf{w}\mathbf{0}^f$ solution**, or a “weak” **$\mathbf{0}^f$ solution**, if it is a feasible vector $\bar{X}_{\tilde{\beta}}$, $\tilde{\beta} \in (0, \beta^*]$, in which for each v_i , *except at most one*, say $v_j \in \mathcal{V}$, $|NZ(\bar{X}) \cap \mathcal{S}_i| = 1$, $v_i \in \mathcal{V} \setminus \{v_j\}$ and $|NZ(\bar{X}) \cap \mathcal{S}_j| = 2$. A solution \bar{X} is called a **$\mathbf{0}^*$ solution** if it is an optimal **$\mathbf{0}^f$ solution**. Let $\mathbf{w}\mathbf{0}^*$ be an optimal **$\mathbf{w}\mathbf{0}^f$ solution**. The following claim holds for every feasible solution of Program (5.14).

CLAIM 5.4. Let \mathcal{L} be an irreducible system with a feasible solution \bar{X}_β , then for every entity v_i there exists an affector $\mathcal{A}_{k_i} \in \mathcal{S}_i$ such that $X_\beta(\mathcal{A}_{k_i}) > 0$, or in other words, $\mathcal{S}_i \cap NZ(\bar{X}_\beta) \neq \emptyset$.

Proof. For clarity of presentation, we begin by considering the case where the system \mathcal{L} is a strongly irreducible system, and in the full version of the paper extend the proof to any irreducible system. Note that by Eq. (5.16) and (5.17), any feasible solution \bar{X}_β satisfies $\bar{X}_\beta \geq 0$, and $\|\bar{X}_\beta\|_1 = 1$. It therefore follows that there exists at least one affector \mathcal{A}_{k_p} such that $X_\beta(\mathcal{A}_{k_p}) > 0$. If $\mathcal{A}_{k_p} \in \mathcal{S}_i$, then we are done. Otherwise, let v_p be such that $\mathcal{A}_{k_p} \in \mathcal{S}_p$ (since no affector is redundant, such v_p must exist). Let $D = SCG_{\mathcal{L}}$, and let $BFS(D, v_p)$ be the BFS tree of D rooted at v_p . Define $L_\ell(D) = LAYER_\ell(BFS(D, v_p))$ to be the ℓ^{th} level of $BFS(D, v_p)$. Formally, $L_\ell(D) = \{v_k \mid d(v_p, v_k) = \ell\}$. Let d_0 be the *depth* of D . We prove by induction on the level ℓ that the claim holds for every $v_i \in L_\ell$. For the base of the induction, consider $L_0(D) = \{v_p\}$. By choice of \mathcal{A}_{k_p} and v_p , $X_\beta(\mathcal{A}_{k_p}) > 0$ and $\mathcal{A}_{k_p} \in \mathcal{S}_p$. Next, assume the claim holds for every level $L_q(D)$, for $q \leq \ell - 1$, and consider $L_\ell(D)$, for $\ell \leq d_0$. By the inductive hypothesis, for every $v_j \in L_{\ell-1}$, $\mathcal{S}_j \cap NZ(\bar{X}_\beta) \neq \emptyset$, that is, there exists an affector $\mathcal{A}_{k_j} \in \mathcal{S}_i$ such that $X_\beta(\mathcal{A}_{k_j}) > 0$. By definition of the graph D , for every $v_i \in L_\ell(D)$ there exists a predecessor $v_j \in L_{\ell-1}(D)$ such that $\mathcal{S}_j \subseteq \mathcal{R}_i$. It therefore follows that $\mathcal{A}_{k_j} \in \mathcal{R}_i$, hence the total repression on v_i satisfies $T^-(\bar{X}_\beta)_i > 0$. By Eq. (4.8), as \bar{X}_β is a feasible solution, it also holds that $T^+(\bar{X}_\beta)_i > 0$. Hence there must exist an affector $\mathcal{A}_{k_i} \in \mathcal{S}_i$ such that

$X_\beta(\mathcal{A}_{k_i}) > 0$, or, $\mathcal{S}_i \cap NZ(\bar{X}_\beta) \neq \emptyset$, as required. This completes the proof of the claim for strongly irreducible systems. The proof for any irreducible system \mathcal{L} is deferred to the full version. \square

We end this section by showing that every vertex $\bar{X} \in V(\mathcal{P}(\beta))$ is a $\mathbf{w}\mathbf{0}^f$ solution.

CLAIM 5.5. *If the system of Program (5.14) is irreducible, then every $\bar{X} \in V(\mathcal{P}(\beta))$ is a $\mathbf{w}\mathbf{0}^f$ solution.*

Proof. By Claim 5.3, for every $\bar{X} \in V(\mathcal{P}(\beta))$, $|NZ(\bar{X})| \leq n + 1$. By Claim 5.4, for every i , $|NZ(\bar{X}) \cap \mathcal{S}_i| \geq 1$. Therefore there exists at most one entity v_i such that $|NZ(\bar{X}) \cap \mathcal{S}_i| = 2$, i.e., the solution is $\mathbf{w}\mathbf{0}^f$. \square

$\mathbf{0}^*$ solutions. In the previous section we established the fact that every vertex $\bar{X} \in V(\mathcal{P}(\beta))$ corresponds to an $\mathbf{w}\mathbf{0}^f$ solution. In particular, this statement holds for $\beta = \beta^*$. By the feasibility of the system for β^* , the corresponding polytope is non-empty and bounded (and each of its vertices is a $\mathbf{w}\mathbf{0}^*$ solution), hence there exist $\mathbf{w}\mathbf{0}^*$ solutions for the problem. The goal of this subsection is to establish the existence of a $\mathbf{0}^*$ solution for the problem. In particular, we show that every optimal $\bar{X} \in V(\mathcal{P}(\beta^*))$ solution is in fact a $\mathbf{0}^*$ solution.

Throughout we consider Program (5.14) for $\beta = \beta^*$, i.e., the optimal value of Program (4.5). We begin by showing that for β^* , the set of n SR Inequalities (Eq. (5.15)) corresponding to $\mathcal{M}^- \cdot \bar{X}^* = 1/\beta^* \cdot \mathcal{M}^+ \cdot \bar{X}^*$ hold with equality for every optimal solution \bar{X}^* , including an \bar{X}^* that is not a $\mathbf{w}\mathbf{0}^*$ solution.

LEMMA 5.2. *If $\mathcal{L} = \langle \mathcal{M}^+, \mathcal{M}^- \rangle$, then $\mathcal{M}^- \cdot \bar{X}^* = 1/\beta^*(\mathcal{L}) \cdot \mathcal{M}^+ \cdot \bar{X}^*$, for every optimal solution \bar{X}^* .*

Proof. By Claim 5.4, every entity v_i has at least one supporter in $NZ(\bar{X}^*)$. Select, for every i , one such supporter $\mathcal{A}_{k_i} \in \mathcal{S}_i \cap NZ(\bar{X}^*)$. Let $\mathbf{S}^* = \{\mathcal{A}_{k_i} \mid 1 \leq i \leq n\}$. By definition, $\mathbf{S}^* \subseteq NZ(\bar{X}^*)$, and since the sets \mathcal{S}_i are disjoint, \mathbf{S}^* is a complete selection (i.e, for every v_i , $|\mathcal{S}_i \cap \mathbf{S}^*| = 1$). Therefore $\mathcal{L}^* = \mathcal{L}(\mathbf{S}^*)$ is a square irreducible system. Let $D^* = \mathcal{CG}_{\mathcal{L}^*}$ be the constraint graph of \mathcal{L}^* . By definition, D^* is strongly connected, and in addition, every edge $e(v_i, v_j) \in E(D^*)$ corresponds to an active effector in \bar{X}^* , i.e., $\mathcal{S}_i \cap \mathcal{R}_j \cap NZ(\bar{X}^*) \neq \emptyset$. To prove this lemma, we establish the existence of a spanning subgraph of the constraint graph, that has the following properties: (a) it is irreducible (strongly connected), (b) every directed edge in this graph is “explained” by an active supporter in \bar{X}^* , where by ”explained”, we mean that there exists an active effector that can be associated with

the directed edge, so the edge remains even if the set of effectors considered is the set of all effectors with positive entry in \bar{X}^* . In other words, even considering only the active effectors in \bar{X}^* (and discarding the others), the constraint graph D^* is guaranteed to be strongly connected (due to Claim 5.4), and moreover, every directed edge is due to some effector with positive entry in \bar{X}^* . Therefore, if we consider an edge (u, v) in D^* by reducing the power of the active supporter of u which, by the definition of D^* , is a repressor of v , v 's inequality can be turned into a strict inequality. This reduction makes sense only because we consider active effectors.

Let $R_i(\bar{X}^*) = 1/\beta^* \cdot T^+(\bar{X}^*)_i - T^-(\bar{X}^*)_i$ be the residual amount of the i 'th SR constraint of Eq. (4.8) (hence $R_i(\bar{X}^*) > 0$ implies strict inequality on the i th constraint with \bar{X}^*).

Assume, toward contradiction, that there exists at least one entity v_{k_0} for which the corresponding SR constraint of Eq. (4.8) holds with strict inequality. In what follows, we gradually construct a new assignment \bar{X}^{**} that achieves strict inequality in Eq. (4.8) for all $v_i \in \mathcal{V}$. Clearly, if all SR constraints of Eq. (4.8) are satisfied with strict inequality, then there exists some $\beta^{**} > \beta^*$ that satisfies all the constraints and we end with a contradiction to the optimality of $\beta^*(\mathcal{L})$.

To construct \bar{X}^{**} , we trace paths of influence in the strongly connected (and active) constraint graph D^* . Let $BFS(D^*, v_{k_0})$ be the BFS tree of D^* rooted at v_{k_0} . Define

$$L_j(D^*) = LAYER_j(BFS(D^*, v_{k_0}))$$

to be the j 'th level of $BFS(D^*, v_{k_0})$. Formally, $L_j(D^*) = \{v_y \mid d(v_y, v_{k_0}) = j\}$. Let $Q_t = \bigcup_{i=0}^t L_i(D^*)$. Let $\mathbf{S}_t \subseteq \mathbf{S}^*$ be the partial selection restricted to entities in Q_{t-1} . I.e., $|\mathbf{S}_t| = |Q_{t-1}|$ and for every $v_i \in Q_{t-1}$, $|\mathbf{S}_t \cap \mathcal{S}_i| = 1$.

The process of constructing \bar{X}^{**} consists of d steps, where d is the depth of $BFS(D^*, v_{k_0})$. At step t , we are given \bar{X}_{t-1} and use it to construct \bar{X}_t . Essentially, \bar{X}_t should satisfy the following properties.

(P1) The set of SR inequalities corresponding to Q_{t-1} entities hold with strict inequality with \bar{X}_t . I.e., for every $v_i \in Q_{t-1}$,

$$1/\beta^*(\mathcal{L}) \cdot T^+(\bar{X}_t, \mathcal{L})_i > T^-(\bar{X}_t, \mathcal{L})_i.$$

(P2) \bar{X}_t is an optimal solution, i.e., it satisfies Program (4.5) with $\beta^*(\mathcal{L})$.

(P3) $X_t(\mathcal{A}_j) = X^*(\mathcal{A}_j)$ for every $\mathcal{A}_j \notin \mathbf{S}_t$ and $X_t(\mathcal{A}_j) < X_{t-1}(\mathcal{A}_j)$ for every $\mathcal{A}_j \in \mathbf{S}_t$.

Let us now describe the construction process in more detail. Let $\bar{X}_0 = \bar{X}^*$. Consider step $t = 1$ and recall that v_{k_0} 's SR constraint holds with strict inequality. Let \mathcal{A}_{j_0} be the active supporter of v_{k_0} , i.e., $\mathcal{A}_{j_0} \in \mathcal{S}_{k_0} \cap \mathbf{S}^*$. Then it is possible to reduce a bit the value of its active supporter \mathcal{A}_{j_0} in \bar{X}^* while still maintaining feasibility. Making this change in \bar{X}_0 yields \bar{X}_1 . Formally, let $X_1(\mathcal{A}_{j_0}) = X^*(\mathcal{A}_{j_0}) - \min\{X^*(\mathcal{A}_{j_0}), R_{k_0}(\bar{X}^*)\}/2$ and leave the rest of the entries unchanged, i.e., $X_1(\mathcal{A}_k) = X^*(\mathcal{A}_k)$ for every other $k \neq j_0$. We now show that properties (P1)-(P3) are satisfied for $t = 0, 1$ and then proceed to consider the construction of \bar{X}_t , for $t > 1$. Since $L_0(D^*) = \{v_{k_0}\}$, and $Q_{-1} = \emptyset$, the solution \bar{X}_0 satisfies (P1)-(P3). Next, consider \bar{X}_1 . By the irreducibility of the system (in particular, see Cor. 4.1), since only \mathcal{A}_{k_0} was reduced in \bar{X}_1 (compared to \bar{X}^*), only the k_0^{th} constraint could have been damaged (i.e., become unsatisfied). Yet, it is easy to verify that the constraint of v_{k_0} still holds with strict inequality for \bar{X}_1 . Properties (P1)-(P3) are satisfied.

Next, we describe the general construction step. Assume that we are given \bar{X}_k for step $k \leq t$ and that the properties (P1)-(P3) hold for each $k \leq t$. We now describe the construction of \bar{X}_{t+1} and then show that it satisfies the desired properties. We begin by showing that the set of SR inequalities of $L_t(D^*)$ nodes (Eq. (4.8)) hold with strict inequality with \bar{X}_t .

CLAIM 5.6. $T^-(\bar{X}_t)_i < 1/\beta^* \cdot T^+(\bar{X}_t)_i$ for every entity $v_j \in L_t(D^*)$.

Proof. Consider some $v_j \in L_t(D^*)$. By definition of $L_t(D^*)$, there exists an entity $v_i \in L_{t-1}(D^*)$ such that $e(i, j) \in E(D^*)$. Since $v_i \in Q_{t-1}$ and \mathbf{S}_t is a selection for Q_{t-1} , a supporter $\mathcal{A}_{i_t} \in \mathbf{S}_t \cap \mathcal{S}_i$ is guaranteed to exist. Observe that $\mathcal{A}_{i_t} \in \mathcal{R}_j$ (by the definition of D^* , $e(v_i, v_j) \in E(D^*)$ implies that $(\mathcal{S}_i \cap \mathbf{S}_t) \in \mathcal{R}_j$). Finally, note that by property (P3), $X_t(\mathcal{A}_{i_t}) < X_{t-1}(\mathcal{A}_{i_t})$ and that $X_t(\mathcal{A}_k) = X^*(\mathcal{A}_k)$ for every $\mathcal{A}_k \in \mathcal{S}_j$. I.e.,

$$(5.18) \quad T^+(\bar{X}_t)_j = T^+(\bar{X}_{t-1})_j \text{ and } T^-(\bar{X}_t)_j < T^-(\bar{X}_{t-1})_j.$$

By the optimality of \bar{X}_{t-1} and \bar{X}_t (property (P3) for step $t - 1$ and t), we have that $R_j(\bar{X}_{t-1}) > 0$ and $R_j(\bar{X}_t) \geq 0$. By Eq. (5.18), $0 \leq R_j(\bar{X}_{t-1}) < R_j(\bar{X}_t)$, which establishes the claim for v_j . The same argument can be applied for every $v_j \in L_t(D^*)$, thus the claim is established. \square

Let Y be the restriction of the selection \mathbf{S}^* to $L_t(D^*)$ nodes. The solution \bar{X}_{t+1} reduces only the entries of Y supporters and the rest of the supporters are as in \bar{X}_t . Recall that by construction, $\mathbf{S}^* \subseteq$

$NZ(\bar{X}^*)$ and therefore also $\mathbf{S}^* \subseteq NZ(\bar{X}_t)$. By Claim 5.6, the constraints of $L_t(D^*)$ nodes hold with strict inequality, and therefore it is possible to reduce a bit the value of their positive supporters while still maintaining the strict inequality (although with a lower residual). Formally, for every $v_k \in L_t(D^*)$, consider its unique supporter in Y , $\mathcal{A}_k \in Y \cap \mathcal{S}_k$. By Claim 5.6, $R_k(\bar{X}_t) > 0$. Set $X_{t+1}(\mathcal{A}_k) = X_t(\mathcal{A}_k) - \min(X_t(\mathcal{A}_k), R_k(\bar{X}_t))/2$. In addition, $X_{t+1}(\mathcal{A}_k) = X_t(\mathcal{A}_k)$ for every other supporter $\mathcal{A}_k \notin Y$.

It remains to show that \bar{X}_{t+1} satisfies the properties (P1)-(P3). (P1) follows by construction. To see (P2), note that since $\mathcal{S}_i \cap \mathcal{S}_j = \emptyset$ for every $v_i, v_j \in \mathcal{V}$, only the constraints of $L_t(D^*)$ nodes might have been violated by the new solution \bar{X}_{t+1} . Formally, $T^+(\bar{X}_{t+1})_i = T^+(\bar{X}_t)_i$ and $T^-(\bar{X}_{t+1})_i \leq T^-(\bar{X}_t)_i$ for every $v_i \notin L_t(D^*)$. Although, for $v_i \in L_t(D^*)$, we get that $T^+(\bar{X}_{t+1})_i < T^+(\bar{X}_t)_i$ (yet $T^-(\bar{X}_{t+1})_i = T^-(\bar{X}_t)_i$), this reduction in the total support of $L_t(D^*)$ nodes was performed in a controlled manner, guaranteeing that the corresponding $L_t(D^*)$ inequalities hold with strict inequality. Finally, (P3) follows immediately. After $d + 1$ steps, by (P1) all inequalities hold with strict inequality (as $Q_d = \mathcal{V}$) with the solution \bar{X}_{d+1} . Thus, it is possible to find some $\beta^{**} > \beta^*(\mathcal{L})$ that would contradict the optimality of β^* . Formally, let $R^* = \min R_i(\bar{X}_{d+1})$. Since $R^* > 0$, we get that \bar{X}_{d+1} is feasible with $\beta^{**} = \beta^*(\mathcal{L}) + R^* > \beta^*(\mathcal{L})$, contradicting the optimality of $\beta^*(\mathcal{L})$. Lemma 5.2 follows. \square

We proceed by considering a vertex of $\bar{X}^* \in V(\mathcal{P}(\beta^*))$. By the previous section, \bar{X}^* corresponds to $\mathbf{0}^*$.

LEMMA 5.3. (a) \bar{X}^* is a $\mathbf{0}^*$ solution. (b) There exists an $F^* \in \mathcal{F}$ such that $r(\mathcal{L}(F^*)) = 1/\beta^*$.

We start with (a) and transform \mathcal{L} into a weak system \mathcal{L}^w . First, if $m = n + 1$, then the system is already weak. Otherwise, without loss of generality, let the i^{th} entry in \bar{X}^* correspond to \mathcal{A}_i where $\mathcal{A}_i = NZ(\bar{X}^*) \cap \mathcal{S}_i$ for $i \in \{1, \dots, n - 1\}$ and the n and $n + 1$ entries correspond to \mathcal{A}_n and \mathcal{A}_{n+1} respectively such that $\{\mathcal{A}_n, \mathcal{A}_{n+1}\} = NZ(\bar{X}^*) \cap \mathcal{S}_n$. It then follows that $X^*(i) \neq 0$ for every $i \in \{1, \dots, n + 1\}$ and $X^*(i) = 0$ for every $i \in \{n + 2, \dots, m\}$. Let $\bar{X}^{**} = (X^*(1), \dots, X^*(n + 1))$. Let $\mathcal{M}_w^+ \in \mathbb{R}^{n \times (n+1)}$ where $\mathcal{M}_w^+(i, j) = \mathcal{M}^+(i, j)$ for every $i \in \{1, \dots, n\}$ every $j \in \{1, \dots, n + 1\}$, and \mathcal{M}_w^- is defined analogously. From now on, we restrict attention to the weak system $\mathcal{L}^w = \langle \mathcal{M}_w^+, \mathcal{M}_w^- \rangle$. This weak system is an almost square system, expect that for the last entity $|\mathcal{S}_n| = 2$. Note that the weak system results from \mathcal{L} by discarding the corresponding entries of $\mathcal{A} \setminus NZ(\bar{X}^*)$. Therefore, $\beta^*(\mathcal{L}) = \beta^*(\mathcal{L}^w)$. Let \mathcal{M}_{n-1}^+

correspond to the upper left $(n-1) \times (n-1)$ submatrix of \mathcal{M}_w^+ . Let \mathcal{M}_n^+ be obtained from \mathcal{M}_w^+ by removing the last $(n+1)^{th}$ column. Finally, \mathcal{M}_{n+1}^+ is obtained from \mathcal{M}_w^+ by removing the n^{th} column. The matrices $\mathcal{M}_{n-1}^-, \mathcal{M}_n^-, \mathcal{M}_{n+1}^-$ are defined analogously.

To study the weak system \mathcal{L}^w , we consider the following three square systems.

$$\begin{aligned}\mathcal{L}_{n-1} &= \langle \mathcal{M}_{n-1}^+, \mathcal{M}_{n-1}^- \rangle, \\ \mathcal{L}_n &= \langle \mathcal{M}_n^+, \mathcal{M}_n^- \rangle, \\ \mathcal{L}_{n+1} &= \langle \mathcal{M}_{n+1}^+, \mathcal{M}_{n+1}^- \rangle.\end{aligned}$$

Note that a feasible solution \bar{X}_{n+i} for the system \mathcal{L}_{n+i} , for $i \in \{0, 1\}$, corresponds to a feasible solution for \mathcal{L}^w by setting $X_w(\mathcal{A}_j) = X_{n+i}(\mathcal{A}_i)$ for every $j \neq n + (1-i)$ and $X_w(\mathcal{A}_{n+(1-i)}) = 0$. For ease of notation, let $P_n(\lambda) = P(Z(\mathcal{L}_n), \lambda)$, $P_{n+1}(\lambda) = P(Z(\mathcal{L}_{n+1}), \lambda)$ and $P_{n-1}(\lambda) = P(Z(\mathcal{L}_{n-1}), \lambda)$, where P is the characteristic polynomial defined in Eq. (3.3). Let β_{n-1}^* , β_n^* , and β_{n+1}^* be the optimal values of Program (4.5) for systems \mathcal{L}_{n-1} , \mathcal{L}_n , and \mathcal{L}_{n+1} , respectively. Let $\lambda^* = 1/\beta^*$ and $\lambda_{n+i}^* = 1/\beta_{n+i}^*$, for $i \in \{-1, 0, 1\}$.

CLAIM 5.7. $\max\{\beta_n^*, \beta_{n+1}^*\} \leq \beta^* < \beta_{n-1}^*$.

Proof. The left inequality follows as any optimal solution \bar{X}^* for \mathcal{L}_n (respectively, \mathcal{L}_{n+1}) can be achieved in the weak system \mathcal{L}^w by setting $X^*(\mathcal{A}_n) = 0$ (resp., $X^*(\mathcal{A}_{n+1}) = 0$). To see that the right inequality is strict, observe that in any solution \bar{X} for \mathcal{L}^w , the two supporters \mathcal{A}_n and \mathcal{A}_{n+1} of v_n satisfy that $X(\mathcal{A}_n) + X(\mathcal{A}_{n+1}) > 0$ by Claim 5.4. Without loss of generality, assume that $X(\mathcal{A}_n) > 0$. Then by Obs. 5.1(a) and the irreducibility of \mathcal{L}^w , v_n is strongly connected to the rest of the graph for every selection of one of its two supporters. It follows that v_n has an outgoing edge $e_{n,j} \in E$ in the constraint graph $\mathcal{CG}_{\mathcal{L}}(V, E)$, i.e., there exists some entity v_j , $j \in [1, n-1]$, such that $\mathcal{A}_n \in \mathcal{R}_j$. Since \mathcal{A}_n does not appear in \mathcal{L}_{n-1} , the total repression on v_j in \mathcal{L}^w (i.e., $(\mathcal{M}_w^- \cdot X^*)_j$) is strictly greater than in \mathcal{L}_{n-1} (i.e., $(\mathcal{M}_{n-1}^- \cdot X^*)_j$). \square

Our goal in this section is to show that the optimal β^* value for \mathcal{L}^w can be achieved by setting either $X^*(\mathcal{A}_n) = 0$ or $X^*(\mathcal{A}_{n+1}) = 0$, essentially showing that the optimal $w\mathbf{0}^*$ solution corresponds to a $\mathbf{0}^*$ solution. This is formalized in the following theorem.

THEOREM 5.1. $\beta^* = \max\{\beta_n^*, \beta_{n+1}^*\}$.

The following observation holds for every $i \in \{-1, 0, 1\}$ and follows immediately by the definitions of feasibility and irreducibility and the PF Theorem 3.1.

OBSERVATION 5.5. (1) $\lambda_{n+i} > 0$ is the maximal eigenvalue of $Z(\mathcal{L}_{n+i})$.

(2) For an irreducible system \mathcal{L} , $\lambda_{n+i} = 1/\beta_{n+i}$.

(3) If the system is feasible then $\lambda_{n+i} > 0$.

For a square system $\mathcal{L} \in \mathfrak{L}^s$, let W^1 be a modified form of the matrix Z , defined as follows.

$$W^1(\mathcal{L}, \beta) = Z(\mathcal{L}) - 1/\beta \cdot I \quad \text{for } \beta \in (0, \beta^*].$$

More explicitly,

$$W^1(\mathcal{L}, \beta)_{i,j} = \begin{cases} -1/\beta, & \text{if } i = j; \\ -g(v_i, \mathcal{A}_j)/g(i, i), & \text{otherwise.} \end{cases}$$

Clearly, $W^1(\mathcal{L}, \beta)$ cannot be defined for a nonsquare system $\mathcal{L} \notin \mathfrak{L}^s$. Instead, a generalization W^2 of W^1 for any (nonsquare) $m \geq n$ system \mathcal{L} is given by

$$W^2(\mathcal{L}, \beta) = \mathcal{M}^- - 1/\beta \cdot \mathcal{M}^+, \quad \text{for } \beta \in (0, \beta^*],$$

or explicitly,

$$W^2(\mathcal{L}, \beta)_{i,j} = \begin{cases} -g(i, i)/\beta, & \text{if } i = j; \\ -g(v_i, \mathcal{A}_j), & \text{otherwise.} \end{cases}$$

Note that if \bar{X}_β is a feasible solution for \mathcal{L} , then $W^2(\mathcal{L}, \beta) \cdot \bar{X}_\beta \leq 0$. If $\mathcal{L} \in \mathfrak{L}^s$, it also holds that $W^1(\mathcal{L}, \beta) \cdot \bar{X}_\beta \leq 0$.

For $\mathcal{L} \in \mathfrak{L}^s$, where both $W^1(\mathcal{L}, \beta)$ and $W^2(\mathcal{L}, \beta)$ are well-defined, the following connection becomes useful in our later argument. Recall that $P(Z(\mathcal{L}), t)$ is the characteristic polynomial of $Z(\mathcal{L})$ (see Eq. (3.3)).

OBSERVATION 5.6. For a square system \mathcal{L} ,

$$\det(W^2(\mathcal{L}, \beta)) = P(Z(\mathcal{L}), 1/\beta) \cdot \prod_{i=1}^n g(i, i).$$

Proof. The observation follows immediately by noting that $W^1(\mathcal{L}, \beta)_{i,j} = W^2(\mathcal{L}, \beta)_{i,j} \cdot g(i, i)$ for every i and j , and by Eq. (3.3). \square

The next equality plays a key role in our analysis.

LEMMA 5.4.
$$\frac{g(n, n) \cdot X^*(n) \cdot P_n(\lambda^*)}{P_{n-1}(\lambda^*)} + \frac{g(n, n+1) \cdot X^*(n+1) \cdot P_{n+1}(\lambda^*)}{P_{n-1}(\lambda^*)} = 0.$$

Our work plan from this point on is as follows. We first define a range of ‘candidate’ values for β^* . Essentially, our interest is in *real* positive β^* . Recall that $Z(\mathcal{L}^w)$, $Z(\mathcal{L}_n)$ and $Z(\mathcal{L}_{n+1})$ are nonnegative irreducible matrices and therefore Theorem 3.1 can be applied throughout the analysis. Without loss of generality, assume that $\beta_n^* \geq \beta_{n+1}^*$ (and thus $\lambda_n^* \leq \lambda_{n+1}^*$) and




Figure 2: Real positive roots of $P_{n+1}(\lambda)$, $P_n(\lambda)$, and $P_{n-1}(\lambda)$.

let $Range_{\beta^*} = (\beta_n^*, \beta_{n-1}^*) \subseteq \mathbb{R}_{>0}$. Let the corresponding range of λ^* be

$$(5.19) \quad Range_{\lambda^*} = (\lambda_{n-1}^*, \lambda_n^*) = (1/\beta_{n-1}^*, 1/\beta_n^*).$$

To complete the proof for Thm. 5.1 we assume, towards contradiction, that $\beta^* > \beta_n^*$. According to Claim 5.7 and the fact that $\beta^* \neq \beta_n^*$, it then follows that $\beta^* \in Range_{\beta^*}$. Note that since $Range_{\beta^*} \subseteq \mathbb{R}_{>0}$, also $Range_{\lambda^*} \subseteq \mathbb{R}_{>0}$. In other words, since we look for an optimal $\beta^* \in \mathbb{R}_{>0}$, the corresponding λ that interests us is real and positive as well. This is important mainly in the context of nonnegative irreducible matrices $Z(\mathcal{L}')$ for $\mathcal{L}' \in \mathcal{L}^s$. In contrast to nonnegative primitive matrices (where $h = 1$) for irreducible matrices, such as $Z(\mathcal{L}')$, by Thm. 3.1 there are $h \geq 1$ eigenvalues, $\lambda_i \in EigVal(\mathcal{L}')$, for which $|\lambda_i| = r(\mathcal{L}')$. However, note that only one of these, namely, $r(\mathcal{L}')$, might belong to $Range_{\lambda^*} \subseteq \mathbb{R}_{>0}$. (This follows as every other such λ_i is either real but negative or with a nonzero complex component).

Fix $j \in \{-1, 0, 1\}$ and let k_{n+j} be the number of real and positive eigenvalues of $Z(\mathcal{L}_{n+j})$. Let $0 < \lambda_{n+j}^1 \leq \lambda_{n+j}^2 \leq \dots \leq \lambda_{n+j}^{k_{n+j}}$ be the ordered set of real and positive eigenvalues for $Z(\mathcal{L}_{n+j})$, i.e., real positive roots of $P_{n+j}(\lambda)$. Note that $\lambda_{n+j}^{n+k_j} = \lambda_{n+j}^*$. By Theorem 3.1, we have that

- (a) $\lambda_{n+j}^* \in \mathbb{R}_{>0}$, and
- (b) $\lambda_{n+j}^* > |\lambda_{n+j}^p|$, $p \in \{1, \dots, k_{n+j} - 1\}$, for every $j \in \{-1, 0, 1\}$.

We proceed by showing that the potential range for λ^* , namely, $Range_{\lambda^*}$, can contain no root of $P_n(\lambda)$ and $P_{n+1}(\lambda)$. Since $Range_{\lambda^*}$ is real and positive, it is sufficient to consider only real and positive roots of $P_n(\lambda)$ and $P_{n+1}(\lambda)$ (or real and positive eigenvalues of $Z(\mathcal{L}_n)$ and $Z(\mathcal{L}_{n+1})$).

CLAIM 5.8. $\lambda_n^{p_1}, \lambda_{n+1}^{p_2} \notin Range_{\lambda^*}$ for every real $\lambda_n^{p_1}, \lambda_{n+1}^{p_2}$, for $p_1 < k_n, p_2 < k_{n+1}$.

Proof. Note that $Z(\mathcal{L}_{n-1})$ is the principal $(n-1)$ minor of both $Z(\mathcal{L}_n)$ and $Z(\mathcal{L}_{n+1})$. We now use the separation

theorem for nonnegative matrices, due to Hall and T. A. Porsching [15], which is an extension to the Cauchy Interlacing Theorem for symmetric matrices. In particular, the separation theorem implies in our context that $\lambda_n^{p_1}, \lambda_{n+1}^{p_2} \leq \lambda_{n-1}^*$ for every $p_1 < k_n$ and $p_2 < k_{n+1}$, concluding by Eq. (5.19) that $\lambda_n^p, \lambda_{n+1}^p \notin Range_{\lambda^*}$. \square

We proceed by showing that $P_n(\lambda)$ and $P_{n+1}(\lambda)$ have the same sign in $Range_{\lambda^*}$. See Fig. 2 for a schematic description of the system.

CLAIM 5.9. $\text{sign}(P_n(\lambda)) = \text{sign}(P_{n+1}(\lambda))$ for every $\lambda \in Range_{\lambda^*}$.

Proof. Fix $i \in \{0, 1\}$. By Claim 5.8, P_{n+i} has no roots in the range $Range_{\lambda^*}$, so $\text{sign}(P_{n+i}(\lambda_1)) = \text{sign}(P_{n+i}(\lambda_2))$ for every $\lambda_1, \lambda_2 \in Range_{\lambda^*}$. Also note that for a fixed $i \in \{0, 1\}$, $\text{sign}(P_{n+i}(\lambda_1)) = \text{sign}(P_{n+i}(\lambda_2))$, for every $\lambda_1, \lambda_2 > \lambda_{n+i}^*$. There are two crucial observations. First, as $P_n(\lambda)$ and $P_{n+1}(\lambda)$ correspond to a characteristic polynomial of an $n \times n$ matrix, they have the same leading coefficient and therefore $\text{sign}(P_n(\lambda)) = \text{sign}(P_{n+1}(\lambda))$ for $\lambda > \lambda_{n+1}^*$ (recall that we assume that $\lambda_{n+1}^* \geq \lambda_n^*$). Next, due to the PF Theorem, the maximal roots of $P_n(\lambda)$ and $P_{n+1}(\lambda)$ are of multiplicity one and therefore the polynomial necessarily changes its sign when passing through its maximal root. Recall that λ_n^* (respectively, λ_{n+1}^*) is the maximal real positive root of $P_n(\lambda)$, (resp., $P_{n+1}(\lambda)$). Assume, toward contradiction, that $\text{sign}(P_n(\lambda)) \neq \text{sign}(P_{n+1}(\lambda))$ for some $\lambda \in Range_{\lambda^*}$. Then $\text{sign}(P_n(\lambda_1)) \neq \text{sign}(P_n(\lambda_2))$ for $\lambda_1 > \lambda_n^*$ and $\lambda_2 \in Range_{\lambda^*}$ also $\text{sign}(P_{n+1}(\lambda_1)) \neq \text{sign}(P_{n+1}(\lambda_2))$ for $\lambda_1 > \lambda_{n+1}^*$ and $\lambda_2 \in Range_{\lambda^*}$. (This holds since when encountering a root of multiplicity one, the sign necessarily flips). In particular, this implies that $\text{sign}(P_n(\lambda)) \neq \text{sign}(P_{n+1}(\lambda))$ for every $\lambda \geq \lambda_{n+1}^*$, in contradiction to the fact that $\text{sign}(P_n(\lambda)) = \text{sign}(P_{n+1}(\lambda))$ for every $\lambda > \lambda_{n+1}^*$. The claim follows. \square

We now complete the proof of Theorem 5.1.

Proof. Due to Thm. 3.1, we have that $\lambda_n^* = 1/\beta_n^*, \lambda_{n+1}^* = 1/\beta_{n+1}^*$ and $\lambda_{n-1}^* = 1/\beta_{n-1}^*$. It therefore holds that $P_{n-1}(\lambda) \neq 0$ for every $\lambda \in Range_{\lambda^*}$. We can now apply safely Lemma 5.4 and Claim 5.9 and get that $\text{sign}(X^*(n)) \neq \text{sign}(X^*(n+1))$. Since $X^*(n)$ and $X^*(n+1)$ are nonnegative, it follows that either $X^*(n) = 0$ or $X^*(n+1) = 0$. Assume, to the contrary, that $\beta^* > \beta_n^*$. Then $\beta^* \in Range_{\beta^*}$, and therefore $\text{sign}(X^*(n)) \neq \text{sign}(X^*(n+1))$. This contradicts the fact that $\overline{X^*}$ is nonnegative. We conclude that $\beta^* = \beta_n^*$. \square

We complete the geometric characterization of the generalized PF Theorem by noting that every vertex of $V(\mathcal{P}(\beta^*))$ is a $\mathbf{0}^*$ solution, thus establishing Lemma 5.3.

LEMMA 5.5. *Every vertex $\bar{X} \in V(\mathcal{P}(\beta^*))$ is a $\mathbf{0}^*$ solution.*

Proof. [Thm. 4.1] Let F^* be the selection such that $r(\mathcal{L}) = r(\mathcal{L}(F^*))$. Note that by the irreducibility of \mathcal{L} , the square system $\mathcal{L}(F^*)$ is irreducible as well and therefore the PF Theorem for irreducible matrices can be applied. In particular, by Thm. 3.1, it follows that $r(\mathcal{L}(F^*)) \in \mathbb{R}_{>0}$ and that $\bar{\mathbf{P}}(\mathcal{L}(F^*)) > 0$. Therefore, by Eq. (4.10) and (4.11), Claims (Q1)-(Q3) of Thm. 4.1 follow.

We now turn to claim (Q4) of the theorem. Note that for a symmetric system, in which $g(i, j_1) = g(i, j_2)$ for every $\mathcal{A}_{j_1}, \mathcal{A}_{j_2} \in \mathcal{S}_k$ and every $k, i \in [1, n]$, the system is invariant to the selection matrix and therefore $r(\mathcal{L}(F_1)) = r(\mathcal{L}(F_2))$ for every $F_1, F_2 \in \mathcal{F}$.

Finally, it remains to consider claim (Q5) of the theorem. Note that the optimization problem specified by Program (4.5) is an alternative formulation to the generalized Collatz-Wielandt formula given in (Q5). We now show that $r(\mathcal{L})$ (respectively, $\bar{\mathbf{P}}(\mathcal{L})$) is the optimum value (resp., point) of Program (4.5). By Lemma 5.3, there exists an optimal point \bar{X}^* for Program (4.5) which is a $\mathbf{0}^*$ solution. Note that a $\mathbf{0}^*$ solution corresponds to a unique hidden square system, given by $\mathcal{L}^* = \mathcal{L}(NZ(\bar{X}^*))$ (\mathcal{L}^* is square since $|NZ(\bar{X}^*)| = n$). Therefore, by Thm. 3.2 and Lemma 5.3(b), we get that

$$(5.20) \quad r(\mathcal{L}^*) = 1/\beta^*(\mathcal{L}^*) = 1/\beta^*(\mathcal{L}).$$

Next, by Observation 4.3(b), we have that $r(\mathcal{L}(F)) \geq r(\mathcal{L})$. It therefore follows that

$$(5.21) \quad r(\mathcal{L}^*) = \min_{F \in \mathcal{F}} r(\mathcal{L}(F)).$$

Combining Eq. (5.20), (5.21) and (4.10), we get that the PF eigenvalue of the system \mathcal{L} satisfies $r(\mathcal{L}) = 1/\beta^*(\mathcal{L})$ as required. Finally, note that by Thm. 3.2, $\bar{\mathbf{P}}(\mathcal{L}^*)$ is the optimal point for Program (4.5) with the square system \mathcal{L}^* . By Eq. (4.11), $\bar{\mathbf{P}}(\mathcal{L})$ is an extension of $\bar{\mathbf{P}}(\mathcal{L}^*)$ with zeros (i.e., a $\mathbf{0}^*$ solution). It can easily be checked that $\bar{\mathbf{P}}(\mathcal{L})$ is a feasible solution for the original system \mathcal{L} with $\beta = \beta^*(\mathcal{L}^*) = \beta^*(\mathcal{L})$, hence it is optimal. Note that by Lemma 5.2, it indeed follows that $\mathcal{M}^- \cdot \bar{\mathbf{P}}(\mathcal{L}) = 1/\beta^*(\mathcal{L}) \cdot \mathcal{M}^+ \cdot \bar{\mathbf{P}}(\mathcal{L})$, for every optimal solution \bar{X}^* . The theorem follows.

Section 6 provides a characterization of systems in which a $\mathbf{0}^*$ solution does not exist.

6 Limitation for the Existence of a $\mathbf{0}^*$ Solution

In this section we provide a characterization of systems in which a $\mathbf{0}^*$ solution does not exist.

Bounded Value Systems. Let X_{\max} be a fixed constant. For a nonnegative vector \bar{X} , let

$$\max(\bar{X}) = \max\{X(j)/X(i) \mid 1 \leq i, j \leq n, X(i) > 0\}.$$

A system \mathcal{L} is called a *bounded power system* if $\max(\bar{X}) \leq X_{\max}$.

LEMMA 6.1. *There exists a bounded power system \mathcal{L} such that no optimal solution \bar{X}^* for \mathcal{L} is a $\mathbf{0}^*$ solution.*

Second eigenvalue maximization. One of the most common applications of the PF Theorem is the existence of the stationary distribution for a transition matrix (representing a random process). The stationary distribution is the eigenvector of the largest eigenvalue of the transition matrix. We remark that if the transition matrix is stochastic, i.e., the sum of each row is 1, then the largest eigenvalue is equal to 1. So this case does not give rise to any optimization problem. However, in many cases we are interested in processes with fast mixing time. Assuming the process is ergodic, the mixing time is determined by the difference between the largest eigenvalue and the second largest eigenvalue. So we can try to solve the following problem. Imagine that there is some rumor that we are interested in spreading over two or more social networks. Each node can be a member of several social networks. We would like to merge all the networks into one large social network in a way that will result in fast mixing time. This problem looks very similar to the one solved in this paper. Indeed, one can use similar techniques and get an approximation. But interestingly, this problem does not have the $\mathbf{0}^*$ solution property, as illustrated in the following example.

Assume we are given n nodes. Consider the $n!$ different social networks that arise by taking, for each permutation $\pi \in S(n)$, the path P_π corresponding to the permutation π . Clearly, the best mixing graph we can get is the complete graph K_n . We can get this graph if each node chooses each permutation with probability $\frac{1}{n!}$. We remind the reader that the mixing time of the graph K_n is 1. On the other hand, any $\mathbf{0}^*$ solution have a mixing time $O(n^2)$. This example shows that in the second largest eigenvalue, the solution is not always a $\mathbf{0}^*$ solution.

7 Computing the Generalized PF Vector

In this section we present a polynomial time algorithm for computing the generalized Perron eigenvector $\bar{\mathbf{P}}(\mathcal{L})$ of an irreducible system \mathcal{L} .

The method. By property (Q5) of Thm. 4.1, computing $\bar{\mathbf{P}}(\mathcal{L})$ is equivalent to finding a $\mathbf{0}^*$ solution for Program (4.5) with $\beta = \beta^*(\mathcal{L})$. For ease of analysis, we assume throughout that the gains are integral, i.e., $g(i, j) \in \mathbb{Z}^+$, for every $i \in \{1, \dots, n\}$ and $j \in \{1, \dots, m\}$. If this does not hold, then the gains can be rounded or scaled to achieve this. Let

$$(7.22) \quad \mathcal{G}_{max}(\mathcal{L}) = \max_{i \in \{1, \dots, n\}, j \in \{1, \dots, m\}} \{|g(i, j)|\},$$

and define T_{LP} as the running time of an LP solver such as interior point algorithm [5] for Program (5.14). Recall that we concern an exact optimal solution for non-convex optimization problem (see Program (4.5)). Using the convex relaxation of Program (5.14), a binary search can be applied for finding an approximate solution up to a predefined accuracy. The main challenge is then to find (a) an optimal solution (and not an approximate solution), and (b) among all the optimal solutions, to find one that is a $\mathbf{0}^*$ solution. Let $F_1, F_2 \in \mathcal{F}$ be two selection matrices for \mathcal{L} . By Thm. 4.1, there exists a selection matrix F^* such that $r(\mathcal{L}) = r(\mathcal{L}(F^*))$ and $\bar{\mathbf{P}}(\mathcal{L})$ is a $\mathbf{0}^*$ solution corresponding to $\bar{\mathbf{P}}(\mathcal{L}(F^*))$ (in addition $\beta^* = 1/r(\mathcal{L}(F^*))$). Our goal then is to find a selection matrix $F^* \in \mathcal{F}$ where $|\mathcal{F}|$ might be exponentially large.

THEOREM 7.1. *Let \mathcal{L} be irreducible system. Then $\bar{\mathbf{P}}(\mathcal{L})$ can be computed in time $O(n^3 \cdot T_{LP} \cdot (\log(n \cdot \mathcal{G}_{max}) + n))$.*

Let

$$(7.23) \quad \Delta_\beta = (n\mathcal{G}_{max})^{-4n^3}.$$

The key observation in this context is the following.

LEMMA 7.1. *Consider a selection matrix $F \in \mathcal{F}$. If $\beta^*(\mathcal{L}) - 1/r(\mathcal{L}(F)) \leq \Delta_\beta$, then $\beta^*(\mathcal{L}) = 1/r(\mathcal{L}(F))$.*

To prove Lemma 7.1, we exploit a lemma of Bugeaud and Mignotte in [6].

Algorithm description. We now describe the algorithm $\text{Compute}\bar{\mathbf{P}}(\mathcal{L})$ for $\bar{\mathbf{P}}(\mathcal{L})$ computation. Consider some partial selection $\mathbf{S}' \subseteq \mathcal{A}$ for $V' \subseteq \mathcal{V}$. For ease of notation, let $\mathcal{L}(\mathbf{S}') = \langle \mathcal{M}^-(\mathbf{S}'), \mathcal{M}^+(\mathbf{S}') \rangle$, where $\mathcal{M}^-(\mathbf{S}') = \mathcal{M}^- \cdot F(\mathbf{S}')$ and $\mathcal{M}^+(\mathbf{S}') = \mathcal{M}^+ \cdot F(\mathbf{S}')$. Consider the Program

$$\begin{aligned} & \max \beta \\ & \text{s.t. } \mathcal{M}^-(\mathbf{S}') \cdot \bar{X} \leq 1/\beta \cdot \mathcal{M}^+(\mathbf{S}') \cdot \bar{X} \\ & \quad \bar{X} \geq \bar{0} \\ & \quad \|\bar{X}\|_1 = 1. \end{aligned}$$

Note that if $\mathbf{S}' = \emptyset$, then the above program is equivalent to Program (4.5), i.e., $\mathcal{L}(\mathbf{S}') = \mathcal{L}$. Define

$$f(\beta, \mathcal{L}(\mathbf{S}')) = \begin{cases} 1, & \text{if there exists an } \bar{X} \text{ such that} \\ & \|\bar{X}\|_1 = 1, \bar{X} \geq \bar{0}, \text{ and} \\ & \mathcal{M}^-(\mathbf{S}') \cdot \bar{X} \leq 1/\beta \cdot \mathcal{M}^+(\mathbf{S}') \cdot \bar{X}, \\ 0, & \text{otherwise.} \end{cases}$$

Note that $f(\beta, \mathcal{L}(\mathbf{S}')) = 1$ iff $\mathcal{L}(\mathbf{S}')$ is feasible for β and that f can be computed in polynomial time using the interior point method.

Algorithm $\text{Compute}\bar{\mathbf{P}}(\mathcal{L})$ is composed of two main phases. In the first phase we find, using binary search, an estimate β^- such that $\beta^*(\mathcal{L}) - \beta^- \leq \Delta_\beta$. In the second phase, we find a hidden square system, $\mathcal{L}(F^*)$, $F^* \in \mathcal{F}$, corresponding to a complete selection vector \mathbf{S}_n of size n for \mathcal{V} . By Lemma 7.1, it follows that $r(\mathcal{L}(F^*)) = 1/\beta^*(\mathcal{L})$.

We now describe the construction of \mathbf{S}_n in more detail. The phase consists of n iterations. On iteration t we obtain a partial selection \mathbf{S}_t for v_1, \dots, v_t such that $f(\beta^-, \mathcal{L}(\mathbf{S}_t)) = 1$. On the final step we achieve the desired \mathbf{S}_n , where $\mathcal{L}(\mathbf{S}_n) \in \mathcal{L}^s$ and $f(\beta^-, \mathcal{L}(\mathbf{S}_n)) = 1$ (therefore also $f(\beta^-, \mathcal{L}(F(\mathbf{S}_n))) = 1$). Initially, \mathbf{S}_0 is empty. On the t 'th iteration, $\mathbf{S}_t = \mathbf{S}_{t-1} \cup \{\mathcal{A}_j\}$ for $\mathcal{A}_j \in \mathcal{S}_t$. Essentially, \mathcal{A}_j is selected such that $f(\beta^-, \mathcal{L}(\mathbf{S}_{t-1} \cup \{\mathcal{A}_j\})) = 1$. We later show (in proof of Thm. 7.1) that such a supporter \mathcal{A}_j exists.

Finally, we use $\bar{\mathbf{P}}(\mathcal{L}(\mathbf{S}_n))$ to construct the Perron vector $\bar{\mathbf{P}}(\mathcal{L})$. This vector contains zeros for the $m - n$ non-selected effectors, and the value of n selected effectors are as in $\bar{\mathbf{P}}(\mathcal{L}(\mathbf{S}_n))$.

To establish Theorem 7.1, we prove the correctness of Algorithm $\text{Compute}\bar{\mathbf{P}}(\mathcal{L})$ and bound its runtime. We begin with two auxiliary claims.

CLAIM 7.1. $\beta^*(\mathcal{L}) \leq \mathcal{G}_{max}$.

CLAIM 7.2. *By the end of phase 1, Alg. $\text{Compute}\bar{\mathbf{P}}(\mathcal{L})$ finds β^- such that $\beta^*(\mathcal{L}) - \beta^- \leq \Delta_\beta$.*

Let $\text{Range}_{\beta^*} = [\beta^-, \beta^+]$. We are now ready to complete the proof of Thm. 7.1.

Proof. [Theorem 7.1] We show that Alg. $\text{Compute}\bar{\mathbf{P}}(\mathcal{L})$ satisfies the requirements of the theorem. Note that at the beginning of phase 2 of Alg. $\text{Compute}\bar{\mathbf{P}}(\mathcal{L})$, the computed value β^- is at most Δ_β apart from β^* . We begin by showing the following.

CLAIM 7.3. *By the end of phase 2, the selection \mathbf{S}_n is such that $r(\mathcal{L}(\mathbf{S}_n)) = 1/\beta^*(\mathcal{L})$.*

Proof. Let \mathbf{S}_t be the partial selection obtained at step t , $\mathcal{L}_t = \mathcal{L}(\mathbf{S}_t)$ be the corresponding system for step t

and $\beta_t = \beta^*(\mathcal{L}_t)$ the optimal solution of Program (4.5) for system \mathcal{L}_t . We claim that \mathbf{S}_t satisfies the following properties for each $t \in \{0, \dots, n\}$:

(P1) \mathbf{S}_t is a partial selection vector of length t , such that $\mathbf{S}_t \sim \mathbf{S}_{t-1}$.

(P2) $\mathcal{L}(\mathbf{S}_t)$ is feasible for β^- .

The proof is by induction. Beginning with $\mathbf{S}_0 = \emptyset$, it is easy to see that (P1) and (P2) are satisfied (since $\mathcal{L}(\mathbf{S}_0) = \mathcal{L}$). Next, assume that (P1) and (P2) hold for \mathbf{S}_i for $i \leq t$ and consider \mathbf{S}_{t+1} . Let $V_t \subseteq \mathcal{V}$ be such that \mathbf{S}_t is a partial selection for V_t (i.e., $|V_t| = |\mathbf{S}_t|$ and for every $v_i \in V_t$, $|\mathcal{S}_i(\mathcal{L}) \cap \mathbf{S}_t| = 1$). Given that \mathbf{S}_t is a selection for nodes v_1, \dots, v_t that satisfies (P1) and (P2), we show that \mathbf{S}_{t+1} satisfies (P1) and (P2) as well.

In particular, it is required to show that there exists at least one supporter of v_{t+1} , namely, $\mathcal{A}_k \in \mathcal{S}_{t+1}(\mathcal{L})$, such that $f(\beta^-, \mathcal{L}(\mathbf{S}_t \cup \{\mathcal{A}_k\})) = 1$. This will imply that step 7(a) always succeeds in expanding \mathbf{S}_t .

By Observation 5.3 and (P2) for step t , the system $\mathcal{L}(\mathbf{S}_t)$ is irreducible with $\beta_t \geq \beta^-$. In addition, note that $\mathcal{F}(\mathcal{L}_t) \subseteq \mathcal{F}(\mathcal{L})$ (as every square system of \mathcal{L}_t is also a square system of \mathcal{L}).

By Theorem 4.1, there exists a square system $\mathcal{L}_t(F_t^*)$, $F_t^* \in \mathcal{F}(\mathcal{L}_t)$, such that $r(\mathcal{L}_t(F_t^*)) = 1/\beta_t$. In addition, $\bar{\mathbf{P}}(\mathcal{L}_t(F_t^*))$ is a feasible solution for Program (5.14) with the system $\mathcal{L}_t(F_t^*)$ and $\beta = \beta_t$.

By Eq. (4.9), the square system $\mathcal{L}_t(F_t^*)$ corresponds to a complete selection \mathbf{S}^{**} , where $|\mathbf{S}^{**}| = n$ and $\mathbf{S}_t \subseteq \mathbf{S}^{**}$, i.e., $\mathcal{L}_t(F_t^*) = \mathcal{L}(\mathbf{S}^{**})$. Observe that by property (Q5) of Thm. 4.1 for the system \mathcal{L}_t , there exists a $\mathbf{0}^*$ solution for Program (5.14) that achieves β_t . This $\mathbf{0}^*$ solution is constructed from the PF eigenvector of $\mathcal{L}_t(\mathbf{S}^{**})$, namely, $\bar{\mathbf{P}}(\mathcal{L}_t(\mathbf{S}^{**}))$.

Let $\mathcal{A}_k \in \mathcal{S}_{t+1}(\mathcal{L}_t) \cap \mathbf{S}^{**}$. Note that by the choice of \mathbf{S}^{**} , such an affector \mathcal{A}_k exists. We now show that $\mathbf{S}_{t+1} = \mathbf{S}_t \cup \{\mathcal{A}_k\}$ satisfies (P2), thus establishing the existence of $\mathcal{A}_k \in \mathcal{S}_{t+1}(\mathcal{L}_t)$ in step 7(a). We show this by constructing a feasible solution $X_{\beta^-}^* \in \mathbb{R}^{m(\mathbf{S}_{t+1})}$ for \mathcal{L}_{t+1} . By the definition of \mathbf{S}^{**} , $f(\beta^-, \mathcal{L}(\mathbf{S}^{**})) = 1$ and therefore there exists a feasible solution $\bar{X}_{\beta^-}^{t+1} \in \mathbb{R}^n$ for $\mathcal{L}(\mathbf{S}^{**})$. Since $\mathbf{S}_{t+1} \subseteq \mathbf{S}^{**}$, it is possible to extend $\bar{X}_{\beta^-}^{t+1} \in \mathbb{R}^n$ to a feasible solution $X_{\beta^-}^*$ for system \mathcal{L}_{t+1} , by setting $X_{\beta^-}^*(\mathcal{A}_q) = \bar{X}_{\beta^-}^{t+1}(\mathcal{A}_q)$ for every $\mathcal{A}_q \in \mathbf{S}^{**}$ and $X_{\beta^-}^*(\mathcal{A}_q) = 0$ otherwise. It is easy to verify that this is indeed a feasible solution for β^- , concluding that $f(\beta^-, \mathcal{L}_{t+1}) = 1$.

So far, we showed that there exists an affector $\mathcal{A}_k \in \mathcal{S}_{t+1}(\mathcal{L}_t)$ such that $f(\beta^-, \mathcal{L}_{t+1}) = 1$. We now show that for any $\mathcal{A}_k \in \mathcal{S}_{t+1}(\mathcal{L}_t)$ such that $f(\beta^-, \mathcal{L}_{t+1}) = 1$, properties (P1) and (P2) are satisfied. This holds

trivially, relying on the criterion for selecting \mathcal{A}_k , since $\mathcal{S}_{t+1}(\mathcal{L}_t) \cap \mathbf{S}_t = \emptyset$.

After n steps, we get that \mathbf{S}_n is a complete selection, $F(\mathbf{S}_n) \in \mathcal{F}(\mathcal{L}_{n-1})$, and therefore by property (P1) for steps $t = 1, \dots, n$, it also holds that $F(\mathbf{S}_n) \in \mathcal{F}(\mathcal{L})$. In addition, by (P2), $f(\beta^-, \mathcal{L}_n) = 1$. Since \mathcal{L}_n is equivalent to $\mathcal{L}(\mathbf{S}_n) \in \mathcal{L}^s$ (obtained by removing the $m-n$ columns corresponding to the effectors not selected by \mathbf{S}_n), it is easy to verify that $f(\beta^-, \mathcal{L}(\mathbf{S}_n)) = 1$. Next, by Thm. 3.2 we have that $1/r(\mathcal{L}(\mathbf{S}_n)) \in \text{Range}_{\beta^*}$.

It remains to show that $1/r(\mathcal{L}(\mathbf{S}_n)) = \beta^*(\mathcal{L})$. By Theorem 4.1, there exists a square system $\mathcal{L}(F^*)$, $F^* \in \mathcal{F}(\mathcal{L})$, such that $r(\mathcal{L}(F^*)) = 1/\beta^*$. Assume, toward contradiction, that $1/r(\mathcal{L}(\mathbf{S}_n)) \neq 1/\beta^*$. Obs. 4.3(b) implies that $r(\mathcal{L}(F^*)) < r(\mathcal{L}(\mathbf{S}_n))$. It therefore follows that $\mathcal{L}(F^*)$ and $\mathcal{L}(\mathbf{S}_n)$ are two non-equivalent hidden square systems of \mathcal{L} such that $1/r(\mathcal{L}(F^*)), 1/r(\mathcal{L}(\mathbf{S}_n)) \in \text{Range}_{\beta^*}$, or, that $1/r(\mathcal{L}(\mathbf{S}_n)) - 1/r(F^*) \leq \Delta_\beta$, in contradiction to Lemma 7.1. This completes the proof of Claim 7.3. \square

By Obs. 4.3(b), $\min_{F \in \mathcal{F}} \{r(\mathcal{L}(F))\} \geq 1/\beta^*(\mathcal{L})$. Therefore, since $r(\mathcal{L}(\mathbf{S}_n)) = 1/\beta^*(\mathcal{L})$, the square system $\mathcal{L}(\mathbf{S}_n)$ constructed in step 7 of the algorithm indeed yields the Perron value (by Eq. (4.10)), hence the correctness of the algorithm is established. Finally we analyze the runtime of the algorithm. Note that there are $O(\log(\beta^*(\mathcal{L})/\Delta_\beta) + n)$ calls for the interior point method (computing $f(\beta^-, \mathcal{L}_i)$), namely, $O(\log(\beta^*(\mathcal{L})/\Delta_\beta))$ calls in the first phase and n calls in the second phase. By plugging Eq. (7.22) in Claim 7.1, Thm. 7.1 follows. \square

8 Applications

We have considered several applications for our generalized PF Theorem. All these examples concern generalizations of well-known applications of the standard PF Theorem. In this section, we illustrate applications for power control in wireless networks, and input-output economic model. (In fact, our initial motivation for the study of generalized PF Theorem arose while studying algorithmic aspects of wireless networks in the SIR model [2, 16, 1].)

8.1 Power control in wireless networks. The rules governing the availability and quality of wireless connections can be described by *physical* or *fading channel* models (cf. [24, 3, 28]). Among those, a commonly studied is the *signal-to-interference ratio (SIR)* model¹. In the SIR model, the energy of a signal fades with the distance to the power of the *path-loss parameter* α .

¹This is a special case of the *signal-to-interference & noise ratio (SINR)* model where the noise is zero.

If the signal strength received by a device divided by the interfering strength of other simultaneous transmissions is above some *reception threshold* β , then the receiver successfully receives the message, otherwise it does not. Formally, let $d(p, q)$ be the Euclidean distance between p and q , and assume that each transmitter t_i transmits with power X_i . At an arbitrary point p , the transmission of station t_i is correctly received if

$$(8.24) \quad \frac{X_i \cdot d(p, t_i)^{-\alpha}}{\sum_{j \neq i} X_j \cdot d(p, t_j)^{-\alpha}} \geq \beta.$$

In the basic setting, known as the SISO (Single Input, Single Output) model, we are given a network of n receivers $\{r_i\}$ and transmitters $\{t_i\}$ embedded in \mathbb{R}^d where each transmitter is assigned to a single receiver. The main question is then is to find the optimal (i.e., largest) β^* and the power assignment \bar{X}^* that achieves it when we consider Eq. (8.24) at each receiver r_i . The larger β , the simpler (and cheaper) is the hardware implementation required to decode messages in a wireless device. In a seminal and elegant work, Zander [37] showed how to compute β^* and \bar{X}^* , which are essentially the PF root and PF vector, if we generate a square matrix A that captures the signal and interference for each station.

The motivation for the general PF Theorem appears when we consider Multiple Input Single Output (MISO) systems. In the MISO setting, a set of multiple synchronized transmitters, located at different places, can transmit at the same time to the same receiver. Formally, for each receiver r_i we have a set of k_i transmitters, to a total of m transmitters. Translating this to the generalized PF Theorem, the n receivers are the entities and the m transmitters are affectors. For each receiver, its supporter set consists of its k_i transmitters and its repressor set contains all other transmitters. The SIR equation at receiver r_i is then:

$$(8.25) \quad \frac{\sum_{\ell \in \mathcal{S}_i} X_\ell \cdot d(r_i, t_\ell)^{-\alpha}}{\sum_{\ell \in \mathcal{R}_i} X_\ell \cdot d(r_i, t_\ell)^{-\alpha}} \geq \beta,$$

where \mathcal{S}_i and \mathcal{R}_i are the sets of supporters and repressors of r_i , respectively. As before, the gain $g(i, j)$ is proportional to $1/d(r_i, t_j)^{-\alpha}$ (where the sign depends on whether t_j is a supporter or repressor of r_i). Using the generalized PF Theorem we can again find the optimal reception threshold β^* and the power assignment \bar{X}^* that achieves it.

An interesting observation is that since our optimal power assignment is a $\mathbf{0}^*$ solution using several transmitters at once for a receiver is not necessary, and will not help to improve β^* , i.e., only the “best” transmitter of each receiver needs to transmit (where “best” is with respect to the entire set of receivers).

Related Work on MISO Power Control. We next highlight the differences between our proposed MISO power-control algorithm and the existing approaches to this problem. The vast literature on power control in MISO and MIMO systems considers mostly the joint optimization of power control with beamforming (which is represented by a precoding and shaping matrix). In the commonly studied *downlink scenario*, a single transmitter with m antennae sends independent information signals to n decentralized receivers. With this formulation, the goal is to find an optimal power vector of length n and a $n \times m$ beamforming matrix. The standard heuristic applied to this problem is an iterative strategy that alternatively repeats a *beamforming* step (i.e., optimizing the beamforming matrix while fixing the powers) and a *power control* step (i.e., optimizing powers while fixing the beamforming matrix) till convergence [7, 8, 9, 30, 33]. In [7], the geometric convergence of such scheme has been established. In addition, [36] formalizes the problem as a conic optimization program that can be solved numerically. In summary, the current algorithms for MIMO power-control (with beamforming) are of numeric and iterative flavor, though with good convergence guarantees. In contrast, the current work considers the simplest MISO setting (without coding techniques) and aims at *characterizing* the mathematical *structure* of the optimum solution. In particular, we establish the fact that the optimal max-min SIR value is an algebraic number (i.e., the root of a characteristic polynomial) and the optimum power vector is a $\mathbf{0}^*$ solution. Equipped with this structure, we design an efficient algorithm which is more accurate than off-the-shelf numeric optimization packages that were usually applied in this context. Needless to say, the structural properties of the optimum solution are of theoretical interest in addition to their applicability.

We note that our results are (somewhat) in contradiction to the well-established fact that MISO and MIMO (Multiple Input Multiple Output) systems, where transmitters transmit in parallel, do improve the capacity of wireless networks, which corresponds to increasing β^* [12]. There are several reasons for this apparent dichotomy, but they are all related to the simplicity of our SIR model. For example, if the ratio between the maximal power to the minimum power is bounded, then our result does not hold any more (as discussed in Section 6). In addition, our model does not capture random noise and small scale fading and scattering [12], which are essential for the benefits of a MIMO system to manifest themselves.

8.2 Input–output economic model. Consider a group of n industries that each produce (output) one




Figure 3: SIR diagram for a power control example, with $\beta^* = 3.5$. (a) An optimal $\mathbf{0}^*$ solution (only t_{11} , t_{21} , t_{31} transmit). All receivers are covered by their corresponding reception zones. (b) A non-optimal $\mathbf{0}^*$ solution (only t_{12} , t_{22} , t_{32} transmit). (c) A non- $\mathbf{0}^*$ solution, where both transmitters collaborate per receiver. Receivers are not covered by their reception zones.

type of commodity, but requires inputs from other industries [23, 26]. Let a_{ij} represent the number of j th industry commodity units that need to be purchased by the i th industry to operate its factory for one time unit divided by the number of commodity units produced by the i th industry in one time unit, where $a_{ij} \geq 0$.

Let X_j represent a unit price of the i th commodity to be determined by the solution. In the following profit model (variant of Leontief's Model [26]), the percentage profit margin of an industry for a time unit is:

$$\beta_i = \text{Profit} = \text{Total income} / \text{Total expenses}.$$

That is, $\beta_i = X_i / \left(\sum_{j=1}^n a_{ij} X_j \right)$. Maximizing the the profit of each industry can be solved via Program (3.4), where β^* is the minimum profit and \bar{X}^* is the optimal pricing.

Consider now a similar model where the i th industry can produce k_i alternative commodities in a time unit and requires inputs from other commodities of industries. The industries are then the entities in the generalized Perron–Frobenius setting, and for each industry, its own commodities are the supporters and input commodities are optional repressors.

The repression gain $\mathcal{M}^-(i, j)$ of industry i and commodity j (produced by some other industry i'), is the number of j th commodity units that are required by the i th industry to produce (i.e., operate) for a one unit of time. Thus, $(\mathcal{M}^- \cdot \bar{X})_i$ is the total expenses of industry i in one time unit.

The supporter gain $\mathcal{M}^+(i, j)$ of industry i to its commodity j is the number of units it can produce in one time unit. Thus, $(\mathcal{M}^+ \cdot \bar{X})_i$ is the total income of industry i in one time unit. Now, similar to the

basic case, β^* is the best minimum percentage profit for an industry and \bar{X}^* is the optimal pricing for the commodities. The existence of a $\mathbf{0}^*$ solution implies that it is sufficient for each industry to charge a nonzero cost for only *one* of its commodities and produce the rest for free.

Finally, we present several open problems and future research directions.

9 Open Problems

Our results concern the generalized eigenpair of a non-square system of dimension $n \times m$, for $m \geq n$. We provide a definition, a geometric and a graph theoretic characterization of this eigenpair, as well as a centralized algorithm for computing it. A natural question for future study is whether there exists an iterative method with a good convergence guarantee for this task, as exists for (the maximal eigenpair of) a square system. In addition, another research direction involves studying the other eigenpairs of a nonsquare irreducible system. In particular, what might be the meaning of the 2nd eigenvalue of this spectrum? Yet another interesting question involves studying the relation of our spectral definitions with existing spectral theories for nonsquare matrices. Specifically, it would be of interest to characterize the relation between the generalized eigenpairs of irreducible systems according to our definition and the eigenpair provided by the SVD approach. Finally, we note that a setting in which $n < m$ might also be of practical use (e.g., for the power control problem in SIMO systems), and therefore deserves exploration.

References

- [1] C. Avin, A. Cohen, Y. Haddad, E. Kantor, Z. Lotker, M. Parter, and D. Peleg. SINR diagram with interference cancellation. In *Proc. 23rd ACM-SIAM SODA*, 502–515, 2012.
- [2] C. Avin, Y. Emek, E. Kantor, Z. Lotker, D. Peleg, and L. Roditty. SINR diagrams: Towards algorithmically usable SINR models of wireless networks. In *Proc. 28th PODC*, 2009.
- [3] U. Black. *Mobile and Wireless Networks*. Prentice Hall, 1996.
- [4] G. Boutry, M. Elad, G.H. Golub, and P. Milanfar. The generalized eigenvalue problem for nonsquare pencils using a minimal perturbation approach. *SIAM J. Matrix Anal. Appl.*, 27(2):582–601, 2005.
- [5] S. Boyd and L. Vandenberghe. *Convex Optimization*. Cambridge University Press, New York, 2004.
- [6] Y. Bugeaud and M. Mignotte. On the distance between roots of integer polynomials. In *Proc. Edinburgh Math. Soc.*, vol. 47, 553–556. Cambridge Univ. Press, 2004.
- [7] D. W. H. Cai, T. Q. S. Quek, and C. W. Tan. A unified analysis of max-min weighted SINR for MIMO downlink system. *IEEE Tr. Signal Process.*, 59, 2011.
- [8] D. W. H. Cai, T. Q. S. Quek, C. W. Tan, and S. H. Low. Max-min weighted SINR in coordinated multicell MIMO downlink. In *Proc. WiOpt*, 2011.
- [9] M. Chiang, P. Hande, T. Lan, and C. W. Tan. Power control in wireless cellular networks. *Foundations and Trends in Networking*, 2(4):381–533, 2007.
- [10] D. Chu and G. H. Golub. On a generalized eigenvalue problem for nonsquare pencils. *SIAM J. Matrix Analysis Applications*, 28(3):770–787, 2006.
- [11] L. Collatz. Einschließungssatz für die charakteristischen zahlen von matrize. *Mathematische Zeitschrift*, 48:221–226, 1942.
- [12] G. J. Foschini and M. J. Gans. On limits of wireless communications in a fading environment when using multiple antennas. *Wireless Personal Communications*, 6:311–335, 1998.
- [13] S. Friedland. On an inverse problem for nonnegative and eventually nonnegative matrices. *Israel J. Math.*, 29:43–60, 1978.
- [14] G.F. Frobenius. Über matrizen aus nicht negativen elementen. *Preussische Akademie der Wissenschaften zu Berlin*, 456–477, 1912.
- [15] C. A. Hall and T. A. Porshing. A separation theorem for nonsymmetric matrices,. *J. Math. Anal. Appl.*, 23:209–212, 1968.
- [16] E. Kantor, Z. Lotker, M. Parter, and D. Peleg. The topology of wireless communication. In *Proc. 43rd ACM STOC* 2011.
- [17] L. G. Khachiyan. A polynomial algorithm in linear programming. *Doklady Akademii Nauk SSSR*, 244:1093–1096, 1979.
- [18] U. Krause. Perron’s stability theorem for non-linear mappings. *J. Math. Econ.*, 15:275 – 282, 1986.
- [19] U. Krause. Concave Perron-Frobenius theory and applications. *Nonlinear Analysis (TMA)*, 47:1457 – 1466, 2001.
- [20] D. Kressner, E. Mengi, I. Nakic, and N. Truhar. Generalized eigenvalue problems with specified eigenvalues. *IMA J. Numer. Anal.*, 2011, submitted.
- [21] O.L. Mangasarian. Perron-Frobenius properties of $Ax = \lambda Bx$. *J. Math. Anal. & Applic.*, 36, 1971.
- [22] V. Mehrmann, R. Nabben, and E. Virnik. On Perron-Frobenius property of matrices having some negative entries. *Lin. Algeb. & its Applic.*, 412:132–153, 2006.
- [23] C.D. Meyer. *Matrix analysis and applied linear algebra: solutions manual*, vol. 2. SIAM, 2000.
- [24] K. Pahlavan and A. Levesque. *Wireless information networks*. Wiley, 1995.
- [25] O. Perron. Grundlagen für eine theorie des Jacobischen kettenbruchalgorithmus. *Mathematische Annalen*, 64:248–263, 1907.
- [26] S.U. Pillai, T. Suel, and S. Cha. The Perron-Frobenius theorem. *IEEE Signal Process. Mag.*, 22:62–75, 2005.
- [27] P. Psarrakos and M. Tsatsomeros. A primer of Perron-Frobenius theory for matrix polynomials. *Lin. Algeb. & its Applic.*, 393:333–351, 2004.
- [28] T. S. Rappaport. *Wireless Communications-Principles and Practice*. Prentice Hall, 1996.
- [29] S.M. Rump. Perron-Frobenius theory for complex matrices. *Lin. Algeb. & its Applic.*, 363:251–273, 2003.
- [30] M. Schubert and H. Boche. Solution of the multiuser downlink beamforming problem with individual SINR constraints. In *IEEE Tr. Vehic. Tech.*, 2004.
- [31] R. Sine. A nonlinear Perron-Frobenius theorem. *Proc. Amer. Math. Soc.*, 109:331–336, 1990.
- [32] C. W. Tan, S. Friedland, and S.H. Low. Nonnegative matrix inequalities and their application to nonconvex power control optimization. *SIAM J. Matrix Analysis Applications*, 32(3):1030–1055, 2011.
- [33] C.W. Tan, M. Chiang, and R. Srikant. Maximizing sum rate and minimizing MSE on multiuser downlink: Optimality, fast algorithms, and equivalence via max-min SINR. *IEEE Tr. Signal Process.*, 59:6127–6143, 2011.
- [34] V.V. Vazirani. The notion of a rational convex program, and an algorithm for the Arrow-Debreu Nash bargaining game. In *Proc. 23rd ACM-SIAM SODA*, 973–992, 2012.
- [35] H. Wielandt. Unzerlegbare, nicht negative matrizen. *Mathematische Zeitschrift*, 52:642–648, 1950.
- [36] A. Wiesel, Y. C. Eldar, and S. Shamai. Linear precoding via conic optimization for fixed MIMO receivers. *IEEE Tr. Signal Process.*, 54:161–176, 2006.
- [37] J. Zander. Performance of optimum transmitter power control in cellular radiosystems. *IEEE Tr. Vehic. Tech.*, 41:57–62, 1992.